



U.S. DEPARTMENT OF
ENERGY

Office of
Science

View from Germantown

Advanced Scientific Computing Research Update

Presented to
ASCAC

by

Barbara Helland, Associate Director

July 29, 2021

Personnel Changes

ASCR's New Facilities Division Director



Ben Brown

On October 11, 2020, Dr. Benjamin Brown became the permanent ASCR Facilities Division Director. Ben joined the Office of Science in 2008 and has over 15 years of experience in the federal government. Since 2017 he has served as Program Manager for ESnet, including the ESnet6 upgrade project. He previously led the design and implementation of the DOE Project Leadership Institute and the Office of Science User Facilities policy and user statistics framework.

Before coming to DOE, Ben worked in the U.S. Senate as a AAAS Congressional Fellow. Ben is an experimental physicist with over a decade of experience conducting research in ultracold atomic physics, optical control of quantum systems, and quantum information science. He received his Ph.D. in optics from the University of Rochester and conducted his doctoral and postdoctoral research at the University of Oxford and NIST, respectively.

New Computer Science Program Manager



Margaret Lentz

Dr. Margaret Lentz joined ASCR as a Program Manager in April 2021. Her previous work in government policy includes coordinating AI efforts across DOE with the Artificial Intelligence & Technology Office, as well as aiding in strategic planning, programming, and budget formulation in NNSA's Defense Programs. Prior to working in policy, Dr. Lentz served as a research scientist specializing in medical imaging, neuroscience, biodefense, and image analysis methods including machine learning and neural networks. She has established and supported imaging research programs in academic and federal research settings, often pertaining to the testing of medical countermeasures for viruses such as HIV, Ebola, MERS and Nipah

Margaret received her PhD in Physical Organic Chemistry from Purdue University and her B.S in Chemistry from Carnegie-Mellon.

ASCR's New Research Division Director



Ceren Susut

On May 23, 2021, Ceren Susut became the permanent ASCR Research Division Director. She joined Office of Science ASCR as a program manager in January 2011. In ASCR, she managed Scientific Discovery through Advanced Computing (SciDAC), and quantum algorithms and quantum computer science portfolios as well as the Exascale Computing Project (ECP). Ceren is also serves as the lead for the SC QIS working group and led the SC-wide National QIS Research Centers effort.

She holds a PhD in chemistry from Georgetown University. Prior to joining DOE, she was a NRC research associate at NIST. In 2015, she was a detailee in the Advanced Cyberinfrastructure Division at NSF.

Sonia Sachs -- Retiring

Sonia Sachs has been a Program Manager for ASCR since March 1, 2010. She will retire on July 31, 2021.



Sonia has a Ph.D. in Electrical Engineering and Computer Science from the University of California at Berkeley. In the last forty years her career included software engineering, consulting, research, research management, SBIR/STTR management, and HPC facilities management. Before joining ASCR, she was at IBM Research, LBNL, UC Berkeley (as a visiting scholar and research manager), her own start-up company, Bell Labs and Apple. More about Sonia is at <https://www.linkedin.com/in/soniasachs>

At ASCR, she managed computer science portfolio in the areas of exascale programming models, programming languages, compilers, runtime systems, operating systems, performance tools, productivity tools, and hardware architectures. For several years, she co-managed the Early Career program in Computer Science, with several of her ECRP PIs receiving PECASE awards. In the last five years she has been managing the Argonne Leadership Computing Facility (ALCF). She has managed SBIR/STTR topics during her ASCR tenure. Since 2019, she has been the SBIR/STTR program portfolio manager. At the same time, she managed the 2019-2020 ALCC program. In early 2021, she was selected co-chair of the NITRD High End Computing (HEC) Inter-agency Working Group.

Her new pursuits in the private sector will leverage her background and experience, particularly the HPC expertise developed at ASCR.

