

Lawrence Berkeley National Laboratory

LBL Publications

Title

Fusion Energy Sciences Network Requirements Review: Mid Cycle Update

Permalink

<https://escholarship.org/uc/item/4w2151rp>

Authors

Zurawski, Jason

Dart, eli

Halfmoon, Michael

et al.

Publication Date

2024-07-26

Copyright Information

This work is made available under the terms of a Creative Commons Attribution-NonCommercial-NoDerivatives License, available at

<https://creativecommons.org/licenses/by-nc-nd/4.0/>

Peer reviewed

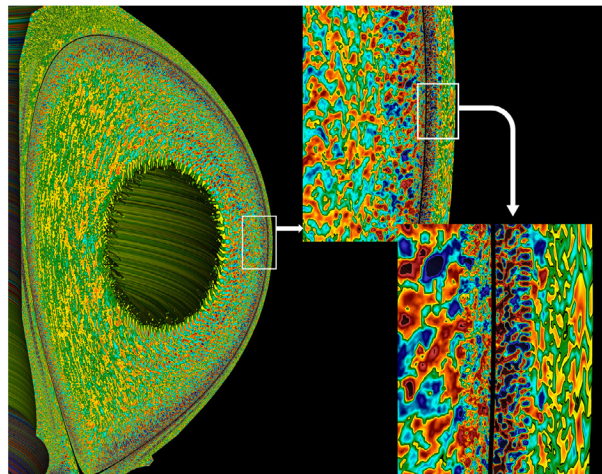




ESnet

ENERGY SCIENCES NETWORK

Fusion Energy Sciences Network Requirements Review

Mid-cycle Update
May 2023



BERKELEY LAB



U.S. DEPARTMENT OF
ENERGY

Office of
Science



ESnet

ENERGY SCIENCES NETWORK

Fusion Energy Sciences Network Requirements Review

Mid-cycle Update

May 2023

Office of High Energy Physics, DOE Office of Science
Energy Sciences Network (ESnet)

ESnet is funded by the US Department of Energy, Office of Science, Office of Advanced Scientific Computing Research. Carol Hawk is the ESnet Program Manager.

ESnet is operated by Lawrence Berkeley National Laboratory (Berkeley Lab), which is operated by the University of California for the US Department of Energy under contract DE-AC02-05CH11231.

This work was supported by the Directors of the Office of Science, Office of Advanced Scientific Computing Research, Facilities Division, and the Office of High Energy Physics.

This is LBNL report number LBNL-2001603

Disclaimer

This document was prepared as an account of work sponsored by the United States Government. While this document is believed to contain correct information, neither the United States Government nor any agency thereof, nor The Regents of the University of California, nor any of their employees, makes any warranty, express or implied, or assumes any legal responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by its trade name, trademark, manufacturer, or otherwise, does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof, or The Regents of the University of California. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof or The Regents of the University of California.

Cover Images:

Upper left: gyrokinetic simulation: MIT PSFC

Upper right: General Atomics

Lower left: Credit © ITER Organization, <http://www.iter.org/>

Lower right: Figure caption: Turbulence in edge plasma: PPPL, ALCF, and OLCF

¹<https://escholarship.org/uc/item/4w2151rp>

Participants and Contributors

Paul Bonoli, *MIT Plasma Science and Fusion Center*

CS Chang, *Princeton Plasma Physics Laboratory*

Michael Churchill, *Princeton Plasma Physics Laboratory*

Doug Curry, *Oak Ridge National Laboratory*

Eli Dart, *Lawrence Berkeley National Laboratory and Energy Sciences Network*

Bill Dorland, *Princeton Plasma Physics Laboratory*

Frederico Fiuza, *SLAC National Accelerator Laboratory*

Mark Foster, *SLAC National Accelerator Laboratory*

Cameron G. R. Geddes, *Lawrence Berkeley National Laboratory*

Richard Gerber, *Lawrence Berkeley National Laboratory and the National Energy Research Scientific Computing Center*

Anthony Gonsalves, *Lawrence Berkeley National Laboratory*

David Green, *Oak Ridge National Laboratory*

Chin Guok, *Lawrence Berkeley National Laboratory and Energy Sciences Network*

Walter Guttenfelder, *Princeton Plasma Physics Laboratory*

Carol Hawk, *Department of Energy Office of Science*

Michael Halfmoon, *Department of Energy Office of Science*

Susan Hicks, *Oak Ridge National Laboratory*

Axel Huebl, *Lawrence Berkeley National Laboratory*

Jerry Hughes, *MIT Plasma Science and Fusion Center*

James Kafader, *Lawrence Berkeley National Laboratory and Energy Sciences Network*

Scott Kampel, *Princeton Plasma Physics Laboratory*

Stan Kaye, *Princeton Plasma Physics Laboratory*

Josh King, *Department of Energy Office of Science*

Cornwall Lau, *Oak Ridge National Laboratory*

Arnold Lumsdaine, *Oak Ridge National Laboratory*

John Mandrekas, *Department of Energy Office of Science*

Orso Meneghini, *General Atomics*

Ken Miller, *Lawrence Berkeley National Laboratory and Energy Sciences Network*

Inder Monga, *Lawrence Berkeley National Laboratory and Energy Sciences Network*

Raffi Nazikian, *General Atomics*

Jeff Nguyen, *General Atomics*

Lieselotte Obst-Huebl, *Lawrence Berkeley National Laboratory*

Francesca Poli, *Princeton Plasma Physics Laboratory*

Juergen Rapp, *Oak Ridge National Laboratory*

Steve Sabbagh, *Columbia University and Princeton Plasma Physics Laboratory*

Brandon Savage, *MIT Plasma Science and Fusion Center*

David Schissel, *General Atomics*

Douglass Schumacher, *The Ohio State University*

Josh Stillerman, *MIT Plasma Science and Fusion Center*

Jean-Luc Vay, *Lawrence Berkeley National Laboratory*

Fatema Bannat Wala, *Lawrence Berkeley National Laboratory and Energy Sciences Network*

Andrew Wiedlea, *Lawrence Berkeley National Laboratory and Energy Sciences Network*

Tong Zhou, *Lawrence Berkeley National Laboratory*

Jason Zurawski, *Lawrence Berkeley National Laboratory and Energy Sciences Network*

Report Editors

Jason Zurawski, *ESnet*: zurawski@es.net

Eli Dart, *ESnet*: dart@es.net

Michael Halfmoon, *Department of Energy Office of Science*: michael.halfmoon@science.doe.gov

Carol Hawk, *Department of Energy Office of Science*:
carol.hawk@science.doe.gov

Josh King, *Department of Energy Office of Science*:
josh.king@science.doe.gov

John Mandrekas, *Department of Energy Office of Science*: john.mandrekas@science.doe.gov

Ken Miller, *ESnet*: ken@es.net

Andrew Wiedlea, *ESnet*: awiedlea@es.net

Table of Contents

Participants and Contributors	III
1 Executive Summary	1
1.1 Summary of Review Findings	2
1.2 Summary of Review Actions	3
2 Requirement Review Overview	4
2.1 Purpose and Process	4
2.2 Structure	4
2.2.1 Background	5
2.2.2 Case Study Methodology	5
2.3 ESnet	6
2.4 About ASCR	7
2.5 About the FES Program	7
3 Review Findings	8
4 Review Actions	11
5 FES Case Study Updates	12
5.1 International Fusion Collaborations	13
5.1.1 Case Study Summary	13
5.1.2 Discussion	13
5.2 Remote Observation and Participation of Fusion Facilities	14
5.2.1 Case Study Summary	14
5.2.2 Discussion	15
5.3 GA: DIII-D National Fusion Facility	16
5.3.1 Case Study Summary	16
5.3.2 Discussion	16
5.4 MIT PSFC	17
5.4.1 Case Study Summary	17
5.4.2 Discussion	18
5.5 PPPL	19
5.5.1 Case Study Summary	19
5.5.2 Discussion	20
5.6 Planning for ITER Operation	21
5.6.1 Case Study Summary	21
5.6.2 Discussion	22
5.7 Public-Private Partnerships in Fusion Research	24
5.7.1 Case Study Summary	24
5.7.2 Discussion	24
5.8 MPEX at ORNL	25
5.8.1 Case Study Summary	25
5.8.2 Discussion	26

5.9 MEC Experiment at SLAC	27
5.9.1 Case Study Summary	27
5.9.2 Discussion	27
5.10 LaserNetUS Program	28
5.10.1 Case Study Summary	28
5.10.2 Berkeley Lab Laser Accelerator (BELLA) Center Update	29
5.10.3 Discussion	30
5.11 Multifacility FES Workflows	31
5.11.1 Case Study Summary	31
5.11.2 Discussion	32
5.12 WDM and FES HPC Activities	33
5.12.1 Case Study Summary	33
5.12.2 Discussion	34
List of Abbreviations	35

1 Executive Summary

The US Department of Energy (DOE) Office of Science (SC) world-class research infrastructure provides the research community with premier observational, experimental, computational, and network capabilities. Each user facility is designed to provide unique capabilities to advance core DOE mission science for its sponsor SC program and to stimulate a rich discovery and innovation ecosystem. Research communities gather and flourish around each user facility, bringing together diverse perspectives. The continual reinvention of the practice of science — as users and staff forge novel approaches expressed in research workflows — unlocks new discoveries and propels scientific progress.

Within this research ecosystem, the high-performance computing (HPC) and networking user facilities stewarded by the SC's Advanced Scientific Computing Research (ASCR) program play a dynamic cross-cutting role, enabling complex workflows demanding high-performance data, networking, and computing solutions. The ASCR facilities enterprise seeks to understand and meet the needs and requirements across SC and DOE domain science programs and priority efforts highlighted by the formal requirements review methodology.

In May 2023, the Energy Sciences Network (ESnet) and the Fusion Energy Sciences program (FES) of the DOE SC organized an interim ESnet requirements review of FES-supported activities to follow up on the work started during the *2021 FES Network Requirements Review*. Preparation for these events included checking back with the key stakeholders: program and facility management, research groups, and technology providers. Each stakeholder group was asked to prepare updates to its previously submitted case study documents, so that ESnet could update the understanding of any changes to the current, near-term, and long-term status, expectations, and processes that will support the science activities of the program.

