HYPERION RESEARCH

HPC User Forum Update

Leadership Computing and the NSF's Computational Ecosystem, Dearborn, Michigan, September 2018

Steve Conway and Bob Sorensen
January 2019

IN THIS UPDATE

The HPC User Forum was established in 1999 to promote the health of the global HPC industry and address issues of common concern to users. In September 2018, Hyperion Research hosted its 70th User Forum in Dearborn, Michigan. This update summarizes a presentation from that meeting given in the session Leadership Computing Initiatives, entitled *Leadership Computing and NSF's Computational Ecosystem* given by Irene Qualters, Senior Science Advisor, NSF.

Qualters' talk highlighted that NSF views HPC as part of a larger computing ecosystem composed of instruments, software, and expertise to provide capacities and capabilities for S&E research not otherwise possible. She also discussed how NSF sees so-called Big Ideas such as quantum computing multi-messenger astrophysics and work at the human-technology frontier as research drivers for NSF's more expansive cyber infrastructures. Finally, new information on the latest NSF-funded HPC, the new Frontera system to be installed at TACC in 2019, was provided.

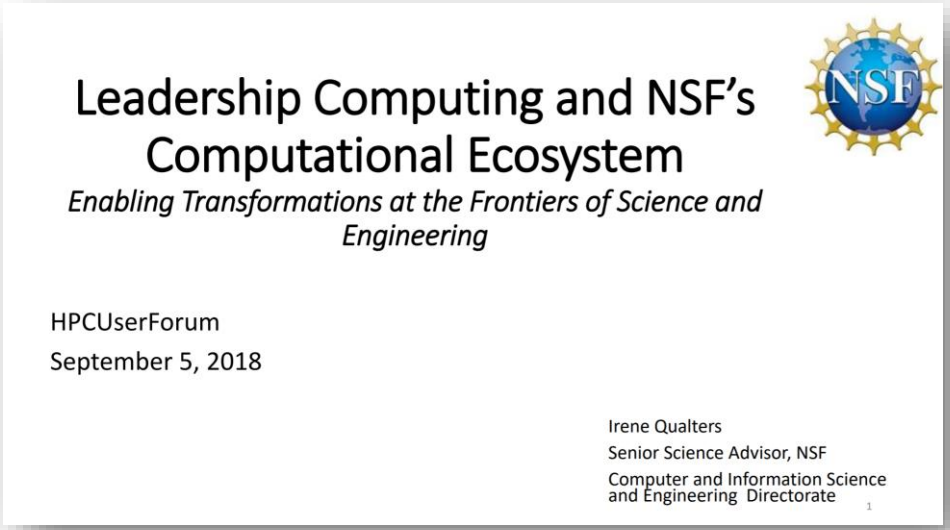

PRESENTATION: LEADERSHIP COMPUTING AND NSF'S COMPUTATIONAL ECOSYSTEM GIVEN BY IRENE QUALTERS, SENIOR SCIENCE ADVISOR, NSF

Qualters began her talk by stating that leadership-class computing, whether it is exascale or not, is a computing ecosystem composed of instruments, software and expertise, and NSF investments are national assets. When the NSF makes investments, they do it in the bigger context of what's happening internationally, both from a science and engineering point of view and looking at what investments other entities are making.

