U.S. DEPARTMENT OF ENERGY | Office of Science

View from Germantown
Advanced Scientific Computing Research

Presented to the
ASCAC

by

Barbara Helland
Associate Director

January 13, 2020

FY 2020 President's Office of Science Budget

FY 2018 Enacted: $6.260B – ASCR $647.0M

FY 2019 Enacted: $6.585B – ASCR $935.5M

FY 2020 Enacted: $7.000B – ASCR $980.0M

Consolidated Appropriations Act 2020
Science

- The Department is directed to provide to the Committees on Appropriations of both Houses of Congress not later than 90 days after enactment of this Act a plan that responds to the findings and recommendations in the Final Report of the Secretary of Energy Advisory Board Task Force on Biomedical Sciences. The plan shall include a reporting of successful collaborations between the Department and the National Institutes of Health to date and plans to expand on those efforts.

- The agreement provides $71,000,000 for Artificial Intelligence and Machine Learning for the six Office of Science programs to apply those capabilities to the Department's mission.

- The agreement provides $195,000,000 for Quantum Information Sciences across the Office of Science programs to advance early-stage fundamental research in this field of science, including $120,000,000 to carry out a basic research program on quantum information science and $75,000,000 for the establishment of up to five National Quantum Information Science Research Centers. To the greatest extent practical, this effort shall be undertaken in coordination with the National Science Foundation and the National Institute of Standards and Technology.

- The agreement provides not less than $10,000,000 and up to $15,000,000 for research in memory advancements for accelerated architectures used to enhance Artificial Intelligence and Machine Learning. The Department is directed to develop a collaborative research program to produce breakthroughs for intelligent memory systems that will enhance the ability of the Department to cost effectively address the largest problems in science while keeping the United States as the leader in semiconductor technologies for advanced computing.

Consolidated Appropriations Act 2020
ASCR

Advanced Scientific Computing Research (ASCR).-Within available funds, $**150,000,000 is for the Argonne Leadership Computing** Facility, **$225,000,000 is for the Oak Ridge Leadership Computing** Facility, **$110,000,000 is for the National Energy Research Scientific** Computing Center at Lawrence Berkeley National Laboratory, and $90,000,000 is for ESnet. Within available funds, not less than $39,000,000 is for Research and Evaluation Prototypes, of which not less than $10,000,000 is for the Computational Science Graduate Fellowship program. The agreement provides **not less than $155,000,000 for Mathematical, Computational, and Computer Sciences Research.**

ASCR FY 2020 President's Request
in thousands

	FY 2018	FY 2019	FY 2020			
	Enacted Approp.	Enacted Approp.	Request	House Mark	Senate Mark	Enacted Approp.
Mathematical, Computational, and Computer Sciences Research						
Applied Mathematics	34,104	28,206	41,500			41,500
Artificial Intelligence and Big Data (Non Add)	*(…)*	*(…)*	*(14,281)*			*(14,281)*
Computer Science	29,508	22,000	38,700			38,700
Artificial Intelligence and Big Data (Non Add)	*(6.402)*	*(2,000)*	*(9,719)*			*(9,719)*
Quantum Information Science (Non Add)	…	(3.000)	(5,000)			(5,000)
Computational Partnerships	49,910	75,667	60,959			69,142
Artificial Intelligence and Big Data (Non Add)	*(3.500)*	*(13,000)*	*(12,000)*			*(12,000)*
Quantum Information Science (Non Add)	*(6,349)*	*(16,214)*	*(16,708)*			*(20,680)*
SBIR/STTR	4,301	4,768	5,347			5,658
Total, Mathematical, Computational, and Computer Sciences Research	117,823	130,641	146,506	155,000	168,944	155,000
High Performance Computing and Network Facilities						
High Performance Production Computing (NERSC)	94,000	104,000	85,000	100,000	115,000	110,000
Leadership Computing Facility at ANL (ALCF)	110,000	140,000	150,000	150,000	165,000	150,000
Exascale (Non Add)	*(110,000)*	*(140,000)*	*(150,000)*	*(150,000)*	*(150,000)*	*(150,000)*
Leadership Computing Facility at ORNL (OLCF)	162,500	199,000	210,000	225,000	235,000	225,000
Exascale (Non Add)	*(62,500)*	*(100,000)*	*(125,000)*	*(125,000)*	*(125,000)*	*(125,000)*
Total, Leadership Computing Facilities	272,500	339,000	360,000	375,000	400,000	375,000
Research and Evaluation Prototypes	24,260	24,452	39,453	25,620	42,000	39,000
CSGF	*(10,000)*	*(10,000)*	*(10,000)*	*(10,000)*	*(12,000)*	*(10,000)*
Quantum Information Science (Non Add)	*(14,260)*	*(14.452)*	*(29,453)*	*(15,620)*	*(30,000)*	*(29,000)*
High Performance Network Facilities and Testbeds (ESnet)	79,000	84,000	80,000	90,000	90,000	90,000
SBIR/STTR	17,417	20,701	21,194	22,185	24,321	22,265
Total, High Performance Computing and Network Facilities	487,177	572,153	585,647	612,805	671,321	636,265
Exascale Computing						
17-SC-20 Office of Science Exascale Computing Project (SC-ECP)	205,000	232,706	188,735	188,735	188,735	188,735
Total, Advanced Scientific Computing Research	**647,000**	**935,500**	**920,888**	**956,540**	**1,029,000**	**980,000**

