DOE's Integrated Research Infrastructure (IRI): Vision, Strategy, and Implementation

Advanced Scientific Computing Advisory Committee Meeting
June 12, 2023

Ben Brown
Director, Facilities Division
Advanced Scientific Computing Research (ASCR)

We are beginning implementation of DOE's Integrated Research Infrastructure (IRI)

Vision: developed and honed in 2020-2021

Strategy: the IRI Framework developed in 2022

Implementation: 2023 efforts and beyond

Coordination between ASCR Facilities and ASCR Research is essential for the success of this long-term endeavor.

DOE is positioned to lead the new era of integrated science within the US Government and the world.

Linking distributed resources is becoming paramount to modern collaborative science.

The challenges of our time call upon DOE and its national laboratories to be an open innovation ecosystem:

Accelerating discovery & innovation

Democratizing access

Drawing new talent

Advancing open science

U.S. DEPARTMENT OF ENERGY | Office of Science

A recent hand-crafted IRI case: COVID vaccine research at LCLS

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.

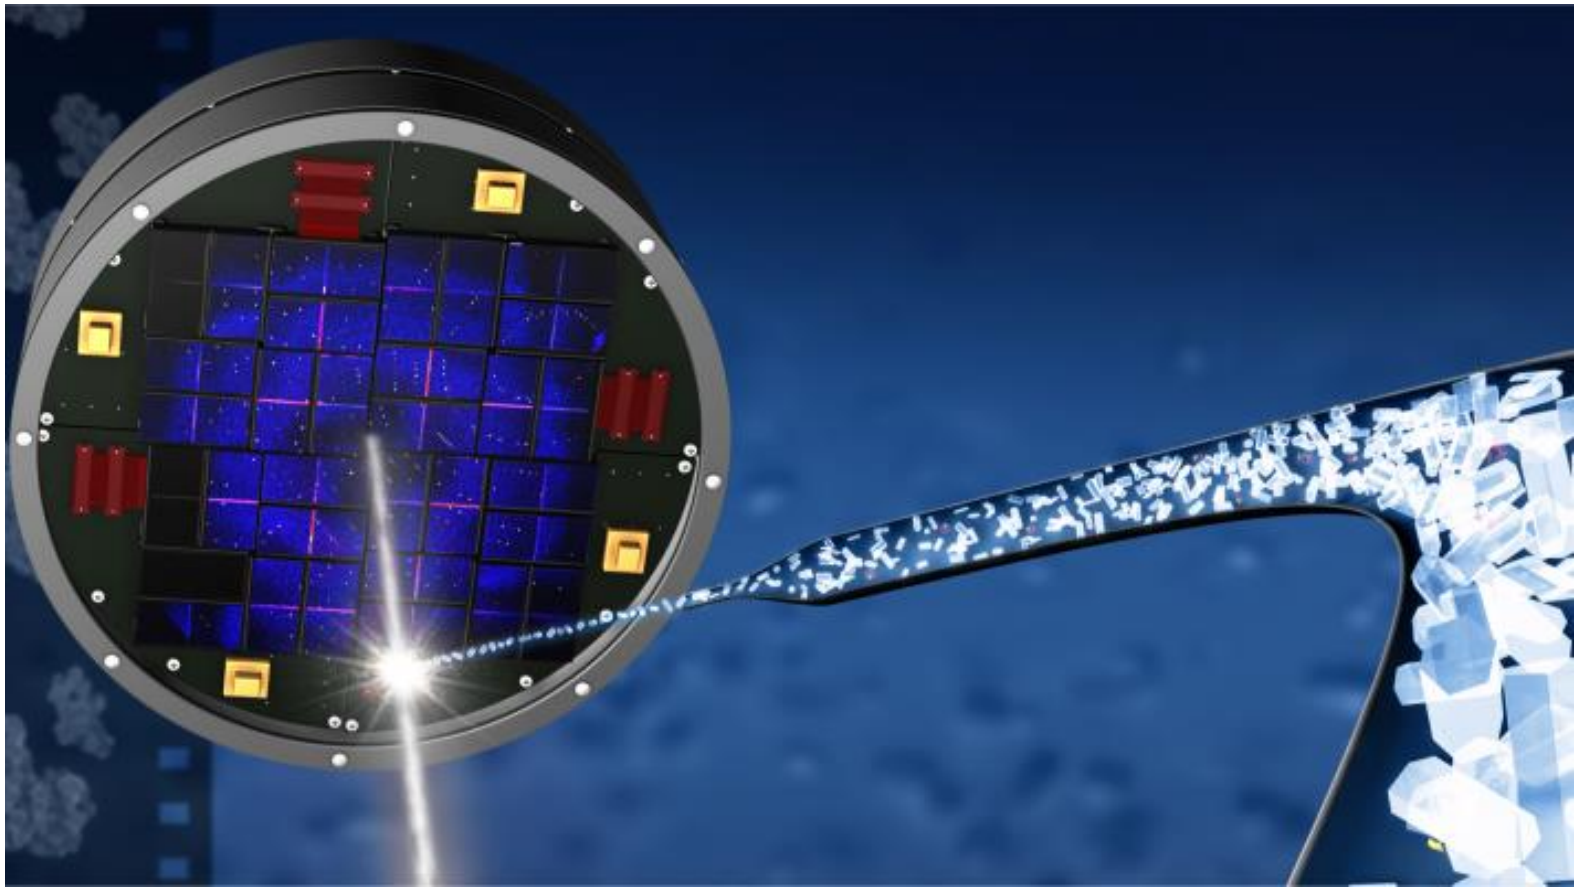

SLAC National Accelerator Laboratory

Linac Coherent Light Source at SLAC

NERSC at LBNL

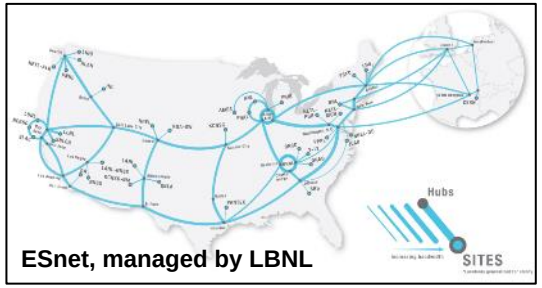

ESnet, managed by LBNL

U.S. DEPARTMENT OF ENERGY | Office of Science

What will future automated IRI look like? Here is but one example…

Compute Reservations via API

Paradigm buster!