FIGURE 1

NSF Leadership-Class Computing

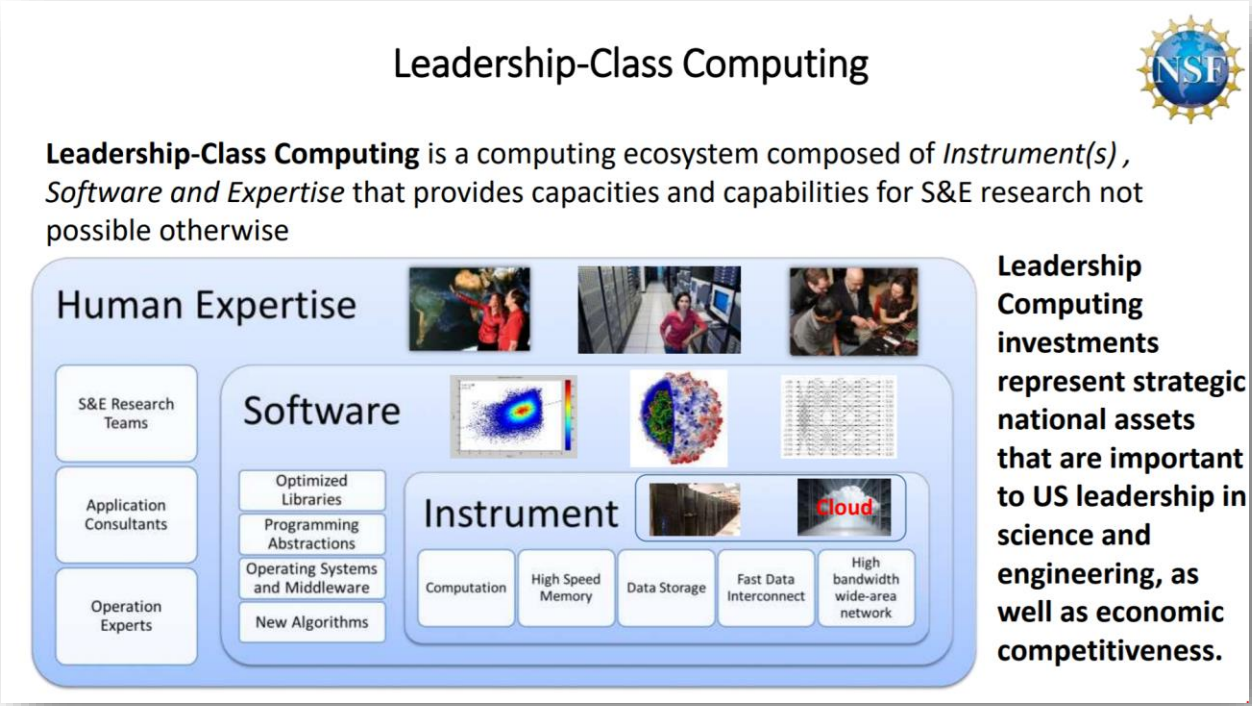

Source: NSF and Hyperion Research, 2018

Qualters stressed that NSF planners look at the entire S&E landscape, beyond the computing sector. She gave a few examples of how this is playing out as the NSF looks to the future. In mid-2015, NSF announced the instantiation of Big Ideas, six research topics (quantum leap, predicting genotype to phenotype, navigating the new arctic, multi-messenger astrophysics, work at the human-technology frontier, and harnessing data) that were selected to help drive NSF's long-term research agenda. These Big Ideas reflect the fact that many of the scientific frontiers that NSF sees today are increasingly convergent from a disciplinary perspective.