This review includes case studies from the following FES user facilities, experiments, and joint collaborative efforts:

- International fusion collaborations.
- Remote observation and participation of fusion facilities.
- General Atomics (GA): DIII-D National Fusion Facility.
- MIT Plasma Science and Fusion Center (PSFC).
- Princeton Plasma Physics Laboratory (PPPL).
- Planning for the International Thermonuclear Experimental Reactor (ITER) operation.
- Public-private partnerships in fusion research.
- Material Plasma Exposure eXperiment (MPEX) at Oak Ridge National Laboratory (ORNL).
- Matter in Extreme Conditions (MEC) Experiment at the SLAC National Accelerator Laboratory (SLAC).
- LaserNetUS program.
- Multifacility FES workflows.
- Whole-device modeling (WDM) and FES HPC activities.

The review participants spanned the following roles:

- Subject-matter experts from the FES activities listed previously.
- ESnet Site Coordinators Committee (ESCC) members from FES activity host institutions, including the following DOE labs and facilities: GA, Lawrence Berkeley National Laboratory (LBNL), MIT PSFC, ORNL, PPPL, and SLAC.
- Networking and/or science engagement leads from the ASCR HPC facilities.

- DOE SC staff spanning both ASCR and FES.
- ESnet staff supporting positions related to facility leadership, scientific engagement, networking, security, software development, and research and development (R&D).

In recent years, the research communities around the SC user facilities have begun experimenting with and demanding solutions directly integrated with HPC and data infrastructure. This rise of integrated-science approaches is well documented, and there is a broad need for integrated computational, data, and networking solutions. In response to these drivers, DOE has developed a vision for an Integrated Research Infrastructure (IRI)¹ 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.

The IRI vision is fundamentally about establishing new data-management and computational paradigms. Within these, DOE SC user facilities and their research communities build bridges across traditional silos to improve existing capabilities and create new possibilities. Implementation of IRI solutions will give researchers simple and powerful tools with which to implement multifacility research data workflows. This work will also extend analysis done on IRI patterns² and discuss ways future FES workflows can benefit from the approach.

1.1 Summary of Review Findings

The review produced several important findings from the case studies and subsequent virtual conversations:

- FES research, development, and operational activities rely heavily on network connectivity, both domestic and international, provided by ESnet. Efforts to upgrade PPPL and GA to 100 Gbps were successful and will reduce network capacity pressure in the coming years.
- Maintaining IPv6 peering across ESnet infrastructure, and with international partners, is critical to the process of science for a number of international FES experiments.
- The FES community has a long history of remote collaboration, which will continue as large international efforts (such as ITER, which features more than 30 countries in collaboration) come into operation.
- The FES community has adopted approaches where computational analysis is often done “closer” to where the experimental data reside rather than transferring data directly. In this paradigm, a user may be sitting at a site with ample local computational resources, but invokes software that runs “remotely” at a location that houses an instrument and dataset.
- The FES community should explore ways to better utilize computational resources that exist at collaborator sites, as well as DOE HPC facilities, as future research depends on the ability to effectively and efficiently utilize computational resources and increasing volumes of data. The advent of IRI will allow for better experiment and computation integration at FES facilities, as well as other possibilities, such as time-dependent workflows being managed across sites.
- Preparing for ITER operation remains an important focus for the FES community. Current timelines indicate that the facility’s first plasma will occur in late 2025, with full operation expected by 2035.
- ITER is anticipating that raw data will flow from the on-site location, through the Marseille facility, and on to international peers. ESnet, along with EU networking partner GÉANT, would be well connected into this facility when it is constructed and could serve as the transit to ensure data flow to US collaborators.

¹ <https://www.osti.gov/biblio/1984466>

² <https://www.osti.gov/biblio/2205078>

- ITER is open to data challenges and will consider doing these in 2024 in the run-up to first plasma. The expected data volumes remain around 2 PB per day, requiring at least 200 Gbps of connectivity to deliver data away from the facility. ESnet networking technology can support this currently, but future storage systems may be a bottleneck.
- FES workflows that span facilities (either experimental site to user, or experimental site to HPC facility) struggle with mechanisms to share and automate the credential exchange required by cybersecurity policies; this typically is required for workflow tools that attempt to migrate data and perform analysis.
- The FES community would rather not see all analysis default to using local computational resources. However, to distribute and manage computational demand, more unification and resource pooling across the FES complex will be needed to allow for fungible operation.
- FES simulation will incorporate the use of artificial intelligence (AI) and machine learning (ML) in the future, as the codes are adapted to run on next-generation machines and at a larger number of facilities.
- The FES community is exploring ways that cloud-provided storage and computation could be integrated into scientific workflows, particularly at facilities that are not able to scale local resources due to cost, lack of space, or lack of expertise to operate long-term storage pools. Investigations are underway to understand the costs and usability for FES workflows.
- DOE HPC allocations for FES are subject to annual renewal. This causes challenges for strategic planning and long-term investments, in a particular computing capability and workflow architecture.

1.2 Summary of Review Actions

Lastly, ESnet will follow up with review participants on a number of high-level actions identified. These items are listed as guidance for future collaboration, and do not reflect formal project timelines. ESnet will review these with FES participations on a yearly basis, until the next requirements review process begins:

- ESnet will continue to work with FES facilities and experiments to evaluate network connectivity, commercial peering options, and deployment of protocols and services that meet scientific requirements.
- ESnet will inform the FES community of opportunities and experiences with supporting networking to cloud computing providers.
- ESnet and GÉANT will discuss network connectivity peering to support ITER in the coming years. ITER is anticipating the construction of a data center in Marseille, France, that will allow for the sharing of experimental data.
- ESnet will collaborate with ITER on data challenges starting in 2024 in the run-up to first plasma. The expected data volumes remain around 2 PB per day, requiring at least 200 Gbps of connectivity to deliver data away from the facility. ESnet networking technology can support this currently, but future storage systems may be a bottleneck.

2 Requirement Review Overview

ESnet and ASCR use the requirements review process to discuss and analyze current and planned science use cases and anticipated data output of a particular program, user facility, or project to inform ESnet's strategic planning, including network operations, capacity upgrades, and other service investments.

2.1 Purpose and Process

The requirements review process, when performed regularly and comprehensively, surveys major science stakeholders' plans and processes to investigate data-management requirements over the next 5–10 years. Questions crafted to explore this space include the following:

- How, and where, will new data be analyzed and used?
- How will the process of doing science change over the next 5–10 years?
- How will changes to the underlying hardware and software technologies influence scientific discovery?

Requirements reviews help ensure that key stakeholders have a common understanding of the issues and the actions that ESnet may need to undertake to offer solutions. The ESnet Science Engagement Team leads the effort and relies on collaboration from other ESnet teams: Software Engineering, Network Engineering, and Network Security. This team meets with each individual program office within the DOE SC every three years, with an intermediate virtual update scheduled between the full review. ESnet collaborates with the relevant program managers to identify the appropriate principal investigators, and their information technology partners, to participate in the review process. ESnet organizes, convenes, executes, and shares the outcomes of the review with all stakeholders.

Requirements reviews are a critical part of a process to understand and analyze current and planned science use cases across the DOE SC. This is done by eliciting and documenting the anticipated data outputs and workflows of a particular program, user facility, or project to better inform strategic planning activities. These include, but are not limited to, network operations, capacity upgrades, and other service investments for ESnet as well as a complete and holistic understanding of science drivers and requirements for the program offices.

We achieve these goals by reviewing the case study documents, discussions with authors, and general analysis of the materials. The resulting output is a set of review findings and actions that will guide future interactions between FES, ASCR, and ESnet. These terms are defined as follows:

- **Findings:** key facts or observations gleaned from the entire review process that highlight specific challenges, particularly those shared among multiple case studies.
- **Actions:** potential strategic or tactical activities, investments, or opportunities that can be evaluated and potentially pursued to address the challenges laid out in the findings.

2.2 Structure

The requirements review process is hybrid, and relies on a combination of asynchronous and synchronous activities to understand specific facility and experimental use cases. The review is a highly conversational process through which all participants gain shared insight into the salient data-management challenges of the subject program/facility/project. Requirements reviews help ensure that key stakeholders have a common understanding of the issues and the potential actions that can be implemented in the coming years.

2.2.1 Background

Through a case study methodology, the review provides ESnet with information about the following:

- Existing and planned data-intensive science experiments and/or user facilities, including the geographical locations of experimental site(s), computing resource(s), data storage, and research collaborator(s).
- For each experiment/facility project, a description of the “process of science,” including the goals of the project and how experiments are performed and/or how the facility is used. This description includes information on the systems and tools used to analyze, transfer, and store the data produced.
- Current and anticipated data output on near- and long-term timescales.
- Timeline(s) for building, operating, and decommissioning of experiments, to the degree these are known.
- Existing and planned network resources, usage, and “pain points” or bottlenecks in transferring or productively using the data produced by the science.

2.2.2 Case Study Methodology

The case study template and methodology are designed to provide stakeholders with the following information:

- Identification and analysis of any data-management gaps and/or network bottlenecks that are barriers to achieving the scientific goals.
- A forecast of capacity/bandwidth needs by area of science, particularly in geographic regions where data production/consumption is anticipated to increase or decrease.
- A survey of the data-management needs, challenges, and capability gaps that could inform strategic investments in solutions.

The case study format seeks a network-centric narrative describing the science, instruments, and facilities currently used or anticipated for future programs; the network services needed; and how the network will be used over three timescales: the near term (immediately and up to two years in the future); the medium term (two to five years in the future); and the long term (greater than five years in the future).

The case studies address the following sections with review participants:

Science Background: a brief description of the scientific research performed or supported, the high-level context, goals, stakeholders, and outcomes. The section includes a brief overview of the data life cycle and how scientific components from the target use case are involved.

Collaborators: aims to capture the breadth of the science collaborations involved in an experiment or facility focusing on geographic locations and how datasets are created, shared, computed, and stored.

Instruments and Facilities: description of the instruments and facilities used, including any plans for major upgrades, new facilities, or similar changes. When applicable, descriptions of the instrument or facility’s compute, storage, and network capabilities are included. An overview of the composition of the datasets produced by the instrument or facility (e.g., file size, number of files, number of directories, total dataset size) is also included.