Budget Review

ASCR FY2021 Appropriations

- ▶ The agreement supports ASCR's leadership in emerging areas relevant to the Department's mission, including artificial intelligence and quantum information science. The agreement supports ASCR's pursuit of machine learning tools for scientific applications and its support for the development of algorithms for future deployable quantum computers. The agreement recognizes that a robust research program in applied and computational mathematics and computer science will be critical to continued progress in these areas and is supportive of the Department's efforts to prioritize these programs.
- ▶ High Performance Computing and Network Facilities.-The agreement provides **not less than \$150,000,000** for the Argonne Leadership Computing Facility, **not less than \$225,000,000** for the Oak Ridge Leadership Computing Facility, **not less than \$110,000,000** for the National Energy Research Scientific Computing Center, and **\$90,000,000** for ESnet.
- ▶ Mathematical, Computational, and Computer Sciences Research.-The agreement provides **not less than \$250,000,000** for Mathematical, Computational, and Computer Sciences Research, including **not less than \$10,000,000** for the Computational Science Graduate Fellowship program. The agreement provides **up to \$40,000,000** for the development of AI-optimized emerging memory technology for AI-specialized hardware to drive national competitiveness.

ASCR Enacted Appropriations

in thousands

	FY2020	FY 2021		FY2022
	Enacted Approp.	Request	Enacted Approp	Request
Mathematical, Computational, and Computer Sciences Research				
Applied Mathematics	41,500	53,728	48,570	51,048
Artificial Intelligence and Big Data (Non Add)	(14,281)	(23,473)	(24,330)	(24,330)
Computer Science	38,700	49,605	46,827	49,773
Artificial Intelligence and Big Data (Non Add)	(9,719)	(14,875)	(14,915)	(16,454)
Quantum Information Science (Non Add)	(5,000)	(7,000)	(7,256)	(7,256)
Computational Partnerships	69,142	75,051	76,194	86,029
Artificial Intelligence and Big Data (Non Add)	(12,000)	(17,652)	(17,621)	(18,036)
Quantum Information Science (Non Add)	(20,680)	(18,709)	(19,209)	(19,209)
Research and Evaluation Prototypes	39,000	76,007	88,274	106,112
CSGF	(10,000)	(10,000)	(10,000)	(15,000)
Quantum Information Science (Non Add)	(29,000)	(60,453)	(71,937)	(81,001)
SBIR/STTR	5,658	9,637		
Total, Mathematical, Computational, and Computer Sciences Research	155,000	264,028	259,865¹	292,962¹
High Performance Computing and Network Facilities				
High Performance Production Computing (NERSC)	110,000	85,000	113,786	115,963
Leadership Computing Facility at ANL (ALCF)	150,000	150,000	152,955	159,047
Exascale (Non Add)	(150,000)	(150,000)	(150,000)	(150,000)
Leadership Computing Facility at ORNL (OLCF)	225,000	220,000	228,120	249,066
Exascale (Non Add)	(125,000)	(125,000)	(120,000)	(125,000)
Total, Leadership Computing Facilities	375,000	370,000	381,075	381,075
High Performance Network Facilities and Testbeds (ESnet)	90,000	90,000	91,329	93,961
SBIR/STTR	22,265	20,078		
Total, High Performance Computing and Network Facilities	636,265*	555,078	586,190¹	618,037¹
Exascale Computing				
17-SC-20 Office of Science Exascale Computing Project (SC-ECP)	188,735	168,945	168,945	129,000
Total, Advanced Scientific Computing Research	980,000	988,051	1,015,000	1,040,000

¹ SBIR/STTR contained in each activity

*Includes R & E Prototype funds;

FY2022 Initiatives: Biopreparedness Research Virtual Environment (BRaVE)


- ▶ A virtual platform to rapidly mobilize DOE's bioscience R&D assets in response to national crises
- ▶ Provides distributed DOE laboratory research teams with collaborative frameworks and rapid access to:
 - DOE's High-Performance Computing resources and x-ray and neutron characterization facilities
 - Collaborative Design-Build-Test-Learn workflows that leverage DOE's databases of biological, chemical and materials datasets
 - Medical radioisotopes that may prove crucial to biological emergency responses
- ▶ Supports development of DOE's next-generation analytical technologies, user facility resources, national preparedness for medical isotope processing, and biotechnological capabilities



