A Vision for the ASCR Facilities Enterprise

Meeting of the Advanced Scientific Computing Advisory Committee
September 29, 2021

Benjamin Brown, Ph.D.

Office of Advanced Scientific Computing Research

Office of Science

Dedication

Bob Astheimer

David Skinner

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)

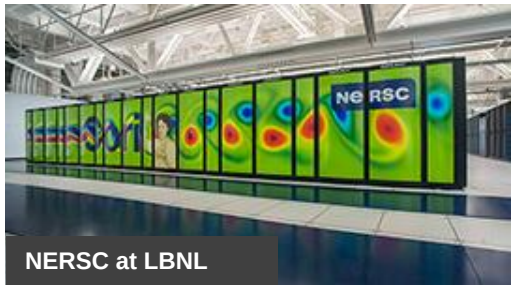

NERSC at LBNL

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

Appreciation: The people of the ASCR Facilities

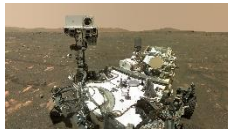

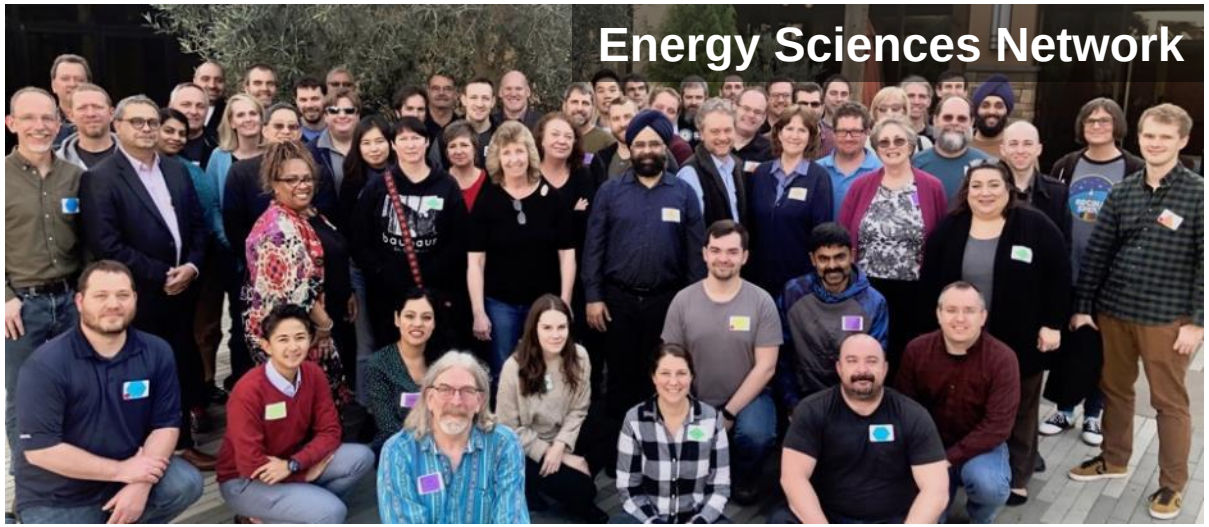

U.S. DEPARTMENT OF ENERGY | Office of Science

Timeline of DOE HPC upgrade projects & ESnet6

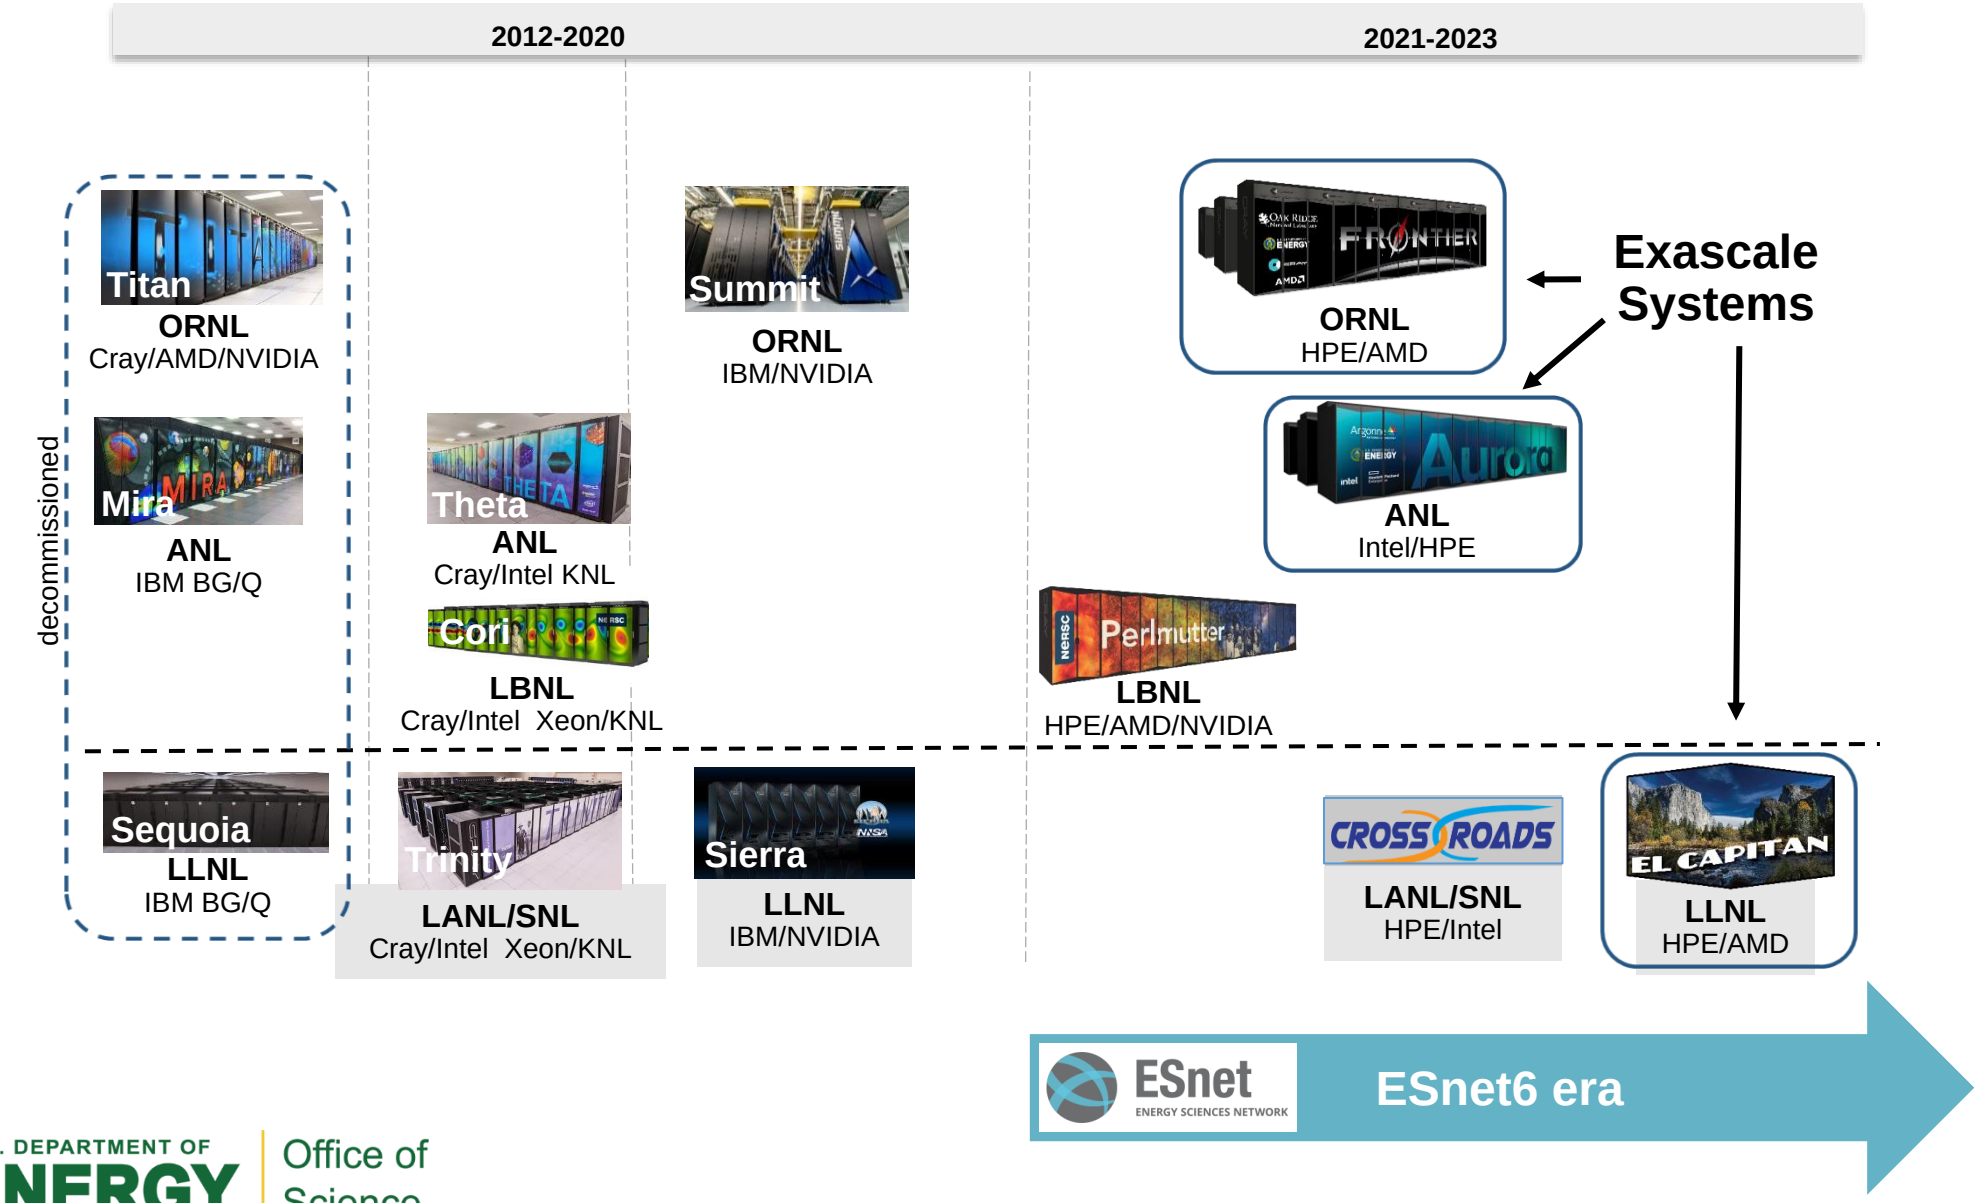

ASCR Facilities: Users statistics, FY 2015-20

Total HPC users per facility, per year

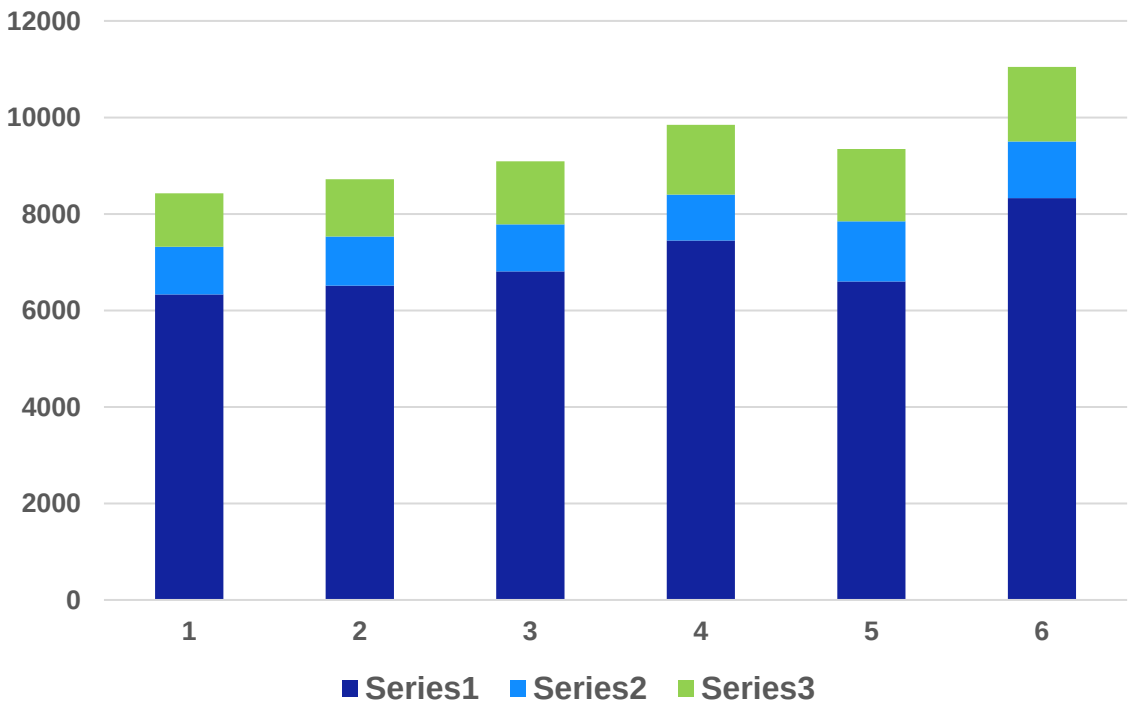

Industrial users

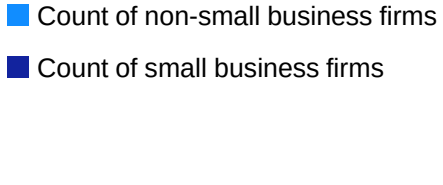

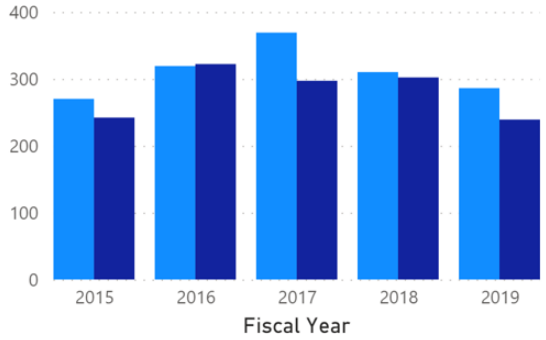

Notable large firms

3M Corporation
Altair Engineering, Inc.
ANSYS Inc.
ARM Inc.
Boeing Company
BP America, Inc.
Caterpillar, Inc.
Cisco Systems
Cray Inc.
Dow Chemical Company
Dresser-Rand Company
Electro-Motive Diesel, Inc.
Engility Corporation
Fiat Chrysler Automobiles US LLC
FM Global
Ford Motor Company
General Atomics
General Electric Company (GE)