OFFICE OF ADVANCED SCIENTIFIC COMPUTING RESEARCH (SC-21) WAS

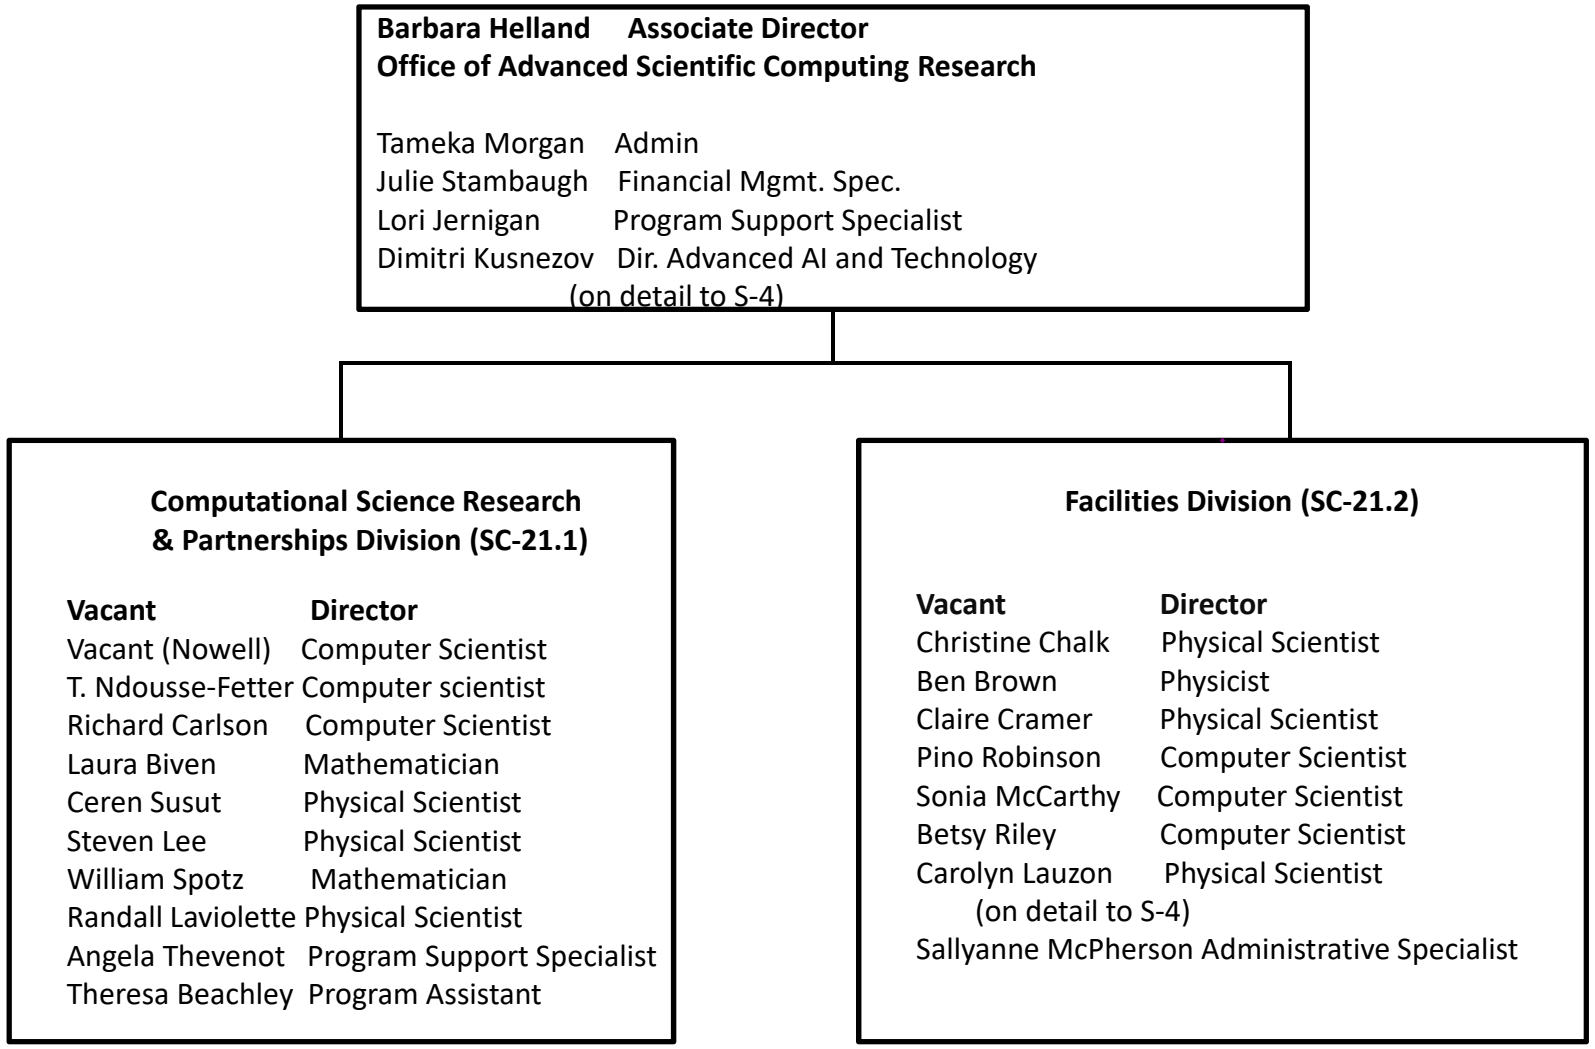

Approved OFFICE OF ADVANCED SCIENTIFIC COMPUTING RESEARCH (SC-31)