Process of Science: documentation on the way in which the instruments and facilities are and will be used for knowledge discovery, emphasizing the role of networking in enabling the science (where applicable). This should include descriptions of the science workflows, methods for data analysis and data reduction, and the integration of experimental data with simulation data or other use cases.

Remote Science Activities: use of any remote instruments or resources for the process of science and how this work affects or may affect the network. This could include any connections to or between instruments, facilities, people, or data at different sites.

Software Infrastructure: discussion of the tools that perform tasks, such as data-source management (local and remote), data-sharing infrastructure, data-movement tools, processing pipelines, collaboration software, etc.

Network and Data Architecture: the network architecture and bandwidth for the facility and/or laboratory and/or campus. The section includes detailed descriptions of the various network layers (local-area network [LAN], metropolitan-area network [MAN], and wide area network [WAN]) capabilities that connect the science experiment/facility/data source to external resources and collaborators.

IRI Readiness: Research communities that utilize DOE SC user facilities are experimenting with and demanding solutions integrated with HPC and data infrastructure. The Integrated Research Infrastructure Architecture Blueprint Activity (IRI-ABA) brought together domain experts from all DOE SC Programs to look for common patterns within diverse workflows across a range of scientific disciplines. This section asks if their workflows can be categorized into the three common patterns:

- Time-sensitive pattern.
- Data integration-intensive pattern.
- Long-term campaign pattern.

Cloud Services: if applicable, cloud services that are in use or planned for use in data analysis, storage, or computing, or other purposes.

Data-Related Resource Constraints: any current or anticipated future constraints that affect productivity, such as insufficient data-transfer performance, insufficient storage system space or performance, difficulty finding or accessing data in community data repositories, or unmet computing needs.

Data Mobility Endpoints: If a facility or experiment has dedicated infrastructure to facilitate data sharing, ESnet is interested in learning more about how it is constructed and maintained. ESnet maintains a set of well-tuned test endpoints and recommends regular testing to evaluate data-transfer capabilities.

Outstanding Issues: an open-ended section where any relevant challenges, barriers, or concerns not discussed elsewhere in the case study can be addressed by ESnet.

2.3 ESnet

ESnet is the high-performance network user facility for the US DOE SC and delivers highly reliable data transport capabilities optimized for the requirements of data-intensive science. In essence, ESnet is the circulatory system that enables the DOE science mission by connecting all its laboratories and facilities in the US and abroad. ESnet is funded and stewarded by the ASCR program and managed and operated by the Scientific Networking Division at LBNL. ESnet is widely regarded as a global leader in the research and education networking community.

ESnet interconnects DOE national laboratories, user facilities, and major experiments so that scientists can use remote instruments and computing resources as well as share data with collaborators, transfer large datasets, and access distributed data repositories. ESnet is specifically built to provide a range of network services tailored to meet the unique requirements of the DOE's data-intensive science.

In short, ESnet's mission is to enable and accelerate scientific discovery by delivering unparalleled network infrastructure, capabilities, and tools. ESnet's vision is summarized by these three points:

1. Scientific progress will be completely unconstrained by the physical location of instruments, people, computational resources, or data.
2. Collaborations at every scale, in every domain, will have the information and tools needed to achieve maximum benefit from scientific facilities, global networks, and emerging network capabilities.
3. ESnet will foster the partnerships and pioneer the technologies necessary to ensure that these transformations occur.

2.4 About ASCR

The mission of the ASCR program is to discover, develop, and deploy computational and networking capabilities to analyze, model, simulate, and predict complex phenomena important to the DOE. A particular challenge of this program is fulfilling the science potential of emerging computing systems and other novel computing architectures, which will require numerous significant modifications to today's tools and techniques to deliver on the promise of exascale science.

To accomplish its mission and address the challenges described previously, the ASCR program is organized into two subprograms:

1. The Mathematical, Computational, and Computer Sciences Research subprogram develops mathematical descriptions, models, methods, and algorithms to describe and understand complex systems, often involving processes that span a wide range of time and/or length scales.
2. The HPC and Network Facilities subprogram delivers forefront computational and networking capabilities and contributes to the development of next-generation capabilities through support of prototypes and test beds.

2.5 About the FES Program

The FES program has two goals: (1) expand the understanding of matter at very high temperatures and densities and (2) build the knowledge needed to develop a fusion energy source. Providing energy from fusion is one of the 14 Grand Challenges for Engineering in the 21st Century,³ and FES is the largest federal government supporter of research addressing the remaining obstacles to overcoming this challenge.

Together with its partner science agencies, FES supports a devoted workforce that has made impressive progress since the first fusion experiments over 60 years ago. Scientists and engineers at DOE national laboratories, at universities, and in private industry make progress each day. With public financial support for this fundamental research, fusion scientists are undertaking fundamental tests of fusion energy's viability using some of the most ambitious energy projects, the most powerful supercomputers, and the fastest networks in the world today.

³ <http://www.engineeringchallenges.org/challenges/fusion.aspx>

3 Review Findings

The requirements review process helps to identify important facts and opportunities from the programs and facilities profiled. These points summarize important information gathered during the review discussions surrounding case studies and the FES program in general:

- FES research, development, and operational activities rely heavily on network connectivity, both domestic and international, provided by ESnet. The coming years will see commissions of new experiments, additions of new collaborators, decreases in time-to-results, and increases in data volume that will place particular emphasis on the reliability and capacity for ESnet's international connections to Europe, and peering relationships with providers that reach other parts of the world (e.g., the Asia-Pacific region, South America, and Africa).
- Maintaining IPv6 peering across ESnet infrastructure, and with international partners, is critical to the process of science for a number of international FES experiments.
- The FES community has a long history of remote collaboration, which will continue as large international efforts (such as ITER, which features more than 30 countries in collaboration) come into operation. Remote use cases require various levels of technology and policy support to be successful, with the core requirement being stable network connectivity to support real-time communication as well as bulk and streaming data movement to share experimental results.
- The FES community has adopted approaches where computational analysis is often done closer to where the experimental data reside rather than transferring data directly. In this paradigm, a user may be sitting at a site with ample local computational resources, but invoke software that runs "remotely" at a location that houses an instrument and dataset. Tools such as MDSplus facilitate this data interaction, and this pattern is expected to remain an important use case to support in the future.
- The FES community should explore ways to better utilize computational resources that exist at collaborator sites, as well as DOE HPC facilities, as future research depends on the ability to effectively and efficiently utilize computational resources and increasing volumes of data. The advent of IRI will allow for better experiment and computation integration at FES facilities, as well as other possibilities, such as time-dependent workflows being managed across sites.
- ESnet connectivity is critical for FES facilities, and backups and capacity augmentations will be required in future years to ensure continuous operation. Efforts to upgrade PPPL and GA to 100 Gbps were successful, and will reduce network capacity pressure in the coming years.
- MIT PSFC's Alcator C-Mod data archive is approximately 150 TB in size and remains heavily accessed by the FES community. Efforts to understand how this archive can be kept active in the coming years are ongoing, as the hardware that provides the archive will require maintenance or augmentation. Upgrading local hardware and software to modernize the portal, or migrating the data to a dedicated facility, remain possibilities. Using cloud services is a possible way to continue to serve the Alcator C-Mod data archive. Concerns remain regarding whether the cloud will be scalable enough to address some of the tools that currently operate on these data, many of which rely on smaller transactions to extract portions of a dataset versus an entire bulk or streaming use case.
- Gyrokinetic simulation will be a major research element during the exascale era of computation. The data produced during runs of this simulation can grow to volumes beyond what current computing storage can handle. As a result of this, effort to reduce data size is required before data can be stored locally or transferred from ASCR HPC centers back to FES facilities.

- The TRANSP tool remains critical to FES analysis and can provide interpretive and predictive simulations of a full tokamak discharge. TRANSP can use both MDSplus and Globus to accomplish computational and data mobility tasks, respectively. As part of the process to define the ITER IMAS, TRANSP will undergo design and development to become compatible with the appropriate Intrusion Detection System (IDS) requirements.
- Preparing for ITER operation remains an important focus for the FES community. Current timelines indicate that the facility's first plasma will occur in late 2025, with full operation expected by 2035.
- ITER contains a number of diagnostic packages, consisting of thousands of data channels, and will eventually produce 2 PB of raw data each day through a gradual increase in capability. ITER will require more than an exabyte of data storage by the mid-2030s, and this estimate does not include the volume of analyzed and simulated data that will also be produced and archived.
- ITER is anticipating that raw data will flow from the on-site location, through the Marseille facility, and on to international peers. ESnet, along with EU networking partner GÉANT, would be well connected into this facility when it is constructed, and could serve as the transit to ensure data flow to US collaborators.
- ITER is open to data challenges and will consider doing these in 2024 in the run-up to first plasma. The expected data volumes remain around 2 PB per day, requiring at least 200 Gbps of connectivity to deliver data away from the facility. ESnet networking technology can support this currently, but future storage systems may be a bottleneck.
- DOE programs that span facilities and communities (e.g., Innovation Network for Fusion Energy [INFUSE], LaserNetUS) do not typically require a data architecture review to facilitate sharing of experimental results, or access to generalized pools of computational resources that can be utilized by participants. Solutions in this space can vary among facilities. While organic approaches have scaled to date, the lack of a cohesive and shared understanding of best practices as data volumes increase will begin to harm productivity. Having access to community-recommended approaches, and potentially more efficient data-transfer hardware and software, would benefit participants and lead to more efficient use of resources over time.
- As an emerging experiment, MPEX will adopt the use of DOE HPC resources for some aspects of the experimental workflow. This is expected to be in the form of the National Energy Research Scientific Computing Center (NERSC) and Oak Ridge Leadership Computing Facility (OLCF), although discussions are ongoing as MPEX is implemented. MPEX could potentially transfer TB to PB volumes of diagnostic data, output from experimental cameras, and simulation workflows to an external DOE HPC facility.
- The MEC Upgrade (MEC-U) facility at SLAC Linac Coherent Light Source (LCLS-II) will have a dedicated infrastructure for reading out detectors, and a shared infrastructure for data reduction, online monitoring, and fast feedback. It will use resources supplied by SLAC or remotely by NERSC. The underlying LCLS-II system, which MEC will take full advantage of, is designed to handle data rates of 100 Gbps and produce 100 PB of data per year.
- The LaserNetUS virtual organization (VO) is loosely coupled, and sites vary in terms of data volume produced and mechanisms to collect, store, and disseminate data to users. The community is working to establish norms around ways that data volumes can be more efficiently managed, stored, shared, and computed.

