

Energy Efficient High Performance Computing Working Group 10/14/14 Meeting Report

INTRODUCTION

The EE HPC WG held a meeting on 10/14/14. 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: December 9th, 9:00-10:00AM Pacific Time

Introductions and Announcements: *Dale Sartor, LBNL*

Conferences Sub-group Update: *Anna Maria Bailey and Marriann Silveira, LLNL*

- **Update on EE HPC WG participation in recent Conferences**

SC14:

Preparation for SC14 is in full swing. Many of you will be able to attend- and that is great! But, there are also a lot of you who won't be able to attend. For those who can't make it, many of the EE HPC WG organized events will be repeated as webinars in upcoming months.

- The 5th Annual EE HPC WG workshop at SC14 will be all day Monday, November 17th.
 - Workshop topics will include:
 - “Considering Energy Efficiency During Procurement”,
 - “Improving Power Measurement Methodologies and Metrics”,
 - “Exploring Opportunities for Tighter Integration of Supercomputing Centers and Electricity Service Providers” and
 - “Control System Challenges and Best Practices.”
 - Robin Goldstone, LLNL has arranged to have Charlie Manese from Facebook as a Guest Speaker. Charlie will discuss how Facebook designs for efficiency and scale and how Facebook contributes those designs to the Open Compute Project.
 - There will be “round tables” at lunch that are organized around interest areas such as energy efficient benchmarks, liquid cooling technology trends and developing a power API.
- The EE HPC WG organized panel on “Challenges with Liquid Cooling – a Look into the Future of HPC Data Centers” is scheduled for Thursday, November 20th from 3:30 to 5:00. This should be a lively discussion with panelists from 5 different companies- IBM, HP, Bull, Eurotech and Megware.
- There will be four EE HPC WG Birds of Feather submissions.
 1. "HPC System and Data Center Energy Efficiency Metrics and Workloads" is scheduled for: Tue Nov 18 @ 12:15pm-1:15p. This Bof will discuss PUE, iTUE as well as FLOPS/Watt in a panel

format. Panelists include Erich Strohmaier from the Top500, Wu Feng from the Green500 and Michael Patterson from The Green Grid.

2. "Design, Commissioning and Controls for Liquid Cooling Infrastructure" is scheduled for: Tue Nov 18 @ 5:30pm-7pm. At this BoF, Josip Loncaric will talk about LANL's recent design experiences and the ASHRAE inlet water temperature guidelines and Lynn Parnell will describe NASA's lessons learned using the EE HPC WG Liquid Cooling Commissioning whitepaper.

3. "Dynamic Power Management for MW-sized Supercomputer Centers" is scheduled for: Wed Nov 19 @ 12:15pm-1:15pm. Some supercomputing centers are starting to employ strategies to dynamically and in real-time control their electricity demand. This BoF seeks those interested in sharing experiences with dynamic power and energy management.

4. "The Green500 List and its Continuing Evolution" is scheduled for: Thu Nov 20 @ 12:15pm-1:15pm. In addition to giving out awards, we are expecting an announcement about changes to the L1 and L2 power measurement methodology at this BoF.

- We will have an exhibitor booth again this year, with thanks to OSISOFT. It will be similar to the one we had last year with posters describing team activities and other team collateral. Please note, we are looking for volunteers to staff the booth. *Please let Natalie know if you are willing to help with staffing the booth for a few hours.*

Other Conferences

- The Demand Response Team has a paper that will be presented and published as part of a Smart Grid Energy Informatics conference. This conference will be held November 13th and 14th in Zurich, Switzerland.
- The EE HPC WG website lists many upcoming Conferences and Workshops that have an HPC Energy Efficiency Focus

Future Conferences: (more details at <http://eehpcwg.lbl.gov/events-and-links>)

Infrastructure Sub-Group Update: William Tschudi, LBNL & Dave Martinez, SNL

EE HPC WG Infrastructure Technical Sessions at SC14:

- As reported in the conferences update, we have a panel at SC14 on "Challenges with Liquid Cooling - a Look Into The Future of HPC Data Centers". The organizers are Torsten Wilde from LRZ, Michael Patterson from Intel and Josip Loncaric from LANL. The panelists include Nic Dube from HP, Ingmar Meijer from IBM, Thomas Blum from Megaware, Paul Arts from Eurotech and Jean-Pierre Panziera from Bull.

