

**Energy Efficient High Performance Computing Working Group
4/11/17 Meeting Report**

INTRODUCTION

The EE HPC WG held a meeting on 4/11/17. This Working Group is composed of members representing major Federal 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 June 13th, 9:00-10:00AM Pacific Time

Introductions and Announcements: *Natalie Bates, EE HPC WG & Anna Maria Bailey, LLNL*

- LRZ has published a paper “Analysis of the Efficiency Characteristics of the First High-Temperature Direct Liquid Cooled Petascale Supercomputer and Its Cooling Infrastructure” in the Journal of Parallel and Distributed Computing. It can be downloaded for free till July 4th:
<https://authors.elsevier.com/a/1V2dq2f6jyvQ80>
- Workshop on Power and Energy Aspects of Computations (PEAC) to be held in conjunction with the 12th International Conference on Parallel Processing and Applied Mathematics - PPAM 2017 in Lublin, Poland, on September 10-13, 2017
- Dale Sartor is organizing a special ½ day workshop on the utility/data center interface as part of the Data Center Dynamics Conference to be held in June in San Francisco.
<http://www.dcdconverged.com/conferences/webscale2017/benefits/energy-smart>
- We hosted a webinar in April on “Energy Efficiency Considerations and HPC Procurement”. Ladina Gilly moderated the panel and presenters included Todd Takken from IBM, Steve Martin from Cray, James Laros from Sandia National Laboratory, Daniel Hackenberg from University of Dresden, and Natalie Bates for Anna Maria Bailey from Lawrence Livermore National Laboratory. The predominant goal for procurement of HPC systems is to identify the optimal solution to both technical and financial targets that maximizes the contribution of that system to the organization's mission. Beyond the acquisition cost of the system, it is also important to consider the total costs of ownership, including the improvements necessary to host the system, the infrastructure that supports its operation, and the significant operational costs associated with that new HPC system. In this webinar, HPC leaders discussed key procurement requirements and lessons learned that can contribute to greater energy efficiency and reduced operational costs.

- We hosted another webinar in June on “Data centre design standards and best practices for public research HPC centres.” Ladina Gilly presented on HPC unique challenges - compared to enterprise data center design standards- that she faced when planning the new CSCS data center. She also interviewed ~20 HPC centers with questions that addressed their experience with enterprise data center design standards. There were marked differences in load per cabinet, raised floor height, raised floor rating, the use of UPS systems, types of cooling technologies and time to build. Her conclusion was that although standards do not cover all requirements of HPC they are nonetheless useful as a starting and reference point. Further, due to the nature of HPC, the definition of a standard is not possible. However, a collection of best practices applied in the peer community can however provide a good knowledge base for the community to work with and exchange and build on.

Conferences Sub-group Update: *Torsten Wilde, Leibniz Supercomputing Centre (LRZ, Germany) & Michael Patterson, Intel*

SC17

SC17 will be held in Denver, Colorado from November 12 through the 17th. The EE HPC WG is planning on having a presence again at SC17. We have already reserved space on the floor for a booth. The Annual EE HPC WG Workshop has been accepted and will probably be held all day on Sunday, November 12th. We may be submitting another panel and, of course, we will be making several BoF submissions. There weren't any Teams that made paper submissions this year.

ISC17

ISC17 will be held in Frankfurt, Germany from June 18th through the 22nd. The EE HPC WG will participate in three ISC17 events; a tutorial, a panel and a BoF. Planning is underway.

On Sunday, 18 June, there is a tutorial called "Introduction to Power Optimization Techniques in HPC". The EE HPC WG system level power methodology for use when running a workload will be covered. This is the methodology that was developed collaboratively with the Green500 and the Top500. The tutorial organizers have also asked that the EE HPC WG present an overview of other relevant WG activities and Torsten Wilde will present.

On Wednesday, 21 June, there will be a panel on Energy Efficiency in HPC. Panelists will include Jonathan Eastep from Intel, Martin Shultz from LLNL, Daniel Hackenberg from University of Dresden and Torsten Wilde from LRZ. This panel will present the state of the art in energy efficiency research from the perspectives of academia, industry and government, and spanning both hardware and software. The aim of the panel will be to stimulate a discussion around the effectiveness of current energy management and energy efficiency techniques that are deployed across the entire system

The BoF is not yet scheduled, but will be a joint EE HPC WG, Green500 and Top500 BoF with an emphasis on hearing feedback from the community about the system level power measurement methodology for use when running a workload.

HPPAC

The High Performance Power Aware Computing Workshop, held in conjunction with the IPDPS conference, will be held in Orlando, Florida on May 29th.

EnaHPC

The Energy Aware High Performance Computing Workshop will be held in conjunction with the International Supercomputing Conference in Frankfurt, Germany. It is scheduled for Thursday, 22 June.

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, ORNL and Dave Martinez, SNL*

LIQUID COOLING CONTROLS TEAM:

Background: The Liquid Cooling Controls Team is defining data inputs for dynamic controls to manage liquid cooling control systems, both with the facility and the HPC system. They have identified a short list of data inputs. These data elements and their characteristics are being described for a particular use case; that of dynamic liquid cooling control. Each data input is described with a name and the unit of measurement. There are also characteristics of the measurement; where it is taken, whether the measurement capability is provided by the facility or the IT system or either one, the frequency with which the measurement is taken and the accuracy of the measurement capabilities. It is the intention of this team to publish the results of this as a guideline for sites to use when designing dynamic controls for liquid cooling systems.

News: Since the last general membership meeting, the Liquid Cooling Controls Team has reviewed and responded to broader feedback on their short list of data inputs. While drafting the guideline, the Team has been struggling with several issues. One issue is whether or not the guideline should include both facility and HPC system liquid cooling controls. There is agreement that the scope includes the facility, but debate about whether or not to include controls that are internal to the HPC system. Another issue is whether or not any of the data inputs should include a level of granularity at the node and component level. Because there are some HPC system liquid cooling systems with control at the node and even the component level, the current thought is that the guideline should include these data elements.

DASHBOARD TEAM:

Background: The Dashboard Team has published guidelines on general recommendations for selecting energy efficiency elements of HPC data center dashboards. A dashboard is a display that is used to provide critical feedback to the users. Carefully selecting the elements to be displayed on the energy dashboard is important, as energy management is a shared responsibility of all stakeholders: operations managers, facilities managers, and system administrators.

News: The Dashboard Team has been renewed and has had their first meeting. There were 10 participants. The Team started to review and approve the goals of the Team. There was discussion about whether or not to expand the scope to include prediction. This discussion will be resumed with the next Team meeting.

TUE TEAM:

Background: The TUE Team has published a paper that defines two new metrics; iTUE and TUE that account for infrastructure elements that are a part of the HPC system (like cooling and power distribution). This is an improvement over PUE.

News: There isn't any news from the TUE Team.

LIQUID COOLED COMMISSIONING TEAM:

Background: The Liquid Cooling Commissioning Team has published a guideline that establishes the fundamentals for liquid-cooled commissioning and develops a comprehensive approach to commissioning liquid-cooled systems. It includes examples from different data center sites and the various approaches utilized in commissioning liquid-cooled solutions.

News: There isn't any news from the Liquid Cooled Commissioning Team.

LIQUID COOLING (W_x) GUIDELINE TEAM:

Background: The original Liquid Cooling Team was dub'ed the W_x Team. This Team worked collaboratively with the American Society of Heating, Refrigerating and Air-Conditioning Engineers Technical Committee 9.9 (ASHRAE TC 9.9) to develop guidelines for warmer liquid-cooling temperatures to guide future supercomputer procurements, and to facilitate the design of warmer temperature cooling systems. ASHRAE TC 9.9 (with EE HPC WG collaboration) authored a book that describes classes of typical infrastructure design and the recommended ranges for water supply temperature in each of those classes.

News: There isn't any news from the W_x Team.

High Voltage and DC- Potential New Team:

Background: For most HPC centers, power goes through multiple conversions in alternating current (AC) with a final conversion to direct current (DC) within power supplies in the HPC system. There may be opportunities for energy efficiency and reduced operational costs with fewer conversions including DC conversion at a higher voltage within the HPC center.

News: There isn't any news.

Systems Sub-group Update: *Natalie Bates, EE HPC WG*

SYSTEM WORKLOAD POWER MEASUREMENT METHODOLOGY:

Background: The Green500, Top500 and the Energy Efficient HPC Working Group have developed a methodology for measuring power while running a workload, such as High Performance Linpack (HPL). The methodology defines three quality levels; essentially a “good”, “better”, “best” with Level 3 having the highest quality. Version 1 was published in 2012. Since then, the EE HPC WG, Green500 and the Top500 have been soliciting, collecting, reviewing and responding to feedback from the broader community. This has resulted in publication of Version 2 in 2015. Version 2 includes some editorial improvements and minor changes, but was mostly focused on fixing some major issues and concerns with V1. Major issues affect the outcome, the way it is done, the intended result. There have been multiple sites that have used the Version 2 Methodology and provided feedback for changes.