- FES workflows that span facilities (either experimental site to user, or experimental site to HPC facility) struggle with mechanisms to share and automate the credential exchange required by cybersecurity policies; this typically is required for workflow tools that attempt to migrate data and perform analysis.
- The FES community would rather not see all analysis default to using local computational resources. However, to distribute and manage computational demand, more unification and resource pooling will be needed across the FES complex to allow for fungible operation. The FES community should explore ways to better utilize computational resources that exist at collaborator sites, as well as DOE HPC facilities, as future research depends on the ability to effectively and efficiently utilize computational resources and increasing volumes of data.
- FES simulation will incorporate the use of AI and ML in the future, as the codes are adapted to run on next-generation machines and at a larger number of facilities.
- The FES community is exploring ways that cloud-provided storage and computation could be integrated into scientific workflows, particularly at facilities that are not able to scale local resources due to cost, lack of space, or lack of expertise to operate long-term storage pools. Investigations are underway to understand the costs and usability for FES workflows.
- DOE HPC allocations for FES are subject to annual renewal. This causes challenges for strategic planning and long-term investments in a particular computing capability or workflow architecture. If renewing an allocation at the same location is not possible, an experiment or facility may experience complications in data and workflow migrating to alternate facilities: adapting software to run on different systems, granting accounts to existing users, and sending most scientific data to another facility. Unified APIs and simplified methods to manage data between DOE HPC facilities could simplify the friction seen in these scenarios. Longer-duration (strategic) allocations of computing at ASCR facilities would allow the FES community to make more effective software investments.

4 Review Actions

ESnet recorded a set of recommendations from the FES and ESnet requirements review that extend ESnet's ongoing support of FES-funded collaborations. Based on the key findings, the review identified several recommendations for FES, ASCR, and ESnet to jointly pursue:

- ESnet will continue to work with FES facilities and experiments to evaluate network connectivity, commercial peering options, and deployment of protocols and services that meet scientific requirements.
- The advent of IRI will allow for better experiment and computation integration at FES facilities, as well as other possibilities, such as time-dependent workflows being managed across sites. ESnet will work with FES community members as they seek new ways to utilize computational resources that exist at collaborator sites, as well as DOE HPC facilities.
- ESnet will inform the FES community of opportunities and experiences with supporting networking to cloud computing providers.
- ESnet will work with GA, MIT PSFC, and PPPL on updates to their connectivity and resiliency plans.
- ESnet will continue to participate with FES community members in the planning for ITER operation. Current timelines indicate that the facility's first plasma will occur in late 2025, with full operation expected by 2035.
- ESnet and GÉANT will discuss network connectivity peering to support ITER in the coming years. ITER is anticipating the construction of a data center in Marseille that will allow for the sharing of experimental data.
- ESnet will collaborate with ITER on data challenges starting in 2024 in the run-up to first plasma. The expected data volumes remain around 2 PB per day, requiring at least 200 Gbps of connectivity to deliver data away from the facility. ESnet networking technology can support this currently, but future storage systems may be a bottleneck.
- ESnet will continue to advise INFUSE and LaserNetUS on ways they can encourage their community to adopt better practices for data mobility and use of remote computing resources. Having access to community-recommended approaches, and potentially more efficient data-transfer hardware and software, would benefit participants and lead to more efficient use of resources over time.

5 FES Case Study Updates

The case studies presented in this document are a written record of the current state of scientific process, and technology integration, for a subset of the projects, facilities, and principal investigators (PIs) funded by the FES program of the DOE SC. These updated case studies were reviewed virtually in May 2023. The case studies profiled and featured in the *2021 FES Network Requirements Review* include the following:

- International fusion collaborations.
- Remote observation and participation of fusion facilities.
- GA: DIII-D National Fusion Facility.
- MIT PSFC.
- PPPL.
- Planning for ITER operation.
- Public-private partnerships in fusion research.
- MPEX at ORNL.
- MEC Experiment at SLAC.
- LaserNetUS program.
- Multifacility FES workflows.
- WDM and FES HPC activities.

5.1 International Fusion Collaborations

International fusion collaborations enable US researchers to explore critical science and technology issues at the frontiers of magnetic fusion research, using the unique capabilities of the most advanced overseas research facilities. This community has a long history of effective collaborative research going back to the 1958 meeting on the Peaceful Uses of Atomic Energy in Geneva. The subsequent years have seen the collaborative environment consistently adopt new technology trends to facilitate information sharing and communication that spans international barriers.

A major emphasis of US international collaborations is superconducting facilities capable of true steady-state operation and large-scale fusion plasmas not currently accessible in the domestic program. Facilities such as Experimental Advanced Superconducting Tokamak (EAST, in China), Korea Superconducting Tokamak Advanced Research (KSTAR, in Korea), WEST (W Environment in Steady-state Tokamak, in Cadarache, France), and W7-X (Greifswald, Germany) offer access to devices that can in principle operate in steady-state. There are also strong collaborations with Joint European Torus (JET) in the UK, which is a more conventional pulsed plasma device.

5.1.1 Case Study Summary

- FES research, development, and operational activities rely heavily on international connectivity provided by ESnet. The coming years will see commissions of new experiments, additions of new collaborators, and increases in data volume that will place particular emphasis on the reliability and capacity for ESnet's international connections to Europe, and peering relationships with providers that reach other parts of the world (e.g., the Asia-Pacific region, South America, and Africa).
- Improvements to existing experiments, and development of new scientific infrastructure, are allowing for longer shot durations in the FES community. Historically a shot may have lasted only seconds, and future patterns indicate it may be possible to extend this to minutes, hours, or even days. Relatedly, the time between these shots can grow smaller, meaning a greater number of experimental results can be gathered during a run along with increasing data volumes for each. This time between shots is critical to the experimental process, placing extreme emphasis on network reliability and performance, particularly when spanning international boundaries.
- The EAST in Hefei, China, is a significant international facility that FES community members, such as GA, utilize. The IPv6 communications protocol is used extensively when communicating with EAST because it affords higher levels of performance. Maintaining IPv6 peering across ESnet infrastructure, and with international partners, is critical to the process of science for these interactions.

5.1.2 Discussion

The international fusion collaborations authors report no significant changes to the case study since the *2021 FES Network Requirements Review*. Many discussions regarding ITER have occurred, and these are outlined in [\[Section 5.6.2\]](#).

5.2 Remote Observation and Participation of Fusion Facilities

From a historical perspective, FES collaboration has been centered on exchanging scientists, ideas, and even data in a highly collaborative effort to advance the science of magnetically confined plasmas. The ITER project, currently under construction in Saint-Paul-lès-Durance France , is a great example of this collaborative spirit where 35 nations have joined together to build the world's largest tokamak.

Over the past decade, fusion research has seen an extension beyond remote scientists participating in experimental operation. Now, in some select cases, fusion researchers will operate and control a remote experiment. Thus, the abbreviation RCR has been extended in some cases from Remote Collaboration Room to Remote Control Room. What follows in this use case is an examination of remote control and operation as opposed to only remote participation, with the realization that this unique capability is possible in only a few select instances.

5.2.1 Case Study Summary

- The FES community has a long history of remote collaboration, which will continue as large international efforts (such as ITER, which features 35 countries in collaboration) come into operation. The community draws distinctions among three major types of remote use cases for their scientific workflows :
 - Remote observation: being able to observe aspects of a running FES experiment/instrument, typically through camera views or observable electronic diagnostics.
 - Remote participation: encapsulates the requirements of the previous category, but adds the ability to communicate with local collaborators to influence direction of experimentation (e.g., modifications that will be made prior to the next shot).
 - Remote control: encapsulates the previous two categories, but affords some level of control over the instrumentation during the experimental process.
- Remote use cases require various levels of technology and policy support to be successful. This comes in the form of either a dedicated environment or known toolsets along with specific information security policies that apply to both the source and users of the end-to-end workflow:
 - It is desirable to make the experience “seamless” so that the process of science is not impeded by technical or policy difficulties; without these considerations in place, the use case will not be successful.
 - Much of the prior work is being done to support the upcoming ITER use case, which will rely on strong international partnerships.
 - Remote use environments are present at the three major facilities to support collaboration: GA, MIT PSFC, and PPPL.
 - PPPL is currently planning for the Princeton Plasma Innovation Center (PPIC), expected in 2027, which will feature dedicated spaces to support remote collaboration.
- The ability to access live data streams from FES experiments will become necessary in the coming years, particularly as experimental facilities more routinely couple to collaborating computing facilities. This multifacility model will require advanced software to link experimental resources to storage and computing via the network infrastructure.
- Improvements to existing experiments, and development of new scientific infrastructure, are allowing for longer shot durations in the FES community. Historically a shot may have lasted only seconds, and future patterns indicate it may be possible to extend this to minutes, hours, or even days. Relatedly, the time between these shots can grow smaller, meaning a greater number of experimental results can be gathered during a run, along with increasing data volumes for

each. These changes to experimental behavior will place more emphasis on networking when remote use cases are present. Collaborators will participate for potentially longer periods of time, and the time between experiments will be critical to guiding next steps. Networks must be stable and predictable and have ample capacity for these needs. The time between shots during a fusion experiment is limited to tens of minutes across the FES facilities, implying that any analysis that can be done must be highly scheduled and responsive, or there is a risk that the output cannot be used to guide future shots. For this reason, many FES experiments rely on local, and instantly available, computational resources and tools versus leveraging other facilities in a coupled model.

- The FES community has adopted approaches where computational analysis is often done “closer” to where the experimental data reside rather than transferring data directly. In this paradigm, a user may be sitting at a site with ample local computational resources, but invoke software that runs remotely at a location that houses an instrument and dataset. Tools such as MDSplus facilitate this data interaction, and this pattern is expected to remain an important use case to support in the future.
- The FES community should explore ways to better utilize computational resources that exist at collaborator sites, as well as DOE HPC facilities, as future research depends on the ability to effectively and efficiently utilize computational resources and increasing volumes of data.

5.2.2 Discussion

Remote use of FES instruments and facilities continues to be strong post pandemic. Travel has restarted, but many within the community have found it possible to achieve the same level of scientific output via the remote control-room approaches put in place when travel was reduced. Overall, the network infrastructure remains important and sufficient to meet these needs.

The EAST facility is maintaining relationships with GA and MIT for the third shift of operations as documented in the case study. The use of remote screen sharing technologies (e.g., nomachine) remains steady, and the use of tools to facilitate computation and data transfer (e.g., MDSplus) has remained constant. Several facilities are working on ways to expose basic observation data through simpler mechanisms, such as dashboards that can be viewed via browser windows, along with APIs that may facilitate programmatic ways to access the same information.

The TRANSP tool continues to be used for remote analysis and simulation workflows where the data volumes can reach multiple MB in size. The development team is experimenting with sending images (e.g., multiple GB in size), as well as other potentially complex datasets. Some bottlenecks have been exposed, but so far these are not directly related to networks. The development team is in contact with the PPPL networking group and will advise on the progress of testing.

5.3 GA: DIII-D National Fusion Facility

GA has been an international leader in magnetic fusion research since the 1950s. The DIII-D National Fusion Facility, operated by GA for the US DOE, is the largest magnetic fusion research facility in the US. DIII-D research has delivered multiple innovations and scientific discoveries that have transformed the prospects for fusion energy.

The DIII-D National Fusion Facility, operated by GA for the US DOE, is a world-leading research facility pioneering the science and innovative techniques that will enable the development of nuclear fusion as an energy source for the next generation. Early tokamak designs, starting in the 1960s, were circular in cross-section, but GA scientists developed the “doublet,” a configuration with an elongated hourglass-shaped plasma cross-section. The Doublet I, II, and III tokamaks in the 1970s and 1980s showed that this approach allowed for a hotter and denser stable plasma. Further research led to a modification of Doublet III in the mid-1980s to DIII-D’s current D-shaped cross-section. Successes with this configuration inspired many other devices to adopt the D-shape, including JET (UK), TCV (Switzerland), ASDEX-U (Germany), JT-60U (Japan), KSTAR (Korea), and EAST (China).

5.3.1 Case Study Summary

- There has been some overall FES community interest in cloud services. Some use cases are easier to approach and could be adapted to a cloud with minimal modifications; others require study to understand the technical costs that would be associated. GA has investigated some cloud providers to manage backup data and some limited analysis use cases.
- Recent advancements by the Globus project at the University of Chicago may allow operation utilizing the IPv6 protocol. If possible, this would open an opportunity for GA to consider use of this tool for data mobility in its ongoing collaboration with the EAST in Hefei, China.
- ESnet connectivity is critical for FES facilities, and backups and capacity augmentations will be required in the future years to ensure continuous operation. GA has a 10 Gbps WAN connection to ESnet, and a 1 Gbps WAN backup connection through a commercial provider. Recent events, including a fiber cut in June 2021, have severely affected the ability of GA to perform daily operations. Upgrading the backup connection to support 10 Gbps to ESnet is viewed as a critical requirement to science productivity.
- The current 10 Gbps ESnet connection for GA is critical for the facility’s operation and must be maintained at all costs. In addition to the ability to access research and education connected facilities domestically and internationally, commercial peering to enable cloud services that support storage, audio, and video is critical to the process of science.
- The overall operation time of GA’s DIII-D will remain similar for the next five years, and it is anticipated that the rate of acquiring new data will continue to increase. From 2010 to 2020 the total amount of DIII-D data increased by an order of magnitude.

5.3.2 Discussion

The GA primary and secondary ESnet connections have been upgraded to 100 Gbps. This network upgrade has helped to alleviate the congestion the facility was experiencing when performing off-site mission work and backups.

GA is installing a new storage array for its data-transfer devices, and is engaged with ESnet and PPPL on testing data mobility between the facilities. This new hardware will be used for some of the multifacility workflows described in the use case.

5.4 MIT PSFC

The MIT PSFC seeks to provide research and educational opportunities for expanding the scientific understanding of the physics of plasmas, the “fourth state of matter,” and to use that knowledge to develop useful applications. The central focus of PSFC activities has been to create a scientific and engineering base for the development of fusion power. A diverse set of nonfusion plasma research areas and related technologies and applications are also actively pursued at the PSFC.

PSFC researchers study the use of strong magnetic fields to confine plasma at the high temperatures and pressures required for practical fusion energy. This research is conducted using on-site experimental facilities, theory and simulation, and collaboration with researchers at other facilities. PSFC scientists, students, and engineers perform experiments and develop technologies to confine and heat the plasma and to manage the interactions between the plasma and the reactor materials.

5.4.1 Case Study Summary

- FES use of cloud services is still being explored. Some use cases are easier to approach, and could be adapted to a cloud with minimal modifications; others require study to understand the technical costs that would be associated. Alcator C-Mod data, housed at MIT, is being explored as a possible cloud use case. Concerns remain regarding whether the cloud will be scalable enough to address some of the tools that currently operate on these data, many of which rely on smaller transactions to extract portions of a dataset versus an entire bulk or streaming use case.
- MDSplus remains critical to the operation of the FES community, and is widely used and deployed at experimental and analysis facilities. Modifications to the core software have helped FES keep pace with increases in networking capabilities and computational availability.
- FES simulation and theory workflows do not utilize MDSplus, and often rely on other tools that are native to the HPC facilities to accomplish data mobility tasks (e.g., Globus/GridFTP). Not all FES experimental facilities have similar hardware or software capabilities available, which can affect the efficiency of data transfer as part of these workflows.
- ESnet connectivity is critical for FES facilities, and backups and capacity augmentations will be required in future years to ensure continuous operation. MIT PSFC has a 1 G bps ESnet connection, through the MIT campus, but is interested in upgrading due to increased use cases that rely on external connectivity to support remote computing and storage, as well as increased levels of remote observation use cases. Upgrading the ESnet connection implies working with the MIT campus to upgrade LAN and MAN connectivity.
- The FES community is exploring ways that cloud-provided storage and computation could be integrated into scientific workflows, particularly at facilities that are not able to scale local resources due to cost, lack of space, or lack of expertise to operate long-term storage pools. Investigations are underway to understand the costs and usability for FES workflows.
- MIT PSFC’s Alcator C-Mod data archive is approximately 150 TB in size and remains heavily accessed by the FES community. Efforts are ongoing to understand how the archive can be kept active in coming years, as the hardware that provides the archive will require maintenance or augmentation. Upgrading local hardware and software to modernize the portal or migrating the data to a dedicated facility remain possibilities.

5.4.2 Discussion

MIT PSFC is still maintaining the Alcator C-Mod data archive, and has not experienced any major changes in the operation of this since the publication of the case study. Access to the archive remains similar to previous years, with a small increase in download volume. The collaboration with W7-X continues, with the facility downloading datasets frequently.

The MIT PSFC theory group reports no significant changes to the case study, but has observed that the amount of remote work, even after pandemic travel restrictions were lifted, remains high.

MIT PSFC (as well as some others at PPPL and GA) is working with collaborators via Scientific Discovery through Advanced Computing (SciDAC) partnerships. These efforts are not expected to require any increase in network support immediately, but will leverage existing technologies and relationships that span the FES community. These SciDAC projects include the following:

- Center for Advanced Simulation of RF - Plasma - Material Interactions (MIT PSFC).
- Computational Evaluation and Design of Actuators for Core-Edge Integration (PPPL).
- Development of High-Fidelity Simulation Capabilities for ELM-free Design Optimization (PPPL).
- High-fidelity Digital Models for Fusion Pilot Plant Design (PPPL).
- FRONTIERS in Leadership Gyrokinetic Simulation (GA).

5.5 PPPL

The U.S. DOE PPPL is a collaborative national center for fusion energy science, basic science, and advanced technology. PPPL is dedicated to developing the scientific and technological knowledge base for fusion energy as a safe, economical, and environmentally attractive energy source for the world's long-term energy requirements. The laboratory has three major missions:

1. Develop the scientific knowledge and advanced engineering to enable fusion to power the United States and the world.
2. Advance the science of nanoscale fabrication for technologies of tomorrow.
3. Further the development of the scientific understanding of the plasma universe from laboratory to astrophysical scales.

For 70 years, PPPL has been a world leader in magnetic confinement experiments, plasma science, fusion science, and engineering. As the only DOE national laboratory with a FES mission, PPPL aspires to be the nation's premier design center for the realization and construction of future fusion concepts. PPPL also aims to drive the next wave of scientific innovation in plasma nanofabrication technologies to maintain US leadership in this critical industry of the future. Further, Princeton University and PPPL develop the workforce of the future by educating and inspiring world-class scientists and engineers to serve the laboratory and national interest.

5.5.1 Case Study Summary

- Gyrokinetic simulation will be a major research element during the exascale era of computation. The data produced during runs of this simulation can grow to volumes beyond what current computing storage can handle. As a result of this, effort to reduce data size is required before data can be stored locally, or transferred from ASCR HPC centers back to PPPL. The following steps are needed to properly support XGC:
 - XGC will limit output to fit within memory regions of current (and future) ASCR HPC resources.
 - PPPL and ASCR HPC facilities will require storage upgrades to offer temporary locations for XGC output. PPPL will double its capacity in the coming years to offer PBs of storage space.
 - PPPL is upgrading its data architecture to install new data-transfer hardware, is adopting Globus as a software package, has upgraded local storage, and will be working with ESnet to increase network capacity.
- XGC can produce a simulation of turbulence transport in an ITER-like plasma for a given equilibrium time slice using ORNL's Summit in two days of run time, but the resulting dataset is approximately 50 PB in size. This volume must be reduced before storage or data transfer, and often only a small portion (typically 1–10 TB) can be sent back to PPPL.
 - Future machines are expected to produce data that can approach 300 PB in size.
 - Full data transfer for volumes this large would require multiple Tbps network connections on ESnet between the ASCR HPC facilities and PPPL.
 - Approaches to optimize bulk data transfer, and streaming, will be required even for reduced datasets.
- XGC is exploring ways to leverage cloud storage as part of the experimental workflow. Due to the relative performance, as well as the volume and potential costs, cloud storage is not expected to replace local resources, but could be used to facilitate data backups, or use cases

that require sharing. Additional work in this area could investigate cloud computing for multidataset analysis.

- The TRANSP tool remains critical to FES analysis, and can provide interpretive and predictive simulations of a full tokamak discharge. TRANSP can use both MDSplus and Globus to accomplish computational and data mobility tasks, respectively. As part of the process to define the ITER IMAS, TRANSP will undergo design and development to become compatible with the appropriate IDS requirements. This marks an early step for the FES community to adopt universal standards for cataloging tokamak data standards.
- The ECP-WDM code, once complete, will undergo a period of distributed community analysis. These simulation data will need to be available for a minimum of five years to provide data for developing fusion surrogate models and digital twins.
- PPPL networking requirements have steadily increased over the years as the facility has taken more active roles in existing global FES experiments, such as KSTAR, and prepares for the future requirements of ITER. PPPL currently connects through Mid-Atlantic GigaPop for Internet2 (MAGPI), and has upgraded its local networking environment to accept a 100 Gbps WAN connection from ESnet. It is pursuing a primary ESnet 100 Gbps connection, and would also like to pursue a backup connection through diverse paths and providers.
- PPPL has a number of use cases that leverage the Google Cloud Platform for storage of data and the execution of software codes; the cloud-based storage may take several TB of space in the coming years. The current usage patterns for the data are not intense, but these may grow as AI/ML-informed simulations are added to workflows. The usage can come from domestic and international partners.
- PPPL has migrated some data analysis tasks into cloud storage and is exploring others as it prepares for upgrades to NSTX-U and the affiliated computational and software requirements. There is an effort to provide container-based versions of tools (e.g., TRANSP) as an alternative to running on PPPL computing resources.

5.5.2 Discussion

PPPL has completed upgrades to its internal core network, resulting in 100 Gbps capacity between the ESnet demarcation and the core routers, switches, and firewalls. This now facilitates a full 100 Gbps path to the two ESnet connections (i.e., Washington, DC, and New York), as well as a path to the Princeton campus. A next iteration in the network design will involve increasing resiliency through a second ESnet router. PPPL will start this conversation with ESnet later in 2024 .

PPPL is in the process of upgrading its campus perfSONAR testing infrastructure and DTNs. All these servers will be 100 Gbps capable when complete, with a 10 Gbps backup connection. The hardware has been specified and delivered; some configuration and testing will be required. ESnet will be available to assist PPPL with this as needed. One area of need will be designing DTN “clusters” that use Globus.

Design of the PPPL PPIC continues, but there are no updates on the timeline since the case study was published.

With the network and server upgrades, PPPL will focus on the performance to the international partners mentioned in [Section 5.1.1], with KSTAR being a high priority due to the amount of data transferred on a routine basis. ITER targets are further in the future, but PPPL staff are staying engaged in the planning process.

5.6 Planning for ITER Operation

The ITER tokamak is the most ambitious fusion experiment ever undertaken. ITER is a magnetic confinement device that heats hydrogen isotopes to temperatures up to 100 million degrees, forming a plasma and forcing nuclei to fuse to create fusion energy. ITER brings together 35 nations and 7 major partners (China, the European Union, India, Japan, Korea, Russia, and the United States) to collaborate on building the world's largest tokamak, designed to achieve sustained high-fusion power (500 MW, 500–550 seconds) by the mid-2030s, and to potentially achieve full steady-state operation thereafter. ITER is in Saint-Paul-lès-Durance, France, only a 350 km drive from CERN, the location of another major global scientific collaboration with significant US participation on the Large Hadron Collider (LHC).

5.6.1 Case Study Summary

- Preparing for ITER operation remains an important focus for the FES community. Current timelines indicate that the facility's first plasma will occur in December 2025, with full operation expected by 2035. The facility will have periods of reduced operation through 2032, and full operation is expected by 2035.
- As ITER is designed, there will be new requirements to facilitate access and sharing of experimental data. These include fast networks, the ability to use resources during time windows of experimentation, and adaptable software that can run in multiple locations and access data where they are located. Some of this discussion may involve the ITPA.
- The ITER software stack, which encompasses many aspects of the complete workflow, including data mobility to support analysis activities, is still being designed and subject to participation by several FES community members.
- Given the extensive experience developed by ESnet in meeting US networking needs for large collaborative and international scientific projects, ASCR, ESnet, and FES should perform a formal assessment of the ITER data analysis and network requirements well in advance of ITER first plasma. The ITER Organization (IO) should be engaged in this assessment, given that IO decisions soon may have important implications for the US and other ITER members regarding the timeliness of data access and the quality of remote participation.
- ITER contains over 50 major diagnostic packages, consisting of thousands of data channels, and will eventually produce 2 PB of raw data each day through a gradual increase in capability. ITER will require more than an exabyte of data storage by the mid-2030s. This estimate does not include the volume of analyzed and simulated data that will also be produced and archived. ITER will commence operations with much less data production per day (~ 20 TB) during the first phase of plasma operation (engineering commissioning, first plasma, and engineering operations) planned for 2026.
- ITER peak data production rates are not fully known as of the *2021 FES Network Requirements Review*. However, aggregate estimates of a 20 TB/day data production rate have been made for the engineering operations phase.
- In present fusion facilities, a typical experiment is a collection of similar discharges executed over a single day or partial day, with each discharge typically lasting < 10 s. Initially, discharges in ITER will be of similar duration per pulse, but with the goal of reaching 500 s. by the mid-2030s. However, unlike existing experiments, ITER may run experiments over multiple days.
- Development and implementation of the policies and infrastructure that support data sharing are crucial needs for the FES community in preparation for ITER experimentation. Having access to those data in a timely manner is critical to advancing R&D activities, as well as remote participation in ITER operation.

- Networking to support ITER remains undecided and opaque; this includes aspects of domestic connectivity within France as well as international connectivity to support distributed collaborators. Options for connectivity could involve the French NREN RENATER⁴ or directly connecting to the pan-European REN GÉANT⁵
- As the FES community prepares for ITER, the multifacility use case will become more important as the ITER data volumes will far exceed the storage and processing capacity of any of the major FES facilities. Integration with DOE HPC facilities is critical. Exploring Science DMZ architectures at all FES facilities will be required to ensure that a baseline for data mobility can be achieved.
- The ITER computing and data-management model is still under development but is expected to consist of a main data center located at the instrument, and some set of policies and technology that will be adopted to manage distributed data dissemination to partners around the world. ITER data management will require coordination from the US FES community to ensure efficient and equitable access.
- ITER data rates are still projected to be 50 Gbps (400 Gbps) at peak operation. The ITER timeline, as of the *2021 FES Network Requirements Review*, was as follows:
 - First plasma: December 2025.
 - Additional commissioning and construction: through December 2028.
 - Pre-fusion power operations (Phase 1): December 2028 through January 2030.
 - Pre-fusion power operations (Phase 2): June 2032 through March 2034.
 - Nuclear assembly: 2035.
 - Regular operations: December 2035.

5.6.2 Discussion

ITER planning has been active since the *2021 FES Network Requirements Review*. Some deviations from what was discussed previously have occurred, but many things remain the same. One notable change to the overall workflow is the ongoing design of ITER data centers that will be constructed on site, as well as closer to the location of the international networking exchange point (Marseille, France), and how this will impact the overall flow of science data.

ITER is anticipating that raw data will flow from the on-site location, through the Marseille facility, and on to international peers. ESnet, along with EU networking partner GÉANT, would be well connected into this facility when it is constructed, and could serve as the transit to ensure data flow to US collaborators.

ITER is open to data challenges and will consider doing these in 2024 in the run-up to first plasma. The expected data volumes remain around 2 PB per day, requiring at least 200 Gbps of connectivity to deliver data away from the facility. ESnet networking technology can support this currently, but future storage systems may be a bottleneck.

If ITER were to follow the data challenge model of HL-LHC, it would need to follow a model that allowed a gradual ramp up toward the 2 PB a day target. For example, that could be aiming for 10% volume with the first challenge (in early 2024), moving toward 50% shortly after, and then 100% before first plasma in 2025. The tools would need to be selected and some sites identified. This is within reason for FES partners, provided they coordinate and develop a plan in the coming years.

⁴ <https://www.renater.fr/en/accueil-english/>

⁵ <https://www.geant.org/About>

The ITER schedule has not changed significantly, but has experienced some compression due to delays. The delays will be made up with significant ramp up between 2025 and 2035, where the data volumes may increase at a steeper rate during the testing and build phases. These assumptions were accurate as of May 2023:

- First plasma: December 2025.
- Regular operations: December 2035.

5.7 Public-Private Partnerships in Fusion Research

DOE FES provides funds for business awards to assist applicants seeking access to the world-class expertise and capabilities available across the US DOE complex. This is one component of the Innovation Network for Fusion Energy (INFUSE), a DOE initiative to provide the fusion industrial community with access to the technical and financial support necessary to move new or advanced fusion technologies toward realization with the assistance of the national laboratories. The objective of INFUSE is to accelerate basic research to develop cost-effective, innovative fusion energy technologies in the private sector.

The INFUSE program will accelerate fusion energy development in the private sector by reducing impediments to collaboration involving the expertise and unique resources available at DOE laboratories. This will ensure the nation's energy, environmental, and security needs by resolving technical, cost, and safety issues for industry.

5.7.1 Case Study Summary

- The INFUSE program features public-private partnerships with non-DOE entities that are funded to perform aspects of FES research. Many of these entities are unfamiliar with mechanisms to interact with DOE SC facilities, including ASCR HPC centers and ESnet.
- DOE programs that span facilities and communities (e.g., INFUSE) do not typically require a data architecture review to facilitate sharing of experimental results; solutions in this space can vary among facilities. While organic approaches have scaled to date, the lack of a cohesive and shared understanding of best practices as data volumes increase will begin to harm productivity. Having access to community-recommended approaches, and potentially more efficient data-transfer hardware and software, would benefit participants and lead to more efficient use of resources over time.
- DOE programs that span facilities and communities (e.g., INFUSE) do not include access to generalized pools of computational resources that participants can utilize. While participants can pursue these resources independently from DOE HPC facilities, this is a secondary step that must be managed independently. Having access to computational resources, and potentially more efficient data-transfer and analysis tools, would benefit participants and lead to more efficient use of resources over time.

5.7.2 Discussion

INFUSE has had leadership changes since the *2021 FES Network Requirements Review*, but this has not changed the program or content of the case study. Overall, more than 30 companies have received awards. None of these awards is anticipated to have significant amounts of research data that will need curation. Some awards focus on simulation and modeling: these will be encouraged to use existing tools at DOE HPC centers.

ESnet gave a talk at the annual virtual INFUSE workshop in December of 2023.

5.8 MPEX at ORNL

MPEX is a next-generation linear plasma device that will support study of the way plasma interacts long term with the components of future fusion reactors. MPEX represents a shift from the historical direction of the plasma-material interaction field, which for many years focused on the effect that materials had on plasma, but not on the effect that plasma had on materials.

MPEX is a linear plasma device to address the challenges of plasma-material interactions for future fusion reactors. This includes the capability to test material samples that have been pre-irradiated with neutrons in fission reactors such as the High Flux Isotope Reactor (HFIR). MPEX will be a steady-state device able to expose components to long pulses up to deuterium ion fluences of $1e+31$ per square meter. MPEX is unique in its capabilities to reach the plasma conditions expected in divertor plasma, utilizing a novel high-power helicon plasma source, as well as an electron heating and ion heating system. In situ diagnostics as well as dedicated surface analysis tools (e.g., FIB/SEM, IBA-NRA) will monitor the plasma-exposed materials in-vacuo to provide information on the surface evolution and hydrogen transport in the material.

5.8.1 Case Study Summary

- The MPEX experiment at ORNL is under design and will be operational by 2027. The MPEX project at ORNL is currently in the DOE 413.3b project phase and has passed CD-1.
- Visible light cameras will be used for measuring the target surface and will produce raw video data streams at 1 Gbps. Up to six cameras can be used at various angles during a run period. These can generate just under 4 TB of raw data frames per hour, or up to 24 TB per hour if all cameras are operating.
- A single infrared camera can be used for measuring surface materials interactions, and it is estimated to produce raw data rates at 9 Gbps or 32 TB per hour .
- The standard short-pulse use case will produce the following:
 - An estimated 50 GB of scientific data per run day, with 100 run days per year. This is an estimated 5 TB of data per year.
 - Approximately 35,000 archived signals for operational data are stored in a relational database. The archived data consume approximately 17 GB/day or 6.2 TB/year.
- A second use case, consisting of a longer pulse (two weeks of continuous operation), has the potential to generate 1 PB of scientific experimental data. The camera rates listed previously will apply as well, but will be limited to the two-week operational period.
- MPEX will expose data via recommended mechanisms that ORNL and OLCF support (e.g., HTTP portals, RSYNC, SCP). Data long-term storage and archiving are expected to be managed at ORNL.
- MPEX is designing experimental workflow, and will approach data handling like other large-scale experiments: saving raw data to archival storage, and generating a system to reduce information to formats that are easy to process and share.
- Data will be produced mainly on MPEX with its installed diagnostics. Collaborators will conduct some post-mortem analysis of material samples in other locations; they will have access to raw and processed data on MPEX and might transfer parts of data for further analysis or processing.

- As an emerging experiment, MPEX will adopt the use of DOE HPC resources for some aspects of the experimental workflow. This is expected to be in the form of NERSC and OLCF, although discussions will be ongoing as MPEX is implemented. MPEX could potentially transfer TB to PB volumes of diagnostic data, output from experimental cameras, and simulation workflows to an external DOE HPC facility.

5.8.2 Discussion

The MPEX at ORNL authors report no significant changes to the case study since the *2021 FES Network Requirements Review*.

5.9 MEC Experiment at SLAC

The MEC experiment, collocated with the LCLS X-ray free-electron laser (XFEL), is at SLAC. The overall scientific goal of the instrument is to deliver ultrashort X-ray pulses to probe the characteristics of matter.

The LCLS XFEL is an open-access user facility at SLAC that delivers ultrashort X-ray pulses nine orders of magnitude brighter than any prior source, able to probe the characteristics of matter with unprecedented spatial and temporal precision. The MEC instrument at LCLS, funded by the DOE SC FES program, combines the XFEL with high-power, short-pulse lasers to produce and study high energy density (HED) plasmas, and to develop the fundamental understanding of plasmas and matter in extreme environments. This has driven a remarkably rich array of high-profile scientific results with applications in fusion energy, isotope production, advanced materials, and medical and nuclear technology.

5.9.1 Case Study Summary

- The LCLS XFEL is an open-access user facility at SLAC that delivers ultrashort X-ray pulses able to probe the characteristics of matter. The MEC instrument at LCLS combines the XFEL with high-power, short-pulse lasers to produce and study HED plasmas.
 - The MEC-U proposes a major upgrade to MEC that would significantly increase the power and repetition rate of the high-intensity laser system to the petawatt level.
 - The CD-1 of the MEC-U was completed Q4 FY2021, and the upgrade has an estimated duration of five years from CD-1 to CD-4. The MEC-U data system is expected to be complete and ready for beam time by June 2026.
 - MEC-U plans to use all existing LCLS-II cyberinfrastructure.
- The MEC-U facility at SLAC LCLS-II will have a dedicated infrastructure for reading out detectors, and a shared infrastructure for data reduction, online monitoring, and fast feedback. It will use resources supplied by SLAC or remotely by NERSC:
 - The underlying LCLS-II system, which MEC will take full advantage of, is designed to handle data rates of 100 Gbps and produce 100 PB of data per year.
 - MEC dataset sizes are highly dependent on the physics case being studied. Based on estimated laser pulses and beam allocations, datasets are expected to be a minimum of 10 GB to a maximum 100 TB, with individual file sizes not exceeding 1 TB. The total number of files per experiment can range from a few hundred to 10,000 with a median of 3,000.
 - MEC data transfer will utilize LCLS systems, with the main data-transfer tools being bbcop and XRootD on-site data-transfer hardware. Other tools are also supported on SLAC's DTNs: scp, sftp, rsync, and a Globus endpoint for data transfers.

5.9.2 Discussion

The MEC at SLAC authors report no significant changes to the case study since the *2021 FES Network Requirements Review*. MEC-U is still under design and still plans to utilize much of the LCLS infrastructure when constructed and operated.

5.10 LaserNetUS Program

LaserNetUS is a program established by DOE FES to help restore the US's once-dominant position in high-intensity laser research. LaserNetUS will provide US scientists increased access to the unique high-intensity laser facilities at 10 institutions: University of Texas at Austin, Ohio State University, Colorado State University, the University of Central Florida, University of Nebraska-Lincoln, University of Rochester, SLAC, LBNL, Lawrence Livermore National Laboratory, and Université du Québec. LaserNetUS provides time to users to run laser-based experiments. The actual amount of data involved during a run is currently relatively small (a few GB is common), and there is little issue with storage or transmission of these data.

5.10.1 Case Study Summary

- LaserNetUS VO is loosely coupled, and sites vary in terms of data volume produced and mechanisms to collect, store, and disseminate data to users.
 - Laser capability dictates factors such as power, pulse length, and number of shots that can be run during an experimental period.
 - Typical shot output is several MB to as much as a GB. An entire experimental run, consisting of tens to hundreds of shots over the course of several days, may approach hundreds of GBs.
 - Managing the data is at the discretion of each site. Typical approaches could be requiring the use of portable media, integration to enterprise cloud storage, or the ability to transfer data from network-enabled portal systems that are on premises.
 - Site users are responsible for data analysis and data reduction, which they do at their home institutions. This includes simulations, which are used to predict the outcome of experiments, or the experimental data are used to guide and benchmark the simulations.
- LaserNetUS does not maintain a suggested set of policies and procedures to address data management and mobility within, or between, facilities.
- LaserNetUS provides time to users to run laser-based experiments utilizing a collection of high-power, short-pulse lasers operated by 10 participating institutions and facilities. These laser systems are often combined with long-pulse “driver” lasers to achieve high density and pressure or with other beams.
 - The actual amount of data involved during a run is small (a few GB is common).
 - Each facility has its own research program that is, to varying degrees, separate from LaserNetUS and data associated with the facilities' local programs.
 - There is no standard approach to handle data mobility, and often facilities rely on nontechnical approaches (e.g., portable media) to transfer research data.
- DOE programs that span facilities and communities (e.g., LaserNetUS) do not include access to generalized pools of computational resources that participants can utilize. While participants can pursue these resources independently from DOE HPC facilities, this is a secondary step that must be managed independently. Having access to computational resources, and potentially more efficient data transfer and analysis tools, would benefit participants and lead to more efficient use of resources over time.
- DOE programs that span facilities and communities (e.g., LaserNetUS) do not typically require a data architecture review to facilitate sharing of experimental results; solutions in this space can vary among facilities. While organic approaches have scaled to date, the lack of a cohesive and

shared understanding of best practices as data volumes increase will begin to harm productivity. Having access to community-recommended approaches, and potentially more efficient data transfer hardware and software, would benefit participants and lead to more efficient use of resources over time.

5.10.2 Berkeley Lab Laser Accelerator (BELLA) Center Update

The laser-plasma accelerator (LPA) is a revolutionary technology already starting to fulfill its promise of making particle accelerators smaller and more affordable and delivering beams with unique properties. The BELLA Center uses some of the world's fastest and most powerful lasers to drive LPAs and to address various scientific needs. These include future high-energy physics colliders, ultrafast photon and particle beams for probing matter, ultra-intense beams for high-energy-density science, novel beams for biological studies, and national security and industry applications. By providing hands-on access to cutting-edge technology, computational tools, and theoretical models, the BELLA Center is also an ideal environment for training the next generation of particle accelerator researchers.

High average power lasers that produce pulses with kHz repetition rates are increasingly available for experiments of laser-particle acceleration and Inertial Fusion Energy (IFE) physics. Commercialization of IFE will require MHz repetition rates.⁶ The increase in data rates and total storage required when going from the current maximum repetition rate of 1–5 Hz to kHz–MHz will be several orders of magnitude and cannot be dealt with using the current methods of on-the-fly analysis and reduction, transport, and storage.

Currently, BELLA experiments at 1 Hz produce about 250 GB of raw data per experimental day. Upgrades to bandwidth, computing power, and data storage will become necessary when increasing the data rate by three orders of magnitude to kHz, as is envisaged with the kBELLA facility currently in the proposal phase. Where fast feedback is needed, on-the-fly (in situ) analysis will require accelerated, local (edge) computing resources near the data source (e.g., FPGA- or GPU-accelerated nodes). AI/ML processing and pre-trained models will become important to reduce and act on data in real time. Some of the processing can be batched and streamed to nearby HPC resources, to be parallel processed (e.g., at NERSC). Low-latency, tens of GBit/s bandwidth between BELLA and NERSC would be required for the initial raw data transport at kHz. A prototypical combination of experimental and HPC facilities in a programmable superfacility for automation is being explored in an upcoming LBNL LDRD.

Advanced laser technologies aimed at high average power and high repetition rates (kHz to MHz) require unprecedented data needs. For example, coherently combined fiber laser systems require control from initially dozens to later hundreds of laser channels in parallel, for which analysis can be parallelized in a similar manner using a superfacility approach for batch processing on HPC. Also considering the high repetition rates of pulsed operation, on-shot pulse characterization for applications will require R&D for high repetition rate diagnostics and reconstruction algorithms, which could benefit from edge computing resources.

⁶ Inertial Fusion Energy Final Report of the 2022 Fusion Energy Sciences Basic Research Needs Workshop (<https://events.bizzabo.com/IFEBRN2022/home>)

5.10.3 Discussion

LaserNetUS has some minor changes since the publication of the *2021 FES Network Requirements Review*. The University of Central Florida has joined the LaserNetUS program, replacing the University of Michigan . The community itself has gained valuable experience since the start of the program and is now midway through cycle 5 of its funding. Some notable changes include the following:

- Establishing a formalized user group to assist experimenters and facilities.
- Planning for semi-regular meetings (some virtual, some in person).
- Forming the following two committees:
 1. Laser Diagnostics: user-run and organized, with a focus on ways to more easily share experimental data.
 2. Computing: also user-run and organized, aiming to help users of all skill levels approach the tasks of modeling, simulation, and analysis using common tools and facilities.

LaserNetUS has no specific charge to investigate the use of ML and AI in its normal operations, but enough users are interested in the topic that it will become part of the community discussion. There will not be formal support, but best common practices can be established. A number of the experimental facilities do not have permanent access to any form of computing, although some have relationships with campus computing resources when these are available. Those that do make HPC available are finding that the amount of data generated via the simulation activity can exceed the experimental data they are collecting.

ESnet can help LaserNetUS with some of the mechanics of data mobility when needed, in particular by recommending tools and configurations that have worked for others.

5.11 Multifacility FES Workflows

A number of pilot use cases and demonstrations have been conducted over the years to couple FES workflows to existing DOE HPC facilities. This experimentation had the modest goals of trying to reduce the number of deployed HPC resources within the FES ecosystem, and utilize higher-performing and better-supported resources. Early efforts identified several areas of improvement, and future goals indicate a desire to continue, provided that some areas of friction can be reduced.

To expand the quality, variety, and quantity of analysis performed for fusion experiments, the use cases here describe workflows to send data generated at experimental machines to remote computing centers in near real time for further analysis/modeling. This use case will focus specifically on the use of remote computing centers for analysis and support during experimental operation, in near real time to aid researchers in providing rapid analysis and shot assessment required to make control-room decisions on the direction of the experiment. The analysis proposed here is in support of the experiments and is distinct from the analysis/simulation of the data, which can come many days or weeks after an experiment is run.

5.11.1 Case Study Summary

- In the FES context, a “Multi/Coupled Facility Workflow” is not considered to be a pairwise operation between two specific entities across a network substrate, as in other use cases (e.g., a Light Source using ESnet to reach an ASCR HPC facility). FES views the multifacility use case as having numerous points:
 - Instrument and local operations staff at once location.
 - Collaborating / participating groups at a number of remote facilities linked via communications tools and remote diagnostics to understand and observe experimental progress.
 - One or more computational and storage facilities where dedicated analysis resources are available for inter-shot diagnostics.
 - Network infrastructure linking all these that carries both communications and data transmission.
- FES workflows that span facilities (either experimental site to user, or experimental site to HPC facility) struggle with mechanisms to share and automate the credential exchange required by cybersecurity policies; this typically is required for workflow tools that attempt to migrate data and perform analysis.
- FES use of cloud services is still being explored. Some use cases are easier to approach and could be adapted to a cloud with minimal modifications; others require study to understand the associated technical costs.
- The FES community would rather not see all analysis default to using local computational resources. However, to distribute and manage computational demand, more unification and resource pooling across the FES complex is needed to allow for fungible operation.
- The ability to support FES workflows at ASCR HPC facilities requires addressing several key areas:
 - Establishing a dedicated pool of computing resources that can be accessed without having to wait in a queue.
 - Enabling a form of system-wide scheduling; namely ensuring that all components of a workflow (computation, storage, networking, and software at all portions of the end-to-end path) are ready when the analysis procedure starts.

- Allowing the worker nodes on an HPC system to retrieve a remote dataset directly for processing.
- Configuring the security of the infrastructure so that it allows for automating the authentication on multiple systems in multiple locations.
- Creating a unified API for all DOE HPC facilities, which supports the multifacility nature of the workflows and allows for workflow portability.
- Creating an intelligent software stack to observe and manage multifacility use cases.
- Ensuring that the networks that link facilities have mechanisms to guarantee performance (latency, bandwidth, etc.), and also allow for the reduction in delays between experimental shots.
- The FES community should explore ways to better utilize computational resources that exist at collaborator sites, as well as at DOE HPC facilities, as future research depends on the ability to effectively and efficiently utilize computational resources and increasing volumes of data.

5.11.2 Discussion

The Multifacility FES Workflows authors report no significant changes to the case study since the *2021 FES Network Requirements Review*. Work has continued on some of the profiled use cases, with some changes becoming necessary due to machines at NERSC being de-commissioned (i.e., Cori), and new resources coming online (i.e., Perlmutter and increased use of SPIN).

5.12 WDM and FES HPC Activities

This case study combines the collective works of several PIs from the FES community. Their work spans the overall field of HPC use and the relationship to FES research as a whole. Due to the overlapping nature of some of the facilities featured in this document, references to prior sections that discuss the overall technical capabilities of a site or project are used.

5.12.1 Case Study Summary

- OMFIT is a modeling and experimental data analysis software used in the FES community. OMFIT will adapt existing workflows to advance modeling approaches that use HPC resources, and will be more widely deployed as the community prepares for ITER. It is expected that OMFIT will expand to allow for the use of more analysis codes, at more locations, with more participants. Improvements to the systems that handle data mobility, and ways to automate authentication and authorization, are expected.
- The process of FES simulation workflows has and will continue to change in the coming years as new codes are developed and more resources become available. The classic style of developing a single code base for a small set of machines is being replaced by models that create ensembles of many codes running on multiple machines. This has also been coupled to research to incorporate a greater number of variables and metrics, adjusting to new time and spatial scales, and overall attempts to create “reduced” data models.
- FES simulation will incorporate the use of AI and ML in the future, as the codes are adapted to run on next-generation machines and at a larger number of facilities.
- PPPL HPC workloads that utilize ASCR facilities routinely are not able to perform at peak efficiency due to limitations. Recent upgrades to the PPPL local network and data architecture are expected to alleviate the problems, but further testing will be needed. Some potential bottlenecks to peak efficiency with data mobility are as follows:
 - Security infrastructure on the PPPL campus was undersized for the expected data volumes and expected capacities. A recent upgrade should enable a higher level of performance.
 - Data transfer hardware was not regularly used. A recent upgrade to deploy purpose-built DTNs will become part of several scientific workflows.
 - Data transfer software was not standardized, with projects using a mixture of tools that could not efficiently utilize the network and hardware. PPPL is moving toward more capable tools (e.g., Globus) for its DTN pool.
 - New use cases that mix bulk data movement, as well as real-time streaming, mean that the network must be responsive to latency as well as bandwidth requirements.
 - Simulations run at ASCR HPC facilities are now generating more data than can be easily stored and transferred using existing capabilities. The upgrades at PPPL, and ongoing upgrades to ASCR HPC facilities, will ensure some mechanisms to scale the requirements into the future as exascale simulations become more common.
- The FES community is exploring ways that cloud-provided storage and computation could be integrated into scientific workflows, particularly at facilities that are not able to scale local resources due to cost, lack of space, or lack of expertise to operate long-term storage pools. Investigations are underway to understand the costs and usability for FES workflows.

- DOE HPC allocations for FES are subject to annual renewal. This causes challenges for strategic planning and long-term investments in a particular computing capability or workflow architecture. If renewing an allocation at the same location is not possible, an experiment or facility may experience complications in data and workflow migrating to alternate facilities: adapting software to run on different systems, granting accounts to existing users, and sending most scientific data to another facility. Unified APIs and simplified methods to manage data between DOE HPC facilities could simplify the friction seen in these scenarios. Longer-duration (strategic) allocations of computing at ASCR facilities would allow the FES community to make more effective software investments.

5.12.2 Discussion

The WDM and FES HPC authors report no significant changes to the case study since the *2021 FES Network Requirements Review*.

List of Abbreviations

AI	artificial intelligence
ASCR	Advanced Scientific Computing Research
BELLA	Berkeley Lab Laser Accelerator
DOE	Department of Energy
DOE SC	Department of Energy Office of Science
DTN	data transfer node
EAST	Experimental Advanced Superconducting Tokamak
ESCC	ESnet Site Coordinators Committee
FES	Fusion Energy Sciences
GA	General Atomics
HED	high energy density
HFIR	High Flux Isotope Reactor
HPC	high-performance computing
IDS	Intrusion Detection System
IFE	Inertial Fusion Energy
INFUSE	Innovation Network for Fusion Energy
IO	ITER Organization
IRI	Integrated Research Infrastructure
ITER	International Thermonuclear Experimental Reactor
JET	Joint European Torus
KSTAR	Korea Superconducting Tokamak Advanced Research
LAN	local-area network
LBNL	Lawrence Livermore National Laboratory
LCