

Energy Efficient High Performance Computing Working Group

04/14/2020 Meeting Report

INTRODUCTION

The Energy Efficient High Performance Computing Working Group (EE HPC WG) held a meeting on 04/14/20. 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, June 9th, 2020 9:00-10:00AM Pacific Time

GENERAL ANNOUNCEMENTS:

We have created a new Mailing List Team. Our mailing list is the same one that we've been using since 2009, which is a word doc. There were originally around 40 people on the list and now it is close to 800 people. Charlotte Seto, Charlotte Seto, as Communication and Collaboration Consultant for the EE HPC WG, will be the technical lead for this Team.

Next Tuesday at this time, Jeff Broughton from Lawrence Berkeley National Lab will be giving the same keynote speech that he gave at the SC19 EE HPC Working Group Workshop. It was a really interesting keynote about Energy Efficiency in an Exascale Era. He covered a broad territory; from the environment, to the infrastructure and all the way down to CPU instruction sets. I encourage you to attend this webinar.

TEAM UPDATES:

Gert Svensson reported on the Procurement Considerations Team. One of the major things we are discussing at this point is whether the document should only focus on the high-performance computer platform itself or whether the focus should include other components, such as procuring scheduler software, or UPS equipment. Our intent is to review and update Section 4, on Usage Cases for Power, Energy and Temperature Management and Control. We will be reaching out to find more content experts to help with this.

Ryan Grant gave a report on Power API. Power API has been working on creating a questionnaire that we want to send out to multiple vendors. The goal of this questionnaire is to understand if we have all the right metadata that attaches to attributes like those that allow you to set frequencies and power caps for different components in the system. This is necessary because setting a frequency on a CPU can mean different things to different vendors. The questionnaire will make sure we know how they work such that the user can use these things in

a portable way, which is the whole goal of the Power API. Certain newer hardware releases have somewhat muddied the waters on what the metadata for a given frequency setting or a piece of hardware may actually do. We're trying to capture all of the subtleties of what may and may not happen. Are those frequencies actually guaranteed? Are they simply suggestions? Are they a mean average? What does it mean to actually set these things? IBM helped us out quite a bit in revising the questions for the questionnaire by filling them out over and over again. So, we already have detailed answers from IBM. After another round of revisions, we will send them out to the rest of our vendor partners. Most of those vendors have been participating in the power API meetings, so they are already familiar with what we are doing. It's much more complicated than you would actually think. There are low-level gotchas and some things that are Trade Secrets that not all the vendors want to admit. The other thing that we've been doing is engaging with the Power Stack engineering effort. So power API is the portable API and interface that is being used inside of Power Stack. Power stack is another effort that is trying to do the engineering for putting together a complete software stack to manage power efficiently and intelligently in large supercomputers. The Power API has been contributing in many different ways with the Power Stack effort in terms of portable APIs to enable that to happen. We don't specify the underlying engineering; for example, the nitty-gritty of how you exchange data from a job scheduler to an individual node. We provide APIs to make sure that the scheduler can communicate with the node.

Michael Ott reported on the Operational Data Analytics (ODA) Team. We have mostly been busy writing up the results from a global site survey on ODA. We did submit a poster on the project to the ISC'20 Conference that was accepted. ISC'20 will be an online event. I'm not sure how that's going to play with the poster. We also submitted an ISC'20 Birds of Feather session, which is on guidelines for HPC data center monitoring and analytic framework development. This BoF was also accepted, but I don't know if they will make it an online event. As I said, we have mostly been working on writing a paper that we are planning to submit to SC20 State of the Practice paper track. The abstract is due tomorrow and the full paper is due next week. The sites that participated in the global survey should have received an early draft of the paper. We are asking for your review to make sure that we got everything right and didn't make any false claims or phrased anything wrong. So if you have any comments, please come back to us. The paper covers the methodology used for the survey, an analysis of the general architecture for ODA, a few use cases and what we've learned overall. I think it is a well written paper and stands a chance of acceptance. That said, the paper acceptance rate for SC is very competitive, so we'll see.