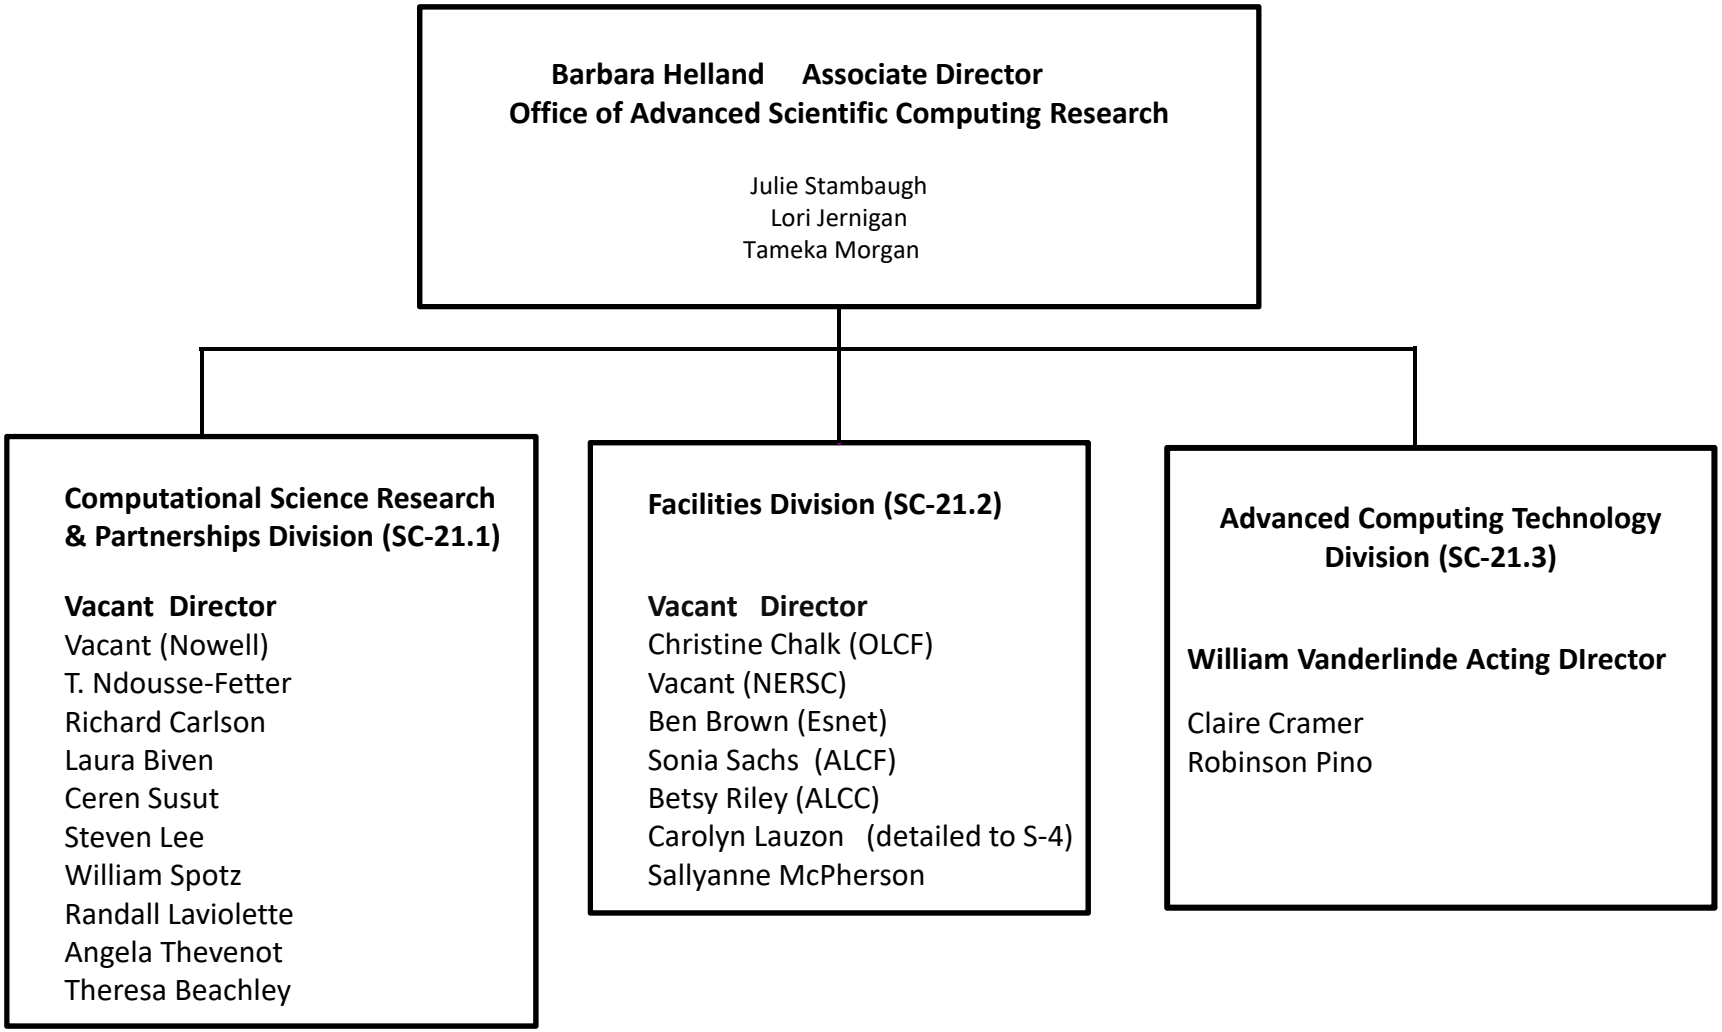

ASCR

AD's view: Co Design

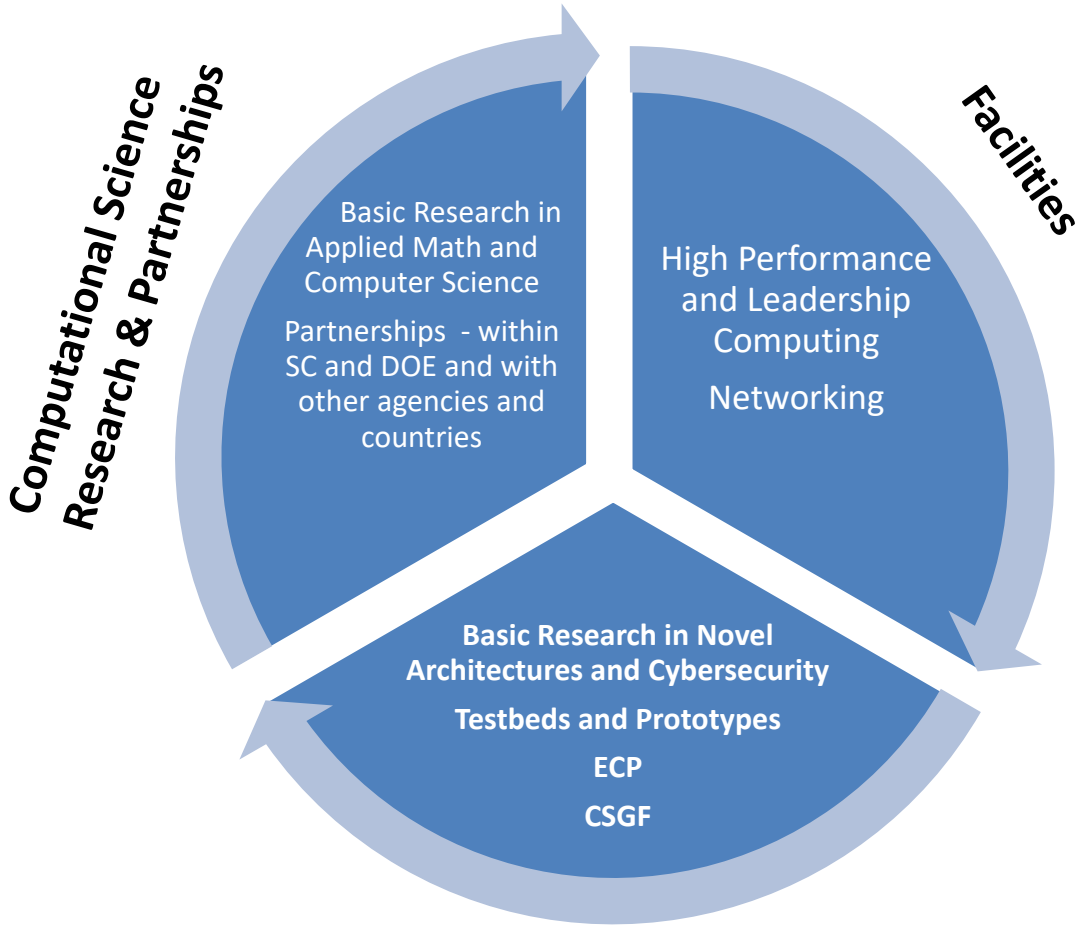

Advanced Computing Technology

Quantum

Quantum Centers Funding Opportunity Announcement

- **Announcement released Friday January 10**
- **Award Size: $10M to $25M per year, 2-5 awards for up to 5 years.**
- **Critical Dates**
  - Preapplications, **required**, Monday, February 10, 2020 at 5 PM Eastern time
  - Pre-application Responses – Tuesday, March 10 2020
  - Full Application Submission – Friday April 10, 2020 at 5PM Eastern time
- **FOA Uniqueness**
  - First large scale effort that crosses the technical breadth of the Office of Science
  - Scope built on extensive community-wide RFI inputs– from technical scope to partnership model to management construct
  - Seamlessly integrates the S&T innovation chain to accelerate progress in QIS R&D
  - Maximizes teaming flexibility and options (TIAs, cooperative agreements, field work authorizations, interagency agreements) to foster direct participation by academics, national/federal labs, and for-profits
  - Leverages other federal agency investments such as NSF's Quantum Leap Challenge Institutes and the NIST Quantum Economic Development Consortium

https://science.osti.gov/grants/FOAs/Open

Artificial Intelligence for Science

- New Projects
- AI for Science Town Halls
- AI charge to ASCAC

Description:

- Lab only call, focus on co-design of learning systems and AI environments that significantly advance the field of AI for public benefit within DOE's Congressionally-authorized mission-space.
- Proposals should present and articulate how implementations and development of AI and machine learning will advance and integrate within emerging computing hardware, where considerations such as power consumption, cost of data movement across the memory hierarchy, and resilience are important to address.

Results:

- 23 letters of intent, 20 full proposals, Evaluated by mail reviews
- 2 Project Awards fully-funded for 3 years. Total amount of awards: $11.1M.