General Motors Corporation
GlobalFoundries
Hitachi America Ltd.
Honda Performance Development, Inc.
IBM Corporation
Intel Inc.
Leidos, Inc.
Lockheed Martin Corporation
Mellanox Technologies, Inc.
Microsoft Corporation
NEC USA
Novartis Vaccines and Diagnostics, Inc.
Novozymes, Inc.
Nvidia Corporation
Orbital ATK
Pfizer, Inc.
Proctor & Gamble Company
Robert Bosch, LLC
Rolls-Royce Corporation
SABIC Americas, Inc.
Samsung Semiconductor Inc.
SAIC
Science Systems and Applications, Inc.
Siemens Corp
SpaceX
Toyota Motor Sales, U.S.A., Inc.
TransTech International
United Technologies Research Center Inc.
UT-Battelle, LLC
Varian Medical Systems
Westinghouse Electric Company LLC

ASCR Facilities: U.S. user institutions, FY 2020

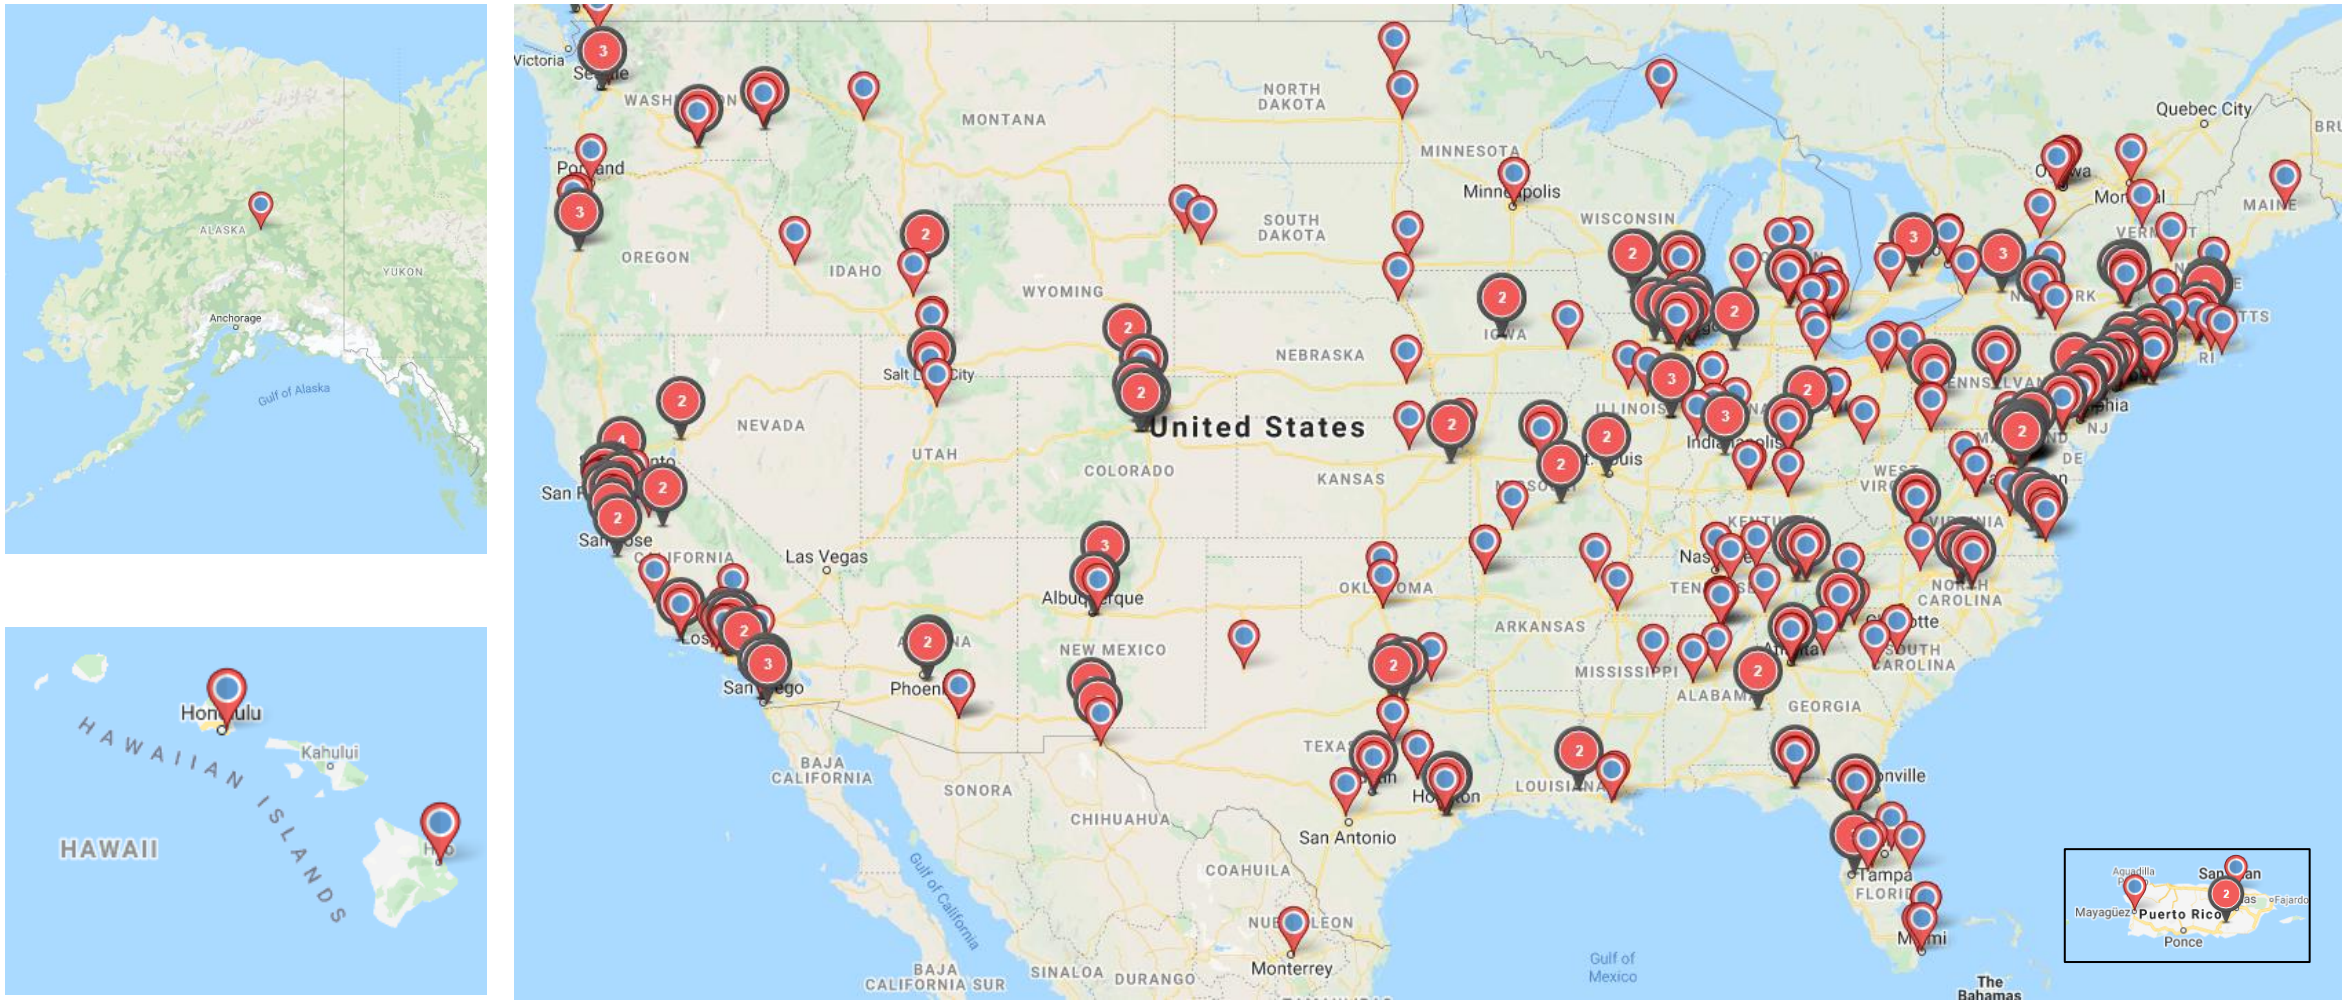

FY 2021
28 scientific user facilities
36,000+ users

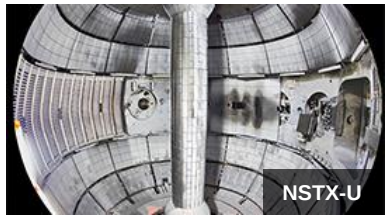

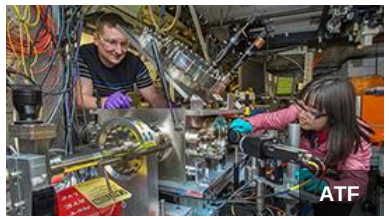

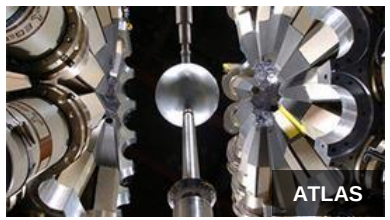

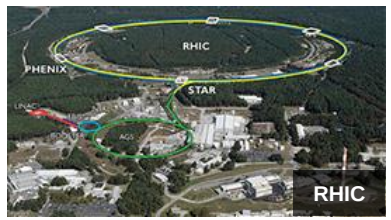

U.S. DEPARTMENT OF ENERGY
Office of Science

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

A significant number of users of experimental and observational SC User Facilities also use the ASCR Facilities

The perimeter arcs are proportional to the number of users who used more than one SC User Facility.