^1H MRI

^3He MRI

^{129}Xe MRI



FY2022 Initiatives: Reaching a New Energy Sciences Workforce (RENEW)



- Outreach
- Listening tours and round tables to:
- Gain understanding about challenges
- Develop evidence-based solutions



- Identify unique Office of Science Lab opportunities
- Partner with Minority Serving Institutions and professional societies
- Implement action plan

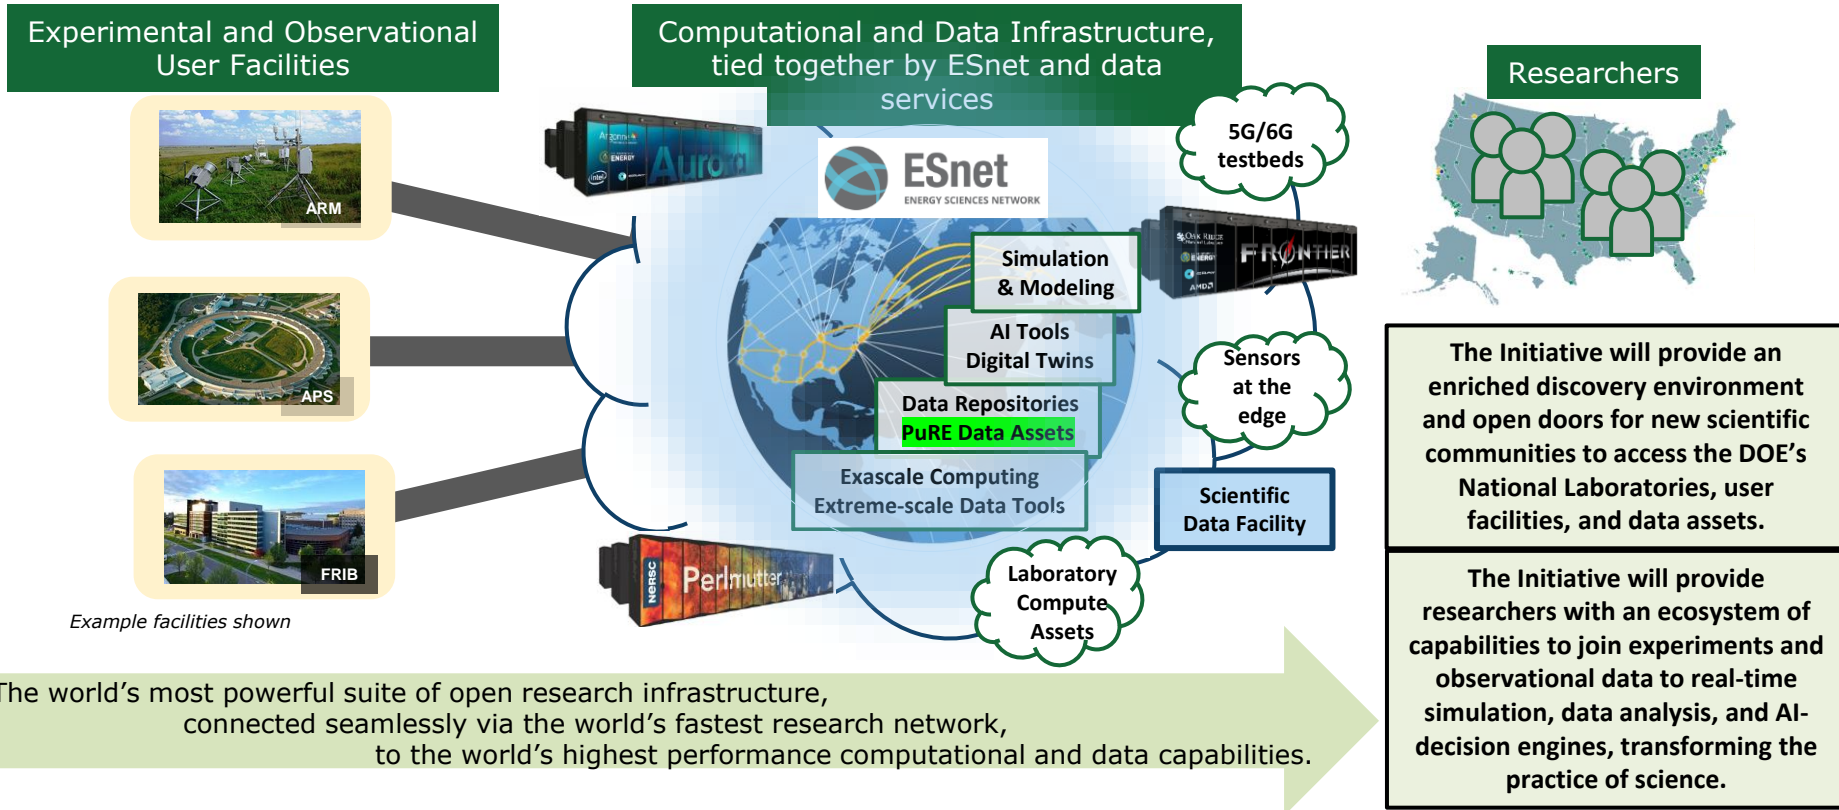
- Competitively support new traineeship awards resulting in:
- “Hands on” experiences, mentoring, enhanced workforce diversity, equity, and inclusion



