

Energy Efficient High Performance Computing Working Group 6/9/15 Meeting Report

INTRODUCTION

The EE HPC WG held a meeting on 6/9/15. 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: August 11th, 9:00-10:00AM Pacific Time

Introductions and Announcements: Anna Maria Bailey and Natalie Bates, LLNL

The EE HPC WG website has had a major facelift and now has a new URL, thanks to Andy Williams, LLNL. Please see the new website – go to <https://eehpcwg.llnl.gov/>. Your feedback is welcome!

Josip Loncaric presented a webinar in late April on LANL's Liquid Cooling Design Considerations. Also, the Liquid Cooling Controls team hosted a webinar on 6/18/15 where system integrator vendors presented their vision and roadmap for HPC liquid cooling controls. The slides and a recording of the webinars are posted on the EE HPC WG website.

An announcement from Chung-Hsing Hsu, ORNL.

The First Workshop on HPC Power Management: Measuring Effectiveness
Tuesday September 22, 2015 8:30am - 5:30pm
Historic Inns of Annapolis, Annapolis, Maryland

This workshop will address the issues and potential solutions for power capacity planning and management in an HPC facility. Two main issues are managing stranded and trapped capacity. Stranded and trapped capacity issues often manifest when facility managers are given highly inaccurate power budget estimations. Power and other physical resources are allocated based on these estimations, which results in overprovisioning and underutilization of these resources. Currently many HPC users are not concerned about power usage effectiveness, but that is certain to change in the near future as the HPC community gears up to meet the exa-scale challenge. Power consumption and management and related infrastructure requirements are considered to be some of the main hurdles to overcome in order to achieve sustained exa-scale operations on future computing platforms and key applications. There is no registration fee for the workshop. However, due to limited space, registration is required.

Conferences Sub-group Update: Marriann Silveira, LLNL

The Silicon Valley Leadership Group held an event in mid-April. Anna Maria Bailey participated in a data center workshop and made a presentation on liquid cooling.

The 6th European workshop on HPC centre infrastructures was held in Stockholm, Sweden in May. There were several members of the EE HPC WG in attendance. Tom Durbin, made a presentation about NCSA's HPC operations, planning, and design issues as well as specific information about the present state of operations of Blue Waters.

SC15:

The EE HPC WG SC15 Workshop was accepted. It will be a full day workshop on Monday, November 16th. We have started planning meetings with the organizing committee. The committee is considering organizing panels that would reach out to HPC segments that are not active in the EE HPC WG. One would be focused on HPC in the oil and gas industry. The other is to invite speakers from Japanese Supercomputing Centers.

The Power Measurement Methodology Team submitted a paper to the SC15 State of the Practice Tract on “Node Power Variability: Implications for System Power Measurement Methodologies”. Natalie will speak to that further in the Systems Sub-Group update.

We made an SC15 panel submission called “Liquid Cooling Promise or Reality.” This will be much like the one we held last year, but the panelists will be users from supercomputing centers instead of vendors this year.

SC15 Birds of Feather submissions are due at the end of July. Last year, we had four Birds of Feather. We haven't yet decided on BoF submissions, but we'll probably make submissions for at least four of them again this year.

ISC15:

The International Supercomputing Conference was held in Frankfurt Germany in mid-July this year.

There are quite a few energy efficiency related sessions again this year.

There was an invited session on “Energy Efficiency and HPC Systems” chaired by Satoshi Matsuoka, Tokyo Institute of Technology. Natalie was one of the speakers of this session and will talk about the latest activities of the EE HPC WG. The other speakers described case studies of deployed energy efficient technologies; one on the #1 Green500 system for November 2014 and the other on a liquid cooling immersion implementation.

There were two Birds of feather. The Power Measurement Methodology Team organized one on “Refining Power Measurement Methodology for Supercomputer-System Benchmarking”. Natalie was the organizer and the speakers were Wu Feng from the Green500, Erich Strohmaier from the Top500 and David Rohr from University of Frankfurt (#1 Green500 system site). Another one was on “Planning for Liquid Cooling in the Data Center.” That included speakers Michael Patterson from Intel, Nic Dube and Tahir Cader from HP, Ingmar Meijer from IBM and Geoff Lyons from CoolIT Systems. There were also 5 papers that discussed energy efficiency. Two of the papers were written by active members of the EE HPC WG; David Rohr and Torsten Wilde. David's paper is a PRACE/ISC award winning paper-congratulations David! All of the papers and abstracts will be listed in the minutes.

Finally, there were three Liquid Cooling Technology vendors who made exhibitor forum presentations.

Other Conferences:

The EE HPC WG website lists many upcoming Conferences and Workshops that have an HPC Energy Efficiency Focus

Infrastructure Sub-Group Update: *Dave Martinez, SNL & David Grant, ORNL*

LIQUID COOLED COMMISSIONING TEAM:

As reported in past meetings, the Liquid Cooling Commissioning Team has been working with ASHRAE to have them publish an updated version of the EE HPC WG Liquid Cooling Commissioning Guidelines. This will first be published as a whitepaper, then included in the next edition of ASHRAE's Liquid Cooling Guidelines for Datacom Equipment Centers.

LIQUID COOLING CONTROLS TEAM:

The Team has been meeting regularly with strong participation. We reported at the April meeting that the team was focused on identifying information/data requirements for control systems from both systems and the building infrastructure. We completed a draft of those requirements and are soliciting feedback.

We are also soliciting case studies. Dave Martinez from Sandia National Lab completed a case study of work in progress that should improve the efficiency of a liquid-cooled HPC cluster named Sky Bridge. They are more tightly integrating and controlling HPC system and building liquid cooling technologies. They hope to prevent wasted energy (heating, etc.) and ultimately create a proactive control of the environment where there will be 'no' delay in response time.