FIGURE 2

NSF Big Ideas as Research Drivers

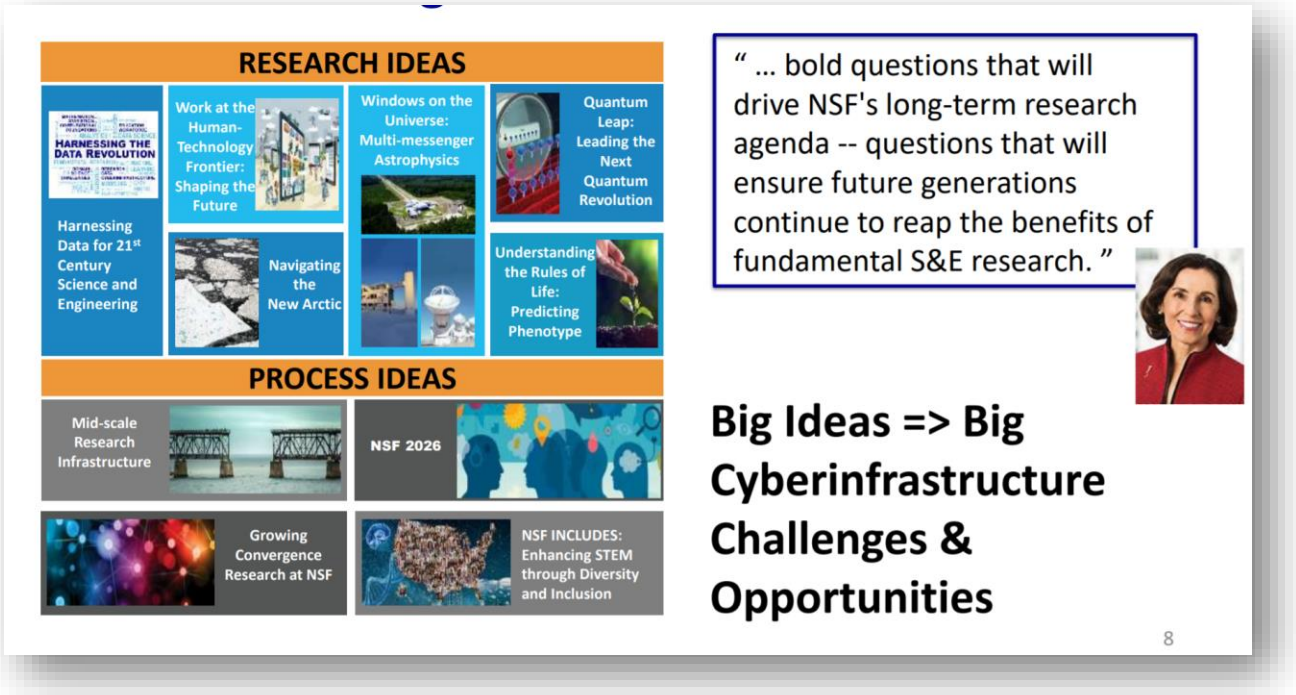

Source: Hyperion Research, 2018

When Qualters sees big research, she sees a large, complex cyberinfrastructure that represent both challenges and opportunities. For example, LIGO (the world's largest gravitational wave observatory) was the single largest investment - from an instrument perspective - that NSF ever made. It cost hundreds of millions of dollars and supported two interferometers to detect gravitational waves. It was a Big Idea, 20 years in the coming, and the first time out it didn't work.

- Ultimately the ability to detect gravity waves (the detection was accomplished in 2017, the Nobel prize was given in 2018) took an enormous amount of computing, collaboration, and software, not just on the application side, to do the detection.
- The effort also required numerous computationally intensive simulations. When the effort was started, Einstein's ten equations had not been solved; no one knew what two black holes colliding looked like, even if they detected it.
- The paper that was finally published had 1,000 people on it from all over the world.

In the meantime, NSF has been developing a next generation leadership-class system, as a two phase system, as well as planning for the rest of the overall computing portfolio to support its research community. For the new two-phase HPC development, NSF issued a solicitation last year looking for at least two- to three-fold performance increase in realized application performance over the current NSF Blue Waters system that would also lead to a phase two system five years out that would be ten times greater application performance.

In August 2018, the National Science Board authorized NSF to approve the acquisition by the University of Texas at Austin of the phase one Frontera system for $60 million, as well as $8 million to begin the phase two preparation.