- Track post-traineeship outcomes
- Assess program effectiveness

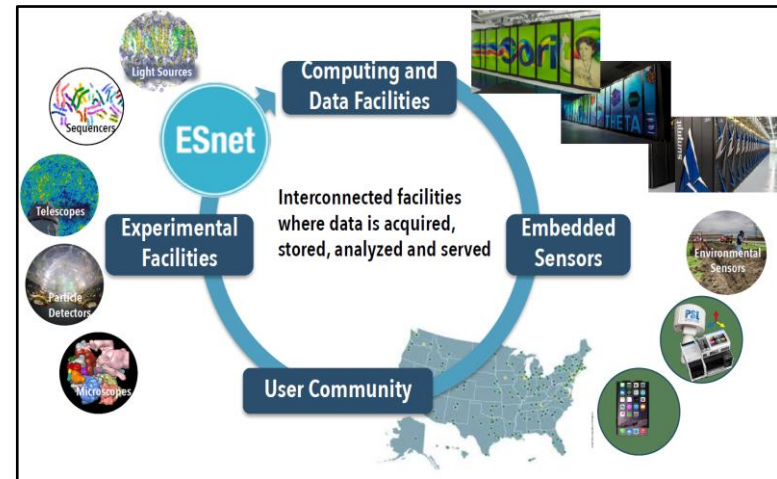
Integrated Computational and Data Infrastructure: A novel discovery environment for the Nation

Creating frictionless access to high performance computing, SC User Facilities, AI tools, and data stores




Scientific HPC Data Facility (HPDF)

- ▶ HPDF will be designed from its inception to support
 - ▶ Data handling
 - ▶ Scientific computing workflow management
 - ▶ High performance, real-time, data processing
- ▶ Provides tools for simulating and managing scientific computing workflows
- ▶ Integrates with edge computing at remote sites to provide real-time access to computing resources that are architecturally diverse than other ASCR facilities
- ▶ Interface smoothly with other ASCR facilities to support acquisition and distribution of data



Data is the Third Pillar of the DOE SC Enterprise



Public Reusable Research (PuRe) Data Resources are:

- data repositories,
- knowledge bases,
- analysis platforms,
- and other activities

that aim to make data **publicly available** in order to advance scientific or technical knowledge.

PuRe Data Resource designations **highlight** and **improve stewardship** of SC-supported community data efforts with strategic impact on the SC mission.