CO-DESIGN OF ADVANCED ARTIFICIAL INTELLIGENCE (AI) SYSTEMS

Predicting behavior of complex systems using multimodal datasets

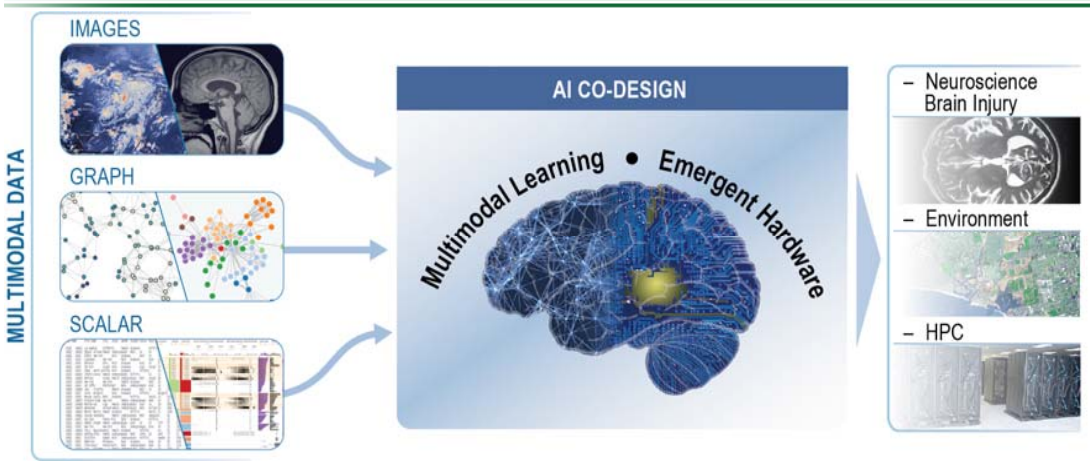

Domain-agnostic framework and tools to advance complex multimodal learning for scientific discovery

Goal: Advancing scientific learning from aggregated multimodal, heterogeneous, variably sized, and unstructured datasets. Improving performance of scientific AI models on emerging hardware

Three scientific use cases: neuroscience, environmental science, and high-performance computing (HPC) operations

Thrust 1: Registration, Segmentation, and Learning from multiple image modalities

Thrust 2: Learning from variably sized graphs containing complex, multimodal attributes

Thrust 3: Learning from mixed-scalar (real-valued and categorical), messy data

Thrust 4: Constructing fusion and co-learning methods that maximizes information in each modality

Thrust 5: Performance (runtime, memory, and storage) of AI/ML algorithms on next-generation computing hardware systems

PI: Jim Brase, LLNL

Argonne NATIONAL LABORATORY · BERKELEY LAB · Lawrence Livermore National Laboratory · UCSF University of California San Francisco

U.S. DEPARTMENT OF ENERGY | Office of Science

Center for co-design of ARtificial Intelligence focused Architectures and Algorithms (ARIAA)

ARIAA is a joint co-design research center that includes PNNL, SNL, and GT, supported by NVIDIA and Qualcomm

Motivation: Codesign of AI/ML accelerators with algorithms and applications will enable the development of AI technologies to suit DOE AI/ML needs

Objectives:

- Co-design novel architectures, algorithms, and programming abstractions to enable AI-based DOE applications
- Understand how AI-focused architectures can accelerate both traditional and AI DOE workloads

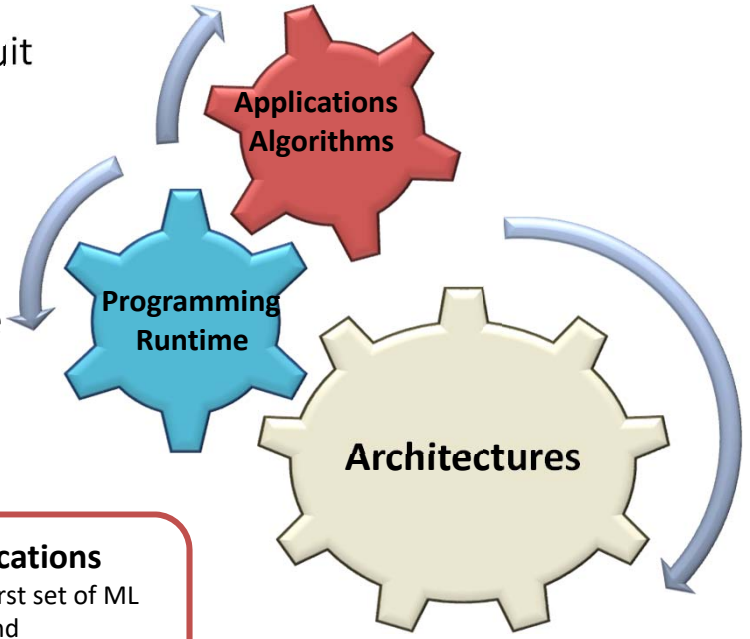

Ongoing Activities

Architecture
- ✓ SST/MAESTRO integration strategy
- ✓ SST accelerator abstractions
- ✓ Initial evaluation of NVIDIA DNN accelerator