FIGURE 3

TACC's Frontera System

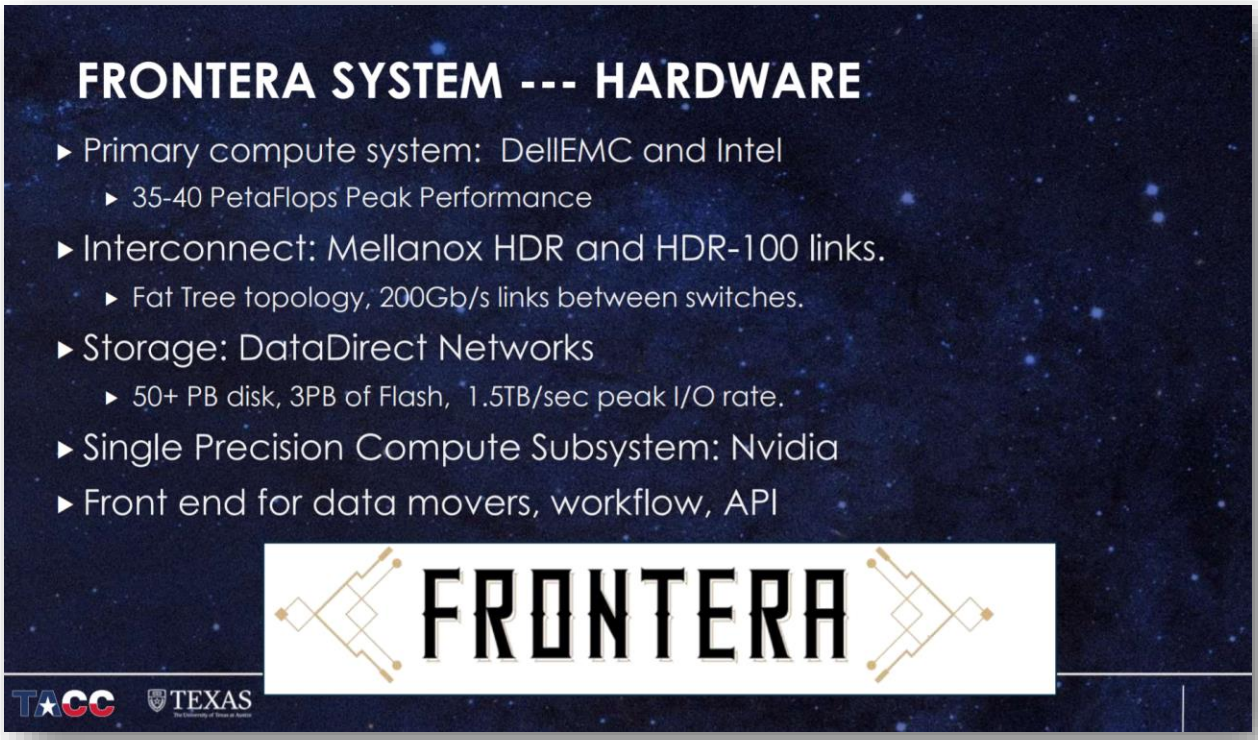

Source: Hyperion Research, 2018

Frontera will be deployed almost immediately, in 2019, and will operate for five years until the phase two system is stood up. Frontera will be supplied by Dell EMC and Intel with a performance range of 25-40 petaflops, a Mellanox HDR interconnect, and 50+ petabytes of DDN disk.

Frontera will interface to other NSF infrastructures for archival systems, public data repositories, and data transfer software. The system will also connect with cloud providers, particularly Microsoft, Amazon, and Google, for the data side.

- The Frontera system will coordinate with NSF's overall scientific ecosystem through partnerships with large-scale instruments, software development teams, and the XSEDE program.

For more information or to view this and other presentations given at HPC User Forums dating back to 2008, visit www.hpcuserforum.com.

About Hyperion Research, LLC

Hyperion Research provides data-driven research, analysis and recommendations for technologies, applications, and markets in high performance computing and emerging technology areas to help organizations worldwide make effective decisions and seize growth opportunities. Research includes market sizing and forecasting, share tracking, segmentation, technology and related trend analysis, and both user & vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). We provide thought leadership and practical guidance for users, vendors and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

Headquarters

365 Summit Avenue
St. Paul, MN 55102
USA
612.812.5798
www.HyperionResearch.com and www.hpcuserforum.com
