U.S. DEPARTMENT OF ENERGY | Office of Science

Ecosystems are the future!

SC21 BOF: Software Engineering and Reuse in Modeling, Simulation, and Data Analytics for Science and Engineering

November 16, 2021

Benjamin Brown, Ph.D.

Director, Facilities Division

Office of Advanced Scientific Computing Research

ASCR Facilities Division: The DOE Headquarters team

Jordan Thomas

Bill Miller
(NSF detailee)

Sash Hier-Majumder

Carol Hawk

Christine Chalk

Argonne NATIONAL LABORATORY

BERKELEY LAB

OAK RIDGE National Laboratory

U.S. DEPARTMENT OF ENERGY | Office of Science

ASCR High Performance Computing and Networking Facilities

World leading capabilities spanning supercomputing, data analysis, data transport & testbeds

Argonne Leadership Computing Facility

Oak Ridge Leadership Computing Facility

Leadership Computing: Extreme-scale resources for the nation

- ALCF and OLCF provide two HPC architectures for technological diversity
- ~3,000 users per year; multiple #1 Top500 rankings over program history
- Emphasis on science and technology applications that use full system capability
- Resources allocated predominantly by competitive merit review
- Current upgrade projects: OLCF-5 **Frontier** (2021) and ALCF-3 **Aurora** (2022)

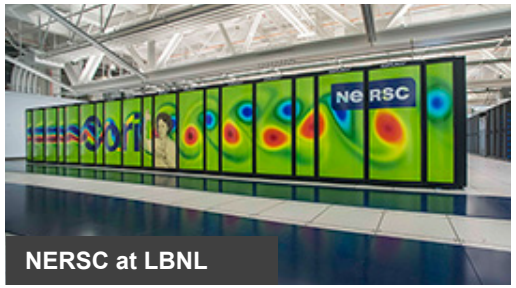

NERSC at LBNL

High Performance Production Computing: A dedicated SC resource

- NERSC's legacy of enabling DOE research with HPC stretches back to 1974
- ~8,000 users per year; NERSC also provides a 200 PB data storage archive
- Emphasis on support for the broadest set of science applications
- Resources allocated predominantly by SC Science Programs to their grantees
- Current upgrade project: NERSC-9 **Perlmutter** (2021)

Advancing U.S. Competitiveness

Every ASCR HPC system procurement includes R&D to drive innovation across the U.S. vendor community.

LCFs constitute a global competitive HPC advantage.


ESnet, managed by LBNL

High Performance Networking: A superhighway for extreme-scale data

- Connects all DOE national laboratories and other DOE sites to global research networks, cloud providers, and the internet
- Many tens of thousands of individual users; ESnet provides DOE the ability to move massive data losslessly
- An open network with high capacity (400+ Gbps), low latency, and innovative services tuned for extreme-scale data
- Transmitted more than one Exabyte (one billion Gigabytes) in the last 12 months; ESnet Testbed enables open R&D
- Current upgrade project: **ESnet6** (2023), a Terabit-scale network with software programmable service orchestration