The panel will be organized with brief presentations from the panelists, followed by questions from the moderators and the audience. The moderator questions will include some that have been previewed by the panelists as well as those that will be 'new' to the panelist. The organizers are well along in the process of defining the moderator questions, but please come to the panel with good questions from the floor.

- We have a Birds of Feather session on "Design, Commissioning and Controls for Liquid Cooling Infrastructure". The organizers for this BoF are Dave Martinez from SNL, Marriann Silveira from

LLNL and Herbert Huber from LRZ. As mentioned in the conferences update, there will be speakers with real-world experiences like Josip Loncaric from LANL and Lynn Parnell from NASA. This BoF is scheduled for one and a half hours, whereas all the other ones are only scheduled for an hour. It will be held on Tuesday evening from 5:30 to 7:00PM. There should be a lot of time for discussion and networking.

- On Tuesday from 12:15 to 1:15, the Metrics and Workloads BoF will host a panel session as well as audience discussion. This BoF cuts across both infrastructure and systems. It will be moderated by Chung-Hsing Hsu, from ORNL. Chung-Hsing is very familiar with computer systems, but has also been a leader in evaluating the iTUE and TUE metrics. He'll ask questions like "What is the appropriate system-verses infrastructure boundary? Why report and compare PUE? Why not TUE? Is real-time iTUE possible today? Is there a better single metric that covers both infrastructure and systems?"
- While on the subject of iTUE and TUE, the Monday workshop will include a session on iTUE and TUE testing results. Ghaleb Abdulla will report on LLNL's recent experiences and Chung-Hsing Hsu will report on ORNL.

LIQUID COOLED COMMISSIONING TEAM: 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. They have been sharing controls designs as well as issues and concerns. So far presentations have been made by LLNL, NCSA, ORNL, SNL, LANL, Argonne NL and LRZ. NCAR is also scheduled to present. Steve Blaine from CH2M gave an interesting presentation to the team on industrial verses commercial controls systems.

Vali Sorell from Syska Hennessy and Bruce Myatt from the Critical Facilities Round Table have taken the lead to outline a whitepaper on HPC controls systems and energy efficiency.

TUE TEAM: As mentioned earlier, both LLNL and ORNL have been testing the iTUE and TUE metrics. iTUE and TUE [Total Power Usage Effectiveness (TUE) and IT Power Usage Effectiveness (iTUE)] account for infrastructure elements that are a part of the HPC system (like cooling and power distribution). TUE is an improvement over PUE as a metric that allows for inter-site comparison. 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. For more information, see a bright talk given by Henry Coles <https://www.brighttalk.com/webcast/679/96847>

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.

There is no new information to report on this activity. Anyone interested in sharing your experiences or testing the ERE metric should contact Natalie.

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

SYSTEM WORKLOAD POWER MEASUREMENT METHODOLOGY: The EE HPC WG along with the Green500, Top500 and Green Grid have developed a standard methodology for measuring system power while running a workload. The ultimate goal is to have broad use of the highest quality energy and power measurement methodology for all of their system workload energy efficiency benchmarking activities.

The methodology defines three quality levels for power measurement; L3 uses the ‘current best’ technique, L1 is what the Green500 has been using as its run rules and L2 is in-between. The team is now focused on doing an analysis that quantifies the variation in measurement results for each of the levels. Our hypothesis is that L1 and L2 variation is substantially greater than the variation seen with L3. We are collecting data in three areas: variation in node power affecting sample size, contribution of the interconnect, differences in power draw during the core run. The intent of this analysis is to recommend changes to the L1 and L2 requirements; increasing the quality and stringency of the requirements. The Leibniz Supercomputing Center, the University of Dresden, France’s CEA Supercomputing Center, Calcul Quebec and the Swiss Supercomputing Center have all contributed data to this analysis.

