

## Energy Efficient High Performance Computing Working Group 4/14/15 Meeting Report

### INTRODUCTION

The EE HPC WG held a meeting on 4/14/15. 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, <http://eehpcwg.lbl.gov>.

*NEXT MEETING: June 9th, 9:00-10:00AM Pacific Time*

### Introductions and Announcements: *Anna Maria Bailey, LLNL*

- Bill Tschudi, Lawrence Berkeley National Laboratory, has announced that he will be retiring this summer and he is not interested in keeping his position as Co-Lead of the EE HPC WG Infrastructure Sub-group after his retirement. It will be almost impossible to fill Bill's exact shoes, as he was one of the founding members of the EE HPC WG and has provided continual leadership since its inception in 2008. Bill wants to immediately pursue a radically different approach to life upon retirement, so he doesn't plan on participating at all in the future with the EE HPC WG. This cold turkey will leave some of us feeling lost, but Bill has been wonderful about communicating his intentions to retire for many months and so has allowed for a graceful succession. We will miss you Bill, and we wish you the best in your next endeavors!
- I am delighted to announce that David Grant from Oak Ridge National Laboratory will replace Bill. The Co-Lead is Dave Martinez, Sandia National Laboratory, who has been in this role since 2011, having assumed it from Otto VanGeet, National Renewable Energy Laboratory. Dave and David will co-chair the Infrastructure Sub-Group. David has been at ORNL for around 6 years. He has recently been the lead mechanical engineer for the Computational Facilities Complex but has been involved with the facilities since the start. David has been actively involved with the EE HPC WG Controls Team and has shown technical and group leadership with that team. Welcome David!
- In the United States, a new federal executive order specifies energy efficiency requirements for agency data centers. <https://www.whitehouse.gov/the-press-office/2015/03/19/executive-order-planning-federal-sustainability-next-decade>
- See LBNL and DOE present along with many others at the Data Center & Computing Conference, Time Square Marriot Marquis, New York City on 15-16 June, 2015. 10 complimentary registration passes are available for end-user EEHPCWG members upon request. Please contact Bruce Myatt [bruce.myatt@comcast.net](mailto:bruce.myatt@comcast.net) for more information. [www.mcdatacenterconf.com](http://www.mcdatacenterconf.com)
- Tapasya Patki and Akhil Langer presented a webinar on dynamic power management in March. Their presentations and a recording of the webinar are posted on the EE HPC WG website (<http://eehpcwg.lbl.gov/resources/discussions-and-webinars>).

- We have additional speakers on dynamic power management lined up for a June webinar, with the exact date still to be determined.

### **Conferences Sub-group Update: *Marriann Silveira, LLNL***

#### **SC15:**

The EE HPC WG SC15 Workshop was accepted. It will be a full day workshop on Monday, November 16<sup>th</sup>. We will start holding planning meetings with the Workshop organizers soon.

We may have a paper submission, but the deadline is very tight for the team. Natalie will speak to that further in the Systems Sub-Group update.

The SC15 Panel Submission is due the 25<sup>th</sup> of April. We are in the process of developing a Liquid Cooling Panel submission again this year.

#### **ISC15:**

The Power Measurement Methodology Team BoF submission was accepted and will be held on Tuesday July 14<sup>th</sup>.

Also on Tuesday, there is an invited session on “Energy Efficiency and HPC Systems” chaired by Satoshi Matsuoka, Tokyo Institute of Technology. Natalie will be one of the speakers of this session and will talk about the latest activities of the EE HPC WG.

Session Abstract: Power is the most constraining factor in the scalability of supercomputers today towards the future, and thus the technologies to build power-efficient machines, as well as infrastructural support to efficiently deliver power and minimizing various overheads such as cooling, is now regarded as crucial. We have three talks in the session regarding the state-of-the-art in power and cooling issues in this session, one presenting the details of the new supercomputer “L-CSC” that became the most power efficient supercomputer on the Green 500, just installed at the Darmstadt GSI Helmholtz-Center, another covering various innovative cooling technologies such as liquid immersive cooling, and finally the up-to-date activities of the Energy Efficient HPC Working Group (EE HPC WG) that is creating various guidelines for efficient datacenter management for supercomputing.