Software
- ✓ Initial integration of MCL with SST and NVIDIA DNN accelerator
- ✓ Programming abstractions for representative kernels

Applications
- ✓ Identify first set of ML kernels and implementations
- ✓ Graph clustering, random forest, all-to-all hashing (streaming), SchNet

PI: Roberto Gioiosa, PNNL

Pacific Northwest NATIONAL LABORATORY | Sandia National Laboratories | Georgia Tech

SciML Priority Research Needs January 2018

SciML Foundational Research Themes

Domain-aware – leverages and respects scientific domain knowledge – physics principles and symmetries, constraints, uncertainties and structure-exploiting models

Interpretable-- explainable and understandable results-- model selection, exploiting structure in high-dimensional data, use of uncertainty quantification with machine learning.

Robust – stable, well-posed & reliable formulations – probabilistic modeling in ML, quantifying well-posedness, reliable hyperparameter estimation

SciML Capabilities

Data-Intensive – scientific inference and data analysis – ML methods for multimodal data, in situ data analysis and to optimally guide data acquisition

Adaptive Simulations – ML hybrid algorithms and models for better scientific computing tools– ML-enabled adaptive algorithms, parameter tuning and multiscale surrogate models

Controlling Complex Systems – automated decision support, adaptivity, resilience, control – exploration of decision space with ML, ML-based resource management and control, optimal decisions for complex systems

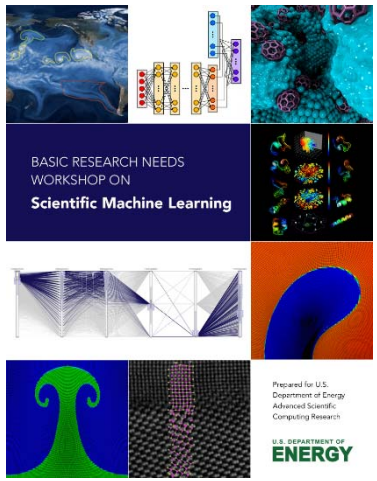

Purpose: Define priority research directions for applied mathematics in scientific machine learning (ML). Identify the challenges and opportunities for increasing the rigor, robustness, and reliability of ML for DOE missions.

Applied Math program has laid the groundwork to harness Machine Learning and Artificial Intelligence for scientific purposes

FY19 Funding Opportunity Announcement
Scientific Machine Learning and Artificial Intelligence: Uncertainty Quantification

Description:

- Foundational research to strengthen the mathematical & statistical basis of validating machine learning & AI predictions from data generated by the Office of Science's user facilities and scientific simulations.
- Nearly 120 Letters of Intent, Over 90 Full Proposals, Evaluated by mail-in reviews.
- 3 Project Awards fully-funded for 2 years. Total amount of awards: $1,900K.

PI Name	Institution	Proposal Title	Objectives
Gunnan Zhang	Oak Ridge Lab	A Stochastic Optimal Control Framework for Quantifying and Reducing Uncertainty in Deep Learning	Quantifying & reducing uncertainties in deep learning by exploiting ties between probabilistic network architectures & optimal control of stochastic dynamical systems.
Jayaraman J. Thiagarajan	Lawrence Livermore Lab	Uncertainty Quantification for Scientific Machine Learning	UQ for deep learning & identifying: Regimes of failure, Confidence intervals, and Out of distribution samples.
Alexander Gorodetsky	University of Michigan	Bayesian Tensor Decompositions for Scalable Supervised Learning of Scientific Data	Embed UQ for improved reliability in supervised learning based on low-rank tensor decompositions.

Report from the Office of Science Data for AI Roundtable

Fagnan, Kjiersten, Nashed, Youssef, Perdue, Gabriel, Ratner, Daniel, Shankar, Arjun, and Yoo, Shinjae. *Data and Models: A Framework for Advancing AI in Science*. United States: N. p., 2019. Web. doi: 10.2172/1579323

ASCR POC: Laura Biven

https://www.osti.gov/biblio/1579323

AI for Science Town Halls: Learning from Exascale

- **Modeled after Exascale Town Halls in 2007**
- **Collect community input on the opportunities and challenges facing the scientific community in the era of convergence of High Performance Computing (HPC) and artificial intelligence (AI) technologies and data**
- **Engage the DOE science community in a series of broad and open discussions about**
  - opportunities that can be realized by advancing and accelerating the development of AI capabilities specifically for science and science use cases,
  - opportunities from the DOE Office of Science as well as selected topics from energy and technology domains and will include approaches combining experiments, traditional modeling and simulation, and machine/deep learning
  - opportunities that include the DOE science user facilities.

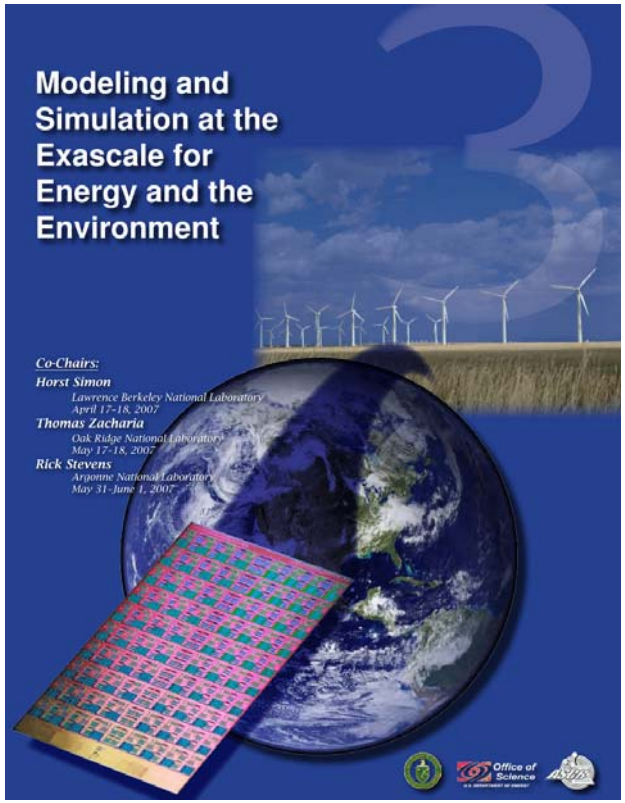

The AI for Science Town Halls statistics

- **Over 1000 registrations across 4 Town Halls**

ANL	July 22-23	357
ORNL	Aug 20-21	330
LBNL	Sept 11-12	349
DC	Oct 22-23	273
Totals	1309	1309

