

Energy Efficient High Performance Computing Working Group 4/12/16 Meeting Report

INTRODUCTION

The EE HPC WG held a meeting on 4/12/16. 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 14th, 9:00-10:00AM Pacific Time

Introductions and Announcements: Natalie Bates, EE HPC WG

- New EE HPC WG Team on Job Scheduling and Resource Management

There is a strong interest in predicting and managing power usage for HPC data centers for many reasons (e.g., more efficient cooling, avoiding high power usage during peak demand periods). This requires a cross-disciplinary approach that addresses different abstraction layers, from the infrastructure, to hardware, to system software and up to applications. The advances at the hardware layer need to be followed by evolution on systems software and middleware in order to provide efficient results. Various techniques have achieved a balance of application performance and power or energy consumption using technologies like Dynamic Voltage and Frequency Scaling (DVFS). There is also research on using resource management and scheduling tools to monitor energy and power consumption and adjust it through coordinated actions. Yet another area of research is investigating run-time operating systems for managing system energy and power. These research efforts are making their way into production environments with implementations at a few leading-edge supercomputing centers.

Lessons learned and best practices are being gained in energy aware resource management and job scheduling software that relies on power monitoring hardware. The Energy Efficient HPC Working Group (EE HPC WG) is initiating a new team to focus on power and energy management with resource management and scheduling tools. The initial deliverable of the team will be a whitepaper that surveys those sites already deployed and/or are exploring the use of this software for HPC power and energy management.

If you are interested in joining this team, please contact Natalie Bates.

Thank you,

Natalie Bates and Anna Maria Bailey

Co-Chairs, Energy Efficient HPC Working Group

- The Annual EE HPC WG Workshop was accepted for SC16. Thank you Mike and Torsten for making this happen.
- On May 12th the WW HPC WG hosted a webinar to give an update on Sandia's Power API, Intel's GEOPM and Redfish.
- The week of April 18th, Natalie presented an update on EE HPC WG activities at the 7th European Workshop on HPC Center Infrastructures.

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

ISC16

- The Demand Response Team has a paper that will be presented by Tapasya Patki at ISC16, which is scheduled for June in Frankfurt Germany. The title is "Supercomputing Centers and Electricity Service Providers: A Geographically Distributed Perspective on Demand Management in Europe and the United States."

HPPAC (Workshop on High Performance Power Aware Computing)

- The Dashboard Team has a paper that was presented by James Laros at the High Performance Power Aware Computing Workshop. The title is "Re-examining HPC Energy Efficiency Dashboard Elements."

SC16

- The Annual EE HPC WG Workshop was accepted for SC16. We also have a booth reserved this year. Panel submissions were due April 24th and we made one on High Voltage and dc energy distribution in HPC data centers. Birds of feather sessions are due in July.

=====

Title: Emerging Realities and New Protocols for Power Distribution: High Voltage and dc

Abstract:

Power requirements for high-performance-computers far surpass those of enterprise-servers, creating substantial loads and challenges. With enterprise racks drawing 4-15kW, their HPC counterparts can draw 40-150kW. Typical servers require 120volts, single-phase to operate, whereas high-performance-computers up the ante to 480volts, three-phase. This allows manageable wire sizes and improved efficiency. Alternating current (ac) is ubiquitous, but is direct current (dc) more efficient, reliable and less expensive? Will this trend accelerate with use of renewables? Many high-performance-computers use high-voltage dc internally; why not in the data center?

The implications for these new power distribution systems demand major shifts for both the infrastructure and operations of the data-center. Are we prepared for this shift? Will dc be compelling enough to drive the eco-system? What needs to be done to bridge the gap and ease the transition?

OTHER CONFERENCES:

See the links and events page on the EE HPC WG website for a list of many other very interesting conferences, workshops and special editions with a focus on Energy Efficiency.

<https://eehpcwg.llnl.gov/pages/events.htm>

Infrastructure Sub-Group Update: *Dave Martinez, SNL*

LIQUID COOLING CONTROLS TEAM:

The Controls Team asked liquid cooling vendors Asetek, CoolIT and Chillydyne to answer a series of questions that had been used for their SC15 BoF panel discussion. This prompted broader ranging discussions about leakage current, heat re-use and the ASHRAE W-levels. The team remains on track to collect case studies of liquid cooling controls, both that use infrastructure and HPC system level controls. It has also developed a list of information requirements for these controls that are being tested against controls implementations at Sandia and Stennis.

DASHBOARD TEAM:

The Dashboard Team has a paper that was presented at the High Performance Power Aware Computing Workshop. The objective of this paper is to provide an update on general recommendations to help select or tailor the energy elements or parameters of an HPC data center infrastructure dashboard. The results of the update largely supported the original recommendations.

TUE TEAM:

There isn't any news from the TUE Team.

LIQUID COOLED COMMISSIONING TEAM:

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

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

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

SYSTEM WORKLOAD POWER MEASUREMENT METHODOLOGY:

This Team activity is suspended pending further experience with implementing the new methodology with submissions to The Green500 and The Top500.

HPC AND GRID INTEGRATION:

This team submitted a paper to ISC High Performance 2016 that was accepted. Tapasya Patki will be presenting the paper.

In this paper, we first present results from a detailed, quantitative survey-based analysis and compare the perspectives of the European grid and SCs to the ones of the United States (US). We then show that contrary to the expectation, SCs in the US are more open toward cooperating and developing demand-management strategies with their ESPs. In order to validate this result and to enable a thorough comparative study, we also conduct a qualitative analysis by interviewing three large-scale, geographically-distributed sites: Oak Ridge National Laboratory (ORNL), Lawrence Livermore National Laboratory (LLNL), and the Leibniz Supercomputing Center (LRZ). We conclude that perspectives on demand management are dependent on the electricity market and pricing in the geographical region and on the degree of control that a particular SC has in terms of power-purchase negotiation.

The team is engaged in research that will generate a paper that looks more closely at the differences in electricity markets both within Europe and between Europe and the United States.

The team has also been working with the Tokyo Institute of Technology and Riken, in Japan, to understand their work on power capping. This has been an important feature for Japanese supercomputing centers because Japan is still severely power constrained as a result of the aftermath of the Tsunami in 2011.

PARTICIPANTS INCLUDED

Name	Organization
Andrea Bartolini	ETH
Natalie Bates	EE HPC WG
Bob Bolz	Aquila
Gary Brown	Adaptive Computing
Anita Cocilova	LLNL
Wade Doll	Cray
Michael Ellsworth	IBM
Steve Harrington	Flometrics
Brent Henderson	HPE
Siddhartha Jana	University of Houston
Sharan Kalwani	Michigan State University
Detlef Labrenz	LRZ
Thomas Leung	GE Global Research
Dave Martinez	SNL
Todor Milkov	SDSC
Scott Milliken	ORNL
Cemil Ozyalcin	CGG
Ben Payne	Laboratory for Physical Sciences – UMD
Ben Radhakrishnan	National University
Suzanne Rivoire	Sonoma State University
Mehdi Sheikhalishahi	University of Calabria
Ram Shetty	Opex Solutions
Ralph Wescott	PNNL
Torsten Wilde	Leibniz Supercomputing Centre (LRZ, Germany)
Foivos Zakkak	Forth Institute