U.S. DEPARTMENT OF ENERGY | Office of Science

The people of the ASCR Facilities

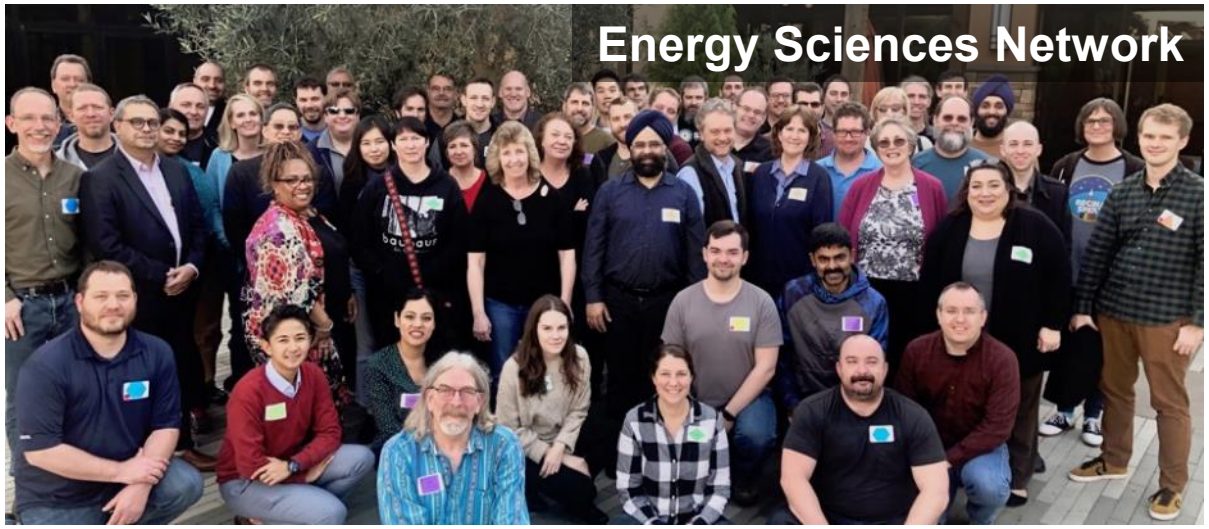

FY 2021
28 scientific user facilities
36,000+ users

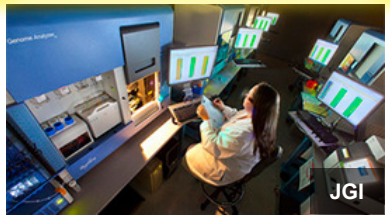

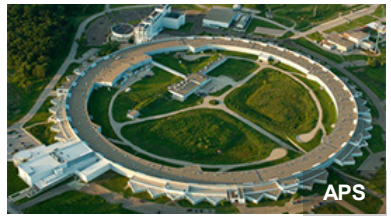

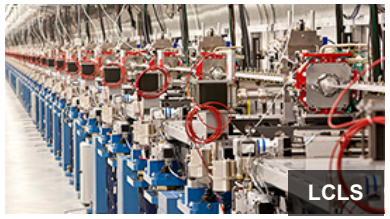

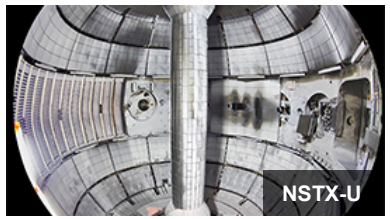

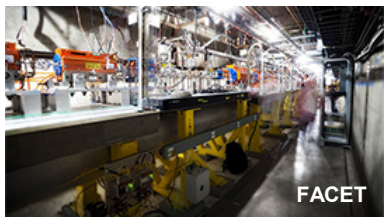

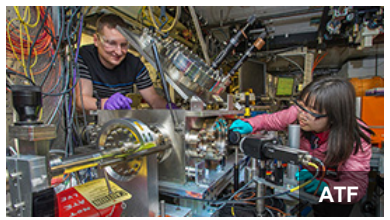

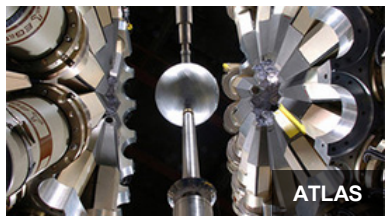

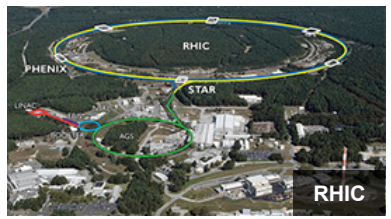

U.S. DEPARTMENT OF ENERGY
Office of Science

Acronym decoder at https://science.osti.gov/User-Facilities

On September 29 I presented my vision for the ASCR Facilities enterprise to the Advanced Scientific Computing Advisory Committee