News: This Team made a submission for a BoF at ISC17. The BoF was accepted, but is not yet scheduled.

AGENDA: With power a first-order design constraint on par with performance, it is important to measure and analyze trends for energy-efficient supercomputing. To raise the awareness of greenness as a first-order design constraint, the Green500 seeks to characterize the energy efficiency of supercomputers with respect to different metrics, workloads, and methodologies. The Green500 list ranks the top 500 supercomputers in the world by energy efficiency. The focus of performance-at-any-cost computer operations has led to the emergence of supercomputers that consume vast amounts of electrical power and produce so much heat that large cooling facilities must be constructed to ensure proper performance. To address this trend, the Green500 list puts a premium on energy-efficient performance for sustainable supercomputing.

This BOF will discuss trends across the Green500 and highlights from the *current* Green500 list. The Green500 and TOP500 have a newly integrated submission process designed to streamline future submissions and to provide a consistent set of data for the historical record. This BOF will also provide a forum for community review of the integrated submission process.

The Green500, Top500 and the Energy Efficient HPC Working Group have developed a methodology for measuring power while running a workload, such as High Performance Linpack (HPL). The methodology defines three quality levels; essentially a “good”, “better”, “best” with Level 3 having the highest quality. This BOF will present case studies from sites that have made submissions that meet the “best” or Level 3 requirements. The BoF will also encourage broader evaluation and feedback on the methodology.

HPC AND GRID INTEGRATION:

Background: This team is analyzing data from 10 major supercomputing centers in Europe and the United States about their electricity contracts and relationships with their electricity service providers. The goal of this study is to determine the structure of the contracts that exist between supercomputing centers and electricity service providers. The team is identifying the influence that this interaction (contractual interaction) may have on demand flexibility. This information is used to understand the degree to which supercomputing centers have, and manage, flexibility

toward grid operation. The information will be further used to identify barriers and opportunities in a demand response-participation context.

News: The Team did confirm our hypothesis that contractual elements – like powerbands – do exist to enable and encourage demand flexibility. However, because of the high HPC system capital depreciation costs, these contractual elements do not translate to dynamic demand response. To analyze all of the contractual data, the Team created a taxonomy for contractual elements that have implications for power and demand measurement. The components of this taxonomy include elements that encourage energy efficiency- like fixed kWh- as well as demand flexibility – like Time of Use and demand pricing. Final analysis of the data is currently in progress and a rough draft of the paper is in progress.

This Team had been targeting submission of a paper to the SC17 technical track, but they have decided to postpone publication until a later date. Perhaps they will submit it to E2SC, which is a workshop held at SC17.

ENERGY AND POWER AWARE JOB SCHEDULING AND RESOURCE MANAGEMENT:

Background: This team has kicked off in high-gear with both great participation as well as immediate work commencing on the initial goal of writing a whitepaper capturing the inventory of sites that are using energy and power aware job scheduling and resource management tools with large-scale implementations in a production and/or pre-production (technology development) environment.

News: Since we last met, the team is has completed interviews with two sites (LANL/SNL for the Trinity machine and STFC in Great Britain) and has two more scheduled for this week (JCAHPC and CINECA). That makes 9 out of 10 possible total sites. The only remaining one is CINES in France.

The team expects to start analyzing the data this month.

RFP CONSIDERATIONS:

Background: The EE HPC WG is maintaining a document that reflects ‘best practices’ for including energy efficiency as an important consideration when writing procurement documents for supercomputer acquisitions.

News: This is a team that has been relatively inactive for several years, but will be called back into action. The immediate next step is to host the webinar on “Energy Efficiency Considerations and HPC Procurement”. As mentioned in the announcements, this webinar originated as a panel discussion that was first held last year at SC16.

PARTICIPANTS INCLUDED

Name	Organization
Natalie Bates	EE HPC WG
Bob Bolz	Aquila
Anita Cocilova	LLNL
Björn Dick	HLRS
Mike Ellsworth	IBM
Parks Fields	LANL
David Grant	ORNL
Daniel Hackenberg	University of Dresden
Glen Hanna	Intel
Sid Jana	Intel
Andree Jacobson	Jacobson Consulting
Greg Koenig	KPMG
Detlef Labrenz	LRZ
Bill Lucnik	ANL
Ian Lumb	Univa
Matthias Maiterth	Intel
Steve Martin	Cray Inc.
Abdelhafid Mazouz	Bull/Atos
Dave Montoya	LANL
Ben Radhakrishnan	National University
Greg Rottman	Army Corps of Engineers
Pamela Tomski	SAS
Torsten Wilde	Leibniz Supercomputing Centre (LRZ, Germany)