

## Energy Efficient High Performance Computing Working Group 02/11/2020 Meeting Report

### INTRODUCTION

The Energy Efficient High Performance Computing Working Group (EE HPC WG) held a meeting on 02/11/20. 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, April 14th, 2020 9:00-10:00AM Pacific Time***

### ANNOUNCEMENTS:

- All of the SC19 EE HPC WG session presentations are posted on the website [https://eehpcwg.llnl.gov/conf\\_sc19.html](https://eehpcwg.llnl.gov/conf_sc19.html).
- Electrical Commissioning Team has been created with technical co-leads Brandon Hong, LLNL and Grant Stewart, LANL.

### TEAM UPDATES:

#### OPERATIONAL DATA ANALYTICS (ODA), Michael Ott (LRZ)

For quite some time, we have been occupied with synthesizing and analyzing the results of a global site survey. The survey is particularly focussed on understanding how sites are really using their data collection systems with specific use case scenarios and to what extent are data streams from HPC platform and infrastructure integrated into a common platform. We presented initial results in the EE HPC WG SC19 workshop. We are preparing a poster for submission to ISC'20, with a deadline of February 23rd. We are also working on an actual paper on the findings of our survey and we are shooting for SC20. There is a small group of us collaborating on writing the paper: Melissa AbdelBaky (LBNL), Norm Bourassa (LBNL) and Torsten Wilde (HPE) are describing specific use cases, myself and Stefan Ceballos (ORNL) are abstracting an architectural model, Woong Shin (ORNL) and Natalie are capturing the methodology used for the survey. Finally, there is also an ODA BoF submission in process for ISC'20 that is a joint effort with the EE HPC WG and HPE.

#### LIQUID COOLING SPECIFICATION, Dale Sartor (LBNL)

We have a large Team of about 60 people that is actually going to meet day after tomorrow. So if you are interested in working on an open liquid-cooled rack specification, please let Natalie or I know and we can add you to the Team meeting. The large Team has sub-teams that work on specific aspects of the specification. The first two were wetted materials and transfer fluids. The wetted materials sub-team has agreed to adopt the OCP (open compute programs) wetted material list. There is a draft of that available and it's going to be finalized later this year. On the transfer fluid specification, we issued a draft late last year prior to SC19. We'll be presenting that to the large Team for their comment and input. Then, we're calling it a milestone that is

done. It will be continuously improved as better information becomes available. Also at the Team meeting later this week, we're going to be choosing another topic to focus on. I think it's going to be the in rack manifold specification. We would be looking for any additional volunteers that would like to work on that specification.

David Sickinger from NREL is leading a parallel effort for this Team. He is developing a complimentary guide which is written in CSI construction institute format and focuses on issues in the design and operations phases of liquid cooling.

The last thing I want to mention is that we're doing a needs analysis for liquid cooling evaluation tools, simulation, etc. We are looking for participants in a focus group that have some knowledge of simulation or modeling or analysis, especially relative to liquid cooling. Where are the holes? What's needed? If we decide that there is a need for an analysis tool, what would be the inputs and the outputs of that kind of a tool? Again, if you're interested in participating in that focus group let Natalie or I know and we will include you.

#### POWERAPI SPECIFICATION, Ryan Grant (Sandia NL)

We've released version 1.0 Power API, which is community specification. Now we're in the process of going back and surveying vendors such that we better understand exactly what it means to set a performance state or a frequency of CPU or other components inside of a system. We have those generic abstracted interfaces, but there are some interesting subtleties that occur. We're going to expose those in our metadata interface. Then everybody who's using the Power API to get power and energy readings on their systems can understand the subtleties between different vendors offerings and ultimately be able to write tools according to those subtleties that would be useful to us. For example, what we have right now with the metadata is the ability to determine whether you get a 1Hz sample out of a system or whether you can get a 1kHz or a 10kHz sample out of a system. That leads us down slightly different paths on what we would do inside of a portable tool. So we want to understand exactly what's going on there. A group of us are writing a survey to send out to all of our different vendors to settle on the most standard forms of what happens when you set a performance state on a given system. Then we can potentially get the new interfaces right in that area.

We've also been planning for our face-to-face this year. We may end up with a dual-location face to face this year; one group in Denver and the other in Japan. The EE HPC State of the Practice workshop is planned to be held as part of IEEE Cluster2020 in Kobe in September. If you want to join for the monthly PowerAPI teleconference or for the face to face, that information will be going out on our mailing list and the mailing list link is also on the EE HPC WG webpage <https://eehpcwg.llnl.gov/power-api.html>

#### COOLING CONTROLS, Chris DePrater (LLNL)

The cooling controls Team is still its is fairly new. We've been meeting a few months and we're going over case studies on different control strategies for cooling controls, whether it be liquid or air cooled. We're going over lessons learned from the different sites and what works, what doesn't work. It's been very informative. We've had three case studies so far. We had one from ENI, one from LLNL and we're currently reviewing one from NCAR so nice work done with that team. The ENI case study was presented at the Data Center Automation, Analytics and Controls

workshop at SC19. There's a reference to it on the website under the bibliography page <https://eehpcwg.llnl.gov/pubs.html>.

