

By Mike Bernhardt
July, 2012

For Lawrence Livermore National Laboratory, the work is just beginning.

LLNL Takes the Lead in Managing the DOE's FastForward Program

Around the end of March, 2012, The U.S. Department of Energy informed the folks at Lawrence Livermore National Laboratory (LLNL) they had been selected to run the DOE's FastForward program.

Now LLNL is certainly not 'running' this on their own – they are for the most part on equal footing with their sister labs (in alphabetical order), Argonne National Laboratory, Lawrence Berkeley National Laboratory, Lawrence Livermore National Laboratory, Los Alamos National Laboratory, Oak Ridge National Laboratory, Pacific Northwest National Laboratory, and Sandia National Laboratories.

They all have parallel and comparable research responsibilities, but for LLNL, those are in addition to the very big task of managing the overall FastForward program.

It was LLNL's responsibility to create the statement of work based on the guidelines they received from DOE headquarters. LLNL has a very experienced procurement team that is well versed in handling this type of assignment, but the timeframe allotted to complete the RFP process and kick off the program was bordering on

ridiculous. They were directed to have the awards ready to announce by the end of June. Two and one half months to prepare the statement of work, get the word out to all the potential bidders, receive the RFP responses, sort through them, analyze the proposals, evaluate the proposals for both technical and business merit, vote, reach a decision, inform the award recipients, and negotiate the contracts and IP provisions. All in two and one half months. And they did it.

It takes strong and wise leadership to manage the interface to multiple brain trusts such as the National Labs, and the coordination and holding to task the most powerful computer manufacturers on the planet.

Meet Terri Quinn, Principal Deputy Department Head, Integrated Computing and Communications at Lawrence Livermore National Laboratory.

Terri started at LLNL in 1984 on the application side writing vector codes. She moved over to high performance computing during the first year of the ASCI program, and had been working on the R&D side as opposed to the operations side.

Now as department head of the LLNL computing center, she's on the same team as some other well known HPC leaders and luminaries including Dona Crawford, Trish Damkroger, and [Michel McCoy](#), recently honored for his pioneering work in HPC by the National Nuclear Security Administration's first Science and Technology Award.



the Exascale Report™

According to Terri, “It was quite a hectic spring and early summer to get this program launched.” One of the most daunting tasks faced by the LLNL team was working through the volumes of Intellectual Property issues that had to be managed. As you can imagine, vendors were being asked to work in very close collaboration, but were, and still are, reluctant to unveil too much about their intellectual property for fear it could derail the competitiveness of their future products.

As the administrator of the FastForward program, Terri realizes the pressure her team is under. “From this point on, the contracts are with Livermore. The lab has a tremendous responsibility and a due diligence to ensure things are done right – that everyone is meeting their specific obligations and accountable for their areas of research.”

And don’t forget those multi-million dollar payments. The Livermore team is responsible for making sure the funds are disbursed appropriately per the contracts.

When it comes down to it, there really hasn’t been much of a challenge in disbursing funds for an exascale stepping stone project up to this point. There really hasn’t been any substantial money allocated to exascale.

Now, teams of NNSA, DOE and lab personnel will need to determine milestones, measure progress, and be the guides for keeping all the parties involved moving in a direction that pushes the envelope of technical computing. Different teams will handle the processor, the memory and the storage requirements.

As already stated in several places, there have been five awards announced to date totaling \$62.5 million:

- AMD, IBM, and Intel Federal were selected for memory.
- AMD, NVIDIA, and Intel Federal were selected for processors.
- And Whamcloud was selected for storage and I/O.

The plan of course is to foster a heavy program of co-design which will require a tremendous amount of cooperation and collaboration. To this point, Quinn commented, “Discussions are still underway on how all these groups will engage. We all understand the need for close collaboration in order to influence future architectures and product designs.”

According to Quinn, “Basically what is happening here is that the government driven initiative is working with the leading manufacturers to say, “we’d like to pay for some things in your future product and architecture design – that you might not have put in on your own – because you presently might not see enough market demand for it or to accelerate technology development because HPC is ready for it now”

So the vendors essentially had to convince the labs that what they had proposed in their RFP responses would ultimately get productized and have a broad market base.

Quinn added, “Keep in mind, this program is only two years. This is not sufficient to get us to exascale. To actually take any of this research all the way down the product development path, these companies will undoubtedly need to ask for more money.”