Charlotte Seto gave a report on the Mailing List Team. The mailing list as it stands today is a hybrid solution with the word doc list as well as google calendar lists. It's not an optimal administrative solution. We created this team which will be meeting for the first time tomorrow to look at administrative platforms and solutions to optimize that system. We have several different requirements that all focus on allowing for better communication that bring all of these incongruent parts together onto a single pane of glass. Given that I haven't met my team members yet, I am excited to know what expertise they bring. We are looking for anyone with specific experience comparing different alternatives. We are comparing Constant Contact,

Google Suite, MailChimp and Causeway. We're also looking for people who have an understanding of potential security and privacy issues related to those four platforms. So if there is anybody on the call who wants to join the team. It should be a fairly short lived team with relatively limited responsibilities, but we'd really appreciate your participation.

Dale Sartor reported on the Liquid Cooling Specification Team. The Liquid Cooling Open Specification Team had divided into two sub-teams; one on transfer fluid and the other on wetted materials. The Transfer Fluid sub-group wrapped up the transfer fluid open specification draft and issued to the overall Liquid Cooling Specification Team. We're still open to comment; that spec will continue to be a living document.. For the time being, though, we're setting that behind us and we're looking to start on an in rack manifold spec. A related topic that I mentioned the last time was that we did do a scoping study for an energy analysis tool. That's been completed and I appreciate the input we got from several people from the working group.

Natalie reported on the Cooling Controls Team. This Team has been developing case studies. We have a couple of them in the works. The first one is completed from ENI and was published last November. NCAR is working on a case study where they designed a very sophisticated cooling control system, but during commissioning, turned it into a much simpler control system. We have another one underway from RSC and the Joint Supercomputing Center from the University of Moscow. Their problem is operating in an extremely cold environment and some of the control challenges they have had operating with such extreme cold temperatures.

Natalie also reported on the Electrical Commissioning Team. This Team is looking at large dynamic voltage power swings and the challenges of commissioning systems that are going to display this kind of behavior. It's the kind of behavior that we haven't really seen too much of in supercomputing, but we anticipate seeing more and more of it in the future. The initial deliverable of this team is just a very short white paper describing what the problem is and making some recommendations as to whether not a commissioning guideline, like what we did for liquid cooling commissioning, makes sense.

Torsten gave an update on the Conferences Sub-Group. Our Annual EE HPC WG Workshop was not accepted by SC this year. We are working with the leadership team and some of the workshop organizers and committee members to discuss alternatives. We are going to try to create a workshop organized by ourselves, much like the European Infrastructure Workshop. We will start in 2021 and it will be located near a major Supercomputing Center in the United States, changing locations from year to year. This year, we will move forward with the EE HPC SOP Workshop to be held as part of Cluster2020. It will be a paper workshop on September 14th. Presenters will be able to participate virtually in the workshop, but there is still planned to be a face-to-face component in Kobe, Japan. A call for papers will be sent in April for this workshop. The submission deadline is July 6th. Both short - 4 page and long - 10 page - papers will be accepted. We highly encourage you to submit a paper.

As you heard earlier from Ryan, Power Stack is an effort to engineer a dynamic power management software stack for HPC. They will have a digital event and currently its scheduled

from June 17th, which is Wednesday to June 19th, which is a Friday. If you are interested in participating please email Sid Jana or visit the Power Stack github.

Finally, as Michael mentioned, ISC20 as a face-to-face event has been cancelled and the organizers have not yet decided upon the schedule and format for a virtual event. So, more details will be forthcoming.

Participants included:

Gert Svennson, KTH
Ryan Grant, Sandia NL
Bill Brantly, AMD
Torsten Wilde, HPE
Shlomo Novotny, Independent
Steve Bruno, DOE HQ
Michael Ott, LRZ
Charlotte Seto, Independent
Benson Muite, Independent
Cosimo Pecchioli, Alfa Laval
Norm Bourassa, LBNL
Dale Sartor, LBNL
Otto VanGeet, NREL
Mike Strevell, LANL
Jason Zeiler, CoolIT
Charlotte Mendes, DOE HQ
Natalie Bates, EE HPC WG
Jim Laros, Sandia NL
Mani Prakash, Intel
Kristi Smith, Alfa Laval