

**Energy Efficient High Performance Computing Working Group
06/11/2019 Meeting Report**

INTRODUCTION

The Energy Efficient High Performance Computing Working Group (EE HPC WG) held a meeting on 6/11/19. 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, August 13th, 2019 9:00-10:00AM Pacific Time

ANNOUNCEMENTS:

Dale Sartor announced the Clean Energy: Data Centers and Green Buildings in Hyderabad, India. The event, to be held on July 15-16, is the third of six workshops to be held in India under the United States Trade and Development Agency in coordination with the Confederation of Indian Industry and the Bureau of Indian Standards (BIS).

Also, Dale announced that the Data Center Dynamics Conference on July 18th in Bangalore will have an energy efficiency focused session.

From April, Andrea announced that there will be a Special Session on Tools for Energy Efficient Computing in Bialystok, Poland from September 8- 11. For more information, see: https://www.ppam.pl/docs/workshops/PPAM_2017_Special_Session_on_EE.pdf

David Grant, ORNL announced that we are planning to kick off the Controls Team again. This time, however, the team will focus more broadly on Cooling Controls and not just Liquid Cooling Controls. We held a webinar on controls that provided the impetus for restarting the team. Computer Scientist Luca Bortot at ENI's Green Data Center has created a control system for the center with extensive instrumentation, automated maintenance routines, and real-time adaptive cooling optimization. The development of a controls algorithm trading off between air and chilled cooling was explained in this webinar. A recording of the webinar and the presentation materials can be found on the EE HPC WG website.

David also announced that ASHRAE TC9.9 has published a new whitepaper on liquid cooling. We also held a webinar with some representatives from ASHRAE TC 9.9 where they described their new white paper on liquid cooling. The technical committee felt like they needed to outline some of the common processes, parts and materials for focus in the use for future water cool design. Some parts in a water-cooled IT system will be specific to the product design such as cold plates, manifolds, arrangement of pipes, pumps, valves and so on. Other parts, such as quick connect hoses, materials and water chemistry fall into more of a common category that

can be used by all current and potential manufacturers of water-cooled IT equipment. The white paper attempts to provide and make available those items that could be classified as common and the material published in the white paper complements the materials published in the Version 2.0 Liquid Cooling Guideline. The white paper also goes into more detail where the committee felt that people were having difficulty with interpretation. The Version 2.0 of the Liquid Cooling Guideline and the whitepaper are available on the ASHRAE TC 9.9 website. The webinar presentation they made and an audio recording of it is also posted on the EE HPC WG website.

GRID INTEGRATION TEAM REPORT:

Grant Stewart, LANL reported on the Grid Integration Team. We're very proud that two papers from the GI Team were accepted at the EE HPC SOP workshop. The paper that's been longest in development and is a longer paper is "An analysis of contracts and relationships between supercomputing centers and electric service providers". This is a very well written paper that surveys 10 sites and goes over some of the things they have in common and gives some insight about strategies and nuances of those contractual relationships. The second paper that we submitted and was accepted is one called "Grid accommodation of dynamic HPC demand." This paper surveys four HPC sites. It goes over some of the electrical engineering concepts that affect how power grids react to power fluctuations from supercomputers. The second part of the paper explores the four sites in more detail as to how the grid is feeding these supercomputers and what other features affect the dynamic power condition.

LIQUID COOLED RACK SPECIFICATION TEAM REPORT:

Dale Sartor, LBNL reported on the Liquid Cooled Rack Specification Team. We are focused on a 'liquid to the chip' cold plate technology rack. We had a draft specification that we started from and we identified nine or so components that we needed to work on to enhance that draft specification. There's a lot of interaction between the components, so some synthesis is also necessary. The first two areas that we've been working on is wetted materials and transfer fluids. The idea is to have an open spec multi-vendor solution especially for clusters of HPCs where we use the same racks for multiple refreshes. You could take servers from multiple vendors that are liquid cooled and basically plug them into the same rack. That's the vision.