The width of a ribbon between two User Facilities is proportional to the number of users that used both of those facilities.

Nearly every non-ASCR Facility has users who leveraged an ASCR HPC Facility. ESnet connects to every SC User Facility.

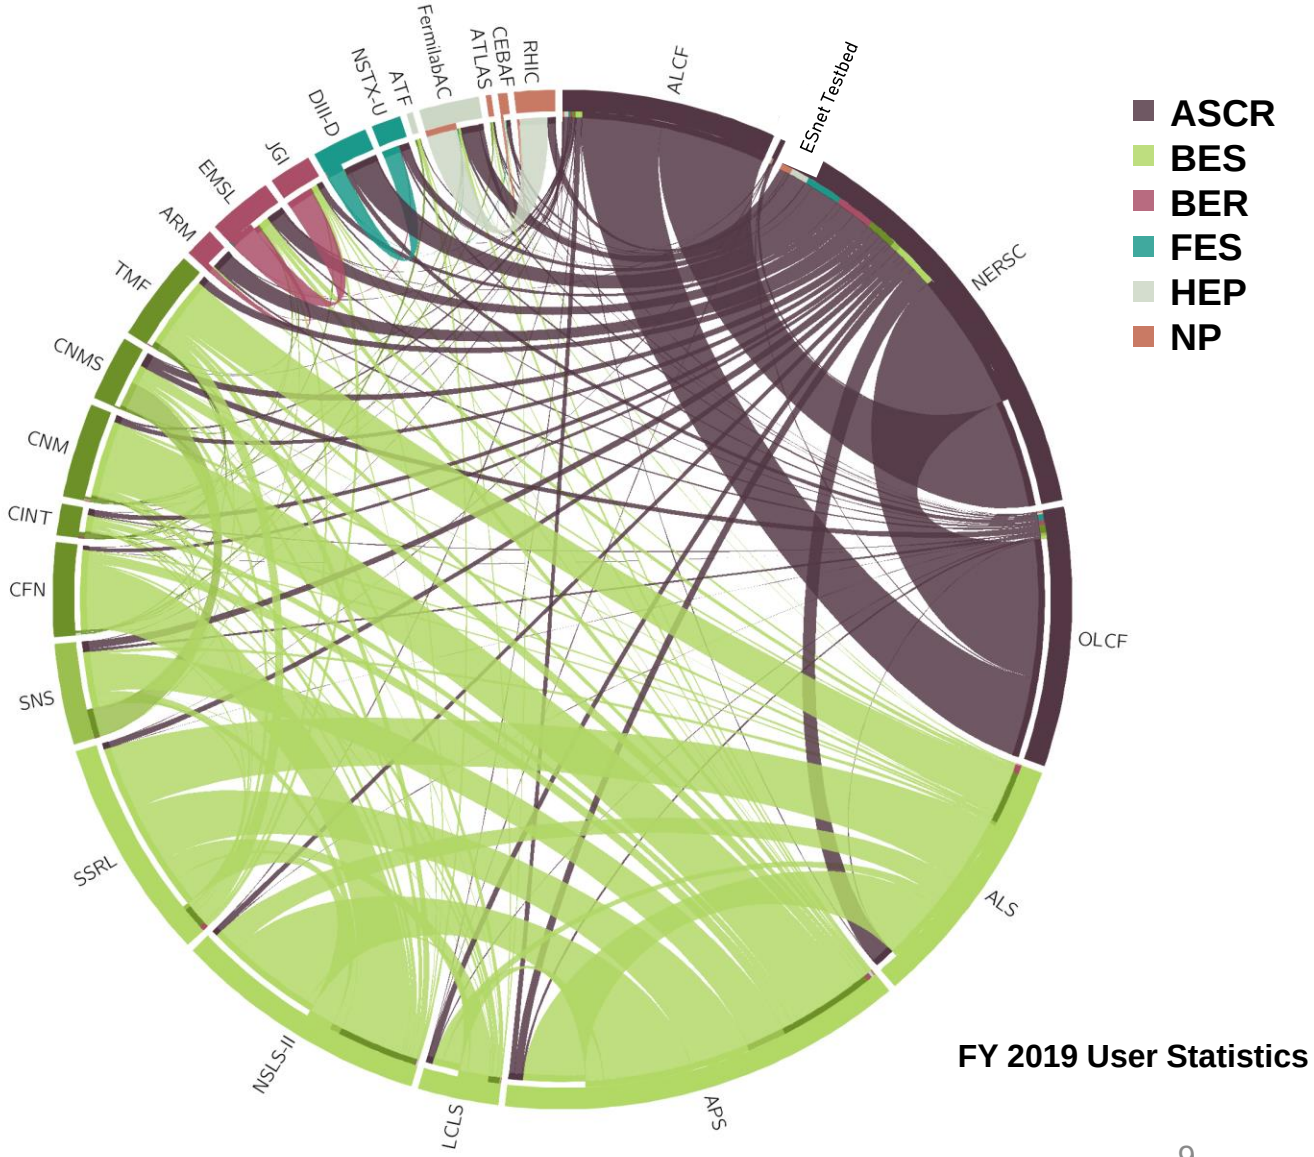

FY 2019 User Statistics

U.S. DEPARTMENT OF ENERGY | Office of Science

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.

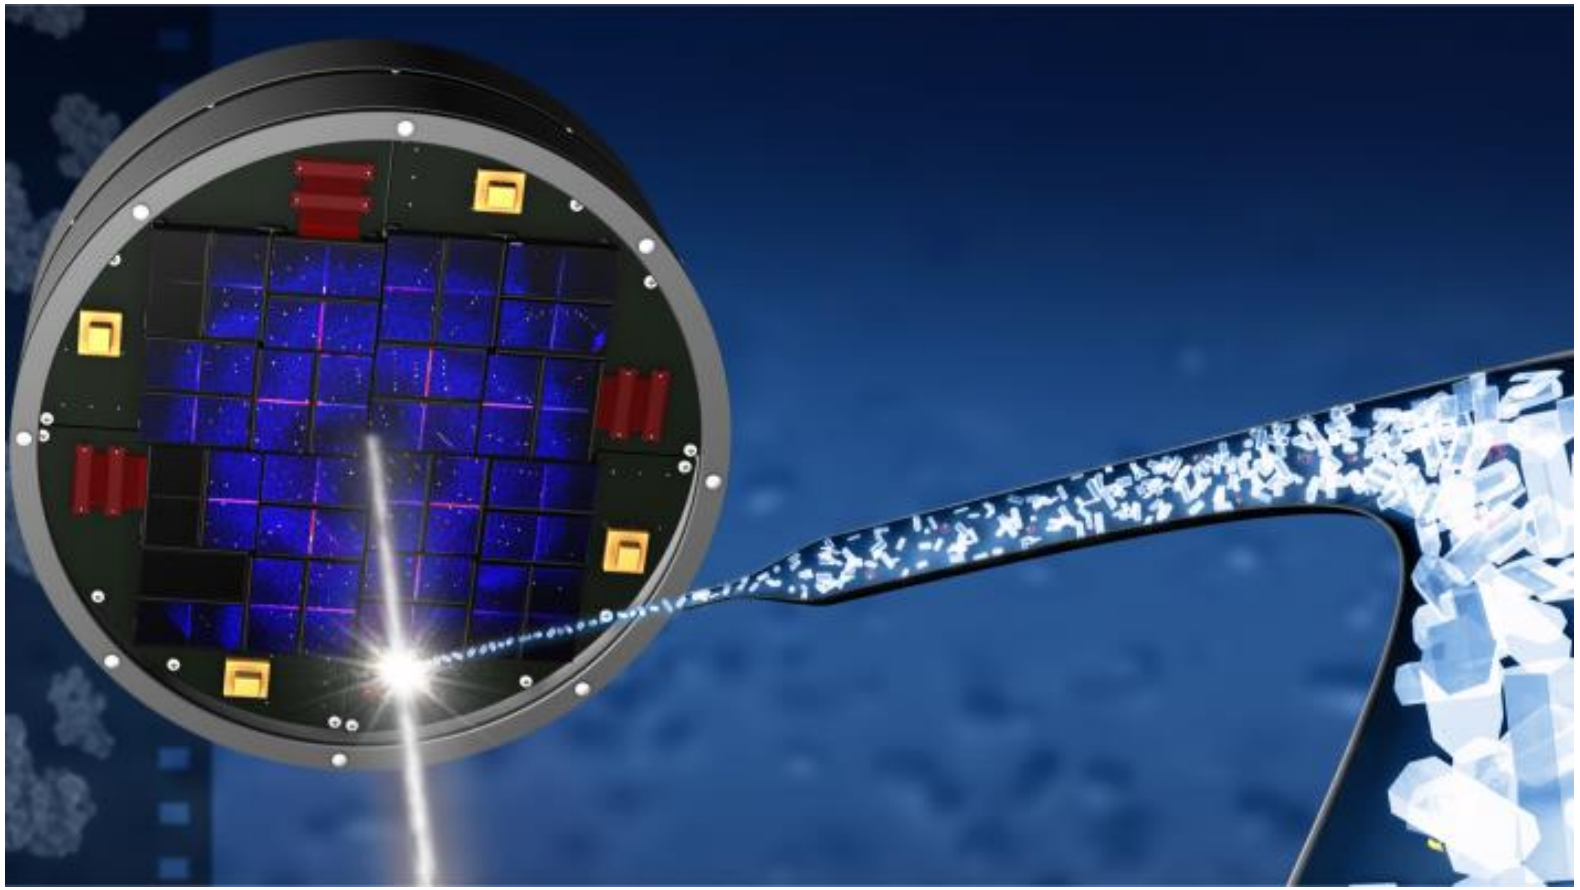

SLAC National Accelerator Laboratory

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.