Another point for all of us to keep in mind is that there really is no exascale program approved at a

the Exascale Report™



Federal level. This is money taken out of the DOE's HPC budget to attempt to influence future

products and architecture design – to start now in paving the path to exascale.

Some people think the timing is ok and that it's not too late to influence product design and development for the end of this decade. Others are still concerned that we should be making much more progress at this point.

"Hopefully, two years out, we'll get some exascale money that we can add to this program or create other programs that will carry this work forward, or add new companies to this because for this pass, we had to leave some technologies out," said Quinn. She went on to elaborate, "Like interconnect, programming models, software, power, packaging, and systems reliability research. That has all been left out. What we did was to start with what we felt were the most critical long-lead items – that would line up with our budget in terms of being able to influence them." Quinn has a determined positive tone to her voice when she responds to a question like, "What happens if there isn't any follow on money or programs to FastForward?"

"We expect that if we only had FastForward and no future money, we would still have influenced some aspect of future design and key components coming to market that would not have changed, and some critical HPC requirements will have been inserted to improve computational capabilities for many applications."

Quinn is enthusiastic and confident. "What I think is really incredible is the opportunity these companies are giving us to actually influence technology. They had a lot of innovative ideas in the proposals, all of which were very high quality, but additionally, they are very open to having us contribute at a really in depth level. We've been able to do this in pieces before, but to this extent, I think it is really an extraordinary opportunity for us

– to not just influence one line of computers, like the Blue Gene line for example, where we've done

co-design with those folks for the past eight years, but in this case, we can have some influence on IBM, Intel, NVIDIA and AMD. It was actually surprising as to how open the companies were about doing this and how willing they were to jump in and work with us on co-design. I think this is creating tremendous opportunities, not just for the DOE people, but for HPC scientific and engineering overall, and it's really going to be exciting to see how this progresses over the next two years and what comes out."



The 23,504-square-meter (253,000-square-foot) Terascale Simulation Facility (TSF) at Lawrence Livermore National Laboratory

So, I'm convinced the FastForward program is in good hands, and I think the elevated discussions around exascale are reigniting a certain U.S. competitive spirit. But the one thing that bothers me with this is the lack of specific funding to research the daunting power issue.

What about the power?

the Exascale Report™

According to Timothy Pricket Morgan in his article “DOE Doles Out Cash to AMD, Whamcloud for Exascale Research”

<http://insidehpc.com/2012/07/12/doe-doles-out-cash-to-amd-whamcloud-for-exascale-research/>

[On a current petaflops-class system today, it costs somewhere between \$5m and \$10m to power and cool the machine today, and extrapolating to an exascale machine using current technology, even with efficiency improvements, you would be in for \$2.5bn a year just to power an exascale beast and you would need something on the order of 1,000 megawatts to power it up. That's 50 nuclear reactors, more or less. The DOE has set a target of a top juice consumption at 20 megawatts for an exascale system.]

The FastForward program is pushing memory, processors, and storage and does emphasize power and energy usage, but only so much can be done with 2 years of R&D investment. Power is the one item that conceivably will keep us from ever achieving a practical exascale system.

Nor does FastForward address research for programming models. And as my esteemed colleague also pointed out so eloquently in his article, “Programming these petaflops machines is a complete bitch, and an exaflops system will be in the range of old battle-axe mother-in-law. Beyond that, you are programming against Death.” Well said Timothy.

We'll close with this. Congratulations to the folks at Lawrence Livermore National Laboratory. Kudos to Terri Quinn and the amazing effort put forth by the LLNL team to date, and our best wishes for an incredibly successful FastForward program to all.

But, before we get too excited, let's not forget - the work has just begun.

Talk to the Emerging Exascale Community

Don't miss out on your opportunity to talk to the emerging Exascale community by submitting a sponsored editorial opinion article.

Talk to the community in your own words. Address the issues you feel most strongly about. Raise your concerns or ideas, or present a challenge to policy makers, funding sources, vendors or end users.

For more information:

<http://theexascalereport.com/pdf/advertorial.pdf>

For subscription information, go to:

<http://theexascalereport.com/user/registTer>

The Exascale Report:

The Voice of the Emerging Exascale Community.

© 2012, The Exascale Report™