U.S. DEPARTMENT OF ENERGY | Office of Science

A View From Germantown

Presented to

Advanced Scientific Computing Research Advisory Committee

Barbara Helland, Associate Director
Advanced Scientific Computing Research

December 12, 2018

Office of Science FY 2019 Appropriations

	FY 2018	FY 2019			
	Enacted Approp.	President's Request	House Mark	Senate Mark	Enacted Approp.
Office Of Science					
Advanced scientific computing research	810,000	899,010	914,500	980,000	935,500
Basic energy sciences	2,090,000	1,850,000	2,129,233	2,193,400	2,166,000
Biological and environmental research	673,000	500,000	673,000	715,000	705,000
Fusion energy sciences	532,111	340,000	590,000	425,000	564,000
High energy physics	908,000	770,000	1,004,510	1,010,000	980,000
Nuclear physics	684,000	600,000	690,000	710,000	690,000
Workforce Development for Teachers and Scientists	19,500	19,000	19,500	24,500	22,500
Science laboratories infrastructure	257,292	126,852	290,147	302,100	232,890
Safeguards and security	103,000	106,110	106,110	106,000	106,110
Program direction	183,000	180,000	183,000	184,000	183,000
Total, Office Of Science	**6,259,903**	**5,390,972**	**6,600,000**	**6,650,000**	**6,585,000**

Budget Language

The following is the only direction provided for ASCR.

- Within available funds, the agreement provides $140,000,000 for the Argonne Leadership Computing Facility, $200,000,000 for the Oak Ridge Leadership Computing Facility, $105,000,000 for the National Energy Research Scientific Computing Center at Lawrence Berkeley National Laboratory, $10,000,000 for the Computational Sciences Graduate Fellowship program, and $85,000,000 for ESnet. The agreement provided $75,667,000 for Computational Partnerships (SciDAC). **Within funds for SciDAC, up to $13,000,000 is to support work on artificial intelligence and big data focused on the development of algorithms and methods to identify new ways of extracting information generated at the Office of Science's large user facilities of validating the use of machine learning in the Office of Science's program's scientific simulations. This is the only funding recommended within the Office of Science that shall be available for this work. The Department is directed to provide to the Committees on Appropriations of both Houses of Congress not later than 90 days after the enactment of this Act a briefing on its plan for implementing this artificial intelligence and big data initiative.**

ASCR FY 2019 Enacted

in thousand

	FY 2018	FY 2019	
	Enacted Approp.	President's Request	Enacted
Mathematical, Computational, and Computer Sciences Research			
Applied Mathematics	30,104	40,316	28,206
Computer Science	29,508	38,296	22,000
Computational Partnership	49,910	62,667	75,667
Artificial Intelligence and Big Data (Non Add)	*(3.500)*	*(…)*	*(13,000)*
SBIR/STTR	4,301	5,352	4,768
Total, Mathematical, Computational, and Computer Sciences Research	117,823	146,631	130,641
High Performance Computing and Network Facilities			
High Performance Production Computing (NERSC)	94,000	80,000	104,000
Leadership Computing Facility at ANL (ALCF)	*110,000*	*140,000*	*140,000*
Exascale	*(…)*	*(100,000)*	*(140,000)*
Leadership Computing Facility at ORNL (OLCF)	*162,500*	*200,000*	*199,000*
Exascale	*(62,500)*	*(100,000)*	*(100,000)*
Total, Leadership Computing Facilities	272,500	340,000	339,000
Research and Evaluation Prototypes	24,260	24,452	24,452
High Performance Network Facilities and Testbeds (ESnet)	79,000	56,435	84,000
SBIR/STTR	17,417	18,786	20,701
Total, High Performance Computing and Network Facilities	487,177	519,673	572,153
Exascale Computing			
17-SC-20 Office of Science Exascale Computing Project (SC-ECP)	205,000	232,706	232,706
Total, Advanced Scientific Computing Research	**647,000**	**899,010**	**935,500**

Darshan earns R&D 100 award as one of the 100 most innovative technologies of 2018