There are several ASCAC reports that point to this complexity

Data | **Methods** | **Software** | **Hardware** | **Platforms**

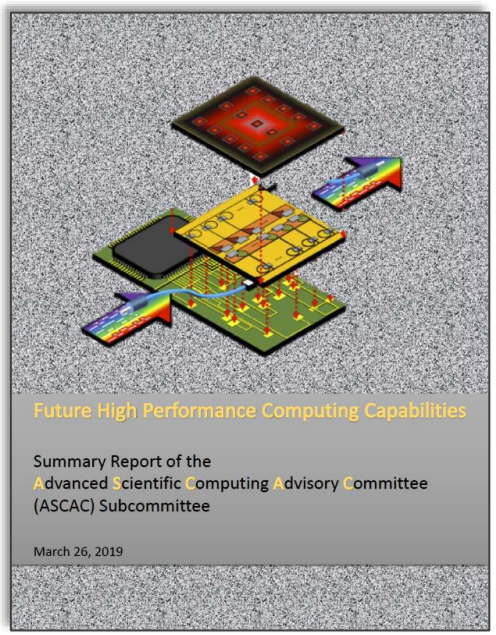

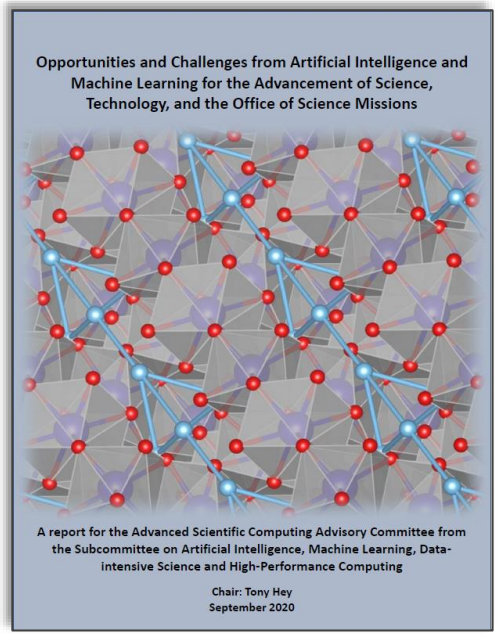

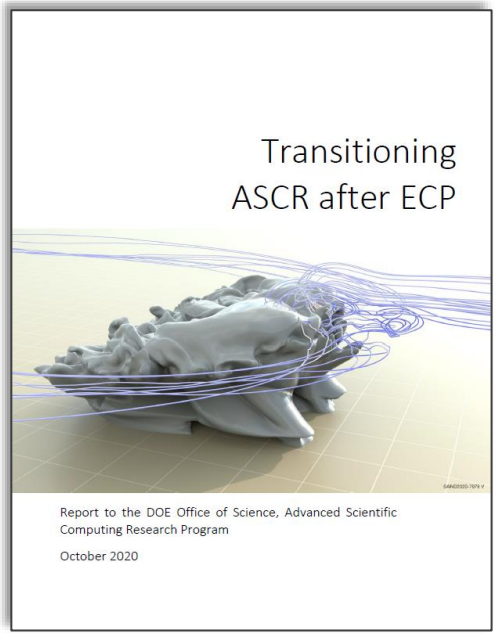

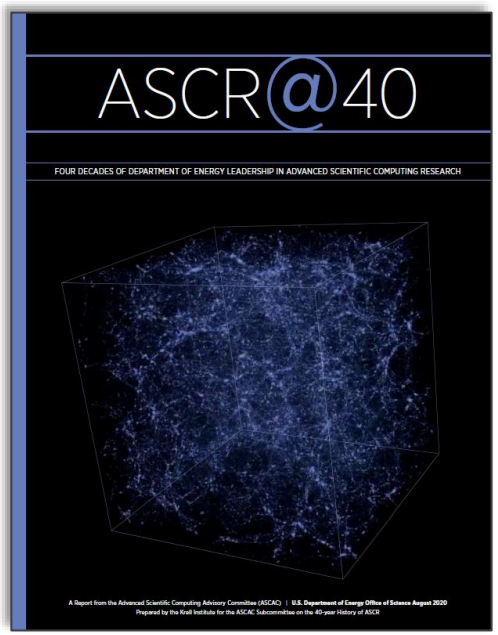

… plus numerous ASCR, SC Program, and community workshop reports of the last few years.

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

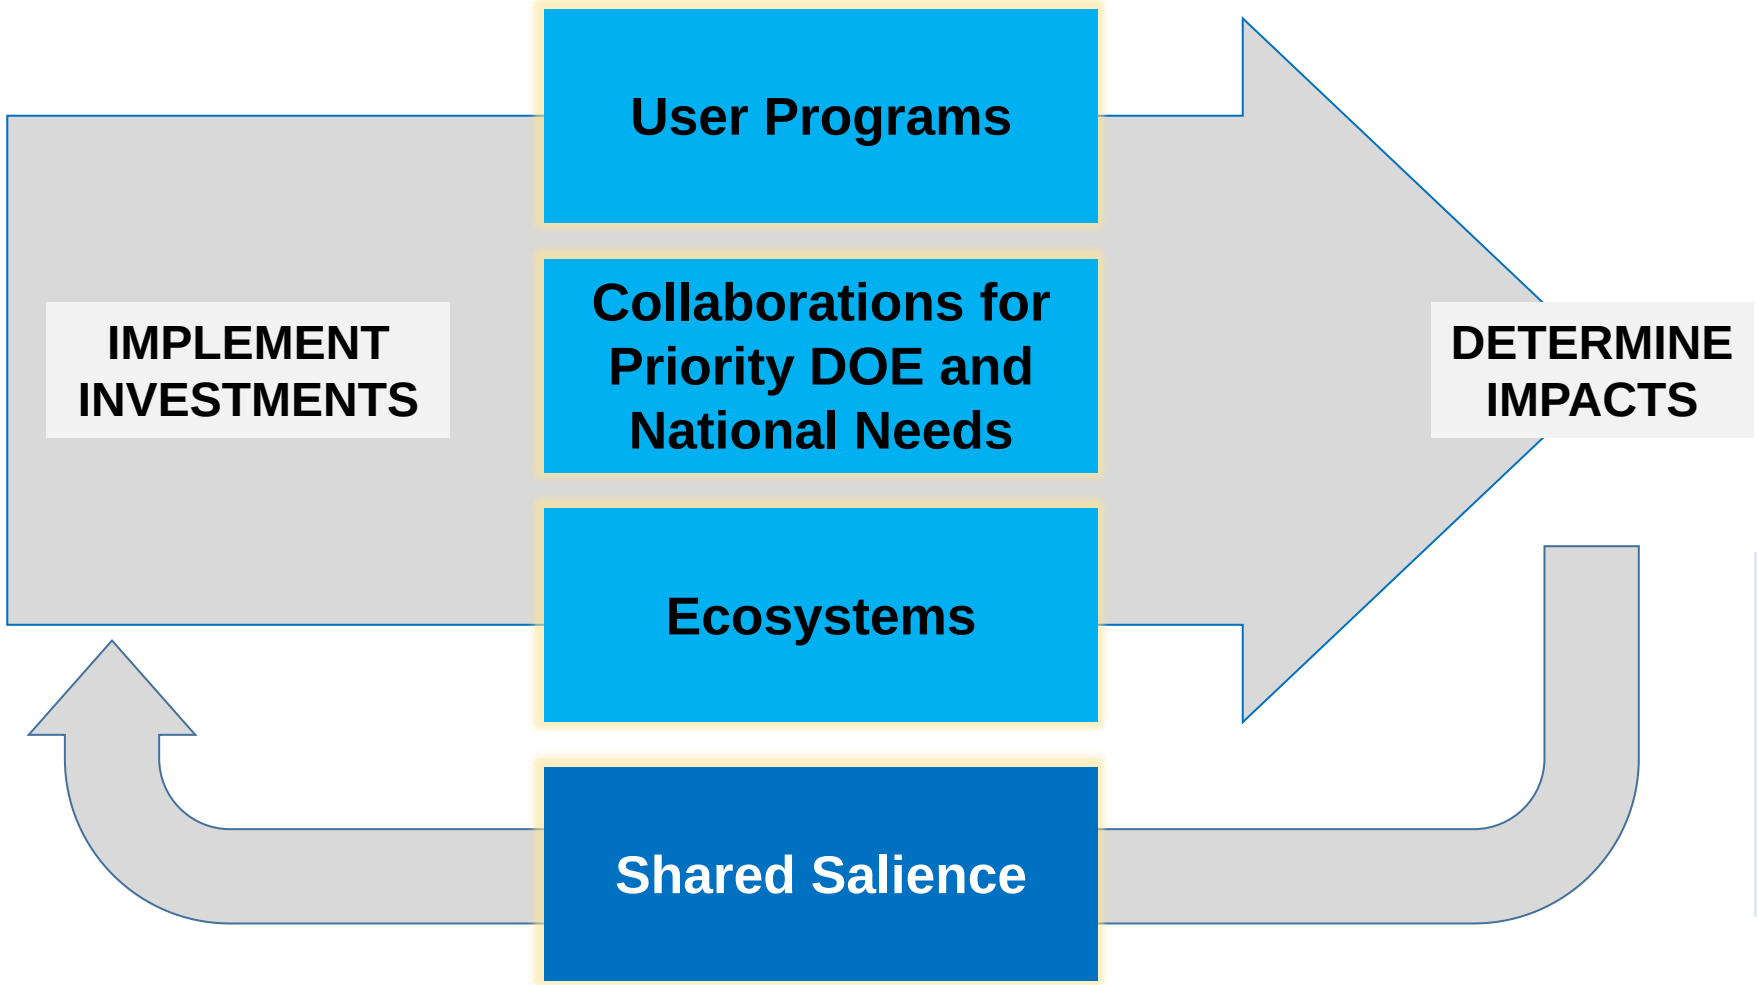

Engage & surface.

Listen & reflect.

Address & implement.

Repeat.

Vision: Clearly articulate and distinguish our core activities

Core activity	Current state	Future state
User Programs	HPC: visible portfolio (INCITE, ALCC, ERCAP) ESnet: visible Site User policy framework	► Continue to broaden the user base. ► Extend the reach of impact to an ever wider range of domains, institutions, and backgrounds.
Collaborations for Priority DOE and National Needs	Collaborations are embedded in core operations and user programs.	► Define the portfolio and make it visible. ► Consider new models (campaigns?) that deliver deep, lasting impact.
Ecosystems	Ecosystems are incipient, but not yet defined	► Define and make visible. Foster stewardship and governance models that provide the foundations for community innovation. ► Inspire participation and sustain careers.
Shared Salience	We have a portfolio of effective, but separate, activities with various stakeholder communities. Coordination is tacit. Synthesis is post facto.	► Make synthesis visible and annual. ► Capture, structure, and aggregate insight “as we go” from a variety of sources.