Tom Durbin from NCSA presented a case study of their system. This study provides verification of savings resulting from implementation of controls that utilize variable frequency drives to minimize energy use in chilled water pumps.

Greg Rottman – formally DoD HPC Modernization Program- presented a paper that could be turned into a case study. This paper describes how they interconnected their chiller with the cooling tower, along with monitoring and control capabilities to allow the system to operate in different modes that leverage the local environmental conditions and minimize power required for chilled water.

TUE TEAM:

The TUE Team has been comparing the iTUE and TUE metrics with some work led by Torsten Wilde at LRZ that he calls DWPE (Data Center Workload Power Efficiency). The TUE Team heard about DWPE and was concerned that there were competing efforts to create a metric that enhances and extends PUE.

After extensive review, discussion and comparison, the Team concluded that DWPE may be most valuable as an evaluative technique or methodology for characterizing the combination of HPC system + workload + at a given site. It was billed as a metric, but it is rather a site tool for evaluation. In any case, the work is very interesting and the paper is freely available. You can download a copy from the SIMOPEK website (<http://www.simopek.de/?q=node/9>).

The TUE Team has approached The Green Grid to have them help champion the iTUE and TUE metrics. There is some interest and it would be great if they helped to drive its implementation.

ENERGY REUSE EFFECTIVENESS:

The Energy Re-use Effectiveness Team in collaboration with The Green Grid has developed a standard metric for measuring the contribution of re-using heat generated by HPC systems for other useful purposes.

If you are interested in participating more actively in any of these efforts or on any of these teams, please contact Natalie.

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

SYSTEM WORKLOAD POWER MEASUREMENT METHODOLOGY:

We have a team working on developing a power measurement methodology while running a workload for a system. We worked with Green500 and Top500 and the GreenGrid to develop a standard methodology with various levels of quality. The lowest quality level measures a small sample of nodes and extrapolates to the entire system. The highest quality level measures the power of the entire system. We have gotten to the point where we are iterating on that methodology. We published it a couple of years ago and are now raising the bar for the lowest level or L1 measurements. We feel that there is too much inaccuracy in the L1 methodology- and some of the assumptions don't hold. In particular, we are requiring a greater minimum sample size for the compute nodes required to be measured. We are also requiring that the average power be measured over the entire workload, not just a fraction of the workload. Finally, we are requiring that the interconnect as well as the compute subsystems are included in the system power. Much of this work has been captured in a paper we wrote and that we submitted to the SC15 State of the Practice. We have also updated the methodology document (available for review and comment at <https://github.com/EEHPCWG/PowerMeasurementMethodology>) which we hope to publish in July in conjunction with the ISC BoF. We are also planning that it would be effective for the Green500 and Top500 lists in November 2015.

HPC AND GRID INTEGRATION:

We have team that is looking at the interaction between supercomputing centers and their electricity service providers. This is an exploratory team and there are two motivating factors behind it. One motivation is that as supercomputing centers continue to increase their power demand and intra-hour power variability, they can actually affect the reliability of their electricity service provider's grid. The other motivation is that as the smart grid continues to develop and electricity service providers develop programs (like demand response) that incentivize their users to participate in grid regulation, supercomputing centers might be in a position to contribute. Some of the strategies that supercomputing centers might use are the same strategies that are currently being used for energy efficiency; fine and coarse grained power management and job scheduling. There might be other programs, like scheduling back-ups to non-peak periods. We looked at the United States supercomputing centers and discovered that, while there is interest in incentive programs, there aren't any sites that are currently actively engaged in electricity grid integration. The team hypothesized that other geographies might be different; particularly those where energy prices are high and there is a greater use of renewable energy (which is a more variable supply than more traditional sources like coal). We selected Europe to replicate the analysis (Germany, France, UK, Switzerland, Sweden and Italy). It ends up that the Europeans are not any further along than the US in terms of grid integration. We think that this is probably due to a few things. First of all, in the US we have four supercomputing centers that are twice the size of all the other sites we looked at in both the US and Europe (LLNL, ORNL, NCSA and LANL). It could be that these extra-large sized sites dominate the results and are harbingers of what we might see at the other, smaller

sites. Also, it may be that the US electricity service provider market may be more advanced than the European market in offering incentive programs.

PROCUREMENT CONSIDERATIONS:

The procurement team has written a document that recommends considerations for energy efficiency items to include in procurement discussions and documents. This document is updated yearly; we just finished the 2014 version and posted it on the EE HPC WG website. We are now starting to work on 2015. Some of the work being done on the controls team will feed into this document. We are looking at metrics like Energy to Solution. We will look at power capping. The process leverages ‘best practices’; that is procurement documents. We have looked at CORAL and are looking at Genci and ARC.

PARTICIPANTS INCLUDED

Name	Organization
Anna Maria Bailey	LLNL
Natalie Bates	EE HPC WG
Gim Leong Chin	Performance Computing LLP
Anita Cocilova	LLNL
Wade Doll	Cray
Thomas Durbin	NCSA
David Grant	ORNL
Glenn Genslinger	United Technologies Research Center
Rajat Ghosh	Georgia Tech
Daniel Hackenberg	University of Dresden
Siddhartha Jana	University of Houston
Dave Martinez	SNL
Alexander Moskovsky	RSC Group
Lynn Parnell	NASA Goddard
Greg Pautsch	Cray
Marriann Silveira	LLNL
Jon Summers	University of Leeds
Bill Tschudi	LBNL
Torsten Wilde	LRZ
Foivos Zakkak	Forth Institute