Integrated Research Infrastructure

Chin Guok
Chief Technology Officer
Energy Sciences Network
Lawrence Berkeley National Laboratory

11th SIG-NGN Meeting
Prague, Czech Republic
April 20, 2023

ESnet is a Science Mission Network

ESnet provides the high-bandwidth, reliable connections that **link scientists** at national laboratories, universities, and other research institutions, **enabling them to collaborate** on some of the world's most important scientific challenges including **energy**, **climate science**, and the **origins of the universe**. **Funded by the DOE Office of Science**, ESnet is managed and operated by the Scientific Networking Division at Lawrence Berkeley National Laboratory. As a nationwide infrastructure and DOE User Facility, ESnet provides scientists with **access to unique DOE research facilities and computing resources**.

ESnet's Mission is to enable and accelerate scientific discovery by delivering unparalleled network infrastructure, capabilities, and tools.

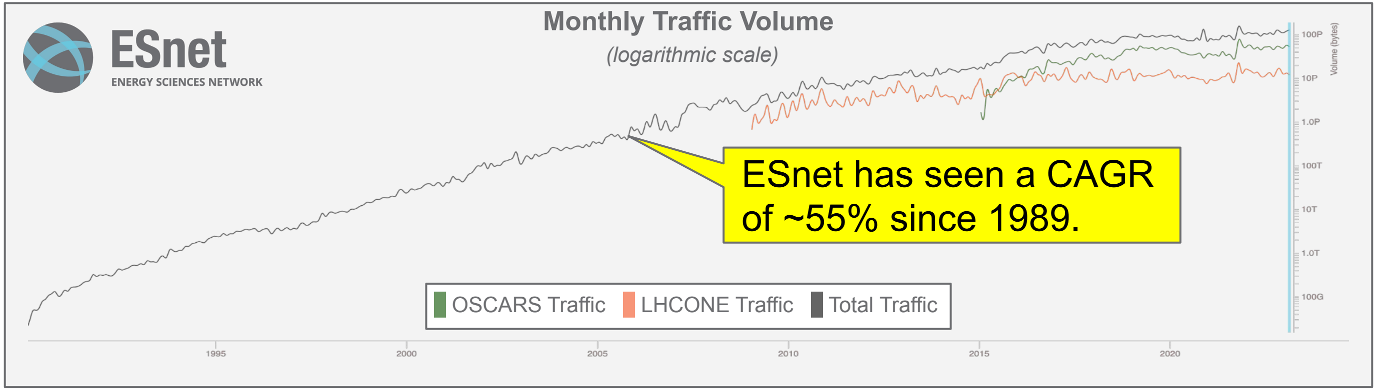

ESnet

DOE Office of Science - Largest supporter of basic research in the physical sciences in the US

The mission of the **Advanced Scientific Computing Research (ASCR)** program is to discover, develop, and deploy computational and networking capabilities to analyze, model, simulate, and predict complex phenomena important to the Department of Energy (DOE).

Basic Energy Sciences (BES) supports fundamental research to understand, predict, and ultimately control matter and energy at the electronic, atomic, and molecular levels in order to provide the foundations for new energy technologies and to support DOE missions in energy, environment, and national security.

The mission of the **Biological and Environmental Research (BER)** program is to support transformative science and scientific user facilities to achieve a predictive understanding of complex biological, earth, and environmental systems for energy and infrastructure security, independence, and prosperity.

The **Fusion Energy Sciences (FES)** program mission is to expand the fundamental understanding of matter at very high temperatures and densities and to build the scientific foundation needed to develop a fusion energy source.

The mission of the **High Energy Physics (HEP)** program is to understand how our universe works at its most fundamental level.

The mission of the **Nuclear Physics (NP)** program is to discover, explore, and understand all forms of nuclear matter.

DOE Office of Science - Uniquely positioned for large scale collaborative science*