Incipient ecosystem: Office of Science User Facilities

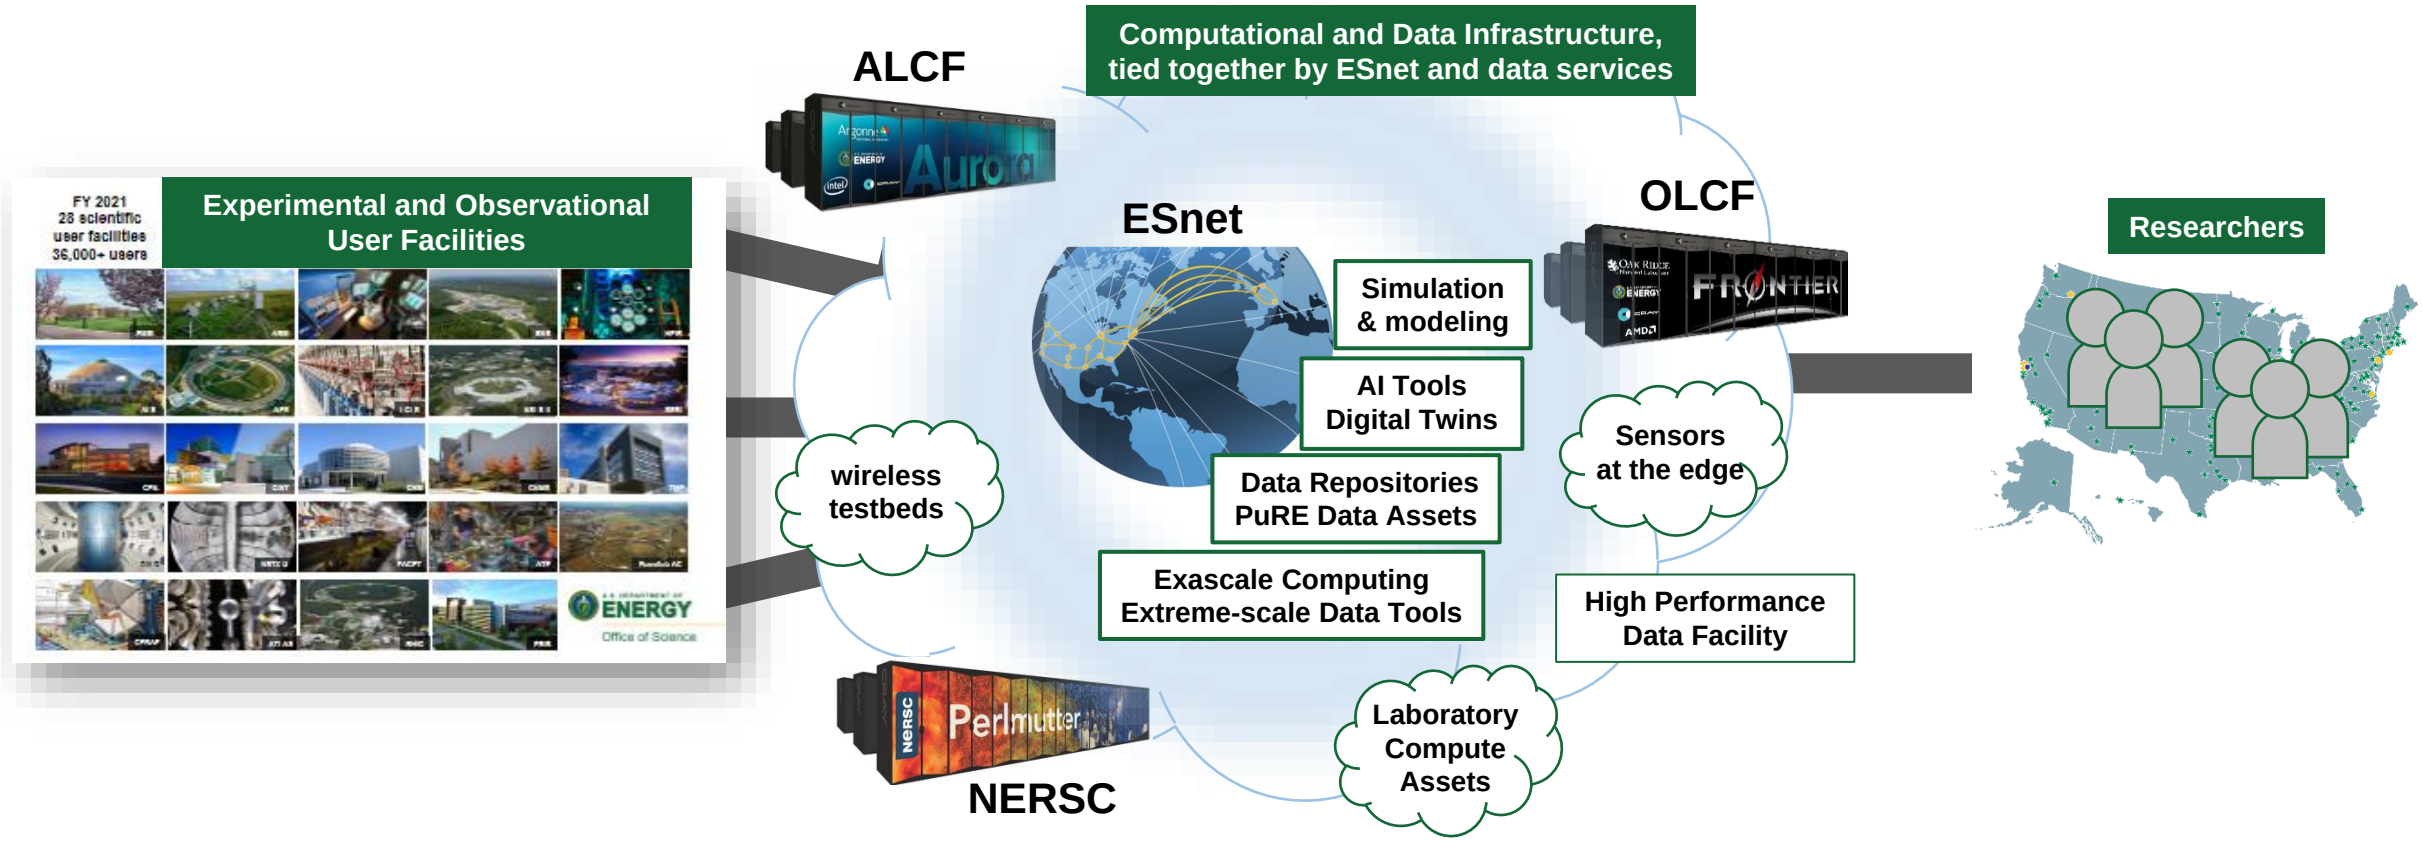

The ASCR Facilities are contemplating the operational implications of integration together and with experimental/observational facilities

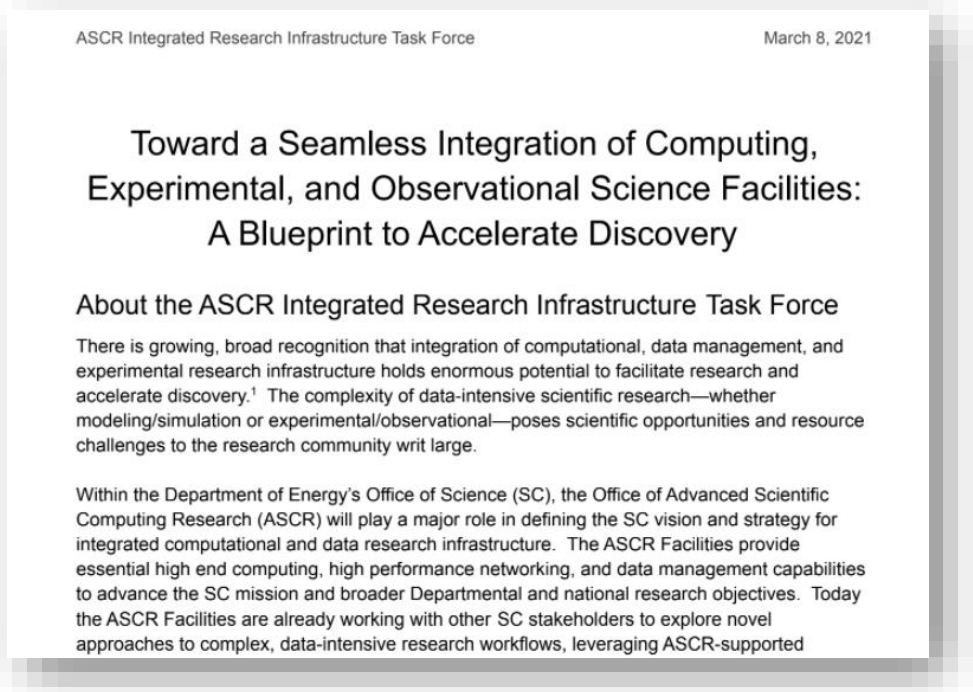

ASCR Integrated Research Infrastructure Task Force March 8, 2021

Toward a Seamless Integration of Computing, Experimental, and Observational Science Facilities: A Blueprint to Accelerate Discovery

About the ASCR Integrated Research Infrastructure Task Force

There is growing, broad recognition that integration of computational, data management, and experimental research infrastructure holds enormous potential to facilitate research and accelerate discovery.[1] The complexity of data-intensive scientific research—whether modeling/simulation or experimental/observational—poses scientific opportunities and resource challenges to the research community writ large.

Within the Department of Energy's Office of Science (SC), the Office of Advanced Scientific Computing Research (ASCR) will play a major role in defining the SC vision and strategy for integrated computational and data research infrastructure. The ASCR Facilities provide essential high end computing, high performance networking, and data management capabilities to advance the SC mission and broader Departmental and national research objectives. Today the ASCR Facilities are already working with other SC stakeholders to explore novel approaches to complex, data-intensive research workflows, leveraging ASCR-supported

Corey Adams
Katie Antypas
Debbie Bard
Shane Canon
Eli Dart
Chin Guok
Ezra Kissel
Eric Lancon
Bronson Messer
Sarp Oral
Jini Ramprakash
Arjun Shankar
Tom Uram

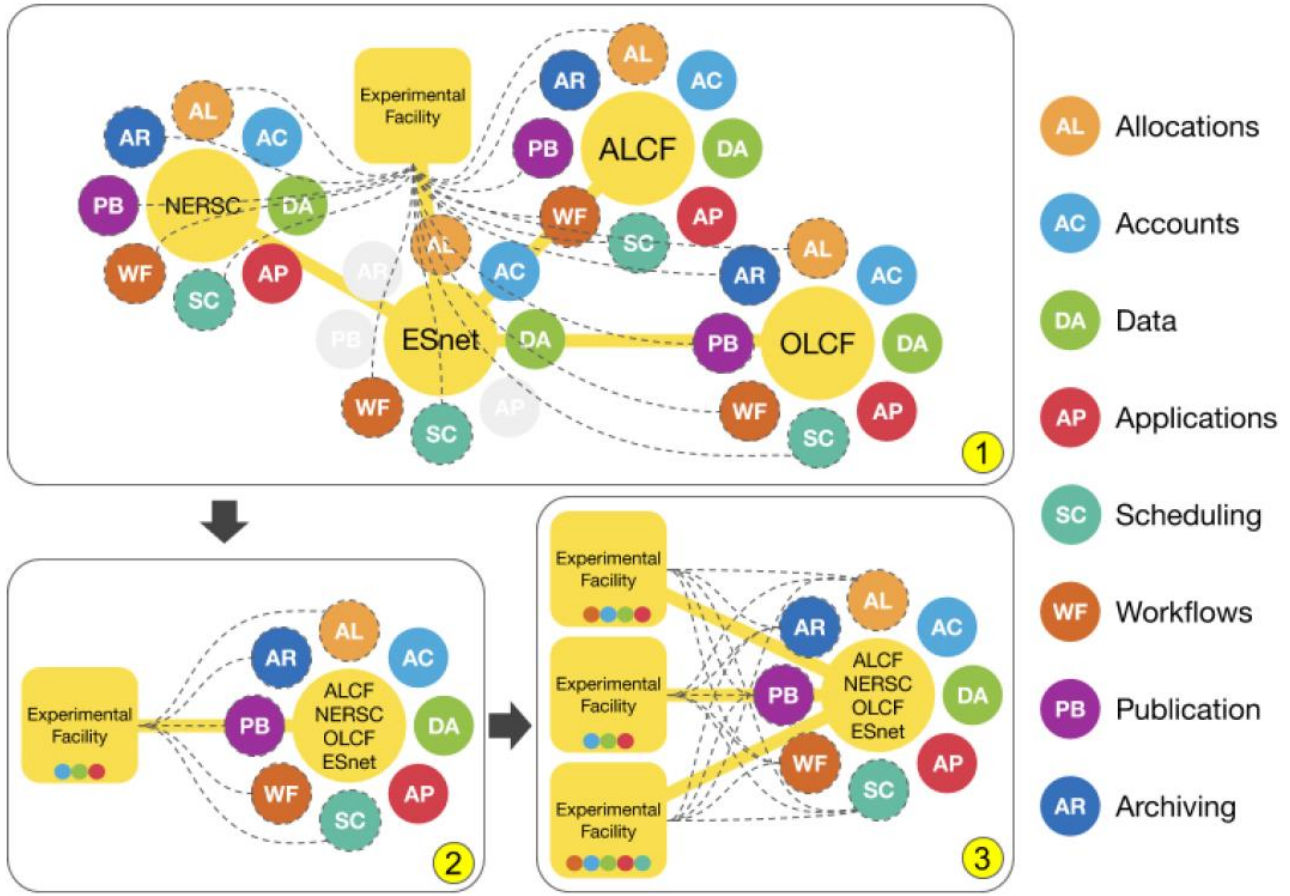

Figure 1. *Depiction of the integration of experimental facilities with computational facilities, across the range of services provided, in contrast with the one-to-one approach required today. 1. Today, an experimental facility must arrange separate bespoke interactions with individual HPC/HPN facilities. 2. A future paradigm with common interfaces could simplify integration of an experimental facility with multiple HPC/HPN facilities. 3. In turn, these common interfaces could support expansion and integration across multiple experimental facilities and HPC/HPN facilities.*