We need to have similar wetted materials in the system for the integration to be successful. The wetted material team has been meeting for a couple of few months and is getting close to a conclusion. We're going to basically adopt the Open Compute Project (OCP) wetted material specification. In general, we're trying to harmonize across other organizations whenever we can. There have been a number of individuals that were interested in an aluminum option. Typically, you can't mix aluminum and copper or steel in the same system for corrosion reasons. These individuals- mostly vendors -were interested in an aluminum-only option so that the entire closed loop would have aluminum or aluminum compatible wetted materials. We're going to defer the establishment of an aluminum-only specification for a future time. It was decided not to try to integrate aluminum in with the current specification.

The transfer fluid team is zeroing in on utilizing the ASHRAE specifications for the transfer fluid. The fluid is water based and we decided not to go the OCP route as they are using a glycol

solution - propylene glycol 25 percent or 50 percent solution. We decided to not have a glycol-based solution and only go with water. It will be treated for corrosion and biocides. We did that for greater compatibility and higher heat transfer. We are hoping that OCP will adopt our standard as an option to their glycol solution.

David Sickinger, NREL added that the Team is also going to put together some best practices on things like transfer fluid and handling practices. Dale added that there's a lot of guidance for operational and maintenance practices that are important that we wanted to have in conjunction with the specification.

POWERAPI TEAM REPORT:

Jim Laros, Sandia NL reported for the PowerAPI Team. The Power API is a holistic specification for power monitoring and control on large-scale systems all the way down to small scale systems. It's basically an abstract device interface so that you can request basic things- like power and energy- to find out more about your system. It has a basic hierarchy that lets you navigate through all the objects in a system. Objects can be anything from the entire chassis, rack nodes and all the way down to individual cores. You can access power and energy consumption and even get things like temperatures. The specification exists and is available on the on the internet. We have a GitHub page and there's a link to it from the EE HPC WG website. The first official community version of the Power API Specification is now available. Please see https://github.com/pwrapi/powerapi_spec/releases for the PDF and full source code for version 0.9. The goal is to have the next version, 1.0, released by mid November.

OPERATIONAL DATA ANALYTICS TEAM REPORT

Norman Bourassa, LBNL/NERSC reported on the Operational Data Analytics (ODA) Team. The most recent activities centered around some case studies to be presented in Kyoto in August at the EE HPC SOP (state of the practice) Workshop which is being held at the International Conference on Parallel Processing (ICPP). We have two from NERSC and one from CINECA. One paper is a description of our OMNI data instrumentation system that we use at NERSC. We basically measure everything on the chance that data can be used in the future. The second paper is more directly related to operational data analytics. It is documenting 5 cases of applying OMNI for operational data analytics. I have not yet seen the paper from CINECA, but I look forward to seeing how they were using ODA to improve their operations with respect to other wider activities. Natalie Bates, EE HPCWG reported that in addition to these three case studies that were submitted to the EE HPC SOP workshop, there will be a poster from the ENI Green Data Center at the ICPP Conference itself. Norm added that for NERSC, the process of writing the case studies also turned out to be enormously useful for communicating what they do to their front office and upper management.

The ODA Team also has a global site survey of sites that are interested in operational data analytics in a production environment on their floor with an intention to integrate both the HPC systems and the facility data. We've collected questionnaire and interview data already from LLNL, NREL, LANL, LBNL, LRZ, RIKEN and CINECA. We have a questionnaire response from NCAR and are expecting one from ORNL as well. The end result of that will be some sort

of a white paper on trends. What are we seeing? Why is this important? We believe that operational data analytics is in the early adopter phase. The case studies and this survey paper will help to communicate to the broader community why this technology is important and how it can help them.

ENERGY AND POWER AWARE JOB SCHEDULING AND RESOURCE MANAGEMENT (EPAJSRM) TEAM REPORT

No news.