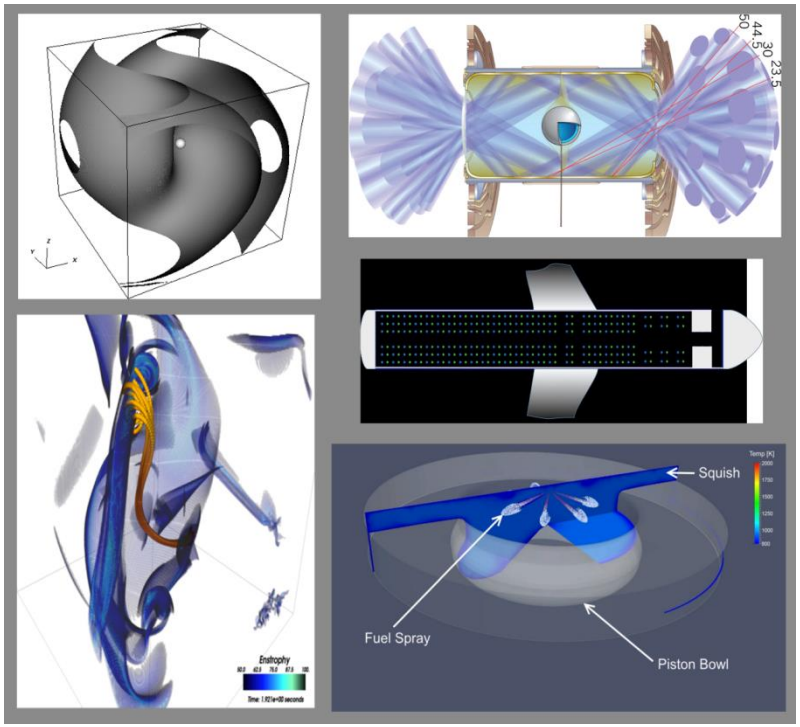

Examples of large-scale scientific applications that have been enhanced with the help of Darshan. Clockwise from the top right: pF3D (fusion physics), VIPRA (infection spread via air traffic), CONVERGE (combustion engine), FLASH (astrophysics), and ATHENA (astrophysics).

R&D 100 award work was performed by the Darshan project team at ANL: Philip Carns, Kevin Harms, Robert Latham, Shane Snyder, and Robert Ross.

Achievement

Darshan was selected by R&D magazine to receive a prestigious R&D 100 award; it honors the top 100 technological innovations of the year in any field.

Significance and Impact

The Darshan software package provides insight into I/O behavior of data intensive scientific computing applications. It has been deployed at every major DOE HPC facility and has fundamentally changed our approach to performance engineering.

Research Details

Darshan has been leveraged to:

- improve scientific applications and facility support methods (SciDAC)
- increase our understanding of emerging technologies at scale (ECP)
- Investigate facility-wide factors that contribute to I/O performance (SSIO)

2018 R&D 100 WINNER

https://www.anl.gov/mcs/article/darshan-wins-rd-100-award-for-2018

MENNDL Wins R&D 100 Award

The Science

ORNL's MENNDL is designed to evaluate, evolve, and optimize neural networks for unique datasets. Scaled across ORNL's Summit supercomputer, MENNDL can test and train thousands of potential networks for a science problem simultaneously, eliminating poor performers and averaging high performers until an optimal network emerges. The process eliminates much of the time-intensive, trial-and-error tuning traditionally required of machine learning experts.

The MENNDL R&D 100 team: pictured left to right are Thomas Karnowski, Derek Rose, Tom Potok, Robert Patton, Seung-Hwan Lim, and Steven Young.

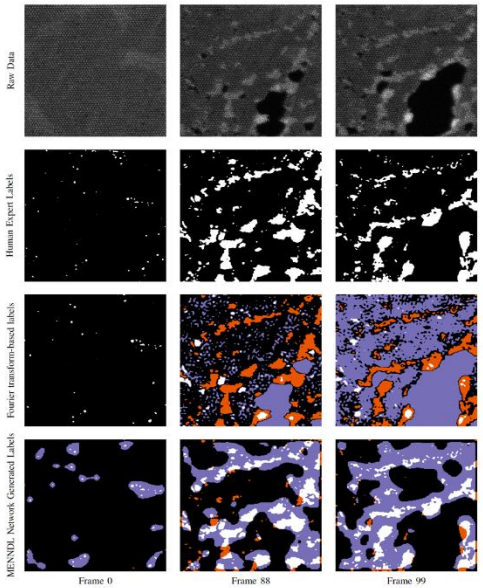

Identification of structural defects in scanning transmission electron microscopy data. First row: raw experimental data. Second row: result of a human expert (semi-)manually labeling atomic defects. Another human expert may label this slightly differently. Third row: result of a Fourier-transform based method tuned to Frame 0. Fourth row: result produced by convolutional neural network. In the third and fourth rows, white indicates correctly identified defects (good), purple pixels identified as defects that were not defects (bad), and orange indicates defects that were not identified (worse). Black indicates regions that contained no defects and were identified as such.