+100 online

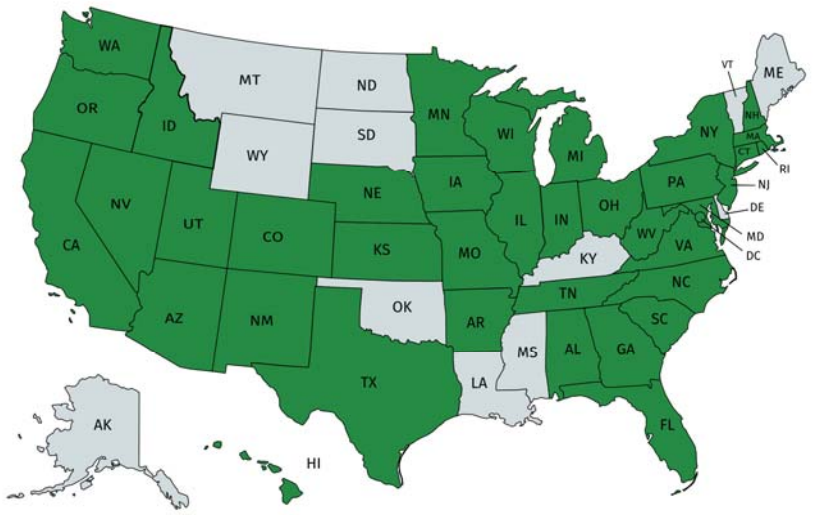

- **All 17 DOE National Laboratories**
- **39 Companies from large and small**
- **Over 90 different universities**
- **6 DOE/SC Offices + EERE and NNSA**

AI for Science Charge

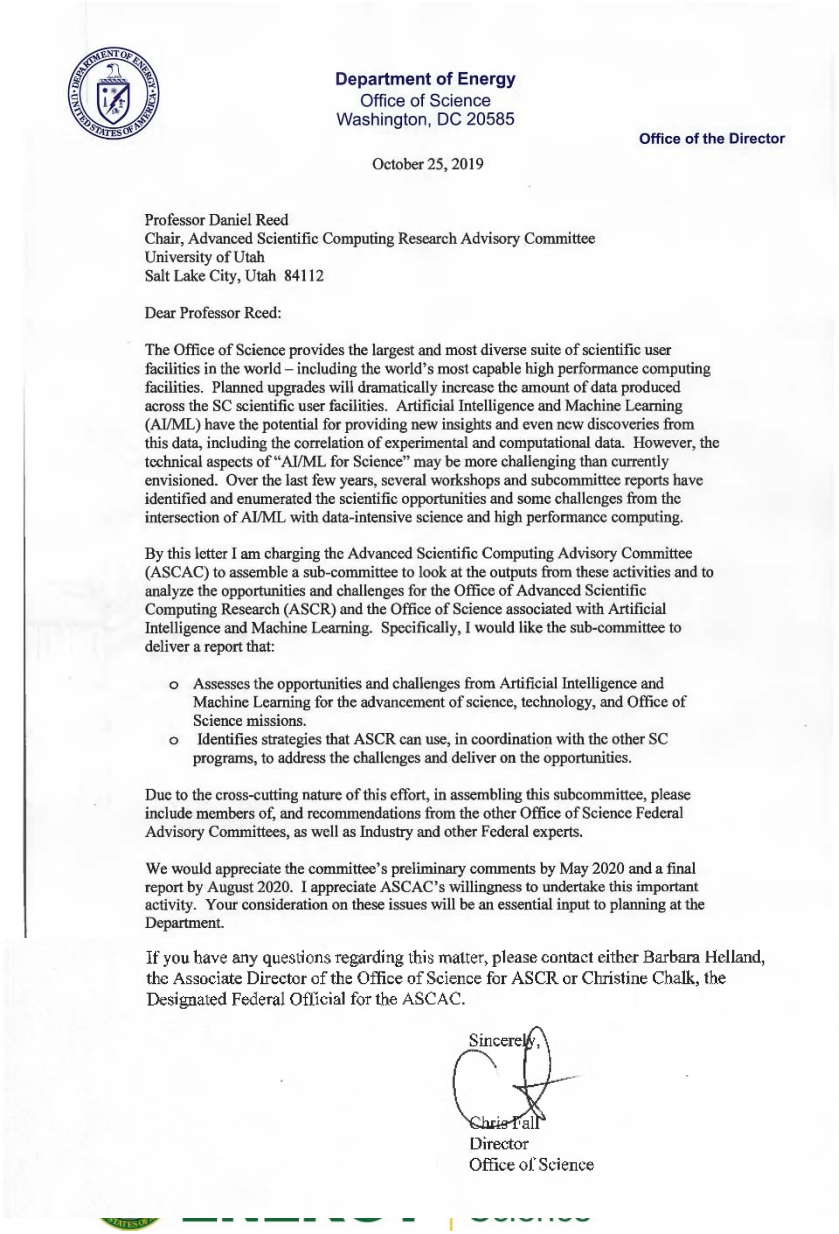

Department of Energy
Office of Science
Washington, DC 20585

Office of the Director

October 25, 2019

Professor Daniel Reed
Chair, Advanced Scientific Computing Research Advisory Committee
University of Utah
Salt Lake City, Utah 84112

Dear Professor Reed:

The Office of Science provides the largest and most diverse suite of scientific user facilities in the world – including the world's most capable high performance computing facilities. Planned upgrades will dramatically increase the amount of data produced across the SC scientific user facilities. Artificial Intelligence and Machine Learning (AI/ML) have the potential for providing new insights and even new discoveries from this data, including the correlation of experimental and computational data. However, the technical aspects of "AI/ML for Science" may be more challenging than currently envisioned. Over the last few years, several workshops and subcommittee reports have identified and enumerated the scientific opportunities and some challenges from the intersection of AI/ML with data-intensive science and high performance computing.