Q: How do we provide real-time compute *without* lowering utilization?
A: Policy-driven compute reservations, requested via API from EOD workflows.

One brief moment in the life of the Superfacility API . . .

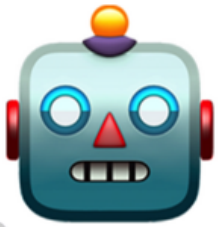

GIVE ME ALL OF THE GPU NODES NEXT WEDNESDAY.

REQUEST DENIED DUE TO GLUTTONY.

MY APOLOGIES. INSTEAD GIVE ME 10 GPU NODES NEXT WEDNESDAY AT 2:00 FOR 4 HOURS.

YOUR RESERVATION IS 8D7A392D. YOU MAY SUBMIT JOBS INTO IT UNTIL NEXT WEDNESDAY AT 5:59.

THANK YOU AND HAVE A NICE HUMAN-DAY.

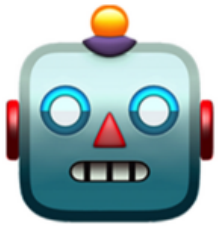

User Workflow Manager

Superfacility API

Linac Coherent Light Source at SLAC

NERSC at LBNL

4

BERKELEY LAB
Bringing Science Solutions to the World

+ Data transport reservations via API

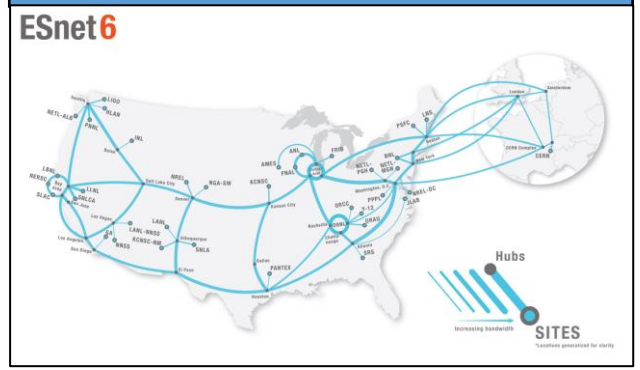

Example from Cory Snavely, NERSC
EOD = experimental and observational data users

DOE's Integrated Research Infrastructure (IRI) Vision:

To empower researchers to meld DOE's world-class research tools, infrastructure, and user facilities seamlessly and securely in novel ways to radically accelerate discovery and innovation