The Impact

MENNDL's unique ability to generate networks and analyze data was recognized in the 2018 R&D 100 Awards. The awards, known as the "Oscars of Invention," honor R&D pioneers and their revolutionary breakthroughs in materials science, biomedicine, consumer products and more from academia, industry and government-sponsored research agencies.

2018 R&D 100 WINNER

ORNL Summit System Overview

System Performance

- Peak of 200 Petaflops (FP_{64}) for modeling & simulation
- Peak of 3.3 ExaOps (FP_{16}) for data analytics and artificial intelligence

The system includes

- 4,608 nodes
- Dual-rail Mellanox EDR InfiniBand network
- 250 PB IBM file system transferring data at 2.5 TB/s

Each node has

- 2 IBM POWER9 processors
- 6 NVIDIA Tesla V100 GPUs
- 608 GB of fast memory

  (96 GB HBM2 + 512 GB DDR4)
- 1.6 TB of NV memory

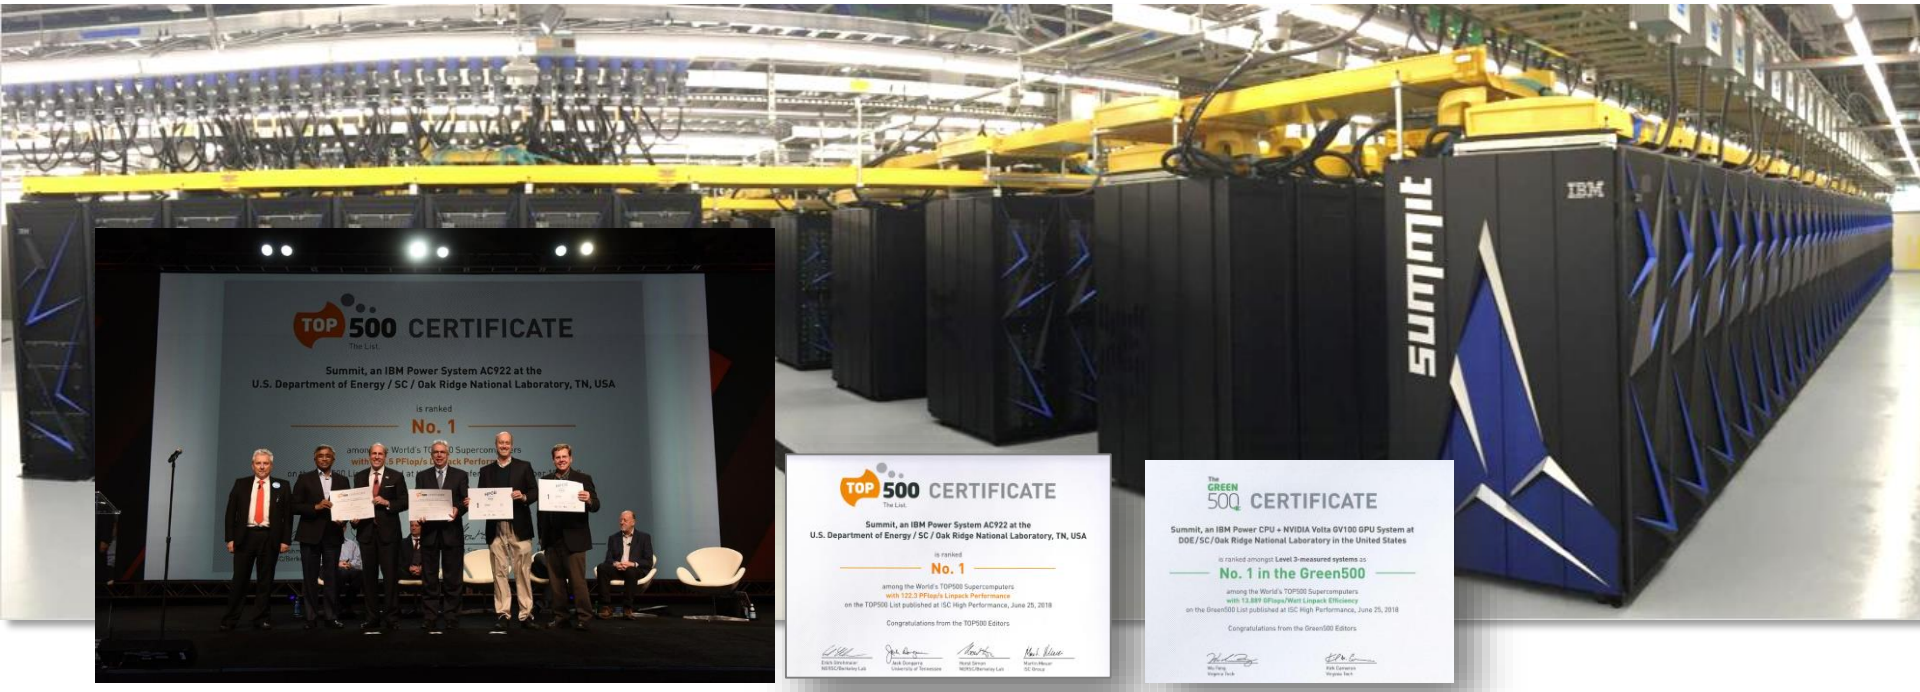

ASCR Facilities Project Update

Achieved in FY2015
OLCF-4 CD 2-3

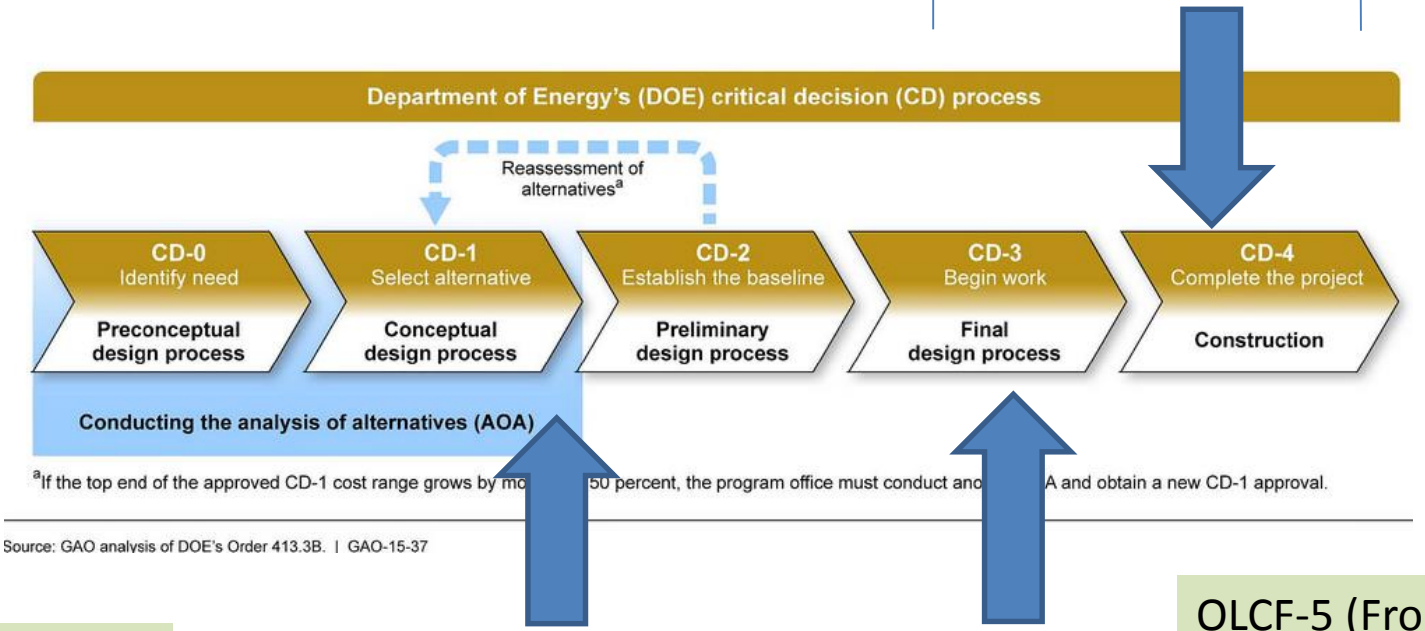

Esnet-6 IPR scheduled for Fall, 2019

Achieved in FY2018
Esnet -- CD 1-3A

OLCF-5 – CD 1-3A

Achieved in FY2018
NERSC- 9 – CD 2-3

ALCF-3 – CD 2-3 –A-21

OLCF-5 (Frontier) CD-2 reviewed planned for February, 2019 and IPR of ALCF-3(Aurora) scheduled for April, 2019

ECP Annual Independent Project Review

- **Conducted by SC's Office of Project Assessment at the request of Mark Anderson, NNSA and Barb Helland, DOE/SC on October 30- November 1**