The results of this analysis will be presented at both the SC14 EE HPC WG Workshop on Monday and the Green500 Birds of Feather on Thursday at noon.

HPC AND GRID INTEGRATION: The Demand Response Team is investigating how HPC centers have, can and should engage more actively with the Grid electricity providers. This is an investigative activity with the ultimate goal of educating the HPC DOE Facility and Operations Managers about HPC and grid integration opportunities and challenges.

The Team has written a paper analyzing data collected from 11 US-based Supercomputing Center sites that are on the Top100 list. This includes LLNL, LANL, ORNL, LBNL, ANL, Purdue, SDSC, NCSA, NOAA, Intel and WPAFB. As mentioned in the Conferences Update, this paper that will be presented at a conference called “Energy Informatics” in mid-November in Zurich. The paper will be presented by Bo Nørregaard Jørgensen from the University of Southern Denmark. Bo, Anders Clausen – also from the University of Southern Denmark - and Rish Ghatikar from LBNL have also published a paper on a similar topic that is being presented at the The International Green Computing Conference in a special session on Data Centers and the Smart Grid.

The team is now focused on extending this work to European-based SC sites that are on the Top50 list. We have received feedback on a similar questionnaire from 8 sites: CEA and EDF from France, LRZ, Juelich and Stuttgart from Germany, CSCS from Switzerland, ECMWF from the UK and KTH from Sweden. We will also be collecting data that will allow for comparing and contrasting electricity markets between countries in Europe and between the US and Europe.

One of the key learnings from this analysis is that some of the tools that are being developed to help with energy efficiency may also be used in the future for electricity grid integration. An example of such a tool is dynamic power management, which will be the subject of discussion for the SC14 BoF on Wednesday mentioned in the Conferences Update.

PROCUREMENT CONSIDERATIONS: The RFP Team has a whitepaper that recommends procurement document requirements that target more energy efficient HPC systems. The intention is to

raise the bar and extend the requirements with a yearly update of the whitepaper. The 2013 whitepaper is posted on the EE HPC WG website.

The Team has almost completed the 2014 update. The 2014 version has four new or enhanced sections: 1) enhanced the measurement section with more detailed definitions and descriptions, 2) added a section on timestamping and clocks, 3) added a section on temperature measurements and 4) enhanced the section on air and liquid cooling. Intel, AMD and Cray all reviewed and gave feedback on a draft of this document.

The 2014 HPC Energy Efficiency Procurement Whitepaper will be posted on the website prior to SC14. The first session of the SC14 EE HPC WG workshop will be a panel on this topic and include panelists from the vendor community as well as Supercomputing Centers.

SW UPDATE: Three efforts continue to develop momentum; these are 1) creating an on-line annotated list with links of energy efficiency workloads/benchmarks 2) promote development of a power measurement API, such as the work being led by Jim Laros from Sandia National Laboratory and 3) share best practices for dynamic power management.

PARTICIPANTS INCLUDED

Name	Organization
Natalie Bates	EE HPC WG
Jee Choi	Georgia Tech
Anders Clausen	University of Southern Denmark
Anita Cocilova	LLNL
Parks Fields	LANL
Jeff Filliez	SDSC
Jay Ford	Liquid Cool Solutions
Kurt Keville	MIT
Detlef Labrenz	LRZ
Rodney Martin	NASA
Steven Martin	Cray Inc.
Dave Martinez	SNL
Florent Parent	Calcul Quebec/Compute Canada

Lynn Parnell	NASA Goddard
Suzanne Rivoire	Sonoma State University
Dale Sartor	LBNL
Marriann Silveira	LLNL
Vali Sorell	Syska Hennessy Group
Hinrich Tobaben	University of Hannover
Bill Tschudi	LBNL
Torsten Wilde	Leibniz Supercomputing Ctr.

EEHPC WG