By this letter I am charging the Advanced Scientific Computing Advisory Committee (ASCAC) to assemble a sub-committee to look at the outputs from these activities and to analyze the opportunities and challenges for the Office of Advanced Scientific Computing Research (ASCR) and the Office of Science associated with Artificial Intelligence and Machine Learning. Specifically, I would like the sub-committee to deliver a report that:

- Assesses the opportunities and challenges from Artificial Intelligence and Machine Learning for the advancement of science, technology, and Office of Science missions.
- Identifies strategies that ASCR can use, in coordination with the other SC programs, to address the challenges and deliver on the opportunities.

Due to the cross-cutting nature of this effort, in assembling this subcommittee, please include members of, and recommendations from the other Office of Science Federal Advisory Committees, as well as Industry and other Federal experts.

We would appreciate the committee's preliminary comments by May 2020 and a final report by August 2020. I appreciate ASCAC's willingness to undertake this important activity. Your consideration on these issues will be an essential input to planning at the Department.

If you have any questions regarding this matter, please contact either Barbara Helland, the Associate Director of the Office of Science for ASCR or Christine Chalk, the Designated Federal Official for the ASCAC.

Sincerely,

Chris Fall
Director
Office of Science

Over the last few years, several workshops and subcommittee reports have identified and enumerated the scientific opportunities and some challenges from the intersection of AI/ML with data-intensive science and high performance computing.

By this letter I am charging the Advanced Scientific Computing Advisory Committee (ASCAC) to assemble a sub-committee to look at the outputs from these activities and to look at the outputs from these activities and to analyze the opportunities and challenges for the Office of Advanced Scientific Computing Research (ASCR) and the Office of Science associated with Artificial Intelligence and Machine Learning. Specifically, I would like the sub-committee to deliver a report that:

- Assesses the opportunities and challenges from Artificial Intelligence and Machine Learning for the advancement of science, technology, and Office of Science missions.
- Identifies strategies that ASCR can use, in coordination with the other SC programs, to address the challenges and deliver on the opportunities.

Due to the cross-cutting nature of this effort, in assembling this subcommittee, please include members of, and recommendations from the other Office of Science Federal Advisory Committees, as well as Industry and other Federal experts.

Key Independent Project Reviews

Exascale Computing Project : CD2/CD3

- **On December 3-5, 2019, the Exascale Computing Project (ECP) underwent an Independent Project Review (IPR) to determine its readiness to seek approval of its project baseline and in ECP's case to move from its R&D phase into final development and testing phase(Critical Decision (CD) 2 and Critical Decision 3). The IPR resulted in only three formal recommendations from the review committee:**
  - Proceed to CD-2/3.
  - Experience with the customized solution Dashboard, including design philosophies, successes, failures and challenges should be captured and archived as a best practice in a form that can be shared with other communities facing similar challenges. Archiving should begin as soon as possible, but before lessons learned begin to fade. Publication should coincide with CD-4.
  - Prior to turning over applications to their science sponsors, in the CD-4 timeframe, the ECP should publish a Community Outreach Document for each application to disseminate the knowledge and experience gained in their development and document their maturity for supporting scientific discovery.
- **The Exascale Computing Project (ECP) has a proposed Total Project Cost (TPC) of $1,812.3 million. The project is jointly funded by SC ($1,326.2 million) and NNSA ($486.1 million). The proposed CD-4 date is September 2024. The project has a cost contingency of $147.6 million (23% of TEC) for SC scope. NNSA cost contingency is controlled at the Program level.**
- **Comments from the review committee:**
  - *Since being recognized as world-class in the October 2018 IPR, the Project Management team has further strengthened. Communication within the project and with external partners is now well defined and appears to be effective.*
  - *Coordination and integration among the Hardware Integration, ST, and Application Development teams has made notable progress since October 2018 and the management team demonstrated support for continuing a culture of team collaboration. Overall, the ECP management is appropriately structured and empowered to ensure the project's success.*

Esnet – CD2/CD3

- **On December 10-12, 2019, ESnet6 underwent an Independent Project Review (IPR) to determine its readiness to seek approval of its project baseline and begin implementation (Critical Decision (CD) 2 and Critical Decision 3). The IPR went very well and resulted in only two formal recommendations from the review committee:**
  - 1) to review the project documentation to eliminate inconsistencies, and
  - 2) to proceed to CD-2/3. In effect, the reviewers identified no impediments to seeking CD-2/3 approval in early 2020.
- **The project's proposed Total Project Cost is $140.4M and a targeted early finish date of January 2023.**
- **Quotes from the closeout report:**
  - **"The project has made tremendous progresses since [the] last [OPA review]: organization, cost & schedule have been finalized."**
  - **"ESnet6 has made excellent progress in incorporating significant number of geographically distributed, additional staff."**
  - **"The Project Management team should be commended for their efforts in addressing prior review recommendations and implementing a robust performance management system."**

MIRA

Department of Energy (DOE) Roadmap to Exascale Systems

An impressive, productive lineup of *accelerated node* systems supporting DOE's mission

Pre-Exascale Systems [Aggregate Linpack (Rmax) = 323 PF!]

First U.S. Exascale Systems

2012	2016	2018	2020	2021-2023

ORNL
Cray/AMD/NVIDIA

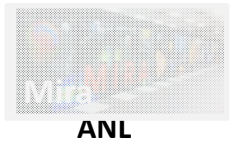

ANL
IBM BG/Q

ANL
Cray/Intel KNL

LBNL
Cray/Intel Xeon/KNL

LLNL
IBM BG/Q

LANL/SNL
Cray/Intel Xeon/KNL

ORNL
IBM/NVIDIA

LLNL
IBM/NVIDIA

LBNL
Cray/AMD/NVIDIA

LANL/SNL
TBD

ORNL
Cray/AMD/AMD

ANL
Intel/Cray

LLNL
Cray

Mira retires after nearly 40 billion core-hours of service

- The ALCF's 10-petaflop IBM Blue Gene/Q, Mira, retired on December 31, 2019 after many productive years of enabling breakthroughs in science and engineering.
- In 7-plus years, Mira delivered 39.6 billion core-hours to more than 800 projects, and supported over 700 full-machine runs.

- When launched, Mira ranked third fastest in the world for speed (Top500) and first in energy efficiency (Green500).
- At retirement, Mira was ranked third on the Graph500, a measure focused on a supercomputer's ability to handle data-intensive applications.