- **In carrying out its charge, the Review Committee should respond to the following questions:**
  - Has the project satisfactorily addressed recommendations from the January 9-11, 2018 Project Review?
  - Is the ECP process to track performance adequately documented? Is this a reasonable and tailored approach to measure project performance?
  - Has adequate technical progress been made since the last IPR, including refinement of KPPs to appropriately reflect ECP goals and mission need, Software Technology's capability assessment, Hardware Integration's facility engagement, and the current status of PathForward projects?
  - Has team adequately identified the information and documentation that needs to be in place for CD-2 approval in Q1 FY2020? Is the Team's plan and schedule to achieve CD-2 approval reasonable?
  - Are the cost and schedule, including contingencies, reasonable for ECP to meet its mission need and preliminary KPPs, especially given the risk of exascale system availability that is beyond the project's control?
  - Is the project on track to meet its major milestones (Level 1) as identified in the PPEP?
  - Is the project being managed appropriately, including effective and agile leadership from the Project Director and his team and establishing and implementing effective tailoring strategies?

ECP Review Panel Findings

- **Overall the Panel recognized that "A world-class leadership team has built project management processes that provide for nimble product development. The project is properly being managed with effective and agile leadership."**

ECP Milestones

ECP PPEP Milestones and Facility System Milestones

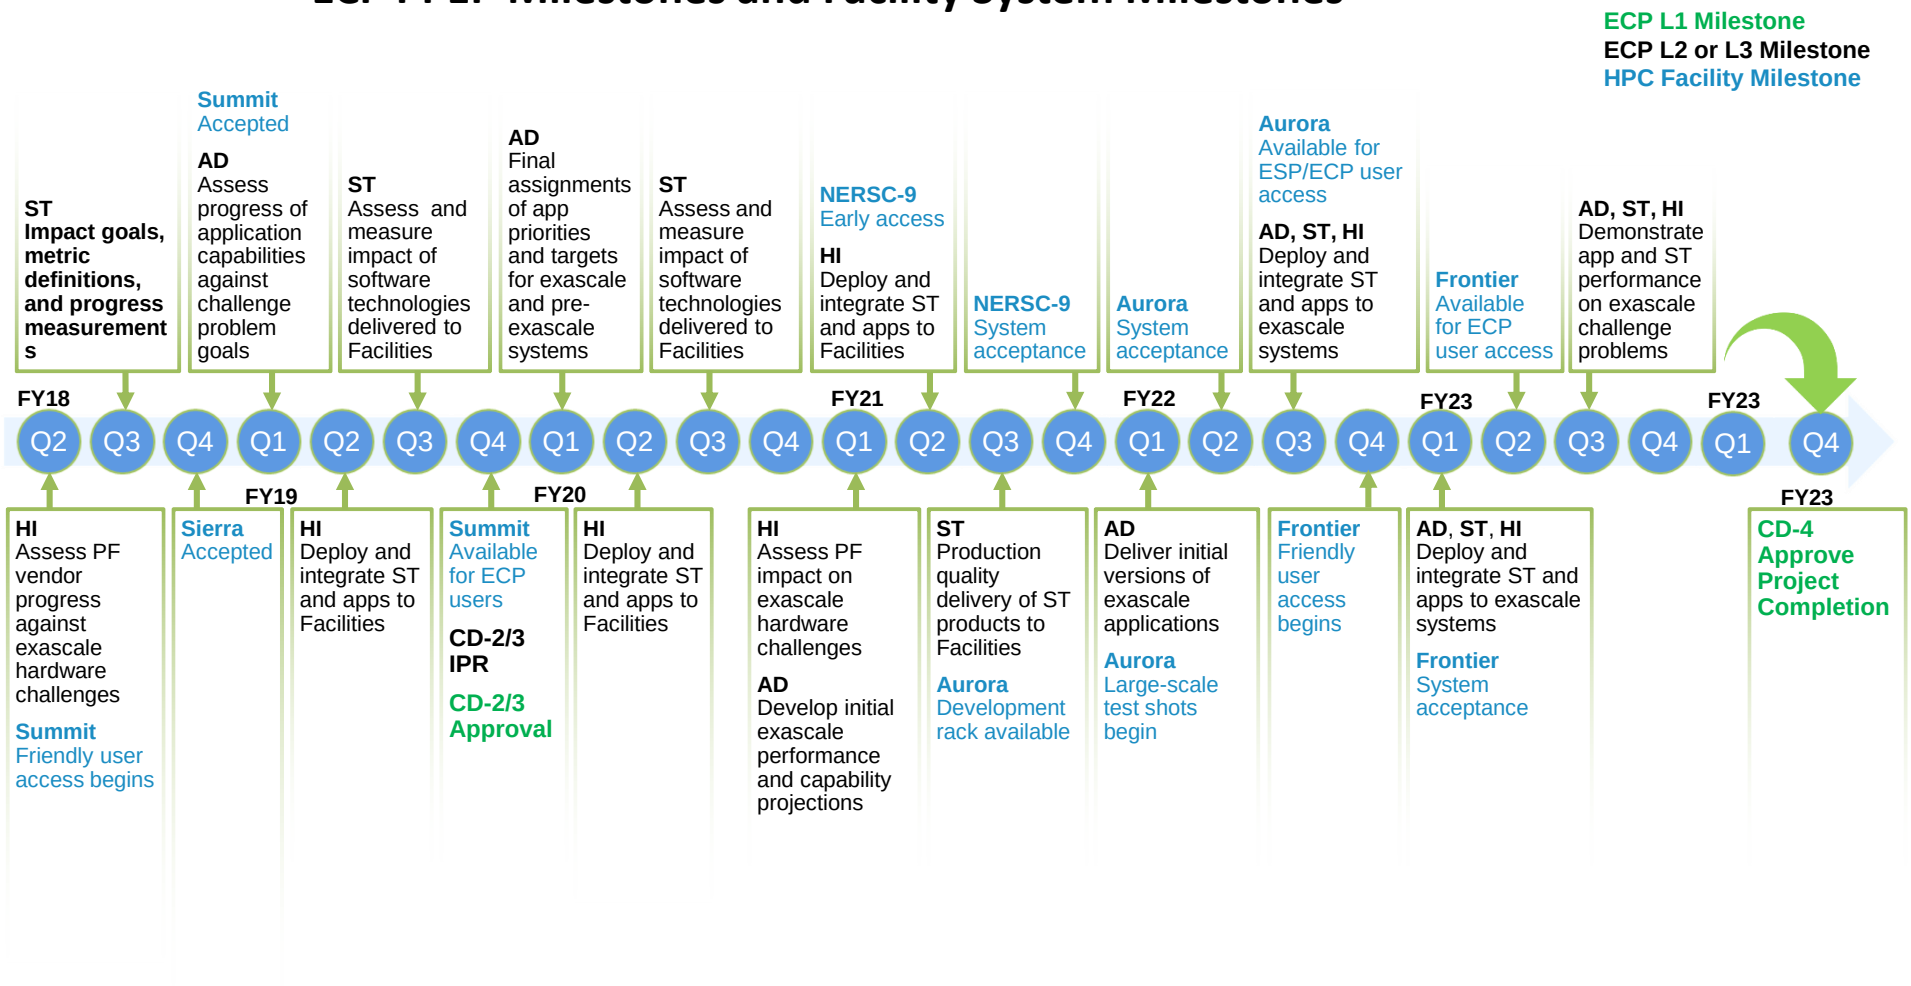

Current Exascale Key Performance Parameters

