

**Energy Efficient High Performance Computing Working Group  
08/13/2019 Meeting Report**

**INTRODUCTION**

The Energy Efficient High Performance Computing Working Group (EE HPC WG) held a meeting on 8/13/19. 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, October 8th, 2019 9:00-10:00AM Pacific Time***

**ANNOUNCEMENTS:**

Natalie Bates reported that Michael Ott from LRZ is the new technical lead for the Operational Data Analytics Team.

She also reported that we are starting a Cooling Controls Team, one that has grown out of the Liquid Cooling Controls Team. Chris DePrater from LLNL and Luca Bortot from ENI Italy are co-technical leads for that Team.

Dale Sartor reported that the Energy Exchange will be hosting a free hands-on tutorial for the DC Pro data center profiling tool. He and Steve Greenberg will also be giving a talk in the overall conference. Energy Exchange is a conference targeting the Federal Energy Management market and there will be a couple of sessions on data centers. <https://www.energy.gov/eere/femp/2019-energy-exchange-pre-event-workshop-series>

He also reported that there's a 7x24 Exchange conference in Phoenix. It's a conference mostly targeting commercial data centers, but there's usually a handful of representatives from super computer centers are like Lawrence Livermore National Laboratory. It is a bit un-usual in that it only has a single session as opposed to a conference divided into simultaneous sessions. We're going to have a session there on liquid cooling and the speakers will be Dave Martinez from Sandia, Otto VanGeet from NREL and myself. October 27 to 30. <https://www.7x24exchange.org/event/7x24-exchange-international-fall-2019-conference/>

Power Stack had a face-to-face prior to ISC as well as several sessions during ISC. It is also planning a face-to-face before SC19. <https://hpcpowerstack.github.io>

The Tenth international Green and Sustainable Computing Conference will be held October 21st to the 24th in Alexandria, Virginia. <https://www.igsc.org/>

The Data Center Automation, Analytics and Controls Workshop will be held on Friday, November 22nd as part of SC19. [daac-general.nsf.gov/daac-2019](https://daac-general.nsf.gov/daac-2019)

## TEAM UPDATES

Liquid Cooling Specification: Dale Sartor reported that David Sickenger from NREL submitted a Birds of Feather proposal for SC19. This would be a sequel to last year's and hopefully it will be accepted. It will be a good forum to further discussion on liquid cooling and open specifications. We are working on an open spec multi-vendor specification for a liquid-cooled rack. We're coordinating that activity with both the Open Compute Project as well as ASHRAE. We divided the effort into ~ nine components and started with two components. These are wetted materials and transfer fluid. Subgroups have been meeting about once a month. The goal is to develop two documents. The first is a specification that could be utilized in a procurement for a multi vendor liquid-cooled rack. The second document is an applications guide that would not only describe some of the nuances of applying the specification but also recommendations and experience relative to the operation of these systems.

Operational Data Analytics (ODA): Michael Ott reported that he had just taken responsibility as the technical lead for the ODA Team and that we were still in the process of coming up to speed on what had been happening with the Team. We are still in the process of a global site survey for Top 100 sized sites that have either integrated or are planning to integrate data from both facilities and HPC systems. The next interview will be with Oak Ridge National Laboratory. We also submitted a Birds of Feather to SC19, which we hope will be accepted. Norm Bourassa reported on the ODA activities at the EE HPC State of the Practice Workshop that was held in Kyoto on August 5th at the International Conference on Parallel Processing.

- Melissa Romanus gave a detailed overview of the Omni system at LBNL's NERSC facility. She described the base data instrumentation system where we pretty much measure everything on the assumption that we don't know what we will need until we need it. We just grab everything and we throw it into our HPSS system and store it permanently so that we can get to it if we need it. She gave a detailed overview of how that works and the type of effort it takes to get the system operational and keep it going. She also described some case study examples of the operational and procurement project benefits that we managed to derive from it.
- Then Walker Johnson and I providing a detailed overview of four case studies using the ODA capabilities of our instrumentation system for specific HVAC improvement examples. We also described a response to the emergency situation of the wildfires in California last year where ODA proved to be extremely useful for us.
- After that Andrea Bartolini from University of Bologna/CINECA provided a detailed description of their MQTT system and its ODA ability to detect anomalies. This is more HPC system data and analysis with operational benefits. It looks at normal compute node activity to detect anomalous behavior. He's presented various aspects of this over several of our functions and he gave a detailed update of their progress.
- Luca Bortot from ENI gave a talk on their HVAC control algorithms and how this could be useful for other sites that are trying to control the use of free cooling systems with supplementary chiller cooling.
- Lastly, Jim Rogers from Oak Ridge National Laboratory provided an update of their rapidly increasing activities of deploying ODA in preparation for the coming Frontier system.

- It's clear that instrumentation and integration of various data packages into a unified action form is gaining traction with hyper-scale system data centers.

Cooling Controls: Luca Bortot reported that the Team will be starting off by writing case studies on cooling controls. ENI and LLNL are the two sites already signed up to write a case study. The Team has been announced, but hasn't yet met.

Procurement Considerations: Gert Svensson reported that the Procurement Team submitted a BoF proposal at SC. The Team has been continuing the work with updating the HPC Energy Efficiency Procurement Considerations document. The idea is to have a document which gives some hints about how you can formulate procurements which are energy efficient. We have been working for quite a long time to get the right level of detail and the correct approach for the information. We have been studying a number of actual procurement documents and have selected requirements from those documents. We worked on finding the correct level of detail for the requirements. We have started with the cooling section. Then we have applied that to the rest of the sections so we got something which is more consistent and easier to read than it was before.

PowerAPI: Ryan Grant reported that the Team has pushed out version 0.9 of the Specification with the intent to have a version 1.0 in time for SC19 in November. So all the efforts at this point are going towards end.

EPA JSRM: Siddhartha Jana reported that the Team completed a third paper; this one a journal paper summarizing details and an analysis of the survey of all the global sites that are using or have plans to use energy and power aware job scheduling and resource management capabilities at scale in a production environment. The goal of the Team is now is to assess and collaborate with the Power Stack community and give guidance and help in reaching out to different HPC sites regarding requirements on how different components of the software stack should interact.

Power Measurement Methodology: Natalie Bates reported that the EE HPC WG, Green500 and Top500 held a BoF at ISC. Buddy Bland from ORNL said a few words about Summit, which is number one on the Top500 and number three on the Green500 Lists. They used the highest quality level3 power measurement methodology for their power submission. We are always looking for level two and level three submissions to the Top500 and Green500 Lists. We have also made a submission for a similar BoF for SC19.

## CONFERENCES UPDATE

Torsten reported on EE HPC WG conference activities.

- At the ISC19 Green500 BoF in June, it was reported that the Top 3 Green500 systems are: 1. Shoubu system B at RIKEN with 17.6 GFlops/Watt power-efficiency during its 1.06 Pflop/s Linpack performance run. It is listed on position 471 in the TOP500, 2. DGX SaturnV Volta at NVIDIA with 15.1 GFlops/Watt power efficiency. It is on position 469 in the TOP500 and 3, Summit at the Oak Ridge National Laboratory (ORNL) in Tennessee. It achieved 14.7 gigaflops/watt and is listed at number one in the TOP500. Interesting to see the smallest and largest systems (as measured by FLOPS) are

together in the Top3 of the Green500. [This is a correction to what was actually reported in the meeting.]

- Norm Bourassa did a thorough report on the ODA aspects of the EE HPC State of the Practice Workshop, which was the lions share of the workshop. There were other papers, including two written by the Electric Grid Integration Team. One was on Supercomputing Center contractual arrangements with Electricity Service Providers. The other was on voltage fluctuations and the local/grid distribution. There was also a paper on modeling the cooling system as well as one on designing an 80MW data center.
- In addition to the EE HPC SOP workshop, the ODA Team had a poster published as part of the ICPP Conference. This was on ENI's ODA capabilities.
- Prior to the EE HPC SOP Workshop in Kyoto, there was a site visit of RIKEN in Kobe. It was attended by people from sites in the United States, Europe as well as Japan. There was an open discussion on how we could improve collaborations between the United States, Europe and Japan. It was a very good exchange with a lot of practical exchange of knowledge. It may be followed by a similar event in September next year in conjunction with Cluster 2020, which will be held in Kobe Japan.
- The main thing that we're working on is organizing for SC19. The EE HPC WG SC19 workshop is scheduled for all day Monday, November 18th. We had our second session discussing a plan for the event. We also have an Exhibitor Booth again this year. There were 6 BoF submissions: PowerAPI, Operational Data Analytics, Liquid Cooling Specification, Green500, Exascale Facility Challenges, and Procurement Considerations. We'll know by the end of the month whether or not they were accepted. There is a panel on HPC AI and Big Data: Computing Under Constraints that we've organized and has been accepted. It which is scheduled for Friday, November 22nd.

***PARTICIPANTS INCLUDED***

<b>Name</b>	<b>Organization</b>
Natalie Bates	EE HPC WG
Torsten Wilde	HPE
Luca Bortot	ENI Energy
Dale Sartor	LBL
Gert Svensson	KTH
Norm Bourassa	LBL
Wade Doll	Cray
Siddhartha Jana	Intel
Andrew Fry	LLNL
Michael Ott	LRZ
Ryan Grant	Sandia NL
Benson Muite	Independent
Cate Berard	US DOE HQ
Mike Mason	LANL
Steve Martin	Cray
Thomas Durbin	University of Illinois