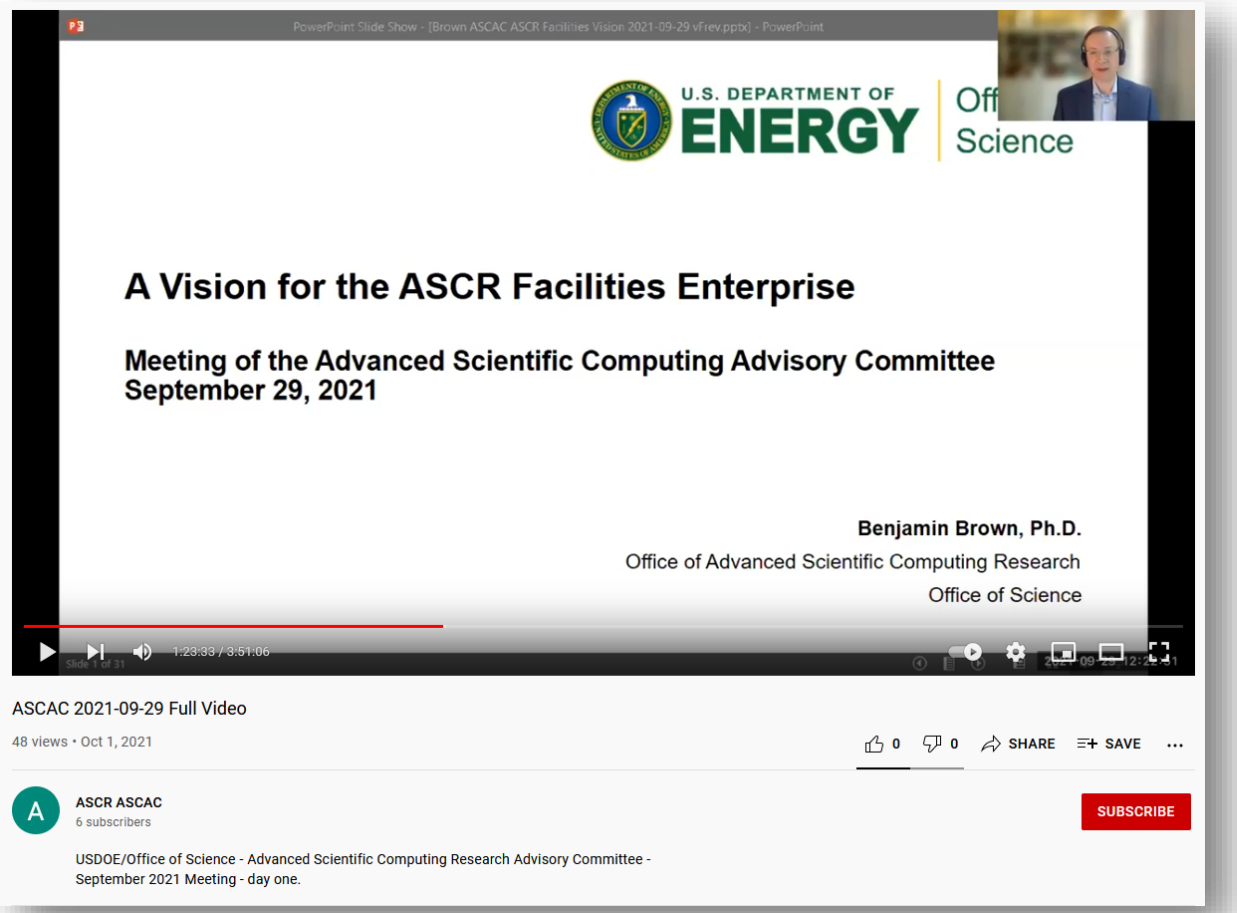

https://youtu.be/ItYuCtS4QH4?t=4971

A complex workflow addressing extraordinary national need

This artist's rendering depicts x-ray crystallography at SLAC's Linac Coherent Light Source. LCLS partnered with NERSC and ESnet to perform real-time image analysis for research of the SARS-CoV-2 virus structure.