PROCUREMENT CONSIDERATIONS TEAM REPORT

Gert Svensson, KTH reported on the Procurement Considerations Team. Since the last public version of this document, we have first started to go through a number of procurement documents we could get and we selected requirements from them which we included in the EE HPC WG Procurement Considerations document. The next step was that we cleaned up the Procurement Considerations document little bit. We defined more rigorously what different classes of requirements really meant and also we structured the document a little bit better. So, we have applied that in the cooling section of the Procurement Considerations document as a first tryout and that has been discussed in length in different groups. I think we have settled on a very good structure of the document now, which has been applied to the cooling section of the document. Now we are working with applying this structure to the entire document and also to write a new section of the document which is one on interfacing with the facility. This describes what kind of impact existing stuff in the facility infrastructure will have on the demands of the procurement. We have made steady progress which is maybe not so fast, but I think the result is going to be good in the end. We are hoping to have something for SC19.

POWER MEASUREMENT METHODOLOGY TEAM REPORT

Natalie reported that the Power Measurement Methodology Team is co-hosting a Birds of Feather at ISC'19 along with the Green500 and the Top500. At SC'19, there were three presentations made from sites that had used the higher quality power measurement methodology. The first person to report was Jim Rogers from Oak Ridge for the Summit supercomputer. So the Oak Ridge facility in which Summit resides has a very high resolution power and energy instrumentation that allows for a level 3 or the highest quality measurement. So, you know is quite simple for them to do a level 3 measurement. They've got it all set up so that it's just part of their normal way of doing business. The second presentation was from AI Bridging Cloud Infrastructure (ABCI), which is an Advanced Industrial Science and Technology (AIST) system in Japan. They were able to make a level two measurement. The third presentation was from a RIKEN system called Shoubu. They were so motivated to make a level three measurement that they used external instrumentation basically clamp sensors and made a level three measurement. So, it's possible to make a level 2 or level 3 measurement with all kinds of instrumentation capabilities.

Just a quick aside. While preparing to do the measurement and collect the data for the RIKEN Shoubu system, they actually realized that they had an opportunity for further optimizing their

cooling subsystem and they added variable frequency pump controllers. Once you start collecting data, you start seeing places where you can optimize and it happened actually during this Green 500 power measurement activity as well.

CONFERENCES REPORT

Torsten Wilde, HPE reported on Conferences activities.

ISC2019 is scheduled for the week of June 16-20 in Frankfurt. The EE HPC WG has a BoF for the Green500. The EE HPC WG has been collaborating with the Green500 and the Top500 to improve the methodology by which power is measured while doing a full-system benchmark run- like MPLinpack. There are other activities at ISC2019 that should be of interest to the EE HPC WG. There is a full day Power Stack Workshop scheduled for Thursday, June 20th. There is a Power Stack project poster. Also, there is a half-day tutorial on GEOPM, which is an open source power management runtime and Intel will host a separate conference room to highlight other vendor plug-ins for GEOPM.

SC2019 is scheduled for the week of November 17-22 in Denver. The annual EE HPC WG workshop has been accepted and is tentatively scheduled for all day on Monday. The Datacenter Automation, Analytics and Control (DAAC) Workshop – which is a paper workshop and very complimentary to EE HPC WG - is scheduled for Friday. The EE HPC WG booth has been secured. The EE HPC WG has a panel on “AI, Big Data and HPC: Computing Under Constraints.” BoF and Poster submission deadlines are the end of July.

The EE HPC State of the Practice Workshop is planned as part of the International Conference for Parallel Processing (ICPP) Conference. It is scheduled for August 5th in Kyoto. This is the first-ever EE HPC WG workshop in the traditional sense with solicited papers from the community. There were 8 papers submitted and 7 accepted. Satoshi Matsuoka will be the keynote speaker.

PARTICIPANTS INCLUDED

Name	Organization
Natalie Bates	EE HPC WG
Torsten Wilde	HPE
Grant Stewart	LANL
Xingfu Wu	Argonne NL
David Strickland	NREL
James Laros	Sandia NL
Steve Bruno	DOE
Luca Bortot	ENI Energy
David Grant	ORNL
Dale Sartor	LBNL
Gert Svensson	KTH
Senthil Kumar	CDAC
Norm Bourassa	LB NL
Steve Leak	LBNL