GRID INTEGRATION, Grant Stewart (LANL)

We've decided to wind down a little bit on our power engineering focus within the grid integration group. We're in a regroup mode right now and exploring a new area that has to do with scheduling policy, metrics and ways to make the computer more compatible or a better partner in the grid environment. There is a lot of different facets of it that we've been talking about, it but it's really hasn't jelled yet.

Dale Sartor question: Has the Grid Integration Team given any consideration to micro grids to provide some greater level of resiliency in our supercomputers? This seem especially relevant in light of the serious power failures that last multiple days if not weeks.

Grant Stewart response: We haven't really talked about that aspect of it. Lots of supercomputing sites have their own generation and micro grids are possible. Microgrids can help a computer ride through a power fluctuation. Microgrids can also help the grid restart in some ways. They can have lots of benefits.

ELECTRICAL COMMISSIONING, Brandon Hong (LLNL) and Grant Stewart (LANL)

The Electrical Commissioning Team is very very new with co-technical leads Brandon Hong and Grant Stewart. We just met one time to discuss the team scope. Almost a decade ago, when liquid cooling became the new reality for large supercomputers, the EE HPC WG created a Liquid Cooling Commissioning Team. Their "Systematic Approach for Universal Commissioning Plan for Liquid Cooled Systems" is available on the website [https://eehpcwg.llnl.gov/infra\\_lcc.html](https://eehpcwg.llnl.gov/infra_lcc.html). Now, in this new age of highly variable power flow, when a new HPC system is starting up and it's in a new facility or the facility has been modified, there can be quite a bit of variability in power flow. It is probably going to generate some uncertainty, maybe trips, and possibly some equipment failures. So, we'd like to collect some ideas from what's been tried and what's working and maybe give some basic guidance for how to develop a commissioning strategy for the electric power system.

Randy Rannow question: Is there any interest in making it cyber-secure?

Grant Stewart response: We have not spent time focused on cyber security aspects of HPC's and their grid environment. I think that's a powerful topic. A lot of people could benefit from some guidance and exploration on what's been working and how the regulations come together.

CONFERENCES SUB-GROUP, Torsten Wilde (HPE)

We have two big conferences; ISC which is in Frankfurt Germany in June and then SC which is in Atlanta, Georgia in November. For ISC, we have two planned BoF submissions; one that Michael already mentioned during the ODA Team update., The other one is a Green500 BoF. We are currently finalizing the workshop proposal for SC which is due by the end of this week and hopefully we will get accepted again. The other thing we are also working on is another paper workshop- the EE HPC State of the Practice (EE HPC SOP) workshop, similar to the one we did last year in Japan. Ryan mentioned this in his PowerAPI report too. It is targeted

for IEEE Cluster 2020, which is in Kobe, Japan on September 14th. I would encourage everyone to submit papers to this workshop.

PROCUREMENT, Natalie Bates

The Energy Efficiency Considerations for HPC Procurement: 2019 document is finalized and posted on the website [https://eehpcwg.llnl.gov/compsys\\_pro.html](https://eehpcwg.llnl.gov/compsys_pro.html). The team is now looking at the next refresh of the document. There are a lot of suggestions and we will be meeting to review and prioritize the options.

POWER MEASUREMENT METHODOLOGY, Natalie Bates

This Team is always encouraging people to use a L2 or L3 (high quality) power measurement methodology when running Linpack for the Top500 and Green500 submissions.

Dale Sartor question: There have been a lot of good ideas brought up today for new teams; things like microgrid and cyber secure grids. Maybe you could describe the process for creating a new EE HPC WG Team?

Natalie response: First, we need a champion who is passionate about a topic that is relevant to the EE HPC WG mission. That person needs to be able to describe why the topic is relevant, who might be interesting and what would be the team deliverable. There also needs to be broader interest; some reasonable number (minimum of 5 committed people) of other people who are interested. Once this is established, the EE HPC WG leadership team reviews the proposal and approves the team. A call for participation announcement describing this is sent to the EE HPC membership and the team is created from there.

### ***PARTICIPANTS INCLUDED***

Natalie Bates, EE HPC WG  
Charlotte Mendez, DOE HQ  
Ryan Grant, Sandia NL  
Steve Martin, HPE  
Daniele Cesarini, CINECA  
Michael Ott, LRZ  
Randy Rannow, Silverdraft Supercomputing  
Torsten Wilde, HPE  
Todd Taken, IBM  
Woong Shin, ORNL  
Dale Sartor, LBNL  
Grant Stewart, LANL  
Dave Martinez, Sandia NL  
Paweł Ciechomski, University of Warsaw  
Alexander Moskovsky, RSC  
Norm Bourassa, LBNL  
Brandon Hong, LLNL

Chris Deprater, LLNL  
Charlotte Seto, EE HPC WG  
Cate Berard, DOE HQ  
Andrew Kolly, DOE HQ  
Phil Naecker, TGS  
Detlef Labrenz, LRZ  
Bruce Myatt, ARUP