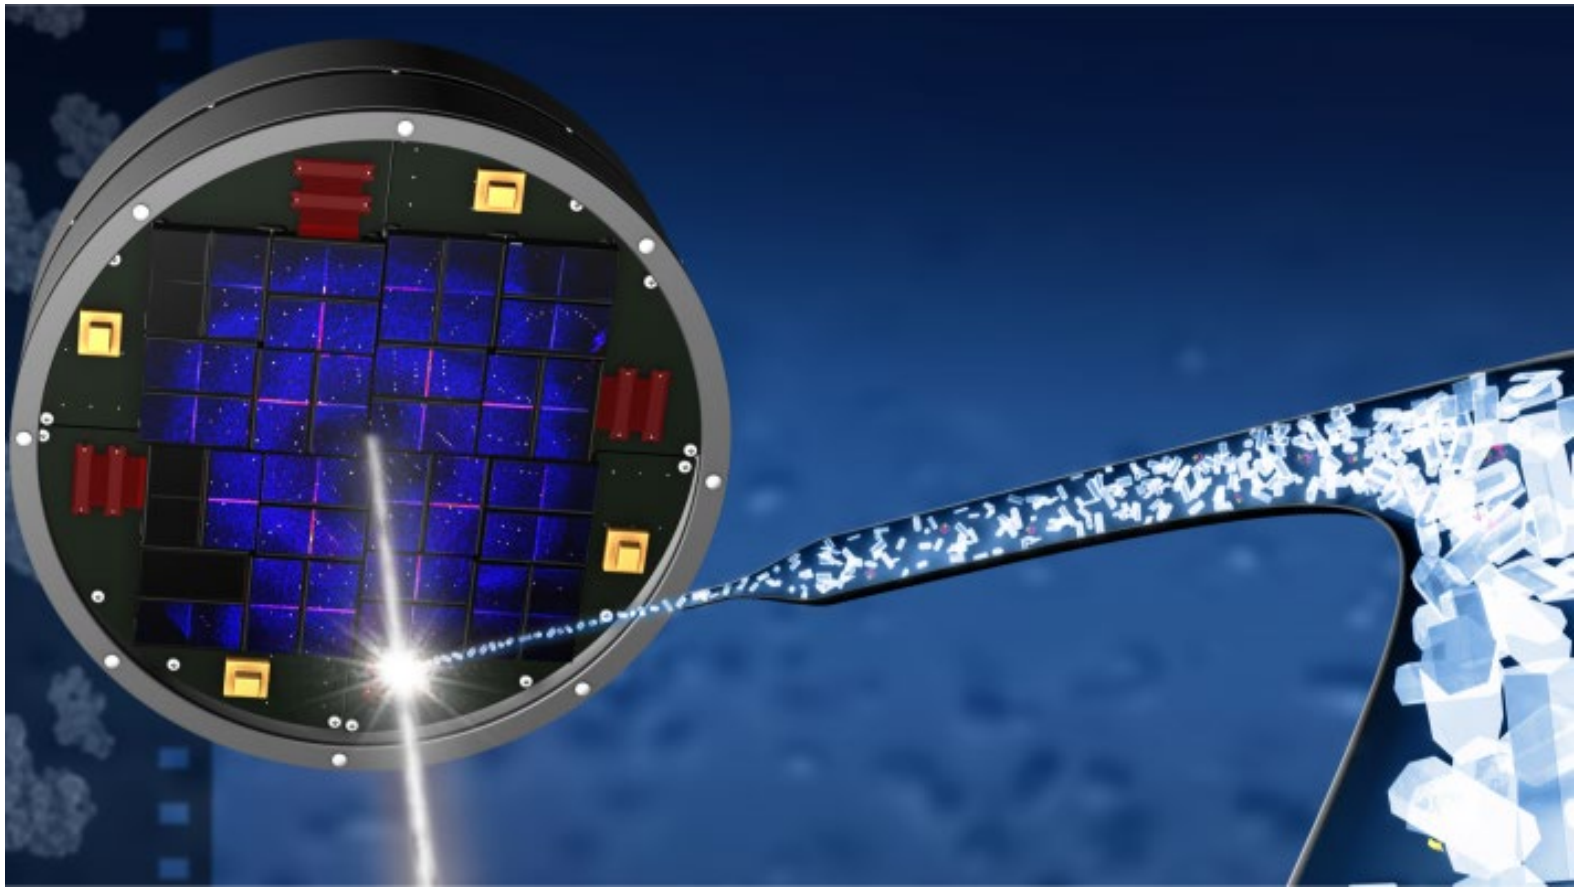

SLAC National Accelerator Laboratory

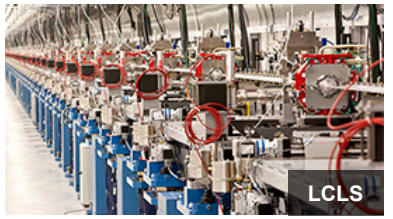

Don't miss Bronson Messer's COVID-19 HPC Consortium talk, Thursday at 11:05am ET.

U.S. DEPARTMENT OF ENERGY | Office of Science

Today we are entering not only the exascale era, but also a new era of complexity in advanced scientific computing.

Today the ASCR Facilities enterprise is contending with new complexity. We are entering a new era of advanced scientific computing.

The practice of science is evolving. Couplings between modeling/simulation, experimental/observational data, advanced algorithms, and AI/ML tools have the power to accelerate discovery and innovation.
Where we once focused on batch jobs and bulk data transfer, we now have complex workflows.

Computing technology is evolving along multiple trajectories. General purpose computing is but one market segment.
Managing risk and opportunity in our hardware choices is increasingly complex.

The ***people*** of the ASCR Facilities enterprise are making extraordinary impacts today; their expertise and efforts are sought by many. And yet many talented individuals do not participate.
Our workforce challenges are significant.

Institutions, programs, and researchers are under pressure to provide/obtain computing and data resources.
Our users, our partners, and we ourselves crave shared clarity of insight and intent.

Our challenge today is to confront this complexity and arrive at a strategy that maximizes the impact of ASCR, Office of Science, and DOE investments—to be greater than the sum of the parts.

Vision for the ASCR Facilities: Thriving together

A complementary system of facilities, each thriving, each possessing agency, collectively driving innovation in advanced scientific computing across DOE and beyond.

- Driving the state-of-the-art with the ASCR research and vendor communities
- Catalyzing discovery and innovation
- Responding to national needs
- Delivering on stakeholder priorities, with balance and equity
- Fostering scientific ecosystems
- Broadening the diversity of individual, institutional, and domain participation
- Demonstrating excellence in project management and operations

… a system in which we (HQ & Facilities) manage enterprise risk and opportunity together and facilitate our stakeholders' abilities to do so effectively.

