10Energy Efficient High Performance Computing Working Group
10/10/17 Meeting Report

INTRODUCTION

The EE HPC WG held a meeting on 10/10/17. 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 December 12th, 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
SC17 was held in Denver, Colorado from November 12 through the 17th. The EE HPC WG was a presence again at SC17.

Eighth Annual Workshop for the Energy Efficient HPC Working Group
Topics include case studies of energy efficient operational lessons learned; the power grid- or "what you need to know about the power grid before adding a 10 MW step-function load generator"; the United States Department of Energy's Path Forward and other Exascale programs; and the software stack's implications for energy efficiency.

Arthur S Buddy Bland from Oak Ridge National Laboratory will be the keynote speaker. Buddy has seen more than 30 years of HPC deployment at ORNL and his keynote will provide insight into operations and energy efficiency for some of the largest supercomputers.

BoF: Total Cost of Ownership and HPC System Procurement
HPC leaders from across the globe will discuss and debate key procurement requirements and lessons learned that can contribute to greater efficiency and reduced operational costs.

BoF: PowerAPI, GEOPM and Redfish: Open Interfaces for Power/Energy Measurement and Control
In this BoF, we discuss three efforts for open-interfaces and frameworks: these are PowerAPI, GEOPM and Redfish. The BoF will introduce each effort and feature an interactive panel discussion with experts from currently implementing and adopting organizations. There will also be a focus on synergy and compatibility between them.

BoF: State of the Practice: Energy and Power Aware Job Scheduling and Resource Management (EPA-JSRM)
This BoF presents results of a global survey of supercomputing centers using Energy and Power Aware JSRM strategies and seeks those interested in sharing experiences with dynamic power and energy management.

BoF: The Green500: Trends in Energy-Efficient Supercomputing
This BoF discusses trends across the Green500 and highlights from the current Green500 list. In addition, the Green500, Top500, and Energy-Efficient HPC Working Group have been working together on improving power-measurement methodology and this BoF presents case studies from sites that have made power submissions that meet the highest quality of measurement methodology.

Panel: Energy Efficiency Gains From Software: Retrospectives and Perspectives
MODERATOR: Daniel Reed
PANELISTS: Satoshi Matsuoka, Thomas Schulthess, Bill Gropp, John Shalf
This panel will explore what HPC software capabilities were most helpful over the past years in improving HPC system energy efficiency? It will then look forward; asking in what layers of the software stack should a priority be put on introducing energy-awareness; e.g., runtime, scheduling, applications? What is needed moving forward? Who is responsible for that forward momentum?

Research Poster: **Global Survey of Energy and Power-aware Job Scheduling and Resource Management in Supercomputing Centers**

Exhibition Floor Booth: EE HPC WG

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***

<u>LIQUID COOLING CONTROLS</u>:
– 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
– Next steps

• Inclusion of data inputs in ASHRAE, Redfish and Power API & EE HPC WG Procurement Considerations Document 2017
• Total Cost of Ownership and HPC System Procurement BoF; Tuesday 12:15
• Redfish, PowerAPI and GEOPM BoF; Tuesday 17:15

– 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:

– 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
    • (4) responses indicate reliability, availability, serviceability (maintainability) extends beyond the systems
– Current activities
    • Soliciting feedback on team creation from major US supercomputing sites
    • Do we create a team on HPC facility maintainability and reliability as it relates to energy efficiency and availability?
– Next steps
    • Waiting for responses
– Help needed
    • Complete the questionnaire
    • 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***

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 List L2/L3 Measurement feedback presented at SC17 Green500 BoF, Wednesday 17:15

– 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
  • EPA JSRM Poster, EPA JSRM BoF; Wednesday 12:15, Software and Energy Efficiency Panel; Friday 10:30.
  • Whitepaper
  • Technical paper
– 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
• Total Cost of Ownership and HPC System Procurement BoF; Tuesday 12:15

– 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

– Next steps
• Analyze questionnaire results and write/publish paper

– Help needed
• Participate in the Dashboard Team and help with analyzing data and writing paper

PARTICIPANTS INCLUDED

Name	Organization
Natalie Bates	EE HPC WG
Anita Cocilova	LLNL
Nic Dube	HP
Anne Elster	Norwegian University of Science and Technology
Michael Garceau	3M
David Grant	ORNL
Jessica Gullbrand	Intel
Jim Laros	SNL
Ronald Lui	IBM
Steve Martin	Cray Inc.
Shlomo Novotny	Independent Contractor
John Peterson	AECOM
Ben Radhakrishnan	National University
Vali Sorell	Sorell Engineering, Inc.
Todd Takken	IBM
Pamela Tomski	SAS