Performance Measure	Threshold	Objective
Performance on scientific and national security applications relative to today's performance	50% of selected applications achieve Figure of merit improvement ≥ 50	100% of selected applications achieve Figure of merit improvement ≥ 50
Broaden the reach of exascale science and mission capability	50% of selected applications can execute their exascale challenge problem[2]	100% of selected applications can execute their exascale challenge problem
Productive and Sustainable High-Performance Computing (HPC) software ecosystem	SW teams meet 50% of their impact goals[3]	SW teams meet 100% of their impact goals
Enrich the HPC Hardware Ecosystem	Vendors meet 80% of all the PathForward milestones	Vendors meet 100% of all the PathForward milestones

[1] Performance is measured by a Figure of Metric that represents the rate of "science work" defined specific to each scientific application and takes into consideration the increased complexity and precision in addition to the speed of solution.

[2] This KPP assess the successful creation of new exascale science and mission capability. An exascale challenge problem is defined for every scientific application in the project. The challenge problem is reviewed annual to ensure it remains both scientifically impactful to the nation and requires exascale-level resources to execute.

[3] This KPP measures progress on the goal to develop a software ecosystem where high-performance applications can be efficiently and effectively designed, developed, tuned, and executed on exascale systems. Each software effort in the project defines 2-4 impact goals, which must be measurable and provide tangible value to the HPC ecosystem.