Research papers have been announced to the authors, but are not yet posted on the website.

Other Conferences:

Natalie Bates gave an invited talk at Supercomputing Frontiers 2015 on the EE HPC WG. Supercomputing Frontiers is Singapore's inaugural conference on trends and innovations in the world of high performance computing. It was held in mid-March. We added ~ 15 new members to the working group as part of our participation in this conference.

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.

### **CONTROLS TEAM:**

There are lessons learned and best practices evolving from implementing and operating supercomputer centers with complex infrastructure systems and the highly variable demands placed upon these systems with today's supercomputers. This team will focus on sharing designs, challenges and best practices for integrated control systems in order to determine if there are universal learnings.

The Team has been meeting regularly with strong participation. We reported at the February meeting that the team was focused on understanding more about control system design considerations and trade-offs between having them located within the computer system verses within the building infrastructure. We are continuing with that line of thinking. We are now identifying liquid cooling control system information/data requirements from both systems and the building infrastructure. This is complimented by case studies. Dave Martinez from Sandia National Lab completed a case study exploring a liquid-cooled HPC cluster named Sky Bridge and the opportunity to merge building controls with the computer cooling controls. Tom Durbin from NCSA will be presenting a case study later in April.

The ASHRAE TC9.9 Committee is writing a book on IT Equipment and asked for review and feedback from the EE HPC WG controls team on a use case of dynamic control of cooling in a data center.

### **TUE TEAM:**

The TUE Team has developed two new metrics; iTUE and TUE that account for infrastructure elements that are a part of the HPC system (like cooling and power distribution). iTUE is not only a metric that is

necessary for calculating TUE, but stands on its own as a metric for a site to use for improving infrastructure energy efficiency.

TUE is an improvement to PUE. PUE is easy to understand, and is comparatively easy to use. It is in common use, so extending its use is thought to be more practical than replacing it.

TUE is positioned to give a more accurate representation of the overall efficiency of the data center with its included IT processing equipment (servers). It primarily allows for consistent results even when moving the location of the air handling and energy storage devices around in the facility / IT chain like from row based to rack located or even directly to the server node itself.

The team has been focused on graphics describing iTUE and TUE for use in presentations and to be posted on the EE HPC WG website. It will also form the basis for a whitepaper.

### **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:**

This team worked with Green500 and Top500 and the GreenGrid to develop a standard methodology for measuring system power while running a workload. The measurement methodology has three levels of measurement 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 included in the system power. We've had a statistician from University of Washington working with us on this team who has helped us with a statistically based analysis for the sample size required. Much of this work has been captured in a paper we wrote and that we hope to submit to the SC15 State of the Practice.

## **HPC AND GRID INTEGRATION:**

We have a 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. It ends up that the Europeans are not any further along than the US in terms of grid integration. We do have quite a few grid experts on the team and so we are diving down into understanding the market and governmental regulations and other environmental factors that might influence the development of methods and programs offered by European electricity service providers. We are specifically looking at Germany, France, UK, Switzerland, Sweden and Italy.

## **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 also looking at metrics like Energy to Solution. We will look at power capping. The process leverages 'best practices'; that is procurement documents such as that used for CORAL.

## ***PARTICIPANTS INCLUDED***

<b>Name</b>	<b>Organization</b>
Anna Maria Bailey	LLNL
Jorge Balcells	Verne Global
Natalie Bates	EE HPC WG
Lloyd Brown	BYU
Anita Cocilova	LLNL
Thomas Durbin	NCSA

David Grant	ORNL
Nathan Gregg	WVU
Detlef Labrenz	LRZ
Piotr Luszczek	University of Tennessee
Dave Martinez	SNL
Andrew Riker	United Technologies, Sikorsky Aircraft
Marriann Silveira	LLNL
Hinrich Tobaben	University of Hannover
Bill Tschudi	LBNL
Robert Voight	Northrop Grumman
Allan Williams	Australian National University