Vision for the ASCR Facilities: How we will thrive together

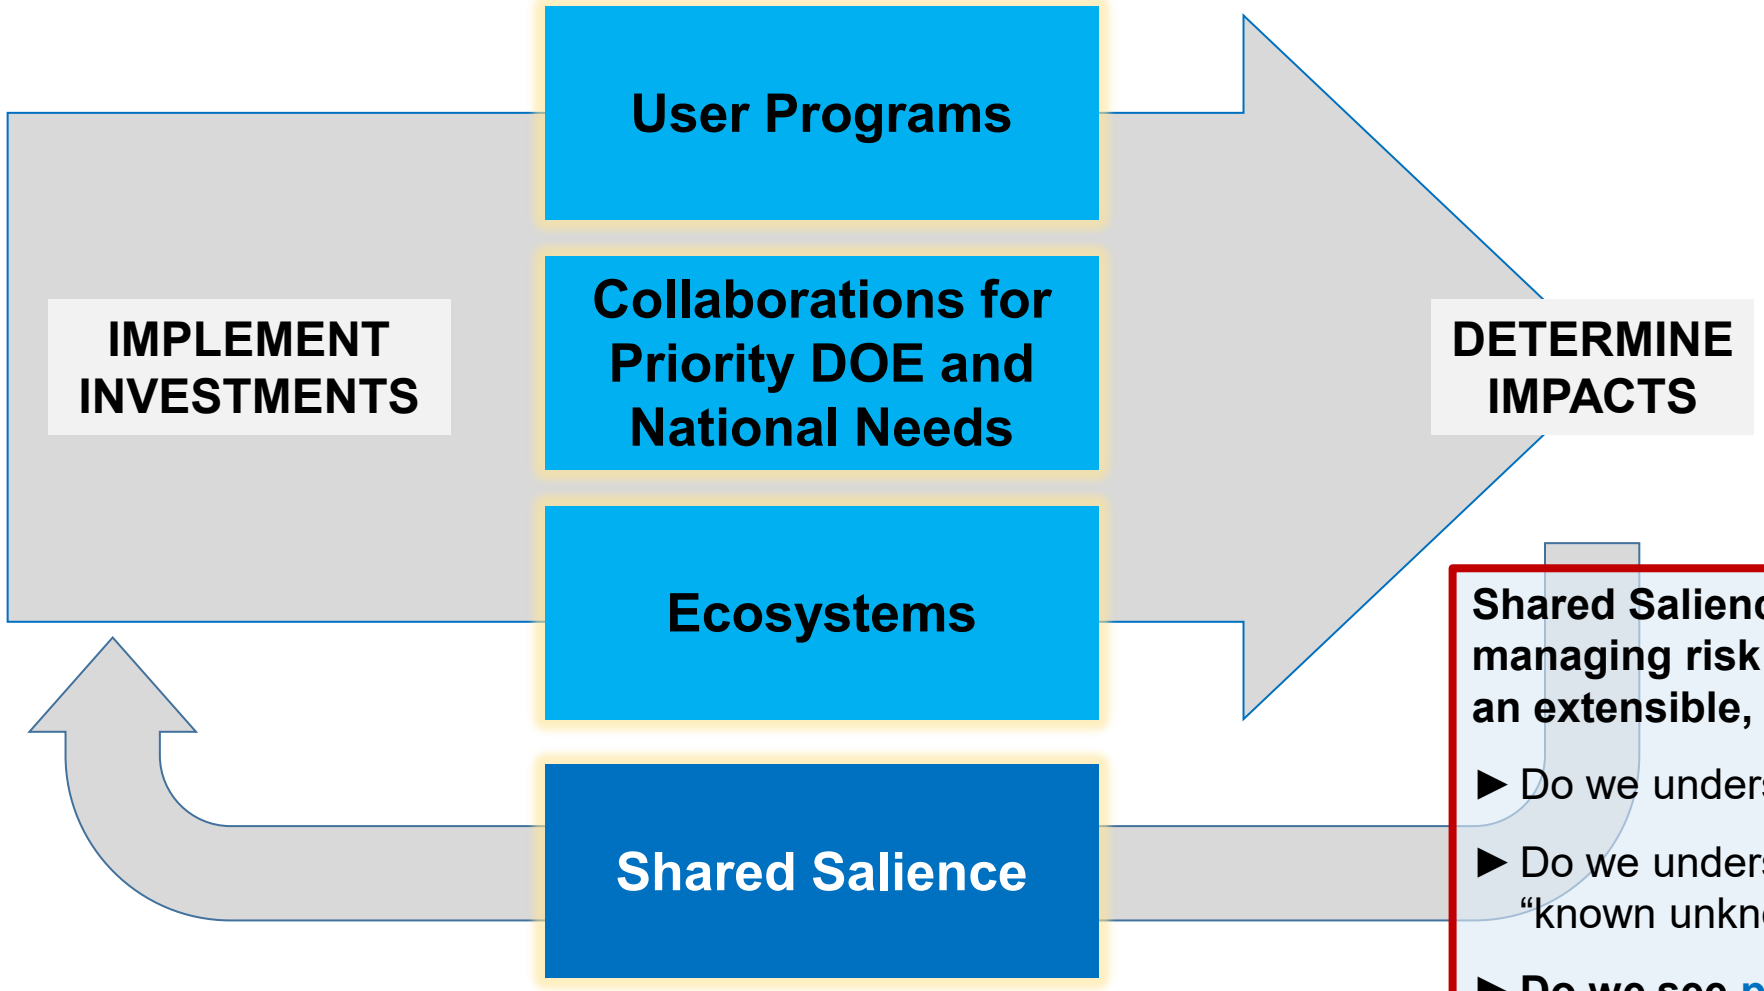

Engage & surface.

Listen & reflect.

Address & implement.

Repeat.

Shared Salience is about all stakeholders managing risk and opportunity, together, with an extensible, systematized methodology.

- ► Do we understand each other's priorities?
- ► Do we understand the "known knowns" and "known unknowns"?
- ► **Do we see patterns to inform investments?**