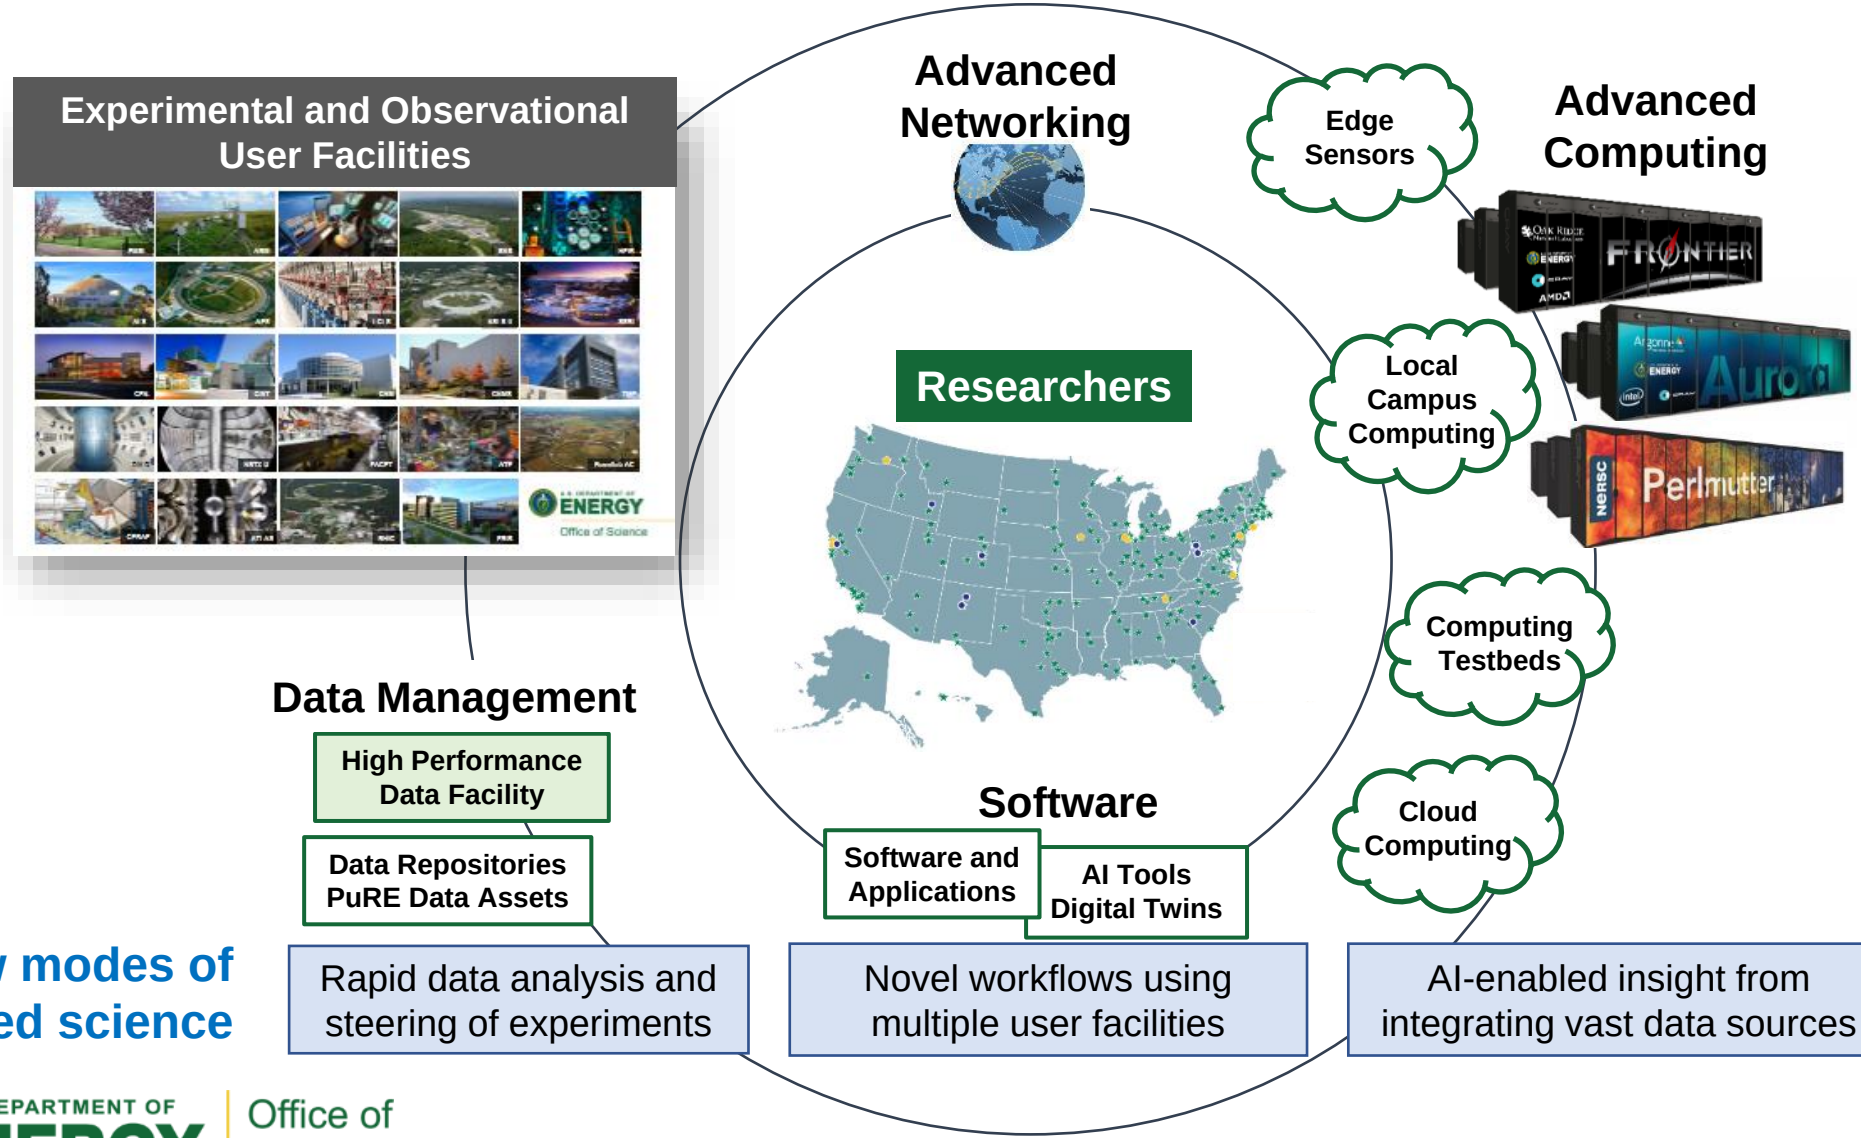

U.S. DEPARTMENT OF ENERGY | Office of Science

DOE's Integrated Research Infrastructure (IRI) Vision:
To empower researchers to meld DOE's world-class research tools, infrastructure, and user facilities seamlessly and securely in novel ways to radically accelerate discovery and innovation