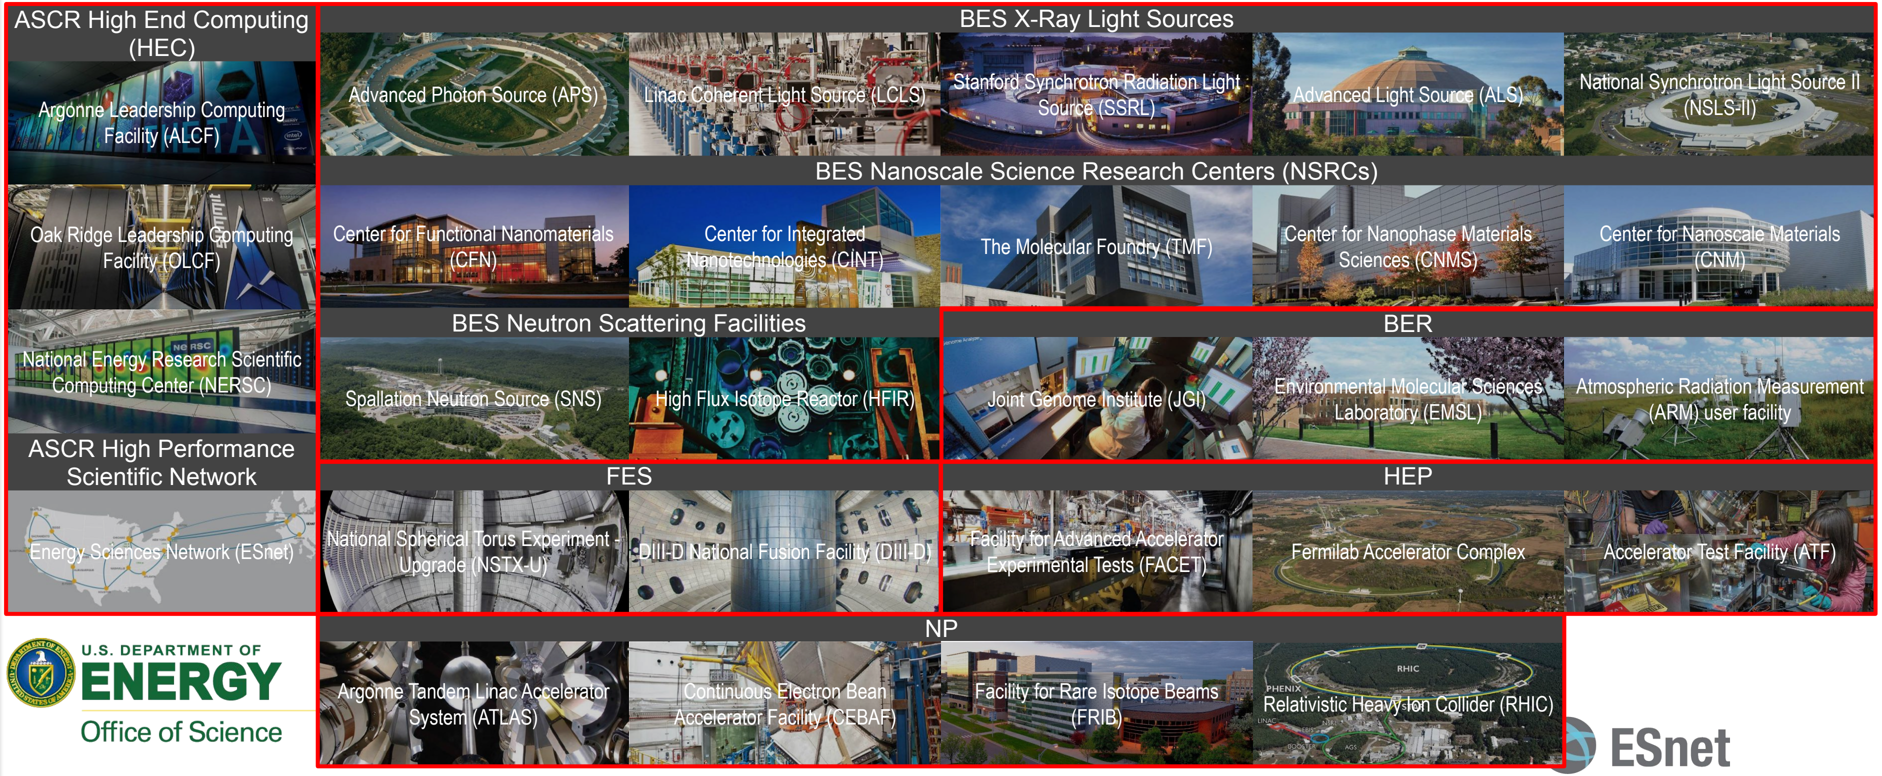

U.S. DEPARTMENT OF ENERGY
Office of Science

ESnet

DOE Office of Science facilities also support other collaborations, e.g., LHC, LSST, etc

DOE Office of Science - Uniquely positioned for large scale collaborative science*

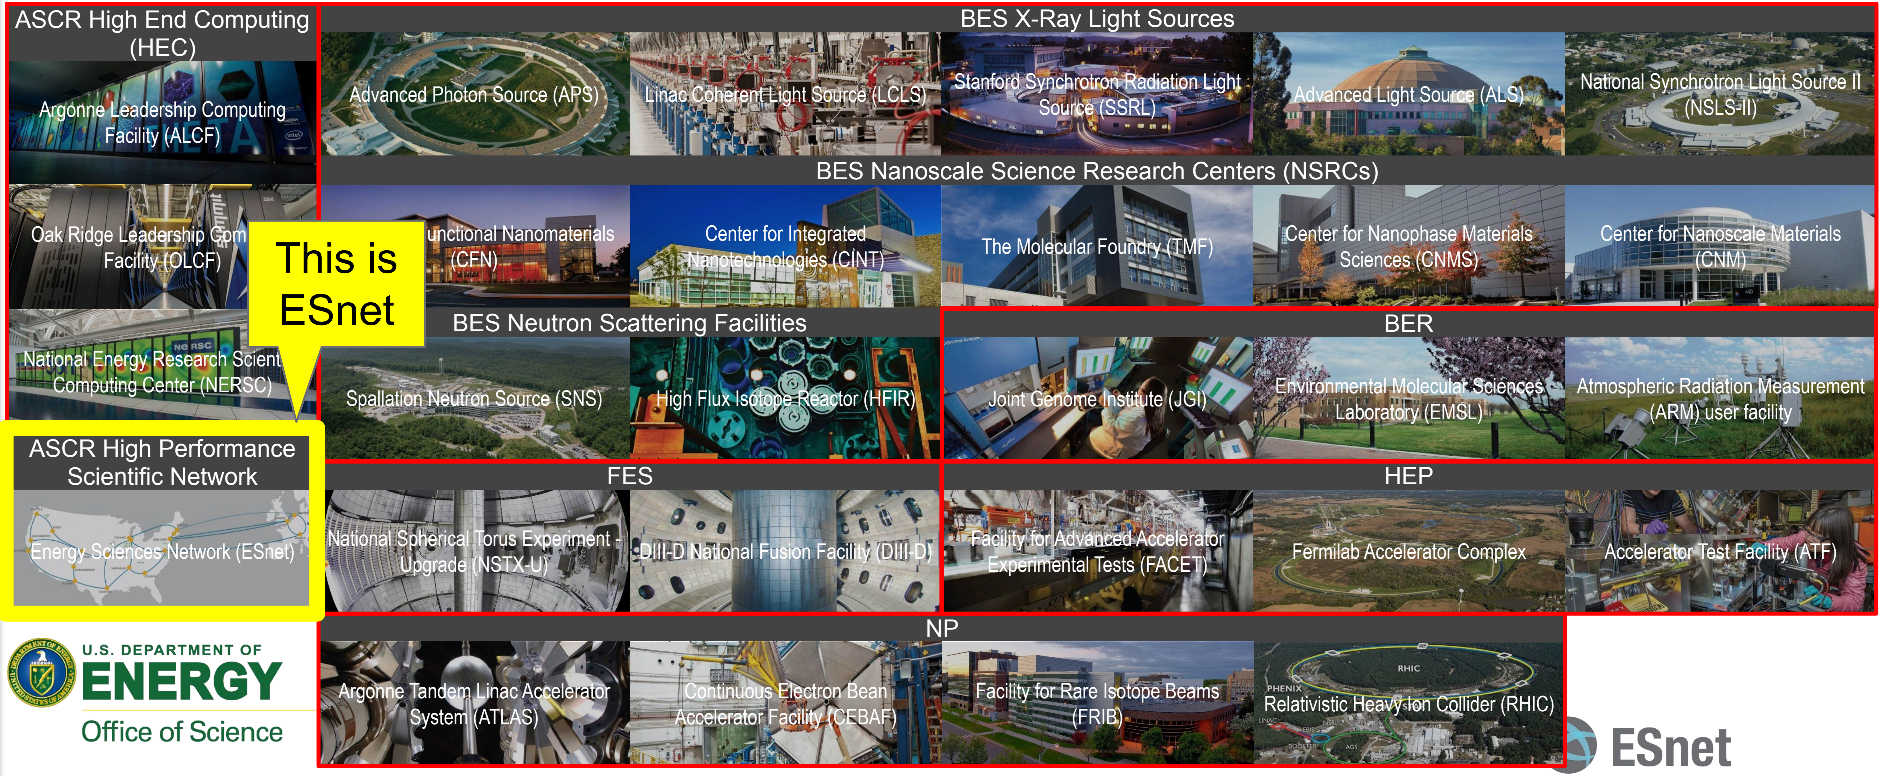

U.S. DEPARTMENT OF ENERGY
Office of Science

ESnet

DOE Office of Science facilities also support other collaborations, e.g., LHC, LSST, etc

What are (some of) the problems we are trying to solve?

- Science projects and facilities lack multi-site, multi-year computing and data storage allocations.
- Users have to establish accounts and identities at each site.
- Projects lack the ability to transparently access large scale datasets across multiple facilities, lack access to distributed data sharing mechanisms, and lack the tools that address the full life-cycle of the data.
- Science projects cannot seamlessly execute and schedule workflows across facilities for resiliency and other purposes.
- Cross-facility governance, policies, and metrics are not yet aligned with the proposed integrated infrastructure.
- Multi-facility workflows and next-generation data analysis techniques don't always have a group of cross-cutting experts to bring these capabilities into science collaborations broadly.

ESnet

Interconnectivity and integration of instrumentation, data and computing have been explicitly recognized as strategic requirements for national R&D

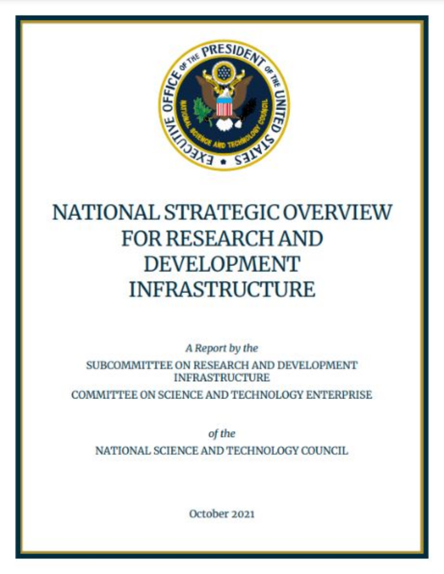
NATIONAL STRATEGIC OVERVIEW FOR RESEARCH AND DEVELOPMENT INFRASTRUCTURE

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

The 2021 National Strategic Overview from the Subcommittee on Research and Development Infrastructure formally redefined "federal R&D Infrastructure" to now include computing, data, and networking facilities, resources and services.

"R&D continues to shift from smaller to bigger science, driven in large part by advances in computing and other research cyberinfrastructure, which interlink[s] research data, analytics, … and experimental instrumentation."

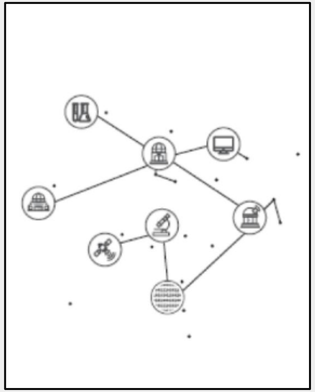

Past

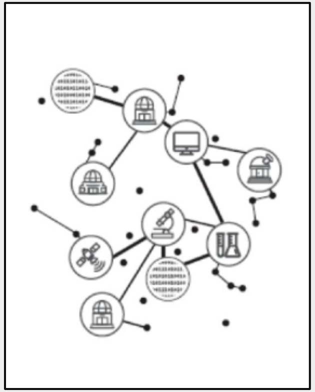

Present

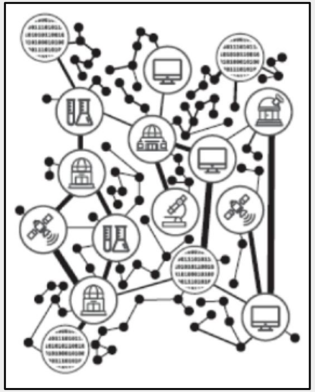

Future

U.S. DEPARTMENT OF ENERGY | Office of Science

ESnet

https://www.whitehouse.gov/wp-content/uploads/2021/10/NSTC-NSO-RDI-_REV_FINAL-10-2021.pdf

Brown et al

The Roundtable report contained key insights regarding the need for distributed and interoperable computing resources

"Facilities may be able to deploy a distributed network of connected and interoperable computing resources that enable all scales of computing, data exploration, and analysis."

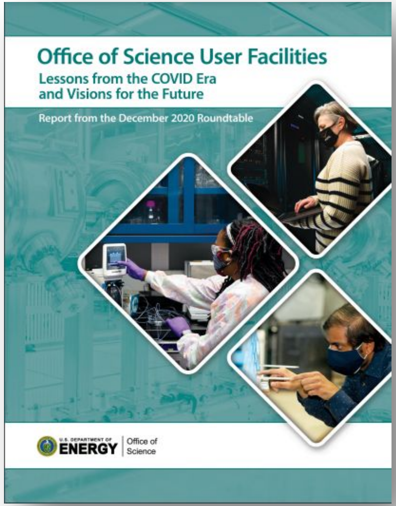

- The pandemic highlighted the importance of **secure remote collaboration and facilitated access to data and computing resources**.
- **Data management and data stewardship** present another critical opportunity. User facilities generate exabytes of unique, irreplaceable data which must be managed, curated and made available for analysis and computation.
- This requires a host of **user-connecting operational approaches and technologies** such as high-quality interfaces, collaboration tools, federated identity management, automation of experiments and workflows, and more.
- "With collaboration among all its user facilities, DOE SC is in a position to facilitate all aspects of the data lifecycle across its facility complex, including simulations, experiment design, data generated at scientific instruments, data analysis, and data archiving for future use.
- "Seamlessly connecting a user with data and computing enables more uniform and egalitarian data exploration and analysis capabilities."

U.S. DEPARTMENT OF ENERGY | Office of Science

ESnet

International efforts in Integrated Research Infrastructure are expanding too

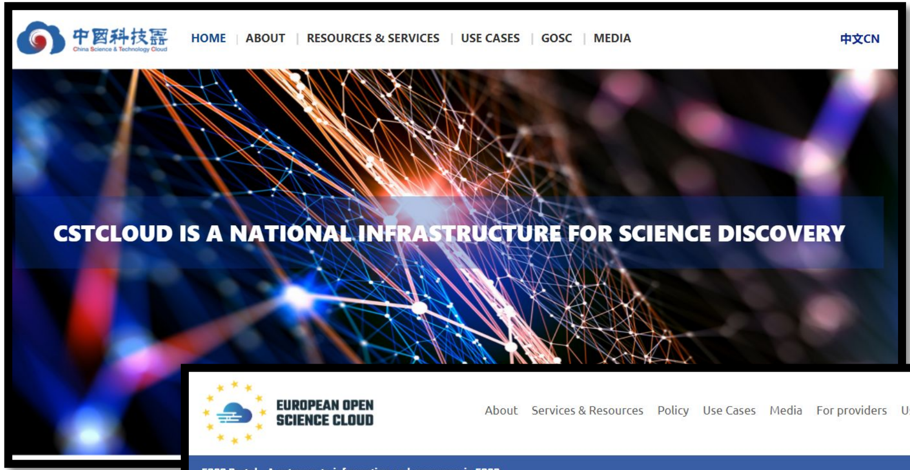

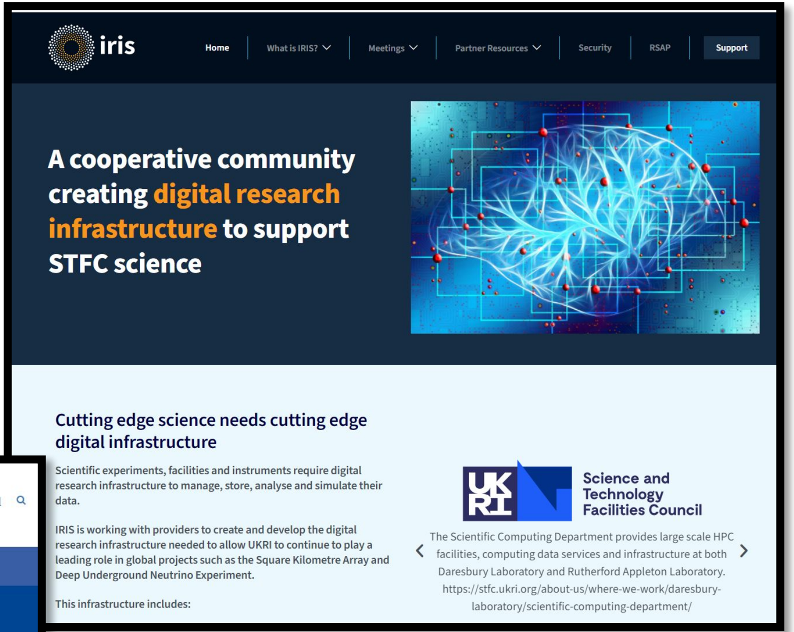

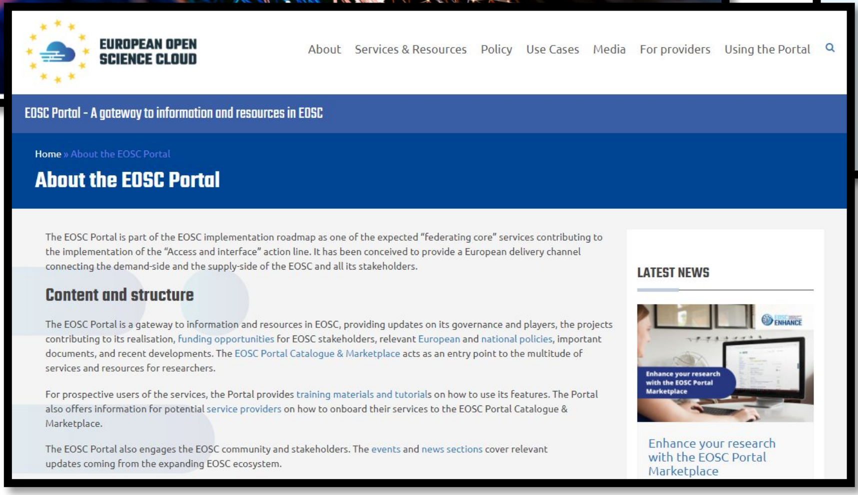

- China Science and Technology Cloud
- European Open Science Cloud
- IRIS UKRI SFTC initiative

ESnet

Brown et al

What does this mean for networks*?

Promoting networks as “first class” resources, similar to instruments, compute and storage, e.g.,

- Accessible
  - Security frameworks for accessing (selected) services
  - APIs to interact with services
- Controllable
  - Resource/service selection/negotiation
  - Service scheduling
- Transparent
  - Resource (general) availability
  - Service (specific) status
- Adaptable
  - Ability to integrating compute and/or storage into the network
  - Rapid prototyping of new services

Networking is an end-to-end service, inter-domain interoperability and service consistency is critical!

ESnet

The vision: A DOE/SC integrated research ecosystem that transforms science via seamless interoperability

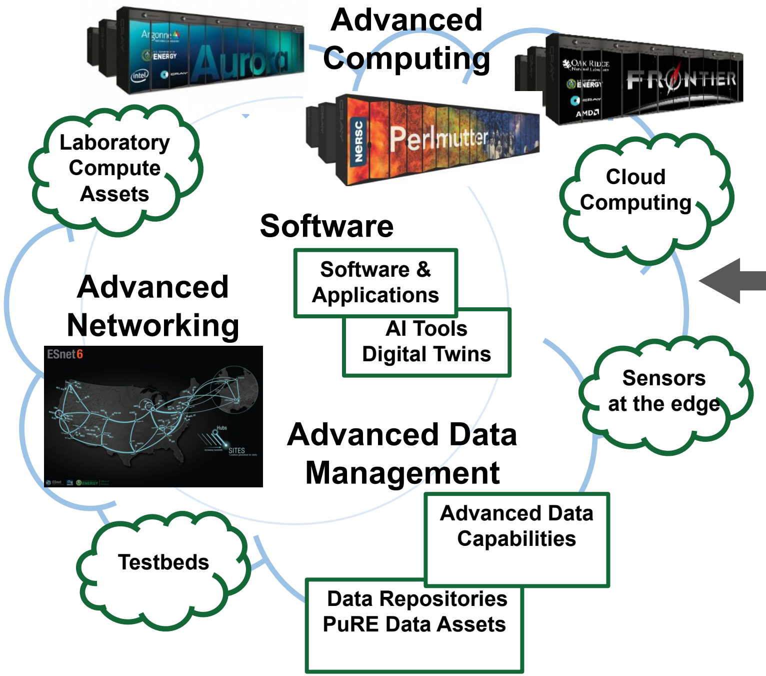

Strategic goal: Broadening participation

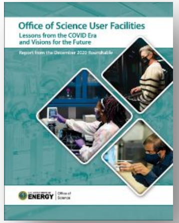

New modes of integrated science

Rapid data analysis and steering of experiments	Novel models for multi-facility allocation/utilization	AI-enabled insight from dynamic, vast multi-modal data	Seamless user interconnectivity via federated IDs

Brown et al

DOE ASCR Integrated Research Infrastructure Vision Statement

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

Addressing national and societal grand challenges and unlocking new opportunities around energy, science, and technology for US competitiveness will require **highly coordinated, collaborative research** and **integrating capabilities across our world-leading facilities**, which currently operate largely independently. We can achieve this vision if the facilities, projects, and science communities have the right **incentives, governance, and operating structure** to enable them to **deliver an integrated research platform** – accelerating time-to-discovery and time-to-innovation.

ESnet

DOE ASCR IRI Task Force contemplated operational models and guiding principles [CY2021]

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 research and other investments. In February 2020, ASCR established the Integrated Research Infrastructure Task Force[2] as a forum for discussion and exploration, with specific focus on the operational opportunities, risks, and challenges that integration poses. In light of the global COVID-19 pandemic, the Task Force conducted its work asynchronously from April through December 2020, meeting via televideo for one hour every other week. The Director of the ASCR Facilities Division facilitated the Task Force, in coordination with the ASCR Facility Directors.

The work of the Task Force began with these questions: Can the group arrive at a shared vision for integrated research infrastructure? If so, what are the core principles that would maximize scientific productivity and optimize infrastructure operations? This paper represents the Task Force's initial answers to these questions and their thoughts on a strategy for world-leading integration capabilities that accelerate discovery across a wide range of science use cases.

B. Brown, C. Adams, K. Antypas, D. Bard, S. Canon, E. Dart, C. Guok, E. Kissel, E. Lancon, B. Messer, S. Oral, J. Ramprakash, A. Shankar, T. Uram, <https://doi.org/10.2172/1863562>

Areas

- AL Allocations
- AC Accounts
- DA Data
- AP Applications
- SC Scheduling
- WF Workflows
- PB Publication
- AR Archiving

Principles

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.

Brown et al

SC Integrated Research Infrastructure Architecture Blueprint Activity (IRI-ABA) [CY2022]

- **Aim:** Produce the **reference conceptual foundations** to inform a coordinated “whole-of-SC” strategy for an integrative research ecosystem.
- **Approach:**
  - Invite DOE experts across the SC User Facilities, SC National Laboratories, and key SC enterprise stakeholders to participate in a series of activities and events:
    - Gather and analyze **integrative use cases** that inclusively span SC programs and user facilities
    - Develop overarching design principles and one or more “architecture blueprints” that will address the chief IRI **design patterns** in an efficient way.
    - Identify **urgent program and lab priorities** and early win opportunities.
- **Intended outcomes:**
  - Produce a shared understanding across SC and DOE of IRI requirements, operational and technical gaps and needed investments, and a common lexicon to describe these.
  - Position SC Programs to contemplate future investment decisions.
  - Explore leveraging existing SC and ASCR resources and services as well as identifying new needs for research and capability gaps for new resources that do not yet exist.
- **Timeline:** February through September 2022.

ESnet

Snapshot of IRI ABA Participants (as of May 2022)

✔ 112 subject matter experts with diverse backgrounds from across the SC lab complex. (Does not include additional 30 science user interviewees who provided input during Sprint 1. The next two slides show that the interviewees were speaking to science use cases that used multiple facilities.)

Interest survey responses …..

"Area that I identify with"

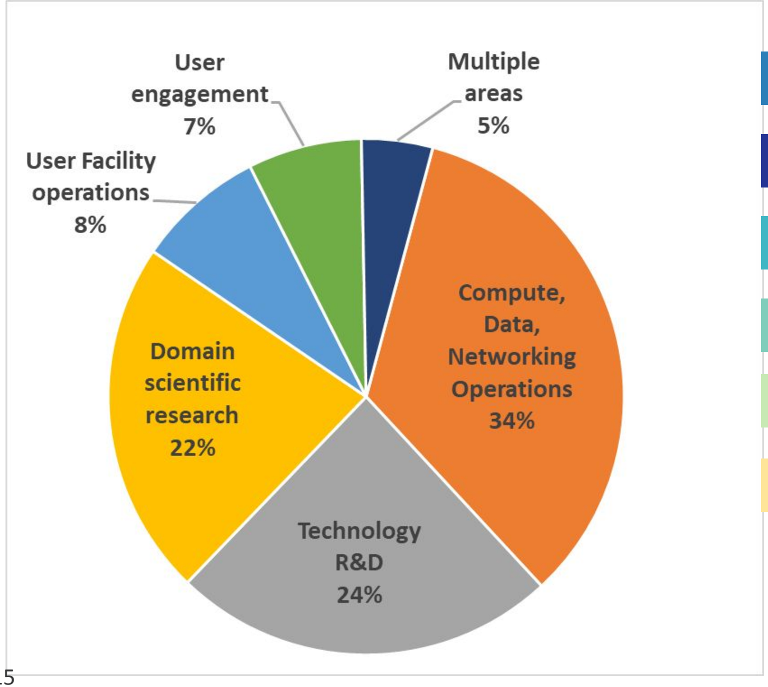

"SC Office sponsoring my work"
(many indicated multiple sponsors)

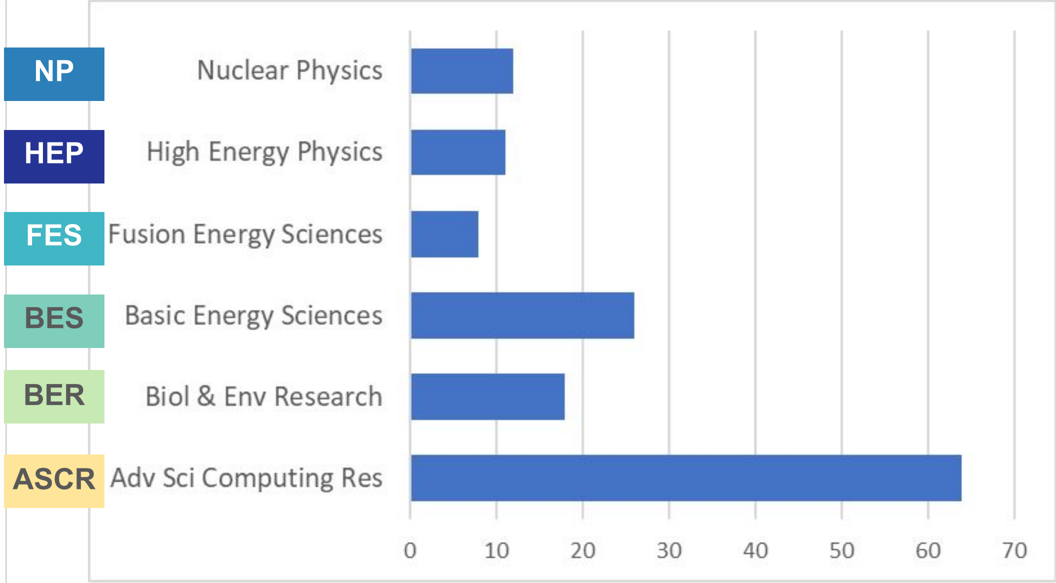

ESnet

User requirements touched on many IRI cases

NP	Large-scale simulations integrating experimental data from DOE and international nuclear physics facilities.
NP	AI/ML incorporated into simulations to drive exploration of parameter space using codes like LAMMPS in the workflow.
NP	GRETA spectrometer online computing pipeline construction - for 1-5 day campaigns.
HEP	Astronomical spectroscopic observations at DESI/KPNO, survey data daily streaming to HPC, target selection pipeline.
HEP	Astronomy data management - DESI workflows - processing 100s of TBs of data utilizing HPC.
HEP	LHC simulation workflows; event generation for integration of WLCG with HPC.
HEP	Event generation for of complex simulation pipeline. Distributed, compute-intensive but not data-intensive.
FES	Heterogeneous data handling and real-time analysis for fusion power experiment steering at DIII-D and JET for fusion.
FES	DIII-D Fusion experiments + simulations, near-real-time analysis to determine/predict plasma conditions.
FES	DIII-D Tokamak plasma physics experiments - diagnostic data collection. Integrating ASCR HPC facilities with the WLCG, for LHC computing.
FES	Fusion workflows: 10s ~45GB per pulse, about 500TB total using custom SQL and NoSQL databases, computing with local and institutional resources, data exposed via APIs also for future ITER workflows.
FES	Data management and real-time analysis for fusion experiments at PPPL.

BES	Beamline data transfer at scale to HPC for real-time analysis and experiment steering, data transfer at scale to HPC (e.g., at SSRL).
BES	Metallurgy, MIDAS x-ray analysis software for APS high-energy X-ray diffraction microscopy beamline data. High data throughput. Pipeline development to NERSC and ALCF for on-demand computing at scale.
BES	Autonomous AI/ML-driven experiments such as for ALS and NSLS-II.
BES	APS 8-ID-I Smal-angle XPCS x-ray spectroscopy, high-frame rate camera, data management workflow including HPC. 0.2 PB of unsparsified data generated/day.
BES	Materials science. High data throughput, facility data transfer from APS to PNNL for analysis.
BES	Light source data processing workflows for large datasets (TBs to PBs), particularly ptychography; using AI/ML to help reduce data quickly.
BES	Data pipelines for high-speed detectors used at lights sources and NSRCs. Workflow development using NERSC.
BES	ML Autonomous materials characterization workflows using data (100s GBs per day) collected at light sources, neutron sources, and NSRCs.
BER	Combining multi-modal data from simulations and experiments for molecular scale imaging workflows (instrument + compute + storage).
BER	Crystallography microscopy and synchrotron light sources workflows and local and HPC computing.
BER	HPC-enabled high throughput sequencing, large scale sequence data analysis, sample and data life cycle, data product development (e.g., at JGI).
BER	Applying AI/ML to EMSL data for analysis.

ESnet

Common recurring sentiments across the user interviews

Data Management

- Users are overwhelmed with large and growing amounts of data to manage, reduce, analyze
  - Users need to move data across facilities and use different systems at different steps of data processing chain
  - Users need bespoke data movement and workflow solutions, and long-duration support for data/metadata.

Automation/AI

- Users need reliable automation & seamless access and try to compensate via human effort.
  - Users need automation, and anticipate AI, but struggle with skills and application of these novel technologies

Heterogeneity

- Users face mismatches between resources, tools and needs
  - Users need heterogeneity in scale and type of resources but have platform fatigue learning many different platforms
  - Users need workflows to be at the center but need software APIs and standardization/uniformity
  - Users have a spectrum of computing needs from elastic computing (matching need to available resources) to urgent computing (near real-time/just-in-time, on-demand)

Ease of Use

- Users find infrastructure hard to use
  - Users encounter a lack of transparency about workflow tools and resources, and many different use policies and cybersecurity barriers
  - Users need infrastructure to be easier to use and be more coordinated across resources and facilities

Workforce Skills Gap

- Users and teams struggle with workforce and training needs
  - Users (and their organizations) struggle with lack of skills, oversubscribed staff, recruiting, and retention
  - Users experience gaps between their working knowledge and skills and those of infrastructure experts.
  - Users need support and expertise in data science.

ESnet

IRI ABA Design Phase

Pattern Blueprints

- Time Sensitive
- Data Integration-Intensive
- Long-Term Campaign

Practice Areas

- Resource Co-Operations
- Cybersecurity & Federated Access
- Workflows, Interfaces & Automation
- Scientific Data Lifecycle
- Portable/Scalable Solutions
- User Experience

Focus Topics

- Overarching IRI Principles
- Blueprint Compare & Contrast
- Governance/Steering

ESnet

IRI ABA Design Phase Pattern Blueprints

Time Sensitive IRI patterns comprise workflows with **time critical/sensitive requirements** (i.e., real-time or near real-time), which can be **motivated by** various factors such as **timely decision making**, **experiment steering**, **virtual proximity**, and **loss of data fidelity**. These time-sensitive workflows involve the integration across multiple facilities and resources.

Data Integration-Intensive are characterized by a need to perform **analysis of data combined from multiple sources** - which can include **data from multiple sites**, **experiments and/or simulations**. This can also include **tracking metadata and provenance** for **reproducible science** and **interactive data analysis**, possibly **at scale**.

Long-Term Campaign patterns is characterized by a need for **sustained access to resources at scale over a longer time** to accomplish a well-defined objective. **Robustness**, **reproducibility**, **and reliability** are important to accomplishing long-term science, and these patterns will likely involve **significant logistical planning**.

ESnet

IRI ABA Design Phase Practice Areas

Resource Co-Operation - **Allocations/provisioning** of multiple heterogeneous **resources across multiple facilities** for large collections of scientific programs must be **aligned in time and planned**.

Cybersecurity & Federated Access - Users require a distributed research Infrastructure with **seamless access and consistent services** while the infrastructure must be operated according to **cybersecurity requirements and policies set at the Federal level**.