<https://www.energy.gov/science/office-science-pure-data-resources>




PuRe Data Resources

<https://science.osti.gov/Initiatives/PuRe-Data/Resources-at-a-Glance>

▶ Initially designated resources:

- ▶ Atmospheric Radiation Measurement Data Center
- ▶ Joint Genome Institute
- ▶ Materials Project
- ▶ National Nuclear Data Center
- ▶ Particle Data Group
- ▶ Systems Biology Knowledgebase (KBase)



ASCAC 07.29.2021

FY2022 House Appropriations Markup Language

- ▶ The Advanced Scientific Computing Research program develops and hosts some of the world's fastest computing and network capabilities to enable science and energy modeling, simulation, and research.
- ▶ *High Performance Computing and Network Facilities.*—The recommendation provides **not less than \$160,000,000 for the Argonne Leadership Computing Facility, \$250,000,000 for the Oak Ridge Leadership Computing Facility, and not less than \$115,000,000 for the National Energy Research Scientific Computing Center** at Lawrence Berkeley National Laboratory. The recommendation includes **not less than \$90,000,000 to support necessary infrastructure upgrades and operations for ESnet.**
- ▶ *Mathematical, Computational, and Computer Sciences Research.*—The recommendation provides **not less than \$250,000,000** for Mathematical, Computational, and Computer Sciences Research, including not less than **\$15,000,000 for computational sciences workforce programs.**
- ▶ The recommendation includes **not less than \$10,000,000 and up to \$40,000,000** for the development of AI-optimized emerging memory technology for AI-specialized hardware allowing for new computing capabilities tailored to the demands of artificial intelligence systems.

FY 2022 Congressional Action: A closer look at the House Mark(\$ in thousands)

	FY2021 Enacted	FY2022 Request	House Mark Language (not less than)	Difference FY2022 Request and House Mark
ASCR	1,015,000	1,040,000	1,025,000	-15,000
ALCF	152,955	159,047	160,000	+953
OLCF	228,120	249,066	250,000	+ 934
NERSC	113,786	115,963	115,000	--
Esnet*	91,329	93,961	90,000	-3,961
CSGF	10,000	15,000	15,000	--
Exascale	168, 945	129,000	129,000	--
Total Research	259,865	292,962	>250,000	
<i>AI-optimized Memory*</i>	<i>\$10,000</i>		<i>\$10,000<\$40,000</i>	<i>>\$10,000<40,000</i>

FY2021 Actions



Early Career Recommended Selections Advanced Scientific Computing Research (ASCR)

PI Name	Proposal Title	Research Area
Ariful Azad, Indiana	Intelligent Primitives for Scalable Graph Analytics and Learning	Applied Math/ AI
Anup Das, Drexel	Architecting the Hardware-Software Interface for Neuromorphic Computers	Computer Science/AI
Manuel Endres, CalTech	Verification of Quantum Devices from Emergent Randomness	Quantum Information Science
Russell Hewett, Virginia Tech	Domain-Decomposition Induced Parallelism for Scientific Deep Learning at Extreme-Scales	Applied Math/ AI
Katherine Isaacs, U of Arizona	Node-to-Code Comparison-Centered Interactive Performance Visualization	Computer Science/ visualization
Aditya Thakur, UC, Davis	AutoNeurify: Automatic Infusion of Learning in HPC Applications	Computer science/ AI
Sheng Di, ANL	Scalable Dynamic Scientific Data Reduction	Computer Science/AI
Drew Kouri, SNL	Adaptive and Fault-Tolerant Algorithms for Data-Driven Optimization, Design, and Learning	Applied Math/ AI
Robert Saye, LBNL	Advanced Numerics for Atomization and Multi-Physics Interface Dynamics	Applied Math/ Multiscale Mathematics
Nathan Tallent, PNNL	Orchestration for Distributed & Data-Intensive Scientific Exploration	Computer science/ Workflows


New ASCR Solicitations

FOA	Description	Call Info	Funding
X-Stack: Programming Models, Runtime Systems, and Tools	Develop innovative techniques to transform applications to exploit parallelism on future systems and verify correct behavior on those systems.	Open: Jan 28 PP due: Mar 15 Closed: Apr 12	Estimated total funding available: \$12M over three years
Entanglement Management and Control in Transparent Optical Quantum Networks	Focus is on entanglement traffic engineering - the scalable management and control of a large pool of quantum entangled states carried over multi-hop multi-user optical quantum networks	Open: Mar. 18 LOI due: Apr 5 Closed: May 17	Estimated total funding available: \$6M over three years
ICDI: Integrated Computational and Data Infrastructure Initiative	Design and deploy a flexible multi-tiered data and computational management AI-enabled architecture that enables a diverse array of on-demand scientific workflows and simulations for SC mission research.	Open: Mar. 19 LOI due: Apr 2 Closed: May 14	Estimated total funding available: \$13.5M over three years
5G Enabled Energy Innovation: Advanced Wireless Networks For Science (Lab only)	Develop new programming framework approaches that will enable effective data collection, movement and analysis to span the continuum from wireless to scientific user facilities, data centers, and high performance computing	Open: Mar. 18 PP due: Apr 20 Closed: May 19	Estimated total funding available: \$6M over three years
Microelectronics Co-Design Research (Lab only)	Multi-disciplinary research to accelerate the advancement of semiconductor/microelectronics technologies in a co-design innovation ecosystem in which materials, chemistries, devices, systems, architectures, algorithms, and software are developed in a closely integrated fashion with BES, FES, HEP and NP	Open: Mar 24 LOI due: Apr 21 Closed: May 19	Estimated total funding: Up to \$54M over three years

New ASCR Solicitations

FOA	Description	Call Info	Funding
Data-Intensive Scientific Machine Learning and Analysis	Develop innovative techniques to transform applications to exploit parallelism on future systems and verify correct behavior on those systems.	Open: Mar 26 PP due: May 3 Closed: May 27	Estimated total funding available: up to \$21 M over three years
EXPRESS – Randomized Algorithms for Extreme-Scale Science	Focus is on entanglement traffic engineering - the scalable management and control of a large pool of quantum entangled states carried over multi-hop multi-user optical quantum networks	Open: Mar. 18 LOI due: Apr 5 Closes: May 17	Estimated total funding available: up to \$8M over two years
Data Reduction	Explore potentially high-impact approaches in the development and use of data reduction techniques and algorithms to facilitate more efficient analysis and use of massive data sets produced by observations, experiments and simulations. (Joint AM/CS call)	Open: Apr. 15 PP due: May 14 Closes: June 4	Estimated total funding available: up to \$10.0M over three years
Advancing the Quantum Internet Backbone for DOE Mission Science (Lab only)	Advance selected priority research objectives and milestones articulated in the report of the DOE Quantum Internet Blueprint Workshop, held in February 2020. This activity joins an array of others initiated by the SC in the quantum information science domain, in alignment with the National Quantum Initiative Act	Open: Apr. 13 Closed: May 28	Estimated total funding available: \$25M over three years
NIH Collaboration: Bridge2AI and Privacy Preserving Artificial Intelligence Research (Lab Only)	Encourage lab collaborations with NIH Bridge2AI program focused on privacy-preserving algorithms for AI datasets	Open: June 7 Closed: July 19	Total funding available: \$1M in FY2021

Office of Science User Facilities Roundtable: Lessons from the COVID Era and Visions for the Future



In December 2020, SC convened a virtual Roundtable of its scientific user facilities to discuss facility challenges and lessons learned during the COVID-19 pandemic as well as facility responses, best practices, and innovations that could be adopted going forward.

Roundtable participants included facility staff, users, and user executive committee chairs. The discussions encompassed topics such as user research and facility operations in virtual and physically distanced contexts; user training and engagement; computation, data, and network resources; and crosscutting issues.

The report contains a chapter devoted to Computing and Data Resources, highlighting that numerous challenges and opportunities exist to improve these resources for the SC User Facilities community.


