Energy Efficient High Performance Computing Working Group
8/14/18 Meeting Report

INTRODUCTION
The Energy Efficient High Performance Computing Working Group (EE HPC WG) held a meeting on 8/14/18. This Working Group is composed of members representing major governmental 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, October 9th, 9:00-10:00AM Pacific Time

Introductions and Announcements: *Natalie Bates, EE HPC WG & Anna Maria Bailey, Lawrence Livermore National Laboratory*

Two new loosely formed groups of interest to the community.

- Lawrence Berkeley National Laboratory and China Institute of Electronics are steering an initiative to create an open-specification for a secondary fluid warm liquid-cooled rack.
- The HPC PowerStack effort will work to identify what power optimization software actors are needed; how they interoperate to achieve stable, synchronized optimization; and how to glue together existing open source projects to engineer a cost-effective but cohesive, cross-platform power stack implementation.

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

ISC18

- ISC18 was held in Frankfurt from June 24th through the 28th.
- The EE HPC WG had a special event presentation, two Birds of Feathers, and participated in a panel as well as a workshop.

Data Center Dynamics

- Dave Martinez and Dale Sartor presented a panel on liquid cooling at the Data Center Dynamics' San Francisco conference in June.

SRMPDS

- The 14th International Workshop on Scheduling and Resource Management for Parallel and Distributed Systems will be held in conjunction w/ ICPP in Eugene, Oregon and Sid Jana has been invited to present a talk on the EE HPC WG EPA JSRM Team results.

7x24 conference

- 7x24 Conference will be held in Phoenix October 21-24, 2018
- Dale Sartor will be moderating a panel on "Lessons Learned in Establishing Metrics, Setting Goals, and Achieving High Data Center Energy Performance". A panel from Digital Realty Trust, Intel, and the Lawrence Livermore, Sandia, and Lawrence Berkeley National Laboratories will share their stories and lessons learned in improving efficiency and resiliency.

International Green and Sustainable Computing Conference

- IGSC will be held October 22-24, 2018 in Pittsburge, PA and the PowerAPI Team is organizing a workshop at this conference

SC18

- Our workshop submission was accepted. We have an organizing committee and held our first meeting in early July.
- There is an SC18 page on our website with content "coming soon"!
- Two papers were submitted to the State of the Practice Tract; one from the Grid Integration Team and the other from the Energy and Power Aware Job Scheduling and Resource Management Team. Both of these were rejected.
- Our panel submission on Software and Energy Efficiency was accepted.
- We made 6 BoF submissions on a variety of topics.

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, Oak Ridge National Laboratory and Dave Martinez, Sandia National Laboratory*

LIQUID COOLING CONTROLS:

The Liquid Cooling Controls Team is focused on improving the efficiency and decreasing the cost of liquid cooling systems. The Team is making it easier to build controls for liquid cooling by working with the HPC system vendor community to deliver the right sensors and telemetry with their IT systems. The Team has created a list of data inputs that are provided by the system vendors or the facility.

There hasn't been much progress with this team since the last time we met, except that Natalie joined as an individual member of the Green Grid – which is the organization working with the Distributed Management Task Force's on providing facility level telemetry to be included in the Redfish API.

LIQUID COOLING STANDARDS:

Lawrence Berkeley National Laboratory and China Institute of Electronics – along with input from major cloud players like Microsoft, Google and Amazon, have drafted an open specification for a liquid cooled rack. Their goal is a specification that would yield a

multivendor solution that could be harmonized across the open standards community (e.g. OCP, Open19, and Scorpio). This open specification for a multivendor supported liquid cooled rack could increase deployment of warm water liquid cooling in HPC clusters and other high power density compute environments.

The EE HPC WG has been working with LBNL to help get out the word on this open specification by jointly preparing an SC18 BoF submission. This BoF will feature a panel of seasoned operations managers from major supercomputing centers to talk about strategies for effectively enabling warm-water cooling, including a discussion on the need for industry standards.