Incipient ecosystem: Office of Science User Facilities

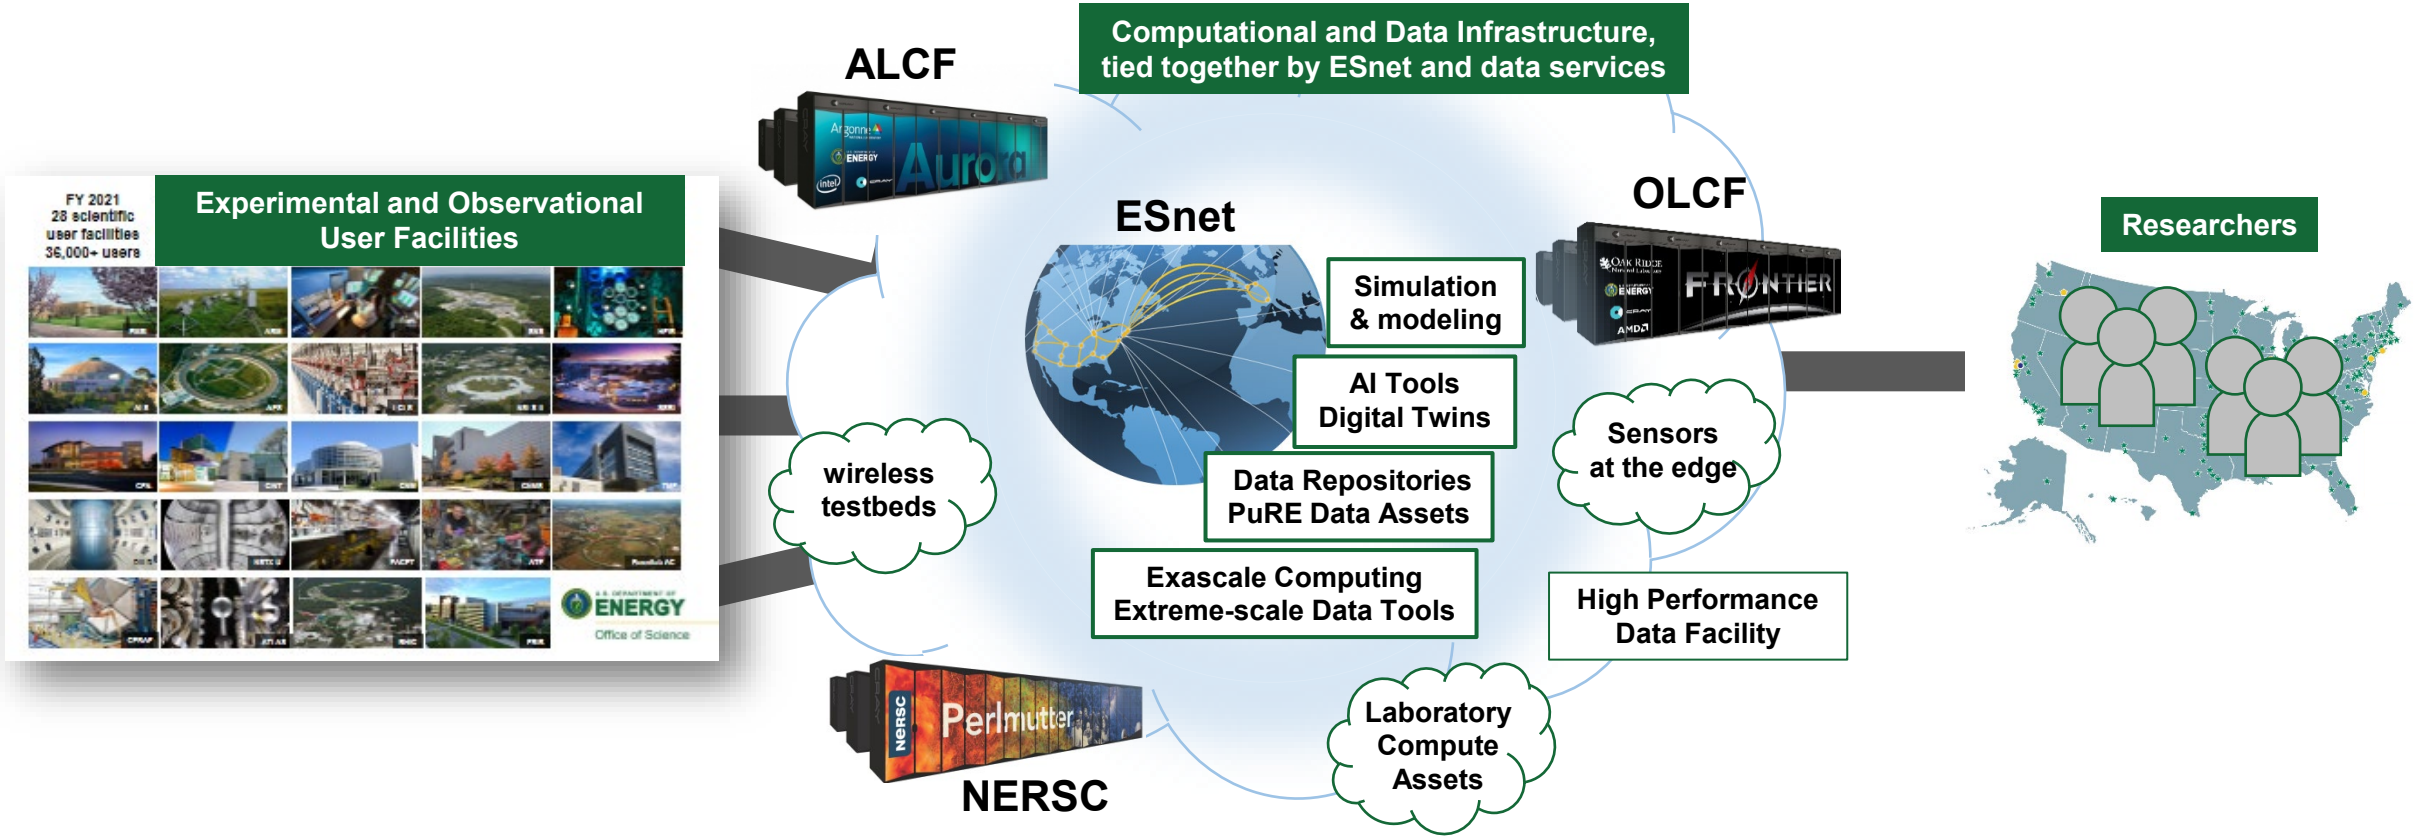

Incipient ecosystem: Scientific software

EXASCALE COMPUTING PROJECT

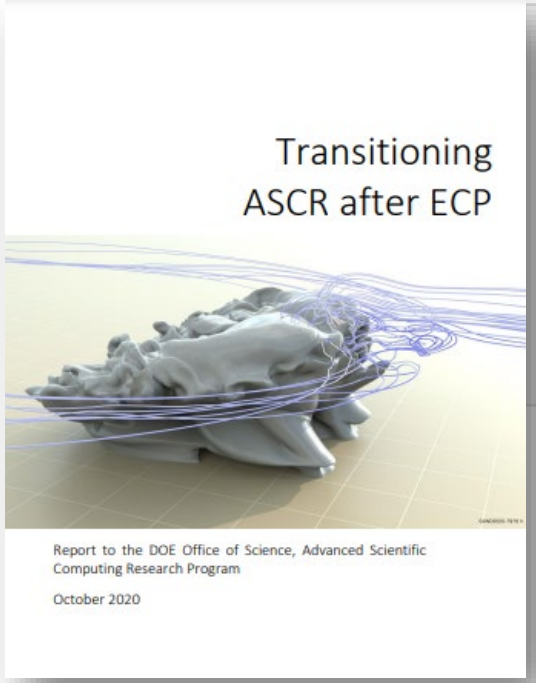

"ECP has created a well-designed software ecosystem for development, curation, and distribution of exascale systems and application software. This ecosystem integrates the fruits of years of basic research in: mathematics, computer science, applications, and systems software.