Interactions with other Agencies

- ▶ NSF – Memorandum of Agreement
 - ▶ Joint Oversight Group (JOG) for Quantum Information Science between NSF and SC
 - ▶ Cooperation in the Funding of NSF Collaborative Research in Computational Neuroscience (CRCNS) -- also NIH

- ▶ NCI – Renewed the Memorandum of Understanding between the National Cancer Institute, the Office of Science and The National Nuclear Security Administration to fund the Joint Design of Advanced Computing Solutions for Cancer (JDACS4C)

ASCR Software-Stewardship Taskforce

- In March 2021, ASCR established an internal taskforce to study the challenges posed by
 - The increasingly complexity of the software ecosystem for scientific and high-performance computing; and
 - The upcoming transition from the Exascale Computing Project (ECP), which has been supporting the development of most of that ecosystem for many years.
 - Led by Hal Finkel and Ben Brown
- The charge of the task force has been substantially shaped by the recommendations in the October 2020 ASCAC report on *Transitioning ASCR after ECP*.



ASCR Software-Stewardship Taskforce

- Nearly all computational science, including computer science and applied mathematics research, depends on building on previous work – *largely gone are the days of starting from a blank slate.*
- The software ecosystem will continue to evolve past ECP, foreseeably growing to encompass additional capabilities for data analytics, AI/ML, and complex workflows.
- Stewardship done well will promote thriving careers in scientific software.
- In light of the ASCAC recommendations and other feedback, the task force recognizes the multi-faceted nature of sustaining the software ecosystem, potentially including:
 - Training on software usage and best practices for development;
 - Support for building and maintaining a diverse, skilled workforce;
 - Infrastructure for software packaging, hosting, testing, and so on;
 - Establishing governance processes and standards to enable effective resource allocation;
 - Collecting information from users and facilities, and providing information for future planning;
 - Software-engineering resources to assist with maintenance activities of key projects; and
 - Providing support for the continued development of key projects.

ASCR Software-Stewardship Taskforce

- The task force has already met with a number of key groups, including:
 - The former ASCAC subcommittee on *Transitioning ASCR after ECP*.
 - ECP leadership and the leadership of the ASCR facilities.
- Meetings with the Computational Research Leadership Council (CRLC) which is composed of Lab Division Directors and other science-funding agencies are upcoming.
- The task force is preparing an RFI to seek broad feedback from the wider community.
- Subject to DOE travel policy, we are planning to attend SC21, please **contact Hal Finkel** to express interest in meeting with us.

Subscribe to ASCR News and Information using GovDelivery

- GovDelivery is a web-based e-mail subscription management system used by DOE and other US government agencies to communicate with the public.
- Signup with GovDelivery to receive notifications regarding ASCR funding announcements, awards, and other news.

On <https://www.energy.gov/science/office-science> (near the bottom of the page):

Email Updates

To receive the latest news and updates about the Office of Science, submit your e-mail address.

- Science Updates**
- Artificial intelligence and machine learning
 - Advanced Scientific Computing Research Program
 - Basic Energy Sciences Program

Check the box for *Advanced Scientific Computing Research Program* in your subscriber preferences.

Some ASCAC Agenda Details

- ▶ EXPANDING PARTICIPATION IN COMPUTATIONAL SCIENCE AND ENGINEERING – *Mary Ann Leung, Sustainable Horizons Institute*
- ▶ DOE COVID-19 S&T RESPONSE: NATIONAL VIRTUAL BIOTECHNOLOGY LABORATORY – *Michelle Buchanan, Office of Science*
- ▶ ASCR ROUNDTABLE DISCUSSIONS ON OPERATING-SYSTEMS RESEARCH -- *Hal Finkel. ASCR*
- ▶ EXASCALE UPDATE – *Katie Antypas, Lawrence Berkeley National Laboratory, Susan Coghlan, Argonne National Laboratory*
- ▶ OVERVIEW OF COLLABORATION BETWEEN THE DOE, NNSA AND NCI – *Emily Greenspan, National Cancer Institute*
- ▶ UPDATE ON FUGAKU– *Satoshi Matsuoka, Director RIKEN Center for Computational Science (R-CCS)*