

**Energy Efficient High Performance Computing Working Group
12/12/17 Meeting Report**

INTRODUCTION

The EE HPC WG held a meeting on 12/12/17. 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, February 13th, 9:00-10:00AM Pacific Time

Introductions and Announcements: *Natalie Bates, EE HPC WG & Anna Maria Bailey, LLNL*

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

SC17

The EE HPC WG had a strong presence at SC17 with a full-day Workshop, four Birds of Feather, one Poster, the Exhibitor Booth and a Panel Session.

All of the sessions were well attended and generated good participation. All of the presentations and some supporting material can be found on the EE HPC WG website (https://eehpcwg.llnl.gov/pages/conf_sc17a.htm) under Conferences and then SC17 Conference.

Inside HPC published a video of the Panel Session and will also publish two more articles- one covering the Birds of Feather on Energy and Power Aware Job Scheduling and Resource Management and the other on Software Stack Energy Efficiency that mentions the Workshop, Poster, Panel and Birds of Feather Sessions. (Special thanks to Sid Jana for his work supporting the Inside HPC coverage.)

Workshop

- Everyone who completed a survey thought the workshop was valuable (14 total respondents).
- The Power Grid and Team Reports were considered very valuable by 85% of the respondents.
- Attendance was comparable to recent prior years with a full room most of the day, mostly all-day participants but some who were attending multiple events and targeted specific sessions.

BoF: Total Cost of Ownership and HPC System Procurement

- Excellent participation from both supercomputing centers and the vendors. We had participants from Supercomputing Centers in the US, Europe and Japan. Vendors included Cray, IBM, HPE, Bull/Atos, Fujitsu, Lenovo and Dell.
- Highlighted the European PRACE work on Total Cost of Ownership as well as the EE HPC WG Procurement Team.

BoF: PowerAPI, GEOPM and Redfish: Open Interfaces for Power/Energy Measurement and Control

- This session has been held over the past few years and this year's presentations demonstrated maturation and collaboration between the three efforts. This year's focus was on user implementations, which are still mostly in the research and early technology development phases.

BoF: State of the Practice: Energy and Power Aware Job Scheduling and Resource Management (EPA-JSRM)

- This was a very well organized BoF with strong audience participation. It was a very effective tool for pulling the EPA JSRM Team together. A whitepaper they generated and reviewed as part of the BoF is also posted on the Conferences website.

BoF: The Green500: Trends in Energy-Efficient Supercomputing

- This BoF is repeated every six months - alternating between ISC and SC. It has a pretty standard format and there weren't any surprises this year. Two sites presented on L2/L3 power measurements. These were Nvidia and the Japan's JCAHPC (Center for Computational Sciences at the University of Tsukuba (CCS) and the Information Technology Center at the University of Tokyo (ITC)).

Panel: Energy Efficiency Gains From Software: Retrospectives and Perspectives

- Although this was the very last session on the very last day of the conference, it was well attended and provided a lively discussion. An audio/video recording of the panel is posted on the EE HPC WG Conferences webpage.

Research Poster: EPA JSRM

- This was the first time that the EE HPC WG had a research poster at any conference. It generated interest and might be a forum that we should consider for the future.

EE HPC WG Booth

- The location was great- right next to a food table on opening night. We had at least 25 people sign up at the booth. It also provides a meet-up and discussion location for members.

ISC18

We are planning to submit a Green500 Birds of Feather as well as a research paper from the Grid Integration Team. The EnaHPC Workshop is also expected to be held at ISC18. ISC18 will be held in Frankfurt in June.

HPPAC

The High Performance Power Aware Computing Workshop call for papers is the end of January. It will be in May in Vancouver BC.

Other Conferences

The EE HPC WG website has a links and events page with many other conferences and workshops listed that have an HPC energy efficiency focus.

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

For all the team reports below, the non-bolded text is a carry-over from prior presentations and **the bolded text is the new material for this meeting.**

LIQUID COOLING CONTROLS:

- Impacts of the work
- Increase ease of deployment for liquid cooling controls
- Improve energy efficiency of the cooling system and reduce costs of cooling
- Deliverables
- “EE HPC WG Liquid Cooling Controls Team Whitepaper
- State of the practice case studies for liquid cooling control systems
- Current activities
- Discussions with ASHRAE, Redfish and Power API on incorporating these data inputs in their recommendations
 - **Steve Martin proposed including data inputs in PowerAPI at the Dec 5th PowerAPI meeting; outcome was to postpone inclusion to a future revision**
- Next steps
- Inclusion of data inputs in ASHRAE, Redfish and Power API & EE HPC WG Procurement Considerations Document 2017
 - **Continue to push for inclusion of data inputs in Redfish**
- Help needed
- Advocates for EE HPC WG to work on Redfish

LIQUID COOLING STANDARDS:

- Impacts of the work
- Encourage liquid-cooled solutions that do not require compressors
- Increase ease of deployment by “standardizing” facility and HPC equipment
- Set the bar for more opportunities to reuse waste heat
- Deliverables
- Wx temperature classes developed, presented, and published (e.g. SC11), and included in ASHRAE TC9.9 Liquid Cooling Guidelines for Datacom Equipment Centers.
- Current activities
- EE HPC WG members providing input and expertise to develop “open” specification for warm water liquid cooled rack with major internet companies
- Next steps
- Continue to communicate and clarify Wx recommendations

- Provide input to liquid cooled rack specification harmonizing U.S. and Chinese standards (e.g., OCP and Scorpio)
 - Help needed
- Input on draft specification including fluid and connectors as well as operating conditions (e.g. temperatures and pressure)

LIQUID COOLING COMMISSIONING:

- Impacts of the work
- Encourage decreased costs and improve energy efficiency with effective liquid cooling commissioning
 - Deliverables
 - “Systematic approach for commissioning liquid cooling infrastructure to support liquid cooled HPC systems”
 - State of the practice case studies for liquid cooling commissioning
 - Current activities
 - Include recommendations in EE HPC WG Procurement Considerations Document 2017
 - Next steps
 - Total Cost of Ownership and HPC System Procurement BoF; Tuesday 12:15
 - ASHRAE TC9.9 to incorporate liquid cooling commissioning in commissioning guideline
 - Help needed
 - Technical expert w/strong technical writing skills to finalize whitepaper with ASHRAE TC9.9 Committee
 - Contribute case studies and lessons learned

RAS AND MAINTAINABILITY (Investigative Team):

- Impacts of the work
- Increase energy and operational efficiency by improving Reliability Availability Serviceability (RAS) and Maintainability beyond the HPC system to facility infrastructure
 - Deliverables
 - Questionnaire created and sent to major US supercomputing sites
 - **Collated initial 5 responses (ORNL, LLNL, Sandia, NREL, Argonne)**
 - (4) responses indicate reliability, availability, serviceability (maintainability) extends beyond the systems
 - Current activities
 - Soliciting feedback on team creation from major US supercomputing sites
 - **(3) more sites to respond (NCSA, LANL, LBNL)**
 - Do we create a team on HPC facility maintainability and reliability as it relates to energy efficiency and availability?
 - Next steps
 - **Waiting for responses from 3 more sites**
 - **Analyze data and make recommendations**
 - Help needed
 - Complete the questionnaire
 - **Join the investigative team**
 - Share best practices/lessons learned

iTUE AND TUE:

- Impacts of the work
 - Combines with PUE to provide a TOTAL view of where the inefficiencies are. Adds a “server PUE”. Precludes miscounting power and cooling losses on the wrong side of the equation.
- Deliverables
 - “TUE, a new energy-efficiency metric applied at ORNL's Jaguar”; Gauss Best Paper Award; ISC13
 - State of the practice case studies for TUE and iTUE
 - Recommended capability in EE HPC WG Procurement Considerations document
- Current activities
 - Evangelizing iTUE/TUE in talks and conferences
- Next steps
 - Develop an iTUE case studies session for next year’s workshop
- Help needed
 - Engage The Green Grid to promote iTUE and TUE
 - Explore your ability to measure or estimate your iTUE and TUE
 - Contribute case studies/lessons learned

Systems Sub-group Update: *Natalie Bates, EE HPC WG and Jim Laros, SNL, Ben Radhakrishnan, National University*

For all the team reports below, the non-bolded text is a carry-over from prior presentations and **the bolded text is the new material for this meeting.**

POWER MEASUREMENT METHODOLOGY:

- Impacts of the work
 - Provides the functions which can monitor and record power consumption of entire system in real time
 - More accurate HPC system architectural trend data for the HPC Community
- Deliverables
 - State of the practice case studies on power measurement methodology
 - “Energy Efficient High Performance Computing Power Measurement Methodology (version 1.0)”; 2012
 - “A power-measurement methodology for large scale, high performance computing”. ACM/SPEC International Conference on Performance Engineering; 2014.
 - “Node Variability in Large-Scale Power Measurements: Perspectives from the Green500, Top500 & EE HPC WG”. SC15; 2015.
 - “Energy Efficient High Performance Computing Power Measurement Methodology (version 2.0 RC 1.0)”; 2016
 - “Submissions Open for Newly Merged TOP500 and Green500”; May 6, 2016; Rich Brueckner; Inside HPC.
- Current activities
 - Solicit and understand feedback on system-level workload power measurement methodology
 - Encourage L2/L3 measurement submissions to Green500/Top500
- Next steps

- **Top500/Green500 Bof at ISC18**

- Help needed
 - Make L2/L3 measurement submissions to Top500/Green500 List
 - Encourage extension of L2/L3 measurement submissions to other benchmarks, e.g., GreenGraph500

ELECTRIC GRID INTEGRATION:

- Impacts of the work
 - Raising awareness of evolving relationship between SCs and their Electricity Service Providers
 - The landscape is changing- get to know your ESP and their partners.
 - Implement contingency planning for power management while minimizing impact to users.
- Deliverables
 - "The Electrical Grid and Supercomputing Centers: An Investigative Analysis of Emerging Opportunities and Challenges"; Energiinformatik; Zurich, Switzerland; 2014.
 - "Supercomputing Centers and Electricity Service Providers: A Geographically Distributed Perspective on Demand Management in Europe and the United States"; ISC16; Frankfurt, Germany; 2016.
- Current activities
 - Writing a paper that examines electricity service contracts in major Supercomputing Centers(SC)
- Next steps
 - Finalize and publish next paper
- Help needed
 - Contribute case studies/lessons learned

ENERGY AND POWER AWARE JOB SCHEDULING AND RESOURCE MANAGEMENT:

- Impacts of the work
 - Share best practices of Energy and Power Aware Job Scheduling and Resource Management (EPA JSRM) and learn from each other
 - Identify opportunities for influencing product development
- Deliverables
 - Interview results from 9 sites that have deployed or are doing technology development (TD) with the intent to deploy large scale EPA JSRM capability in a production environment
- Current activities
 - Analyzing data from survey of large scale EPA JSRM deployments
- Next steps
 - **Write a paper for submission to SC18, possibly other papers as well**
- Help needed
 - Join the team and help with data collection, analysis and writing the paper
 - Identify other potential sites with large scale EPA JSRM in production of TD with an intent to deploy

PROCUREMENT CONSIDERATIONS:

- Impacts of the work
 - Influence product development to drive energy efficient HPC systems
- Deliverables
 - State of the practice case studies for energy efficiency considerations in procurement
 - Component and System Integrator responses to EE HPC WG Procurement Considerations
 - Energy Efficiency Considerations for HPC Procurement Documents: 2014 and Energy Efficiency Considerations for HPC Procurement Documents: 2013
- Current activities
 - Updating procurement considerations document for 2017 with vetted material (e.g., liquid cooling controls)
 - Collaborating with PRACE (Partnership for Advanced Computing in Europe) on TCO and Procurement
- Next steps
 - Publish Energy Efficiency Considerations for HPC Procurement Documents: 2017
 - Start working on Energy Efficiency Considerations for HPC Procurement Documents: 2018
- Help needed
 - Participate on the Procurement Considerations Team and help write documents
 - Share best practice procurement documents

DASHBOARDS:

- Impacts of the work
 - Strive for consensus on HPC center dashboard energy efficiency elements and metrics
- Deliverables
 - "Re-examining HPC Energy Efficiency Dashboard Elements"; 12th Workshop on High Performance Power Aware Computing
 - "General Recommendations for High Performance Computing Data Center Energy Management Dashboard Display"; 9th Workshop on High-Performance Power-Aware Computing Conference
- Current activities
 - Questionnaire - current use of dashboards at major supercomputing centers in USA, Europe & Japan
 - **Analyzed questionnaire results and currently finalizing whitepaper**
 - **LLNL and NERSC have very sophisticated capabilities. Tightly integrated with many data sources. From this questionnaire, they provide the high-bar.**
 - **NERL is the next tier with similar tight integration and multiple data sources, but they say that the system is ad hoc.**
 - **All other sites have isolated (not tightly integrated) solutions.**
 - **There is a desire for more/better capabilities from every site- including LLNL and NERSC.**
- Next steps
 - **Discuss more about how best to share what we have so far.**

- **Identify how LLNL, NERSC and NREL are really using their dashboard tools – analysis, trends, etc. – specific use case scenarios. Some of these could become best practices across data centers.**
- **Equipment Maintenance: Investigate data analytic tools which can do trend analysis on specific data collected over a period of time that will identify potentially-failing equipment before it fails to avoid service downgrade or shut down.**
- **Predictive Analysis of data – here the sky is open and we need again to sit down with some data operators and data scientists to look what they would like to see predicted for the data centers, what data to be collected and what analysis to be done, identify tools to do them (I would certainly think Tableau has a place here and there may other tools). This is taking to the next level.**
- **Artificial Intelligence; this is another use of new technology and its use in data center operation to make it efficient and effective in all its operation given all the data being collected – best optimization of energy and other operational items.**
- **What business metrics can be calculated with the data and/or improve what we have.**

– Help needed

- Participate in the Dashboard Team

PARTICIPANTS INCLUDED

Name	Organization
Natalie Bates	EE HPC WG
Anita Cocilova	LLNL
Thomas Durbin	Independent
David Grant	ORNL
Brandon Hong	LLNL
Jim Laros	SNL
Steve Martin	Cray Inc.
Dave Martinez	SNL
Rick Poole	Idaho National Laboratory
Ben Radhakrishnan	National University
Dale Sartor	LBL
Torsten Wilde	Leibniz Supercomputing Centre (LRZ, Germany)
Xingfu Wu	ANL
Jason Zeiler	Coolit Systems