DASHBOARD TEAM:

The Dashboard team has been evolving into one focused less on dashboards and more on integrated and aggregated data analytics. Ghaleb Abdulla from Lawrence Livermore National Laboratory has developed a questionnaire that is intended to gather information from sites that have implemented – or are planning to implement - aggregated data collection for operational management (including energy management) in a production environment on at least one large-scale system (Top500 sized system) with integration that extends from the facility down to the CPU. It focuses on the data collection process and covers use cases for the data.

The intent behind this survey is to learn from others. There are two audiences: people who are already engaged and way ahead in the process (early adopters) and those who are further behind but interested enough to consider starting down this path. It assumes that a key challenge for using data analytic capabilities is the infrastructure; communication with the sensors and data grabbing. A lot of time is spent trying to go around meters, getting them to talk to each other, programming them the proper way, connecting them to the network, making sure that security is in place to be able to grab the data from a specific building to our server. The questionnaire is currently being trialed with Los Alamos and Argonne National Laboratories.

This Team made a submission for a Birds of Feather at SC18. If accepted, this BoF will present experiences from HPC centers in Europe, Japan and the United States regarding data source integration and analytics for dynamic power and energy management.

LIQUID COOLING COMMISSIONING:

No news from the Liquid Cooling Commissioning Team.

iTUE AND TUE:

No news from the iTUE and TUE Team.

RAS AND MAINTAINABILITY:

No news from the RAS and Maintainability Team.

ASHRAE:

The liquid cooling group within the IT sub-committee of ASHRAE TC9.9 has been working on a white paper primarily intended to supplement the liquid cooling book. It is intended to provide better clarity on operation with CDU vs operation without a CDU. More people want to connect

up to facilities and water systems directly, so what are the implications of doing that. In general, the white paper is intended to provide some design guidelines and best practices on how to interface with and connect up to the IT portion of liquid cooling systems. I'd say there is quite a bit of text in the white paper, it's getting reviewed now. I think we plan on having a version that will be voted out by the TC 9.9 committee by the winter meeting in 2019.

Systems Sub-group Update: *Natalie Bates, EE HPC WG and Jim Laros, Sandia National Laboratory*

ELECTRIC GRID INTEGRATION:

We reported at the last meeting that the Grid Integration team's work on electricity service contracts of supercomputing centers was presented in a Special Interest presentation at ISC'18. The session was well attended and participants stayed after the session was over to continue to ask questions. The Grid Integration Team had also submitted a paper on the same topic to SC18. Unfortunately, the paper was rejected. It has been resubmitted to the Ninth International Green and Sustainable Computing Conference.

The Team has transitioned from its focus on electricity service contracts is currently focused on potential power engineering challenges with future supercomputers and the impacts and nature of voltage fluctuations, dynamic load conditions and grid stiffness. Five sites have volunteered to participate on the team and share their concerns as well as their plans in this area. These are ORNL, UCAR/NCAR, ECMWF, LLNL and LANL.

The team's first deliverable will be a whitepaper that provides a comparative case study of five sites. The focus of the paper is future looking and asks questions about characteristics of grid connection, SC system design, power supplies, metering and the planning relationship of the SC and electricity service provider. The whitepaper may also describe lessons learned, best practices and may suggest possible next steps for this area within the high-performance computing community.

POWER MEASUREMENT METHODOLOGY:

The work of this Team is currently focused on encouraging the HPC community who make submissions to the various benchmark lists, like the Top500, to use a high quality power measurement methodology that captures the power of the entire system.

Oak Ridge National Laboratory provided the highest quality power measurement- Level 3- for their IBMPower Summit supercomputer submission to the Top500/Green500 List. Summit ranked #5 on the Green500 List and #1 on the Top500 List. There were two other systems with L2 submissions. These were from France and Germany; ROMEO HPC Center and Juelich Supercomputing Center, respectively. They were both Bull systems.

A Green500/Top500 and EE HPC WG Birds of Feather was held during ISC18 in June. The power measurement methodology was promoted during this BoF. This Team has made a similar submission for an SC18 BoF.