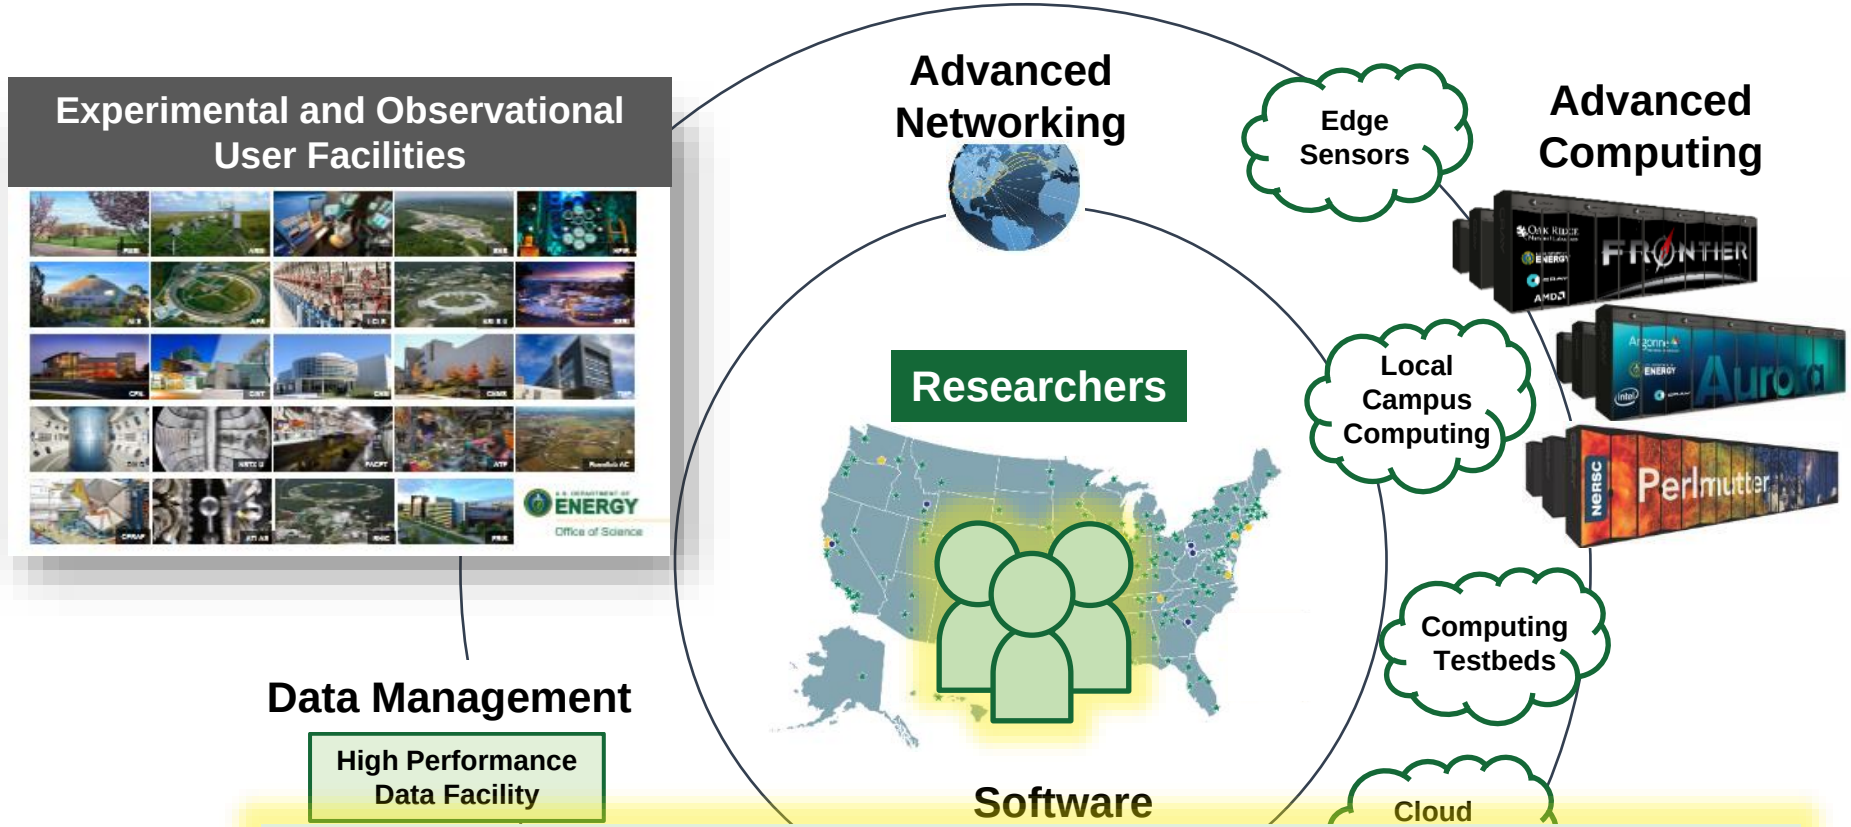

New modes of integrated science

The IRI Vision:
It's about empowering people.
It's about data.

U.S. DEPARTMENT OF ENERGY

IRI Mission Needs & Drivers (a few key ones)

Exascale science era

Experimental data deluge: advances in source/detector technology

Observational data deluge: wireless & edge platforms

AI for Science, Energy, and Security

…… virtually every major research challenge is calling on DOE to be more than the sum of its constituent labs, programs, facilities ……

Timeline of key IRI activities and reports, 2020-23

FY 2021 President's Budget Request includes **Integrated** Computation and Data Infrastructure Initiative

ASCR IRI Task Force launch

ASCR IRI Task Force report

SC IRI Blueprint Activity launch

IRI Blueprint Activity results

FY 2024 PBR advances IRI and the High Performance Data Facility

IRI Blueprint Activity report

IRI Program Development

Jan **2020** Jan **2021** Jan **2022** Jan **2023** Jan **2024**

Integration of instrumentation, data, and computing infrastructure are essential requirements for national R&D objectives

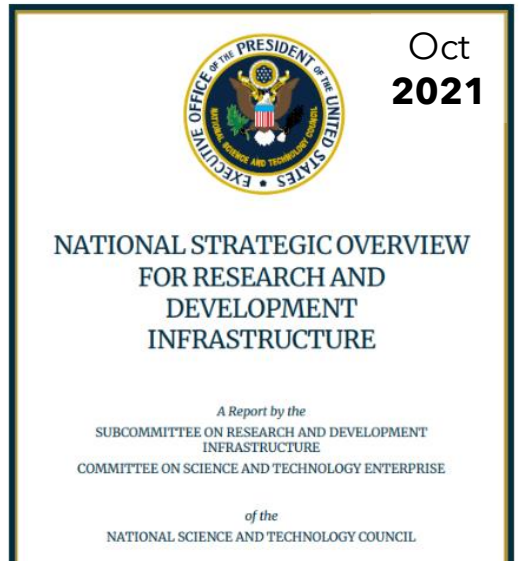

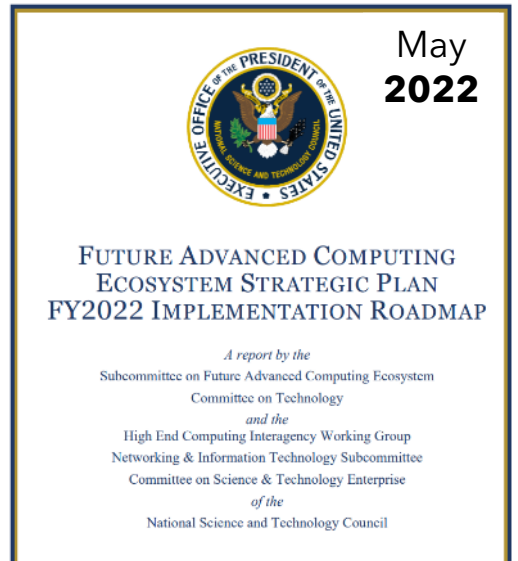

U.S. DEPARTMENT OF ENERGY | Office of Science

The ASCR IRI Task Force aligned the ASCR Facilities around the IRI vision and articulated core IRI design principles

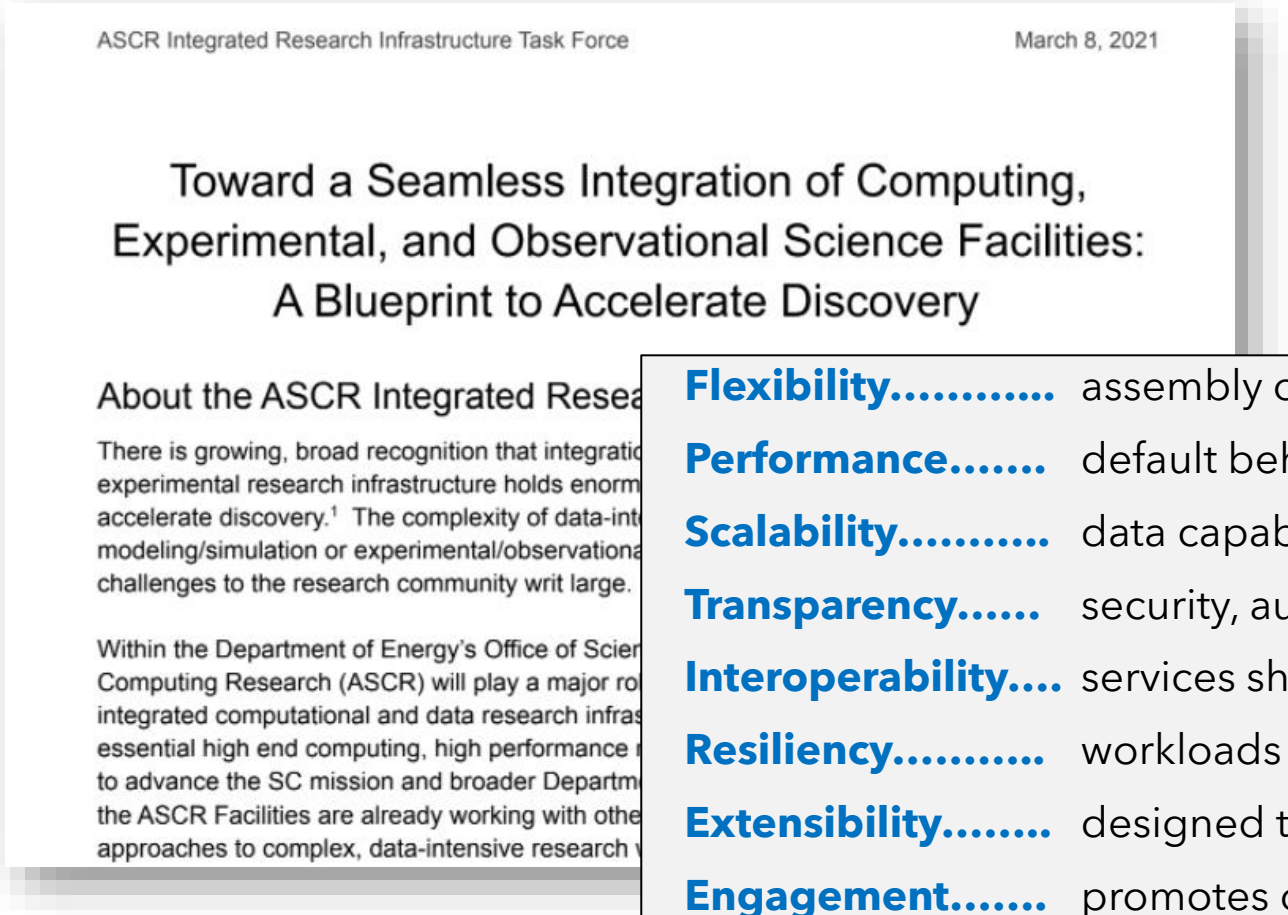

ASCR Integrated Research Infrastructure Task Force March 8, 2021

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

About the ASCR Integrated Resea

There is growing, broad recognition that integratic
experimental research infrastructure holds enorm
accelerate discovery.[1] The complexity of data-int
modeling/simulation or experimental/observationa
challenges to the research community writ large.

Within the Department of Energy's Office of Scier
Computing Research (ASCR) will play a major ro
integrated computational and data research infras
essential high end computing, high performance
to advance the SC mission and broader Departm
the ASCR Facilities are already working with othe
approaches to complex, data-intensive research

Flexibility.....	assembly of resource workflows is facile; complexity is concealed
Performance.....	default behavior is performant, without arcane requirements
Scalability.....	data capabilities without excessive customizations
Transparency.....	security, authentication, authorization should support automation
Interoperability....	services should extend outside the DOE environment
Resiliency.....	workloads are sustained across planned and unplanned events
Extensibility.....	designed to adapt and grow to meet unknown future needs
Engagement.....	promotes co-design, cooperation, partnership
Cybersecurity.....	security for facilities and users is essential

The IRI Blueprint Activity created a framework for IRI implementation

The IRI Framework comprises:

- **IRI Science Patterns (3)** represent integrated science use cases across DOE science domains.
  - **Provide the basis for organizing diverse program requirements into strategic priorities.**

- **IRI Practice Areas (6)** represent critical topics that require close coordination to realize and sustain a thriving IRI ecosystem across DOE institutions.
  - **Provide the basis for organizing the program governance model and cross-cutting efforts.**

Convened over **150 DOE national laboratory experts** from **all 28 SC user facilities** across **13 national laboratories** to consider the **technological, policy, and sociological challenges** to implementing IRI.

The IRI Blueprint Activity created a framework for IRI implementation

IRI Science Patterns (3)

Time-sensitive pattern has *urgency*, requiring real-time or end-to-end performance with high reliability, e.g., for timely decision-making, experiment steering, and virtual proximity.

Data integration-intensive pattern requires combining and analyzing data from multiple sources, e.g., sites, experiments, and/or computational runs.

Long-term campaign pattern requires sustained access to resources over a long period to accomplish a well-defined objective.

Note: the patterns are not mutually exclusive.

IRI Practice Areas (6)

User experience practice will ensure relentless attention to user perspectives and needs through requirements gathering, user-centric (co)-design, continuous feedback, and other means.

Resource co-operations practice is focused on creating new modes of cooperation, collaboration, co-scheduling, and joint planning across facilities and DOE programs.

Cybersecurity and federated access practice is focused on creating novel solutions that enable seamless scientific collaboration within a secure and trusted IRI ecosystem.

Workflows, interfaces, and automation practice is focused on creating novel solutions that facilitate the dynamic assembly of components across facilities into end-to-end IRI pipelines.

Scientific data life cycle practice is focused on ensuring that users can manage their data and metadata across facilities from inception to curation, archiving, dissemination, and publication.

Portable/scalable solutions practice is focused on ensuring that transitions can be made across heterogeneous facilities (portability) and from smaller to larger resources (scalability).

The value propositions of an IRI Program

For the taxpayer:

Achieve greater productivity and avoid duplication of effort.

For the researcher:

Achieve transformational reduction in *time to insight* and *complexity*.

For program/RI/institutional leaders:

- **Achieve greater effectiveness and efficiency in coordinating efforts;**
- **Achieve more nimble solutions than would be possible alone;**
- **Gain leverage with partners who possess like requirements;**
- **Avoid single points of failure; and**
- **Gain access to expertise and shared experience.**

Implementing IRI: Today ASCR is taking the first big steps

1. **Invest in IRI foundational infrastructure**
2. **Bring existing IRI projects into formal coordination**
3. **Deploy an IRI Pathfinding Testbed across the four ASCR Facilities**
4. **Stand up an IRI Program structure (HQ and Field)**

1. Invest in IRI foundational infrastructure

Early focus: Report IRI Science Pattern requirements across all SC Programs and assess IRI readiness
Early focus: Advance the NERSC-10, High Performance Data Facility, and OLCF-6 projects to CD-1

All current and future ASCR Facilities projects are being conceived with IRI requirements in mind.

ASCR Facilities Program Element	Major Project(s)
High Performance Networking	ESnet6 ✓
High Performance Production Computing	NERSC-10, HPDF
Leadership Computing	OLCF-6, ALCF-4

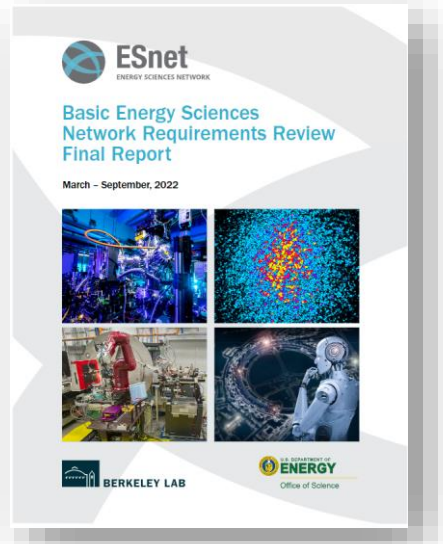

Search for "ESnet Requirements Reviews" and "NERSC-10 Technical Requirements"

IRI in the FY 2024 President's Budget Request

ASCR Facilities

"In FY 2024, the ASCR facilities will continue planning and begin implementation to advance DOE's Integrated Research Infrastructure (IRI) so that researchers can seamlessly and securely meld DOE's unique data, user facilities, and computing resources to accelerate discovery and innovation."

High Performance Data Facility (HPDF) project

"The proposed HPDF will serve as a foundational element in enabling the DOE Integrated Research Infrastructure; will provide crucial resources to Office of Science programs to attack fundamental problems in science and engineering that require nimble shared access to large data sets, increasingly aggregated from multiple sources; will partner and operate in concert with other ASCR Facilities and potentially other DOE laboratory computing resource providers to provide a high availability high performance computing ecosystem for a wide variety of applications; will serve as a 'Hub' enabling 'Spoke' sites to deploy and orchestrate distributed infrastructure to enable high priority DOE mission applications."

See DOE Lab Funding Announcement LAB 23-3020 for more information.

https://science.osti.gov/grants/Lab-Announcements/Open

2. Bring existing IRI projects into formal coordination

Early focus: Move towards a common user experience across the ASCR HPC Facilities
Early focus: Create an IRI HPC allocation

LBNL's Superfacility project
ORNL's INTERSECT initiative
ANL's ALCF-APS Nexus project

These are all separate initiatives with similar integration goals.
Let's now row in the same direction!

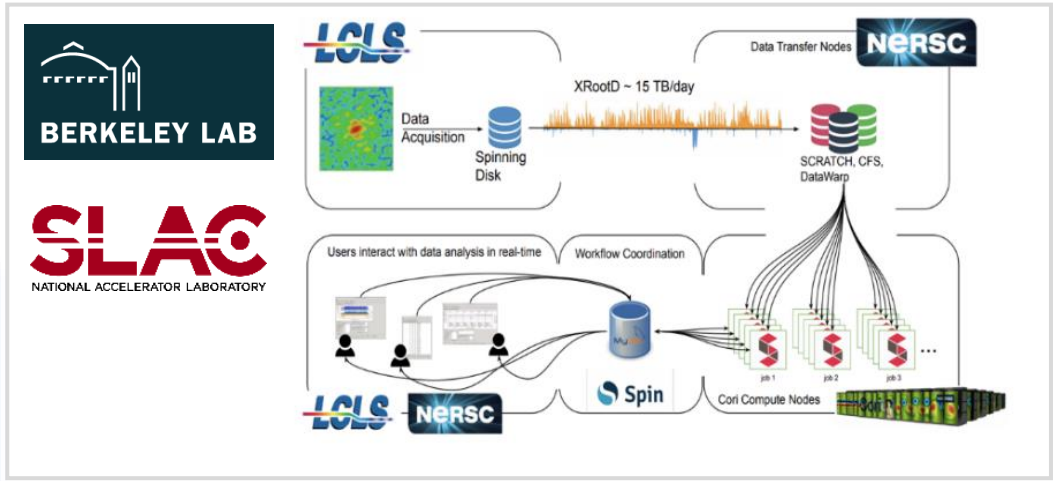

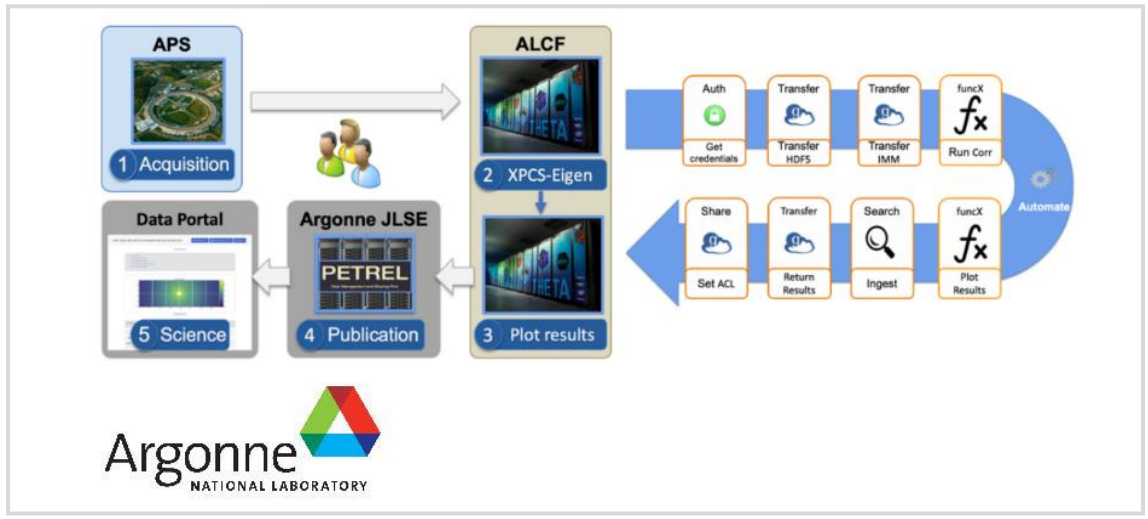

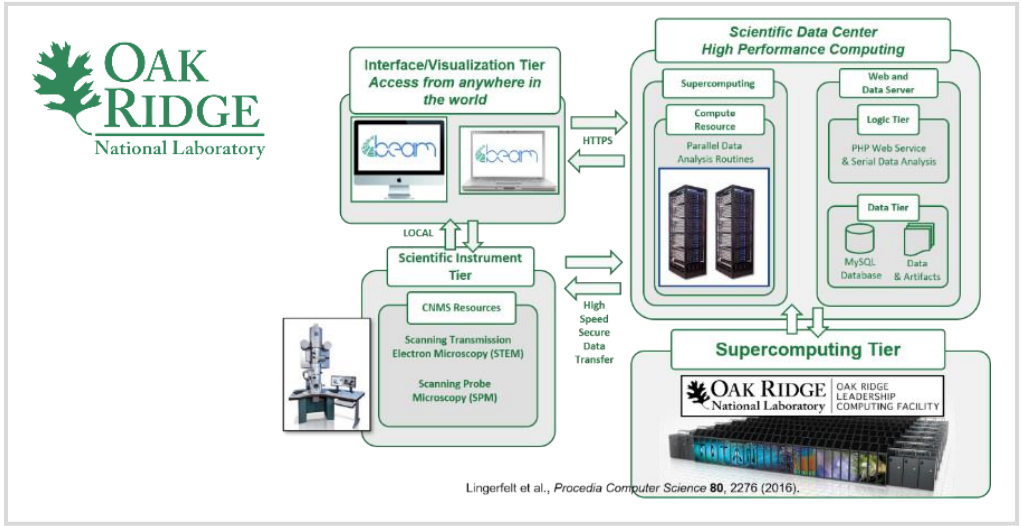

3. Deploy an IRI Pathfinding Testbed across the four ASCR Facilities

Early focus: Initiate projects for the Time-Sensitive and Data-Integration Intensive patterns
Early focus: Engage ASCR Facilities with new IRI-oriented research projects

IRI Pathfinding Testbed

ESnet, OLCF, NERSC, and ALCF are finalizing a jointly-authored concept paper, quoted at right.

Each facility will contribute resources to creating this research environment.

"The proposed IRI Testbed is a progressive design-experiment and test-refine approach proposed to ***establish a shared environment for IRI developers and pilot application users to come together*** *and advance the IRI vision.*

"The goal is to build cyberinfrastructure enabling multiple user facilities to experiment with the design patterns and address the gaps identified in the IRI Architecture Blueprint Activity (ABA) report."

A selection of recent research solicitations with IRI themes:

Advanced Scientific Computing Research for DOE User Facilities
DOE Lab Call 23-3030