In particular, the ecosystem greatly reduces barriers for ASCR fundamental research maturation and impactful delivery at the facilities and with users. Several of our recommendations focus on realizing the potential of this new ecosystem."

The Importance of Stewardship and Sustainability of Research Software in the Office of Science

Anshu Dubey, Mathematics and Computer Science, Argonne National Laboratory
Katherine Riley, Argonne Leadership Computing Facility, Argonne National Laboratory
Nicholas Schwarz, Advanced Photon Source, Argonne National Laboratory
David E. Bernholdt, Computer Science and Mathematics and Oak Ridge Leadership Computing Facility, Oak Ridge National Laboratory
Bronson Messer, Oak Ridge Leadership Computing Facility, Oak Ridge National Laboratory
Mathieu Doucet, Neutron Scattering Division, Oak Ridge National Laboratory
Rama K. Vasudevan, Center for Nanophase Materials Sciences, Oak Ridge National Laboratory
Deborah Agrawal, Computing Research Division, Lawrence Berkeley National Laboratory
Katerina Antypas, National Energy Research Scientific Computing, Lawrence Berkeley National Laboratory
Harinarayan Krishnan, Advanced Light Source/Computing Research Division, Lawrence Berkeley National Laboratory
Edward Balas, Energy Sciences Network, Lawrence Berkeley National Laboratory

August 3, 2021

Software ecosystems are research infrastructure!

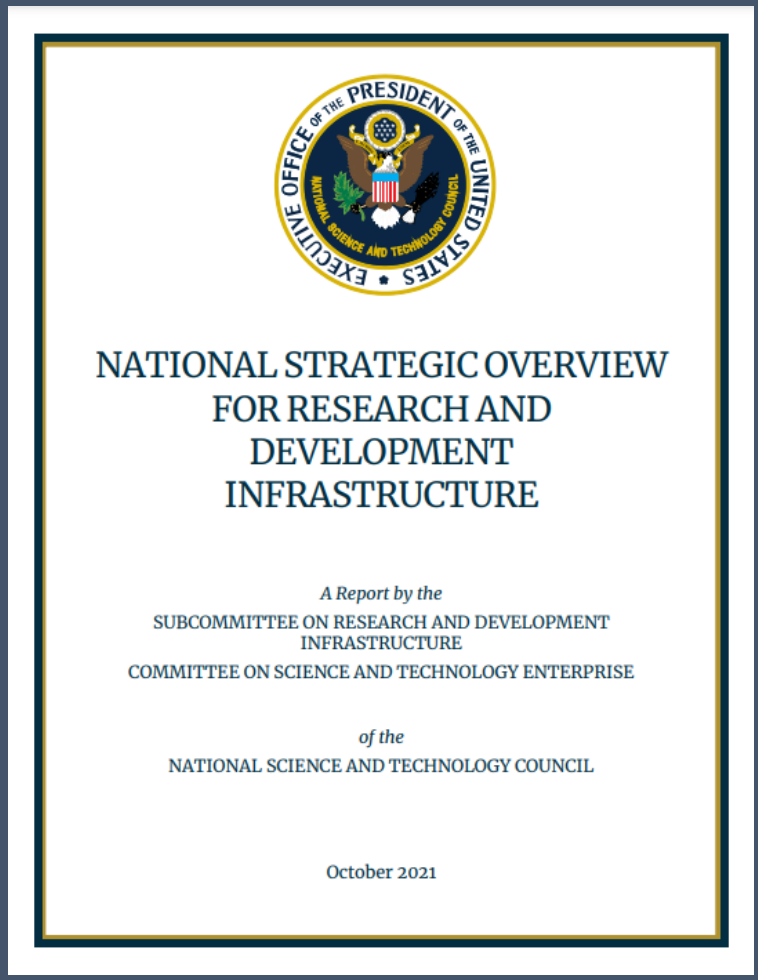

NATIONAL STRATEGIC OVERVIEW FOR RESEARCH AND DEVELOPMENT INFRASTRUCTURE

A Report by the

SUBCOMMITTEE ON RESEARCH AND DEVELOPMENT INFRASTRUCTURE

COMMITTEE ON SCIENCE AND TECHNOLOGY ENTERPRISE

of the

NATIONAL SCIENCE AND TECHNOLOGY COUNCIL

October 2021

ben.brown@science.doe.gov