

Energy Efficient High Performance Computing Working Group
December 14, 2010 Meeting Report

INTRODUCTION

The EE HPC WG held a meeting on December 14, 2010. 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>. Documents from the group can be found at <https://docs.google.com/leaf?id=0BzyTVVVRdMKpNwVjNTI5YTEtMTIIZi00YTA5LTlkMTYtZmY3ZDIyZjJjZmMy&hl=en>.

NEXT MEETING: February 8th, 9:00-10:00AM Pacific Time

Introductions and Announcements: Dale Sartor

The EE HPC WG membership is up to 125, from 105 last month.

Members are invited to participate in a discussion about creating a guideline for the procurement requirements of liquid cooled HPC systems. It would provide an impetus for market transformation, to be achieved as buyers commit to using it. This meeting will be coming up in January.

In the upcoming year, the group will host webinars and talks on some of the following topics. Other ideas are also appreciated.

- NCSA power measurement in Green500 submission (Craig Steffan)
- Power management at LLNL (Anna Maria)
- DC power (UCSD)

Dale also requested information on heat reuse projects either already implemented or planned. Please send any contact information to Natalie Bates

Conferences Sub-group Update: Marriann Silveira & Anna Maria Bailey

Over 100 attendees, about 40 from DOE national labs, participated in the SC'10 workshop "Running a Lean and Productive HPC Center." The conference also featured a Birds-of-a-Feather session with the EE HPC WG. There were approximately 75 attendees, ~ 30 from DOE national labs. The collaboration on compute system metrics was highlighted at the Top 500 and Green500 Birds-of-a-Feather session.

The International Supercomputing Conference (ISC) and the International Green Computing Conference have already called for papers. This group may be making a submission for an ISC tutorial, due February 15th.

Below are topics of interest for technical sessions at upcoming conferences. Anyone interested in helping to prepare and run a session should contact Anna Maria or Marriann.

- Power Management
- Liquid Cooling
- Getting the Data Center ready for Multi-Petascale and Exascale
- HPC System/Rack Power Budget
- Free Cooling and Heat Re-use Case Studies
- DC Power Demonstrations
- Glass House and Regional Strengths and Weaknesses
- Storage and Networking Energy Efficiency
- EE HPC WG BoF

Infrastructure Sub-group Update: *William Tschudi*

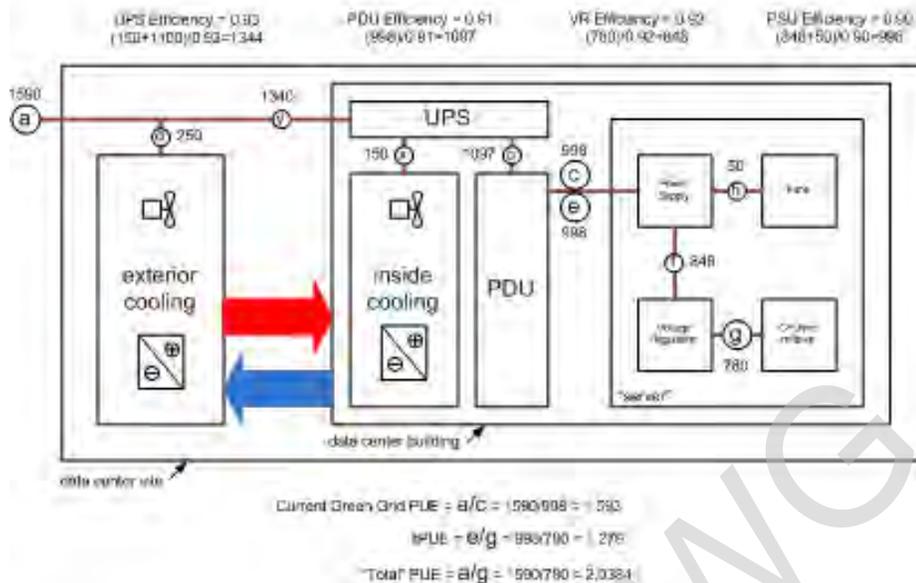
A smaller group consisting of Bill Tschudi, Mike Patterson (Intel), Larry Seibold (SGI), Henry Coles, and Natalie Bates discussed metrics covering situations where parasitic loads shift from IT equipment to the building or vice versa. Several examples were given:

- SGI removes fans in servers and relies on larger infrastructure fans in their container solution
- Yahoo removes infrastructure fans and relies on fans in IT equipment at the new NY data center
- HPC vendors provide rack level power supplies instead of individual server level power supplies

This is an issue since PUE is widely regarded as the metric for data center efficiency. Shifting from IT to infrastructure or vice versa will skew PUE and could give a misleading representation of efficiency.

New metrics need to consider total power and consider all parasitic loads whether they are in the IT equipment or the building. The idea of ITPUE was introduced. The attached diagram was presented by Henry Coles and depicts the idea.

Concept "Total" PUE and itPUE



For now, one could estimate the effects of placing fans or power supplies in another place but it would be difficult to measure. Possibly in the future, the IT equipment could report out on its parasitic loads. Others who are interested are welcome to join the discussion.

Green Grid issued a new metric for carbon utilization. The white paper is available on the Green Grid website. Green Grid also issued a white paper on power management. While their paper is not too applicable for HPC, it basically concludes that most people do not understand power management and don't implement it.

Compute System Metrics Sub-group Update: *Natalie Bates*

The group has completed definitions of terms for metric, methodology and workload; clarified the intended audience and desired result; and started to identify workloads and to iterate between methodology and metric specification. Some outstanding issues that need resolution:

- Can we improve (and standardize?) the power measurement methodology? How do we increase the number of reported (verses derived) power numbers? How accurate and precise are the power measurement reported numbers? Are there discrepancies between the Top500 and Green500 methodologies? What is included in the power measurement (e.g., cooling systems).
- We still need to decide upon the exact metric(s) for the compute subsystem. Do we use classes of systems (e.g., Top50 or Little500)? Are there multiple metrics or a single index?
- Is GUP a good workload for memory?
- What are possible workloads for I/O and storage?