User Experience - **Understanding evolving users' needs and experiences are critical** for technologists' to develop effective IRI solutions.

Workflows, Interfaces & Automation - Users need to **systematically and easily assemble system components** to support IRI science cases in the form of **end-to-end pipelines**.

Scientific Data Lifecycle - Users need to **manage their data (along with meta-data) across facilities from inception, to curation, archiving, dissemination and publication**.

Portable / Scalable Solutions - Users and technologists need to **move/translate their applications across heterogeneous facilities** (be portable) and go from **smaller to larger resources** (be scalable).

IRI ABA Implications for ESnet

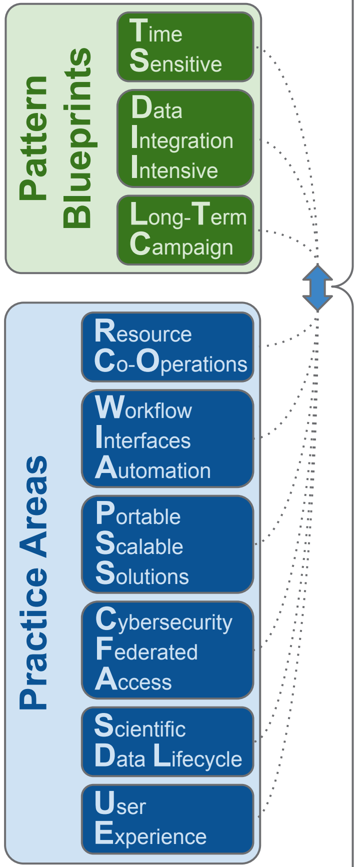

Enable **predictable (end-to-end) network services**, e.g., guaranteed bandwidth/latency/jitter, load-balancing, network resiliency

Provide **high bandwidth and rich connectivity**, e.g., capacity planning, Cloud-connect/peering strategies

Support **application/network interaction**, e.g., availability, provisioning, verification, monitoring

Facilitate **“friction-free” data movement**, e.g., low-impedance architectures, data movement tools

Provide/support **network computational storage** capabilities, e.g., workflow integrated edge compute, in-network compute, in-network data caching

Support **multi-modal network connectivity**, e.g., wireless sensor nets

Advocate for supported **programming constructs**, e.g., orchestration/automation, inter-facility APIs, common (portable) programming and runtime environments, software lifecycle, “standardization”

Collaborate on **common access framework**, e.g., cybersecurity, federated access, resource allocations

Support **resource allocation policies**, e.g., (guaranteed/transferrable) resource allocations, facility metics

Encourage **development and testing environments**, e.g., (federated) testbeds, prototyping collaborations

Facilitate **co-design services**, e.g., design patterns, standard practices

Empower **engagement and partnerships**, e.g., outreach, practice groups, forums

High Performance Data Facility (HPDF)

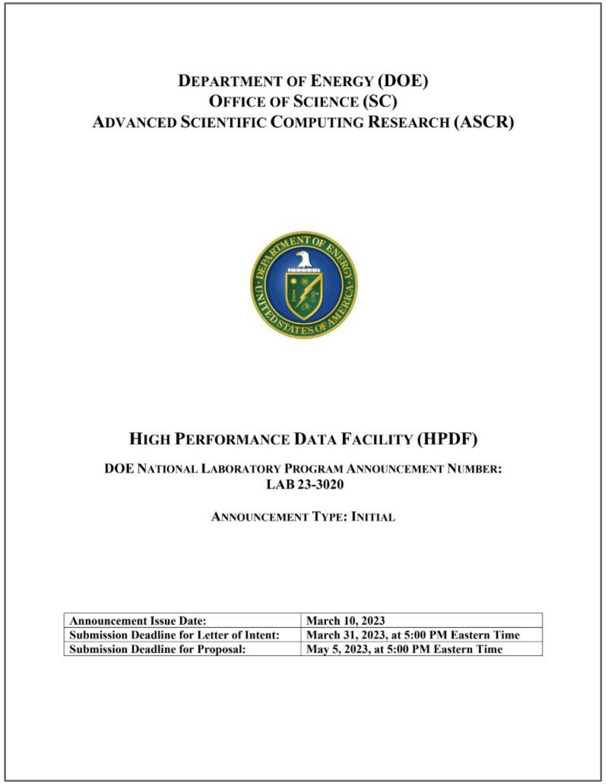

DEPARTMENT OF ENERGY (DOE)
OFFICE OF SCIENCE (SC)
ADVANCED SCIENTIFIC COMPUTING RESEARCH (ASCR)

HIGH PERFORMANCE DATA FACILITY (HPDF)

DOE NATIONAL LABORATORY PROGRAM ANNOUNCEMENT NUMBER:
LAB 23-3020

ANNOUNCEMENT TYPE: INITIAL

Announcement Issue Date:	March 10, 2023
Submission Deadline for Letter of Intent:	March 31, 2023, at 5:00 PM Eastern Time
Submission Deadline for Proposal:	May 5, 2023, at 5:00 PM Eastern Time

“To meet these challenges, SC is advancing the Integrated Research Infrastructure (IRI) vision: DOE will empower researchers to seamlessly and securely meld DOE’s world-class research tools, infrastructure, and user facilities in novel ways to radically accelerate discovery and innovation.”

“The High Performance Data Facility (HPDF) will serve as a foundational element in enabling the DOE Integrated Research Infrastructure.”

“DOE requires a dynamic and scalable data management infrastructure that is network-integrated with the DOE computing ecosystem, with diverse capabilities”

“DOE requires a diversified computing ecosystem that can provide researchers with access to an appropriate computing resource at the appropriate time”

ESnet

Questions…

Chin Guok <chin@es.net>

ESnet