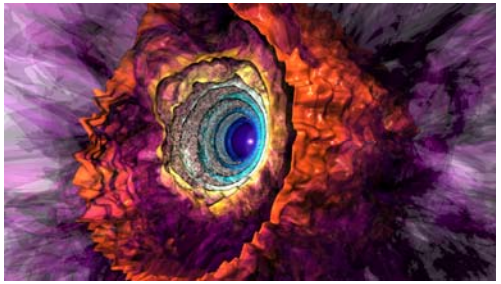

- Directly cooling the chips using pipes carrying water through the node boards (vs. blowing air over the chips) was one of the advances that made Mira so effective and energy efficient.
- Mira helped advance science in many different disciplines and spanned the research of very basic science (origin of life and the evolution of the Universe) to engineering applications (new materials and improving fuel efficiency).

A lasting impact on science…

When Mira came online in 2012, it was 20 times more powerful than Argonne's previous machine Intrepid, giving researchers a tool to perform simulations on unprecedented scales and create more accurate models of everything from combustion engines to blood flow, solving nearly intractable problems in scientific fields ranging from pharmacology to astrophysics.

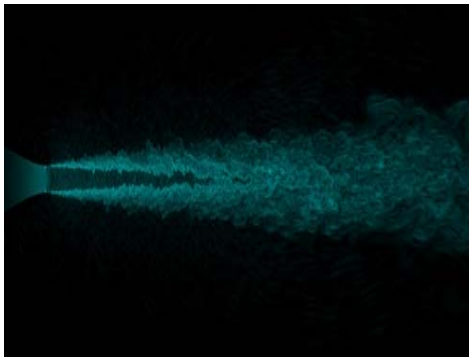

Scientists from GE Global Research used Mira to study the complex behavior of air as it passes through jet exhaust nozzles and over wind turbine blades to inform the design of quieter, **more efficient wind turbines and jet engines**.

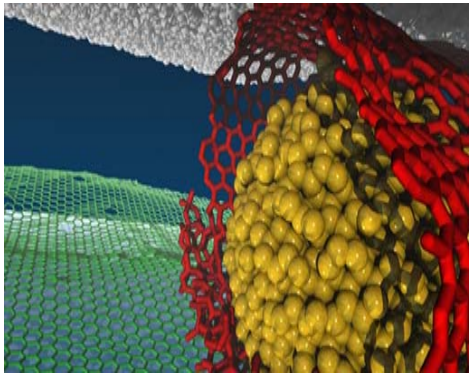

A multidisciplinary team from Argonne used Mira to identify and improve a **new mechanism for eliminating friction**. Their simulations led to the synthesis of a novel material that demonstrated superlubricity (near zero friction) at the macroscale for the first time.

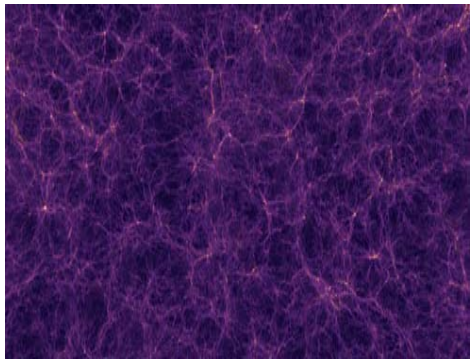

An Argonne-led team used Mira to run some of the **largest and most detailed simulations** ever performed on the **evolution of the universe,** resulting in the first simulations accurate enough to compare with large-scale cosmological surveys.

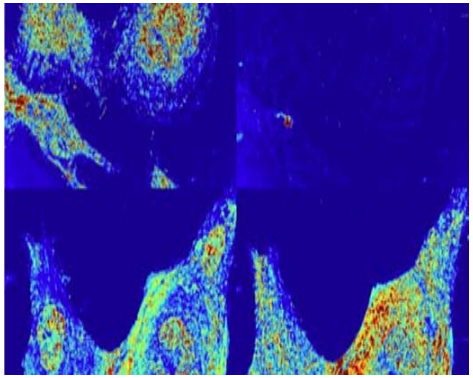

Researchers from Northwestern University tapped Mira to advance the development of an optical microscopy technique that can **predict and quantify cancer risks** at extremely early stages.

Some ASCAC Agenda Details

- **REPORT FROM SUBCOMMITTEE ON EXASCALE TRANSITION --** *Roscoe Giles, Boston University*
- **REPORT FROM SUBCOMMITTEE ON 40th ANNIVERSARY ACCOMPLISHMENTS --** *Bruce Hendrickson, Lawrence Livermore National Laboratory*
- **QUANTUM USER PROGRAM AT THE OLCF --** *Travis Humble, Oak Ridge National Laboratory*
- **DISTINGUISHED SCIENTIST –** *Ian Foster, Argonne National Laboratory*
- **DEPLOYING 5G WIRELESS ACROSS DOE --** *Robinson Pino, ASCR*
- **ARTIFICIAL INTELLIGENCE TECHNOLOGY OFFICE –** *Fred Streitz, AITO*
- **UPDATE ON THE EXASCALE COMPUTING PROJECT –** *Doug Kothe, ORNL and Lori Diachin, LLNL*
- **QUANTUM COMPUTING RESEARCH TESTBEDS –** *Claire Cramer, ASCR*
- **UPDATE ON CURRENT CHARGES**
  - *Artificial Intelligence, Tony Hey ASCAC*
- **ENGAGEMENT WITH EXPERIMENTAL FACILITIES –** *Tom Uram, Argonne National Laboratoroy*
- **FEDERATED IDENTITY MANAGEMENT –** *Rich Carlson, ASCR*

SC's Distinguished Scientist Fellows

- **Newly established award was authorized by the America COMPETES Act**
- **Bestowed on National Laboratory scientists with outstanding records of achievement,**
- **Provides each Fellow with $1 million over three years to be devoted to a project or projects of the Fellow's choosing**
- **Nominated by their individual labs and chosen through a peer review process**
- **Five awardees named on October 68, 2019**
  - Sally Dawson, Brookhaven National Laboratory (HEP)
  - Ian Foster, Argonne National Laboratory (ASCR)
  - Joshua Frieman, Fermi National Accelerator Laboratory (HEP)
  - Barbara Jacak, Lawrence Berkeley National Laboratory (NP)
  - José Rodiguez, Brookhaven National Laboratory (BES)

Ian Foster, SC Distinguished Scientist Fellow

Recognized for his "for trailblazing work in distributed and high performance computing with fundamental and long-lasting impacts on both computer science as a discipline and the practice of computing across the Office of Science."