ASCR Workshops – January, 2019

Sunday	Monday	Tuesday	Wednesday	Thursday	Friday	Saturday
		1	2	3	4	5
6	7	8	9	10	11	12
13	14	15	16	17	18	19
20	21	22	23	24	25	26
27	28 **In Situ Data Management**	29 **In Situ Data Management** **Applied Math PIs Meeting**	30 **Applied Math PIs Meeting**	31 **SC Quantum Information Science PI Kickoff Meeting**	February 1 **SC Quantum Information Science PI Kickoff Meeting**	

Dan Reed will share information from the October, 2018 ASCR-BES-HEP Microelectronics Workshop at 2:30 today

ASCR Workshop on In Situ Data Management

January 28-29, 2019; Bethesda North Marriott
POC: Laura Biven

Organizing Committee

Name	Affiliation
Tom Peterka (Chair)	**ANL**
Debbie Bard	NERSC
Janine Bennett	SNL
Wes Bethel	LBNL
Ron Oldfield	SNL
Line Pouchard	BNL
Christine Sweeney	LANL
Mathew Wolf	ORNL

~ 60 participants

This workshop seeks community input on the development of in situ capabilities for organizing data and enabling the data flow among a wide variety of coordinated tasks for scientific computing.

The workshop considers ISDM beyond the traditional roles of accelerating simulation I/O and visualizing simulation results, to more broadly support future scientific computing needs. In particular, the convergence of simulation, data analysis, and artificial intelligence will require machine learning, data manipulation, creation of data products, assimilation of experimental and observational data, analysis across ensemble members, and, eventually the incorporation of tasks on non-von Neumann architecture.

ASCR Applied Mathematics Principal Investigators Meeting

SAVE-THE-DATE: January 29 – 30, 2019

LOCATION: Rockville Hilton

Goals

Foster community building & brainstorm on research & workforce issues

- Communicate research challenges & accomplishments
- Forum for PIs to learn about ASCR priorities & research opportunities
- Encourage sharing of best practices in carrying out research & training

Format: TBD & similar to PI meeting in September 2017

- Keynote presentations
- Posters – Each project PI presents at least one poster, also poster blitz
- Whitepapers – Forward-looking ideas & basis for breakout discussions
- Breakout sessions – Focused on current & future research directions
- Working Lunches

DOE Quantum Information Science Kick-Off Meeting

DATE: January 30 – February 1, 2019

LOCATION: Bethesda, Maryland

The Kick-off PI meeting

- is being jointly sponsored by the Offices of Advanced Scientific Computing Research (ASCR), Basic Energy Sciences (BES), and High Energy Physics (HEP), and
- will focus on introducing the principal investigators supported by the Office of Science to each other and identifying opportunities to leverage their research and establish collaborations.
- Will include over 200 PIs, co-PIs and significant collaborators from national laboratories, universities and industry from recent QIS and QIS related awards

Format: TBD

- Plenary presentations
- Panel Discussion with SC leadership
- Poster Session

QAT Projects on Enabling Quantum Computing for Scientific Discovery
Overview & Synergistic Activities