Distributed Resilient Systems
DE-FOA-0002902

Biopreparedness Research Virtual Environment (BRaVE)
DOE Lab Call 23-2955

4. Stand up an IRI Program structure (HQ and Field)

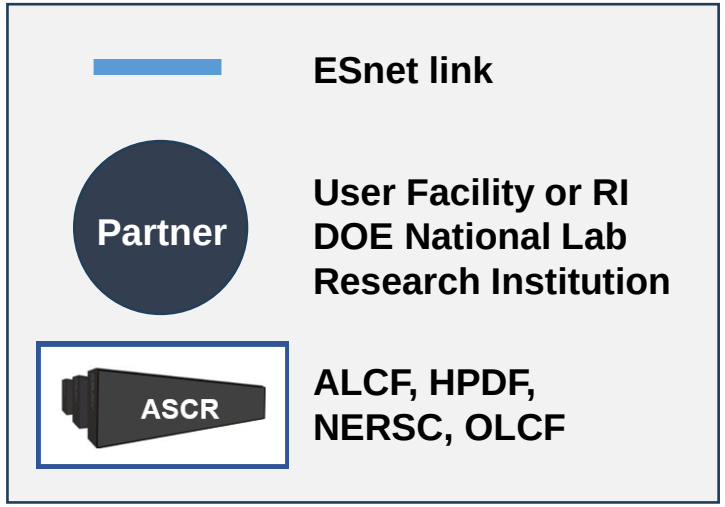

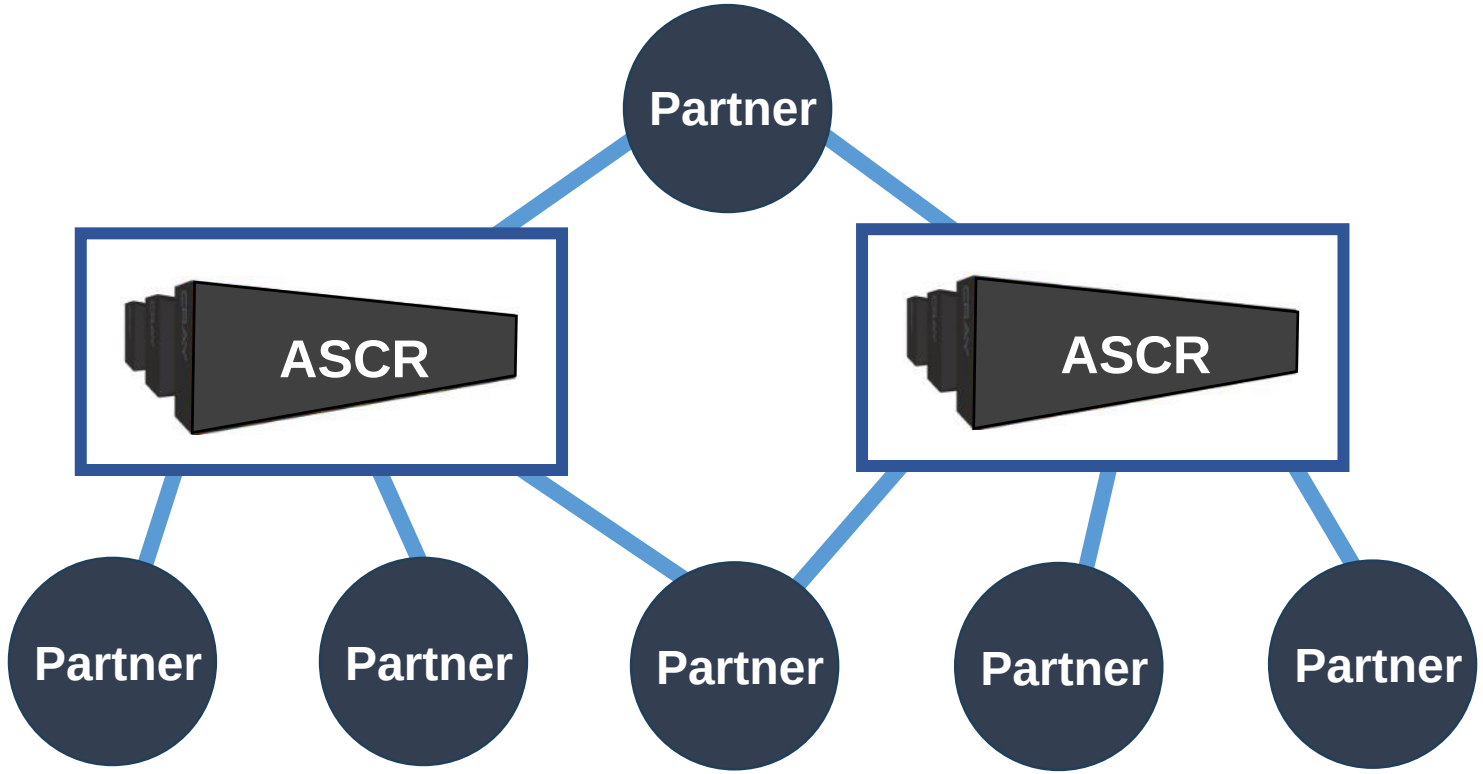

ALCF, HPDF, NERSC, and OLCF each provide unique high performance compute/data resources and each has partners. ESnet provides high performance network connections and IRI data services.

At HQ, IRI Program will be led by ASCR. Governance will span partner institutions, their DOE Program sponsors, and National Laboratory leadership.

Implementing IRI: Today ASCR is taking the first big steps

1. **Invest in IRI foundational infrastructure**
   - Early focus: Report IRI Science Pattern requirements across all SC Programs and assess IRI readiness
   - Early focus: Advance the NERSC-10, High Performance Data Facility, and OLCF-6 projects to CD-1

2. **Bring existing IRI projects into formal coordination**
   - Early focus: Move towards a common user experience across the ASCR HPC Facilities
   - Early focus: Create an IRI HPC allocation

3. **Deploy an IRI Pathfinding Testbed spanning the four ASCR Facilities**
   - Early focus: Initiate projects for the Time-Sensitive and Data-Integration Intensive patterns
   - Early focus: Engage ASCR Facilities with new IRI-oriented research projects

4. **Stand up an IRI Program structure (HQ and Field)**
   - Early focus: Develop a DOE Authentication/Authorization standard
   - Early focus: Devise and revise governance structure with the community

Time-Sensitive Pattern

Data-Integration Intensive Pattern

Long-Term Campaign Pattern

IRI Look Ahead for 2023

Release of **IRI Blueprint Activity** final report

Release of **ESnet Requirements Reviews IRI meta-analysis**

Release of the ASCR Facilities' **IRI Pathfinding Testbed** white paper

Advancing ASCR major projects

Convening event

Acknowledgements

SC Science Programs IRI Coordination Group

BER	Paul Bayer, Jay Hnilo, Resham Kulkarni
BES	Dava Keavney, Tom Russell, Misha Zhernenkov
FES	Josh King, Matt Lanctot
HEP	Jeremy Love, Eric Church
IP	Julie Ezold, Kristian Myhre
NP	Xiaofeng Guo, Paul Mantica, Jim Sowinski

IRI Blueprint Activity Leadership Group

ASCR Colleagues

ASCR Facility Directors

IRI Blueprint Activity participants!

ASCR Facilities IRI Task Force

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

U.S. DEPARTMENT OF ENERGY | Office of Science

The people of the ASCR Facilities: Providing high performance Research Computing, Data, and Networking for DOE and the Nation