U.S. DEPARTMENT OF ENERGY | Office of Science

Incipient ecosystem: Scientific software

EXASCALE COMPUTING PROJECT

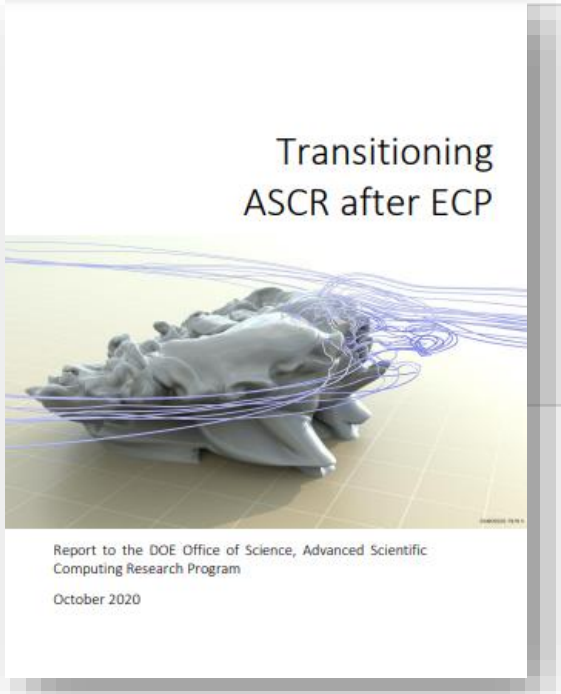

"ECP has created a well-designed software ecosystem for development, curation, and distribution of exascale systems and application software. This ecosystem integrates the fruits of years of basic research in: mathematics, computer science, applications, and systems software.

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

U.S. DEPARTMENT OF ENERGY | Office of Science

The Exascale Requirements Reviews yielded the Exascale Crosscut Report, which synthesized insight, creating shared salience.

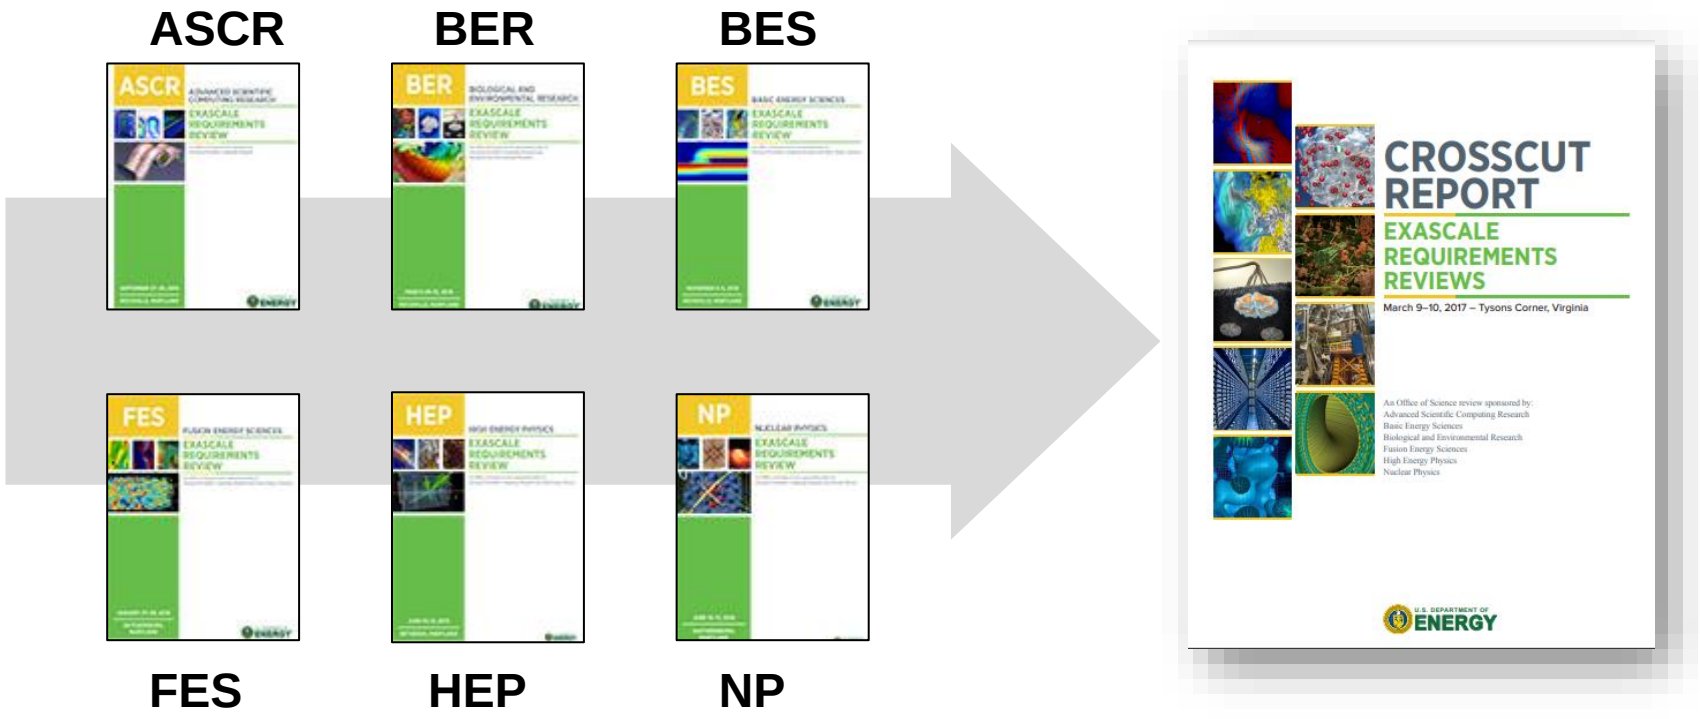

The Crosscut Report, in addition to its HPC focus, yielded deep strategic insight across a range of challenges, including data-intensive workflows, AI/ML, networking, and more.

The Office of Science User Facilities: A unique community

These core principles guide operations and oversight:

Open

The facility is open to all interested potential users without regard to nationality or institutional affiliation.

Accessible

The facility provides resources sufficient for users to conduct work safely and efficiently.

Competitive

Allocation of facility resources is determined by merit review of the proposed work.

Unique

The facility capability does not compete with an available private sector capability.

Free

User fees are not charged for non-proprietary work if the user intends to publish the research results in the open literature. Full cost recovery is required for proprietary work.

Collaborative

The facility supports a formal user organization to represent the users and facilitate sharing of information, forming collaborations, and organizing research efforts among users.

Building a community of practice through shared insight

December 2020 Roundtable

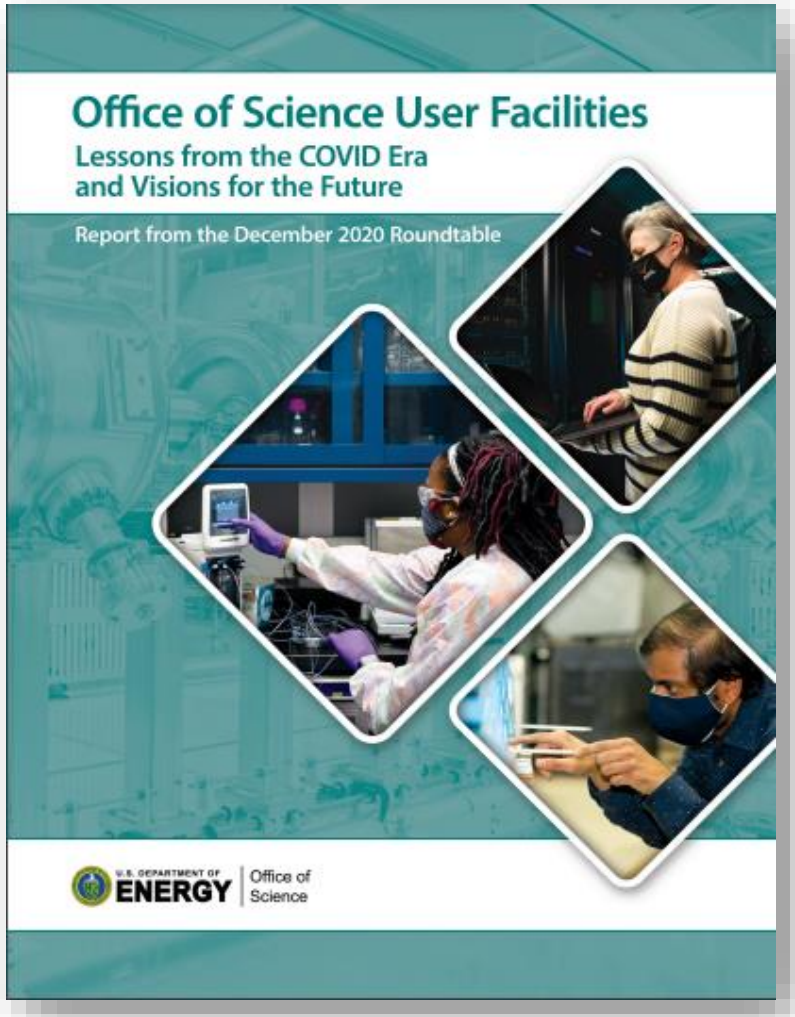

Future All-Facility Meetings

"Facility directors have expressed an interest in hearing about other facilities' challenges and strategies. During the SC roundtable, participants noted and appreciated that, for the first time, all SC user facilities were involved in the same meeting. The exchange of information was beneficial and enlightening, allowing facility directors and staff to better understand the entire DOE system, the challenges faced by other facilities, and the strategies employed to meet those challenges."

https://doi.org/10.2172/1785683

Vision for the ASCR Facilities: How we will thrive together

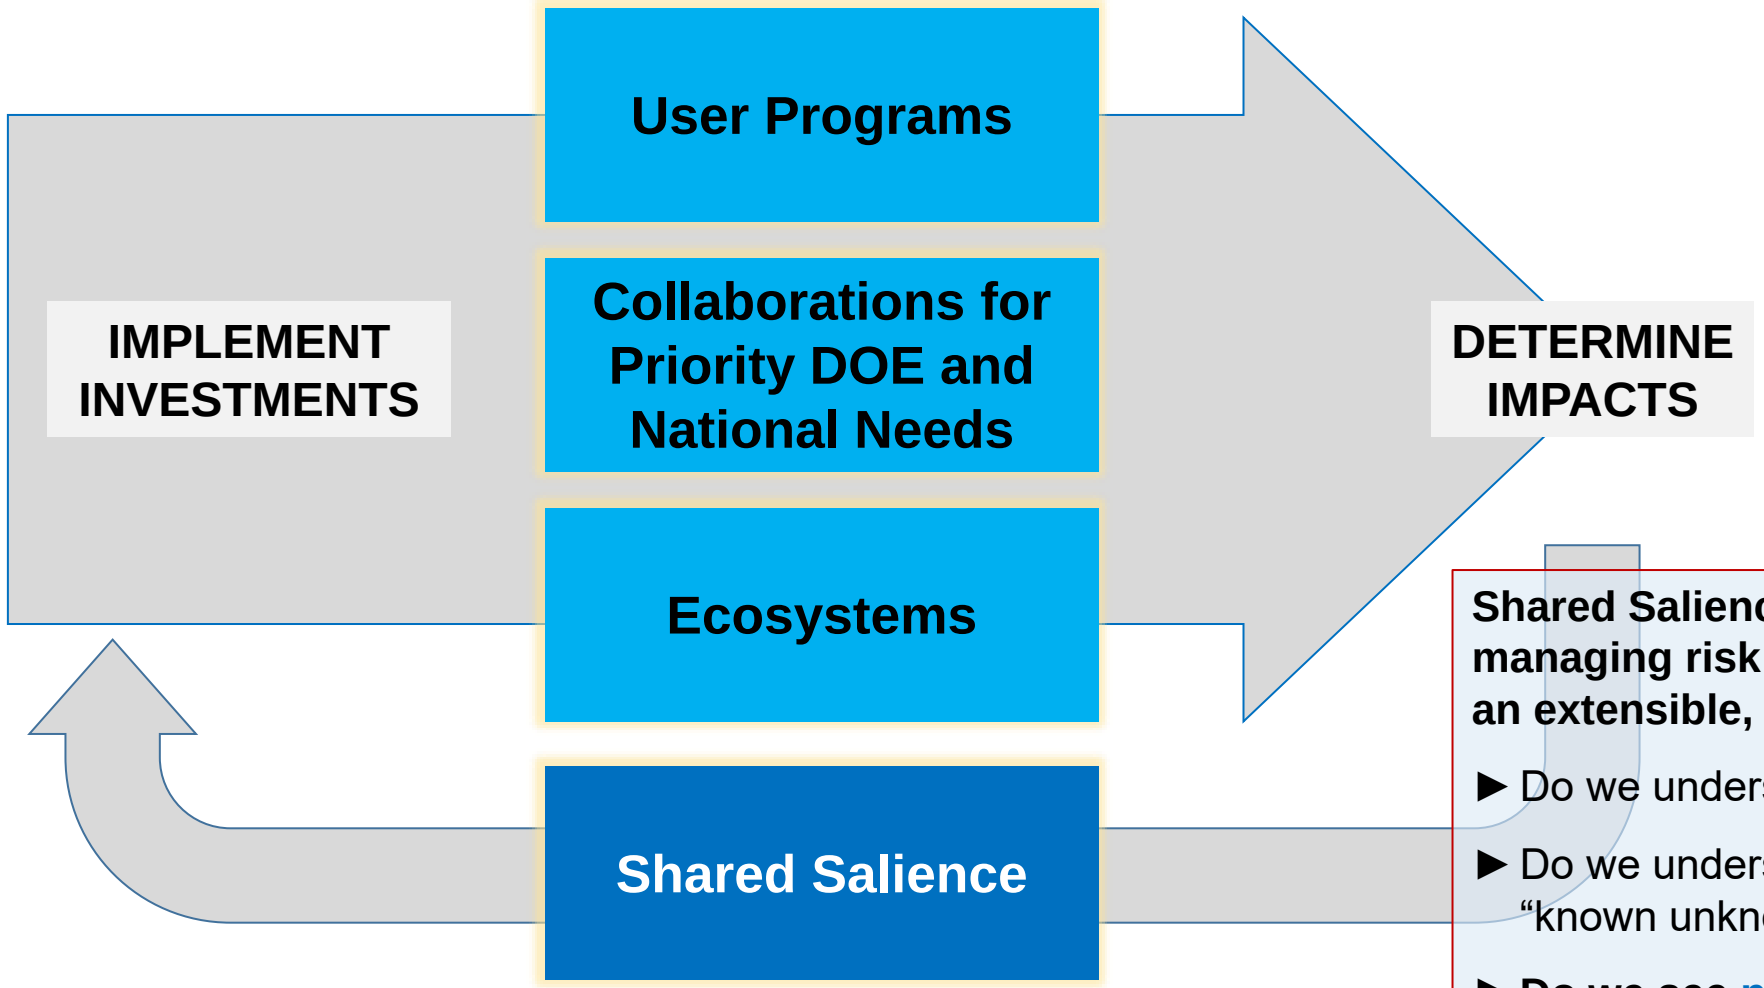

Engage & surface.

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?**

The arc of DOE history
bends towards ecosystems.

DOE originates from the Manhattan Project

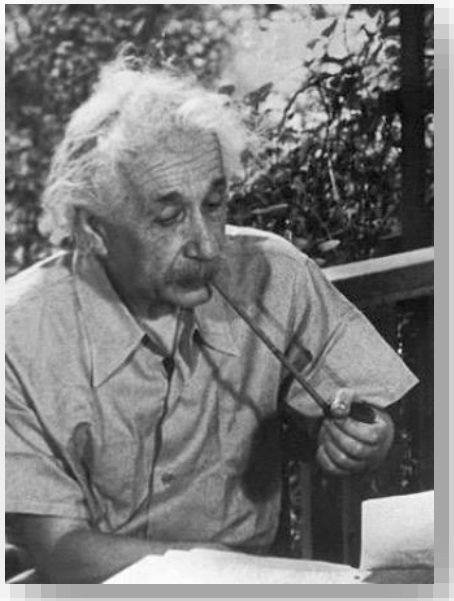

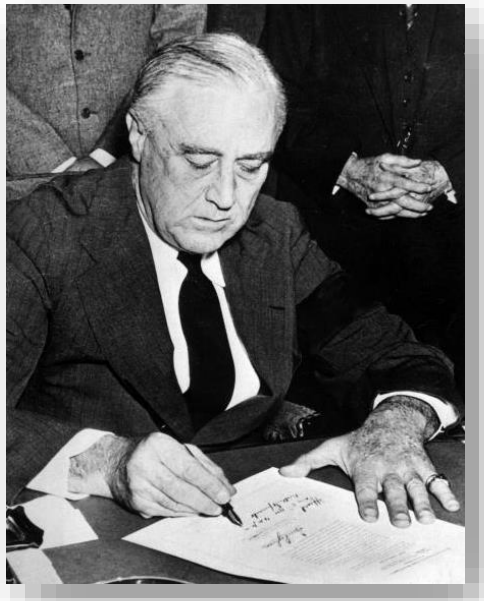

Manhattan Project sites

DOE Era 0

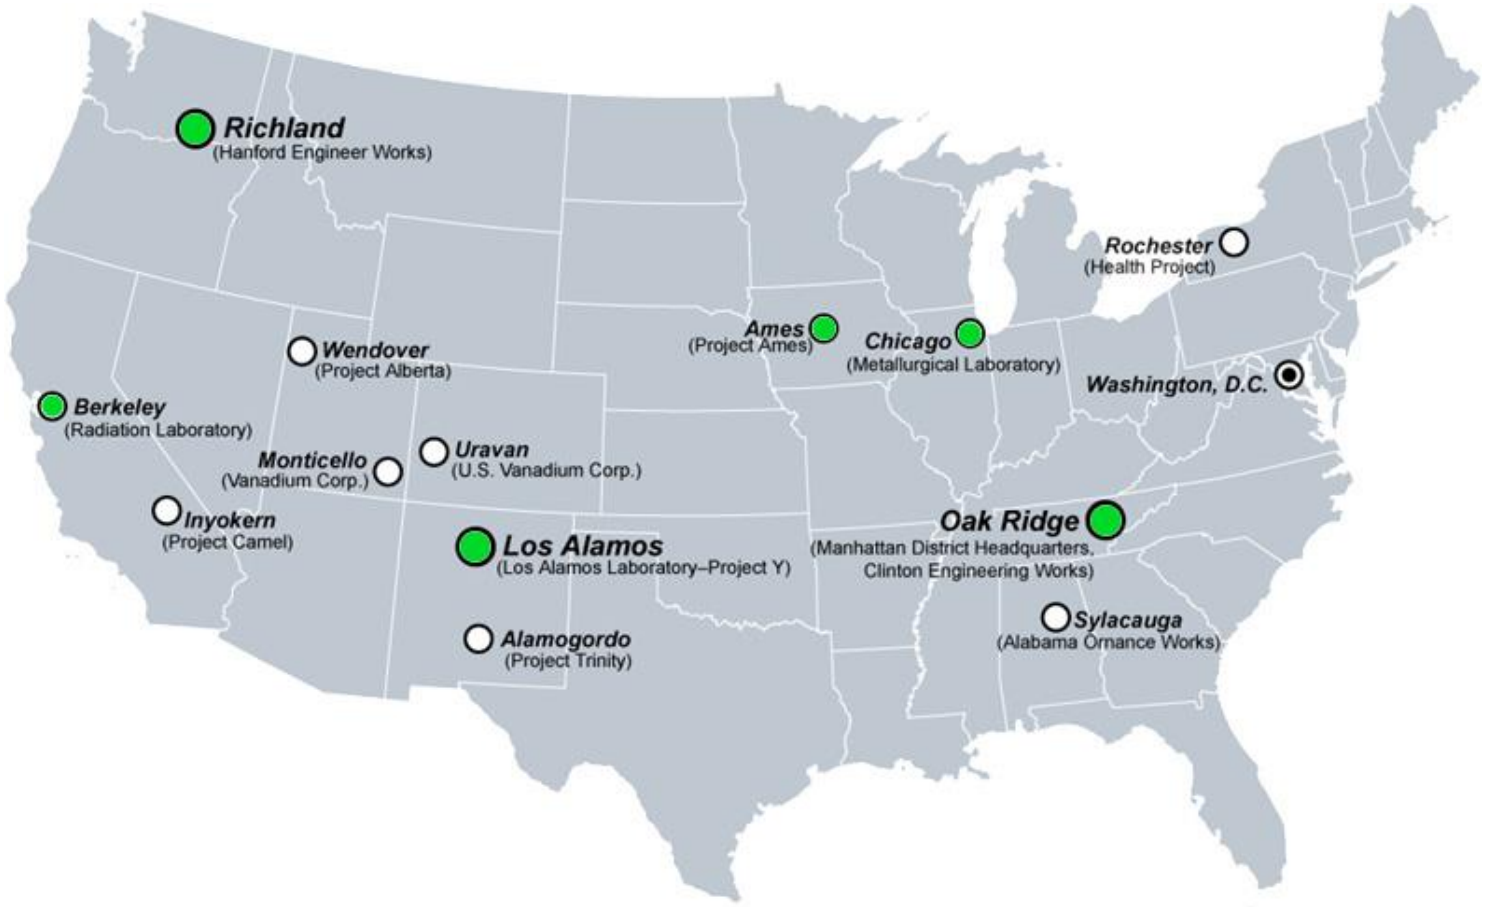

The dawn of the DOE (AEC) national laboratory system and “big science” research infrastructure, built primarily for intramural use

DOE Era 1

The DOE National Laboratories evolve to become more open, exemplified by open-access scientific user facilities

DOE Era 2

Office of Science Laboratories

1. **Ames Laboratory**
   Ames, Iowa
2. **Argonne National Laboratory**
   Argonne, Illinois
3. **Brookhaven National Laboratory**
   Upton, New York
4. **Fermi National Accelerator Laboratory**
   Batavia, Illinois
5. **Lawrence Berkeley National Laboratory**
   Berkeley, California
6. **Oak Ridge National Laboratory**
   Oak Ridge, Tennessee
7. **Pacific Northwest National Laboratory**
   Richland, Washington
8. **Princeton Plasma Physics Laboratory**
   Princeton, New Jersey
9. **SLAC National Accelerator Laboratory**
   Menlo Park, California
10. **Thomas Jefferson National Accelerator Facility**
    Newport News, Virginia

Other DOE Laboratories

1. **Idaho National Laboratory**
   Idaho Falls, Idaho
2. **National Energy Technology Laboratory**
   Morgantown, West Virginia
   Pittsburgh, Pennsylvania
   Albany, Oregon
3. **National Renewable Energy Laboratory**
   Golden, Colorado
4. **Savannah River National Laboratory**
   Aiken, South Carolina

NNSA Laboratories

1. **Lawrence Livermore National Laboratory**
   Livermore, California
2. **Los Alamos National Laboratory**
   Los Alamos, New Mexico
3. **Sandia National Laboratory**
   Albuquerque, New Mexico
   Livermore, California

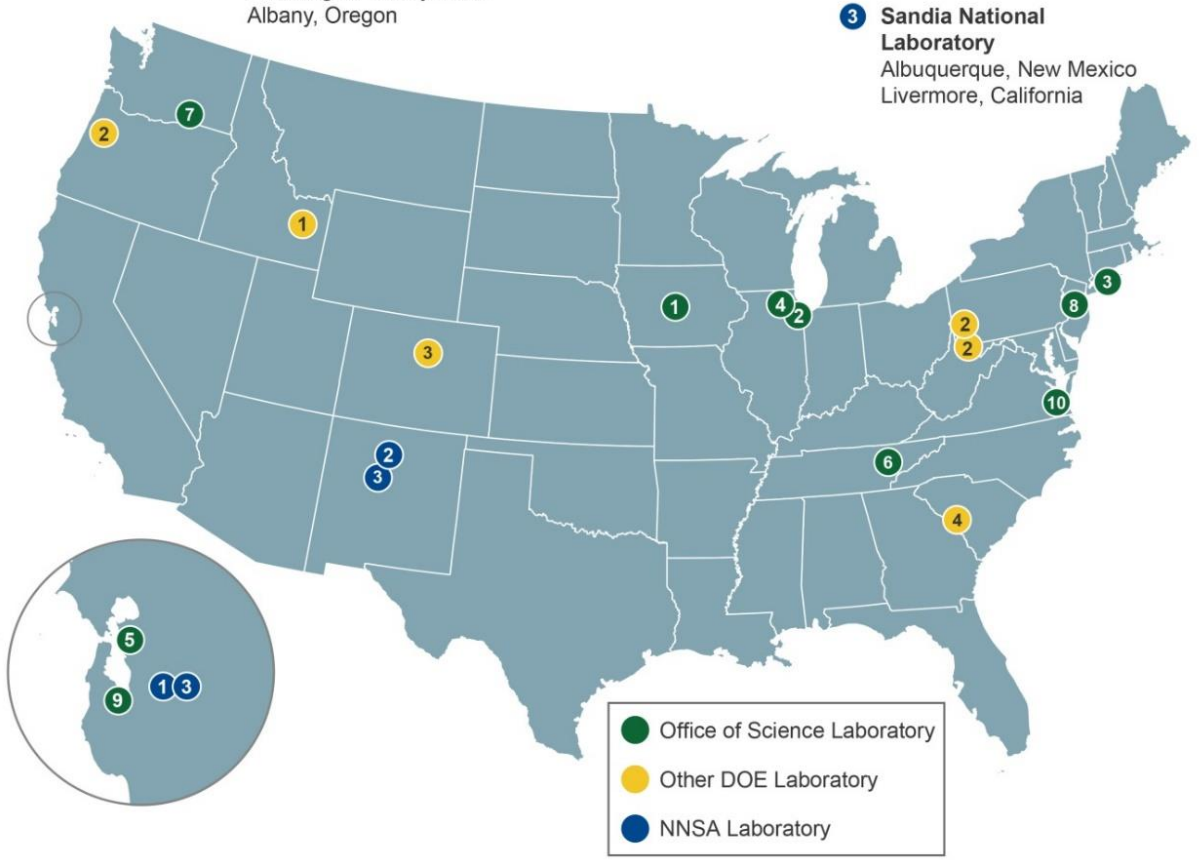

U.S. DEPARTMENT OF ENERGY

ce of
ence

The national laboratories will become an open innovation ecosystem of research platforms, connected together

ECP EXASCALE COMPUTING PROJECT

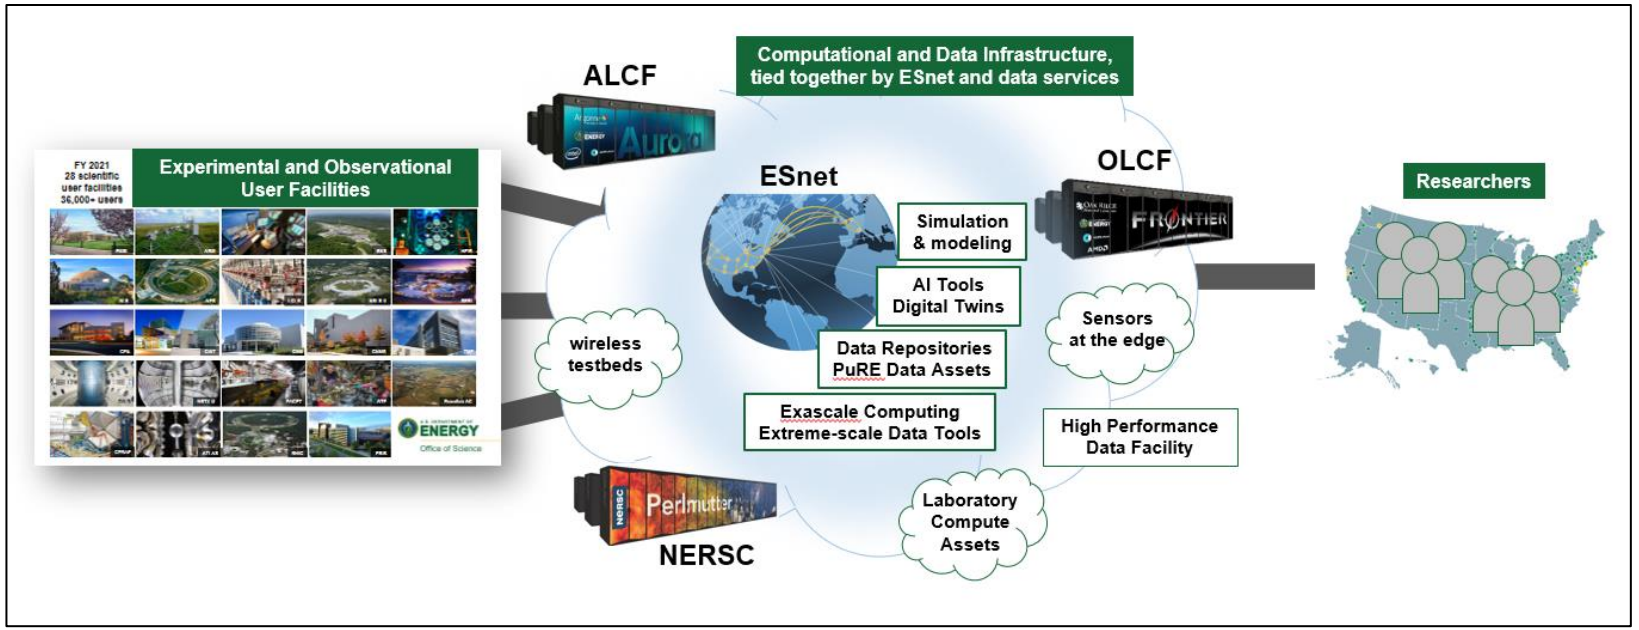

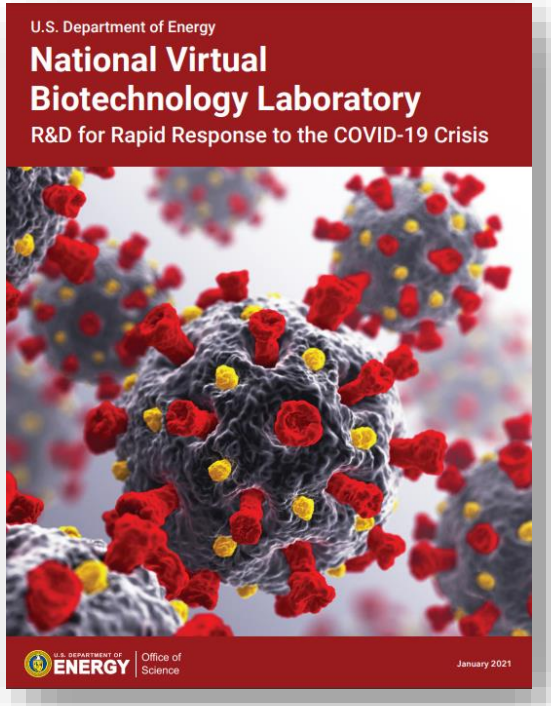

DOE's next Era

U.S. DEPARTMENT OF ENERGY | Office of Science

The ASCR Facilities enterprise will be an engine of insight, innovation, and impact for the next era.

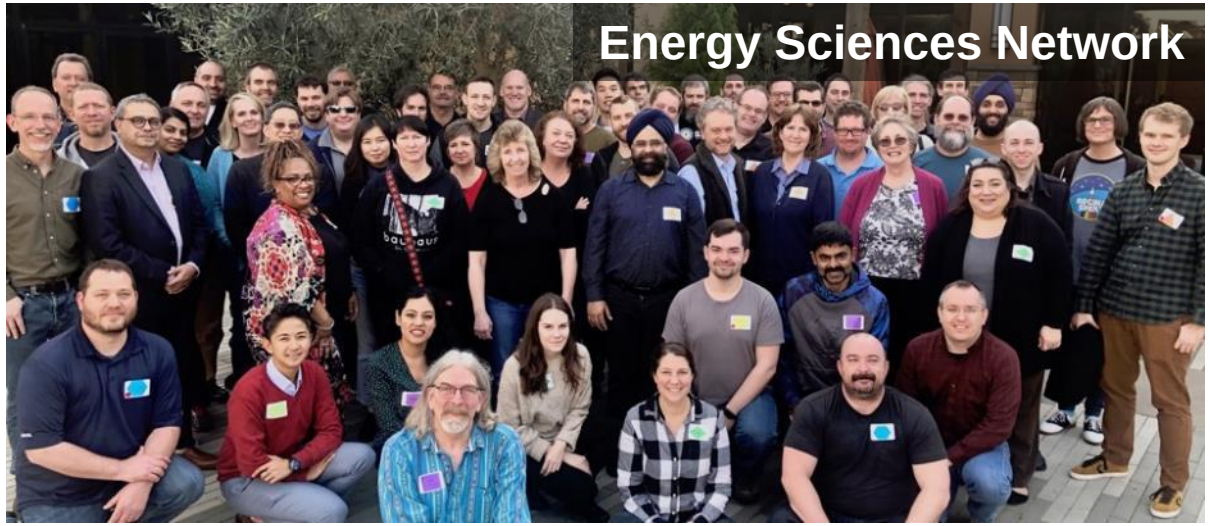

Let’s see how far we can go together.