ENERGY AND POWER AWARE JOB SCHEDULING AND RESOURCE MANAGEMENT:

The Energy and Power Aware Job Scheduling and Resource Management Team held a Birds of Feather session at ISC18 in late June. The BoF generated discussion relatively quickly. The Team had also held a Birds of Feather session at SC17 last November. Comparing the SC and ISC BoFs, the SC18 focus was on power management at the job level whereas at ISC18, the discussion focused more on infrastructure and user participation.

This Team is still wrapping up the first phase of its work, which was to have done a survey of all supercomputing sites with large scale production deployments of energy and/or power aware job scheduling and resource management capabilities. They have written two papers describing their work. One paper was presented and will be published as part of the High Performance Power Aware Computing Workshop held in conjunction with IPDPS in May. The other paper was submitted to SC18, but was rejected. It has been resubmitted to the 2nd International Industry/University Workshop on Data-center Automation, Analytics, and Control, which is part of SC18.

For the next phase, the Team is looking at sites that are interested in using energy and power aware job scheduling and resource management capabilities to allow an application to provide hints and other relevant information to an EPAJSRM job scheduler and/or to notify applications of power management decisions, such as changes in power usage targets and providing awareness of what is going on in the machine that might have made a job run slower.

The Team has made a BoF submission that would explores these capabilities from the perspective of three different sites; LANL, LRZ and STFC.

PROCUREMENT CONSIDERATIONS:

The procurement team has spent the last six months analyzing energy efficiency data from 11 recent large HPC procurements from around the world. We agreed on a structure to our procurement best practices paper that includes an introduction, interfacing with facilities, evaluation, management and future trends. The team is working to merge content from the previous paper and information from analysis of the recent HPC procurements into this structure to produce a new paper. We don't have a tentative completion date, but won't have the paper completed by SC18.

The procurement team is hosting a Bird-of-a-Feather at SC18 to provide a preview to the paper by discussing energy efficiency highlights from a select set of the 9 large HPC procurements. So far we have garnered support from Anna Maria Bailey from LLNL who will speak about CORAL-2, Herbert Huber from LRZ to speak about SuperMUC-3 and SuperMUC-NG, David Martinez from SNL to speak about future trends, and Toshihiro Hanawa from U of Tokyo to

highlight the JCAHPC procurement. After brief presentations, we plan to open up to questions and discussion on energy efficiency for HPC procurements.

POWER API

The Power API Team has also made an SC18 BoF submission in collaboration with the Distributed Management Task Force's Redfish API. In this BOF, we will discuss the Power API and Redfish; APIs for measurement and control of power/energy on large systems. The BOF will introduce newcomers to these efforts, differentiate the goals of the two APIs and discuss inter-operability. An interactive panel discussion with experts from involved organizations will facilitate discussions between both API communities with ample time for audience questions and comments.

The PowerAPI is holding a workshop as part of the Ninth International Green and Sustainable Computing Conference. The aim of the workshop is to bring together researchers and developers to present and discuss innovative algorithms and concepts in the power/energy management of High Performance Computing systems and to create a forum for open and potentially controversial discussions on the future of power management in the Exascale era. Possible workshop topics include innovative algorithms for managing power budgets at system scale, job power management, management and optimization of energy/power at the node level and scalable monitoring methods.

PARTICIPANTS INCLUDED

Name	Organization
Natalie Bates	EE HPC WG
Anita Cocilova	LLNL
Sid Jana	Intel
Greg Koenig	KPMG
Wade Doll	Cray
Detlef Labrenz	LRZ
Ralph Wescott	PNNL
Aniruddha Marathe	LLNL
Matthias Maiterth	Intel
Todor Milkov	SDSC
Michael Ellsworth	IBM
Todd Takken	IBM
Michael Bane	STFC
Steve Martin	Cray Inc.
Dale Sartor	LBNL
Mike Mason	LANL
Ben Radhakrishnan	National University
Kevin Ng	Nvidea

Anna Maria Bailey	LLNL
Tim McCann	HPE
Ghaleb Abdulla	LLNL
Piotr Luszczek	University of Tennessee
Cate Berard	US DOE HQ