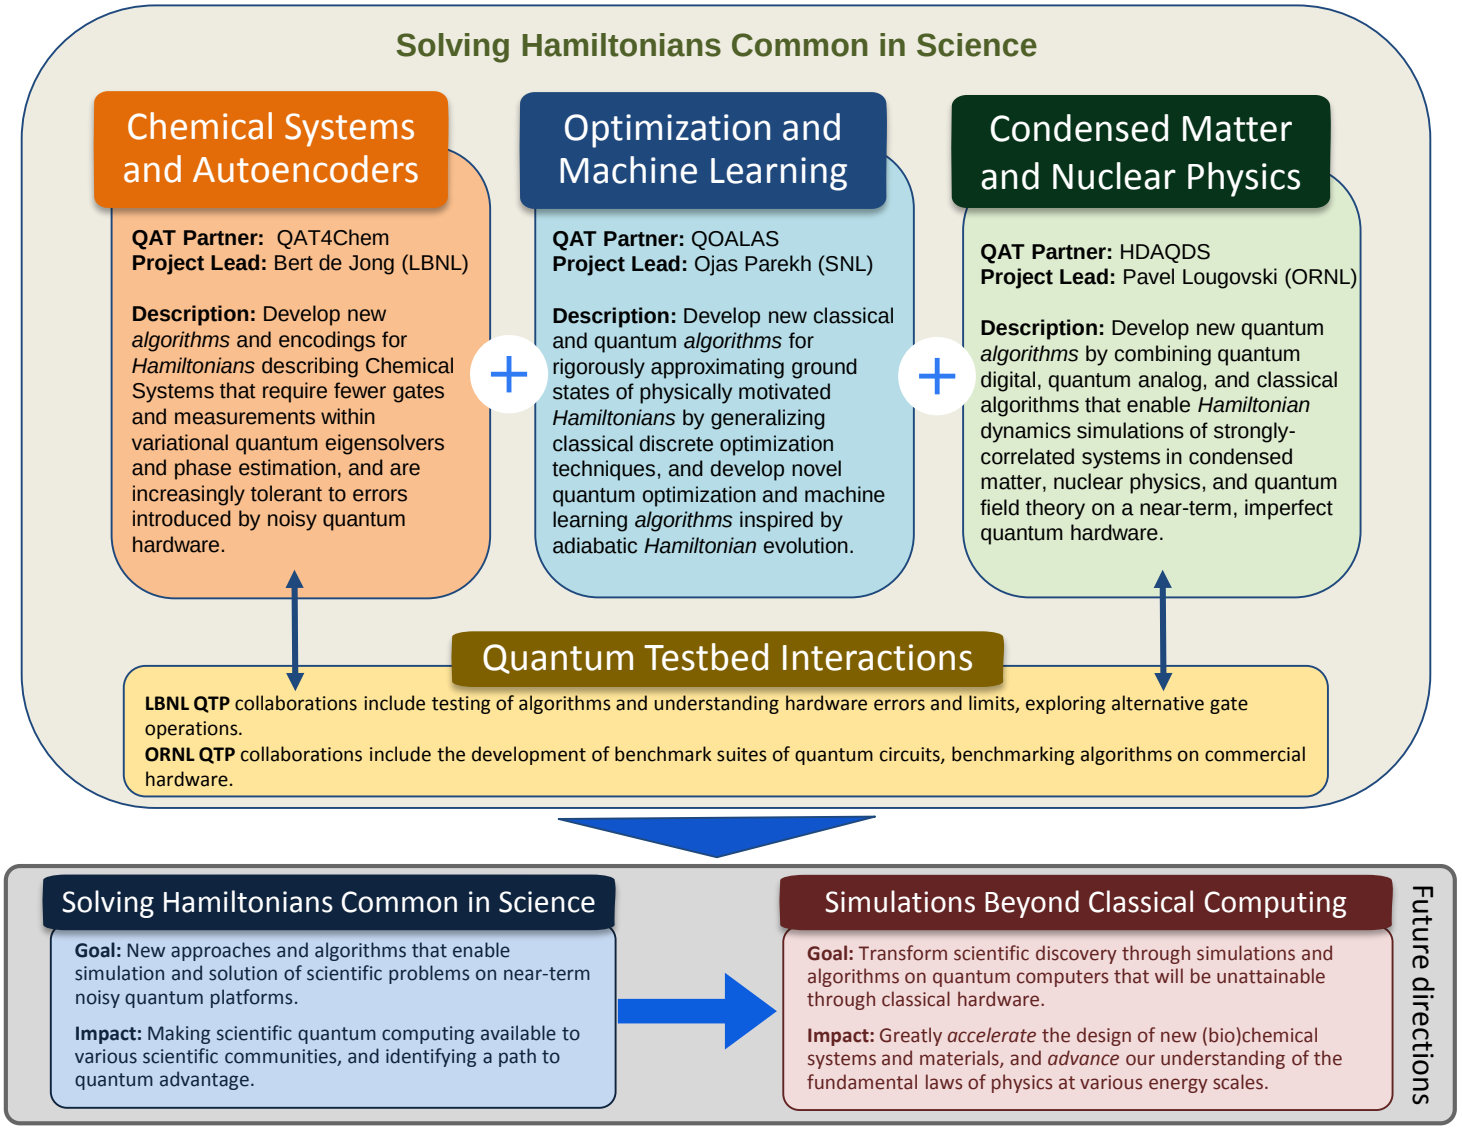

Preview for ASCAC 2019

Sexual Harassment of Women:
Climate, Culture, and Consequences in Academic Sciences, Engineering, and Medicine

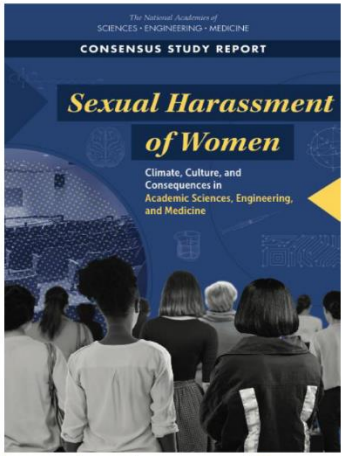

http://nationalacademies.org/SexualHarassment

#ScienceToo

The National Academies of
SCIENCES • ENGINEERING • MEDICINE

Slide from presentation to HEPAP, December 2018

Last, but Certainly not Least

ASCR's Computer Science Program Manager
Dr. Lucy Nowell is retiring December 31, 2018.

We will miss you Lucy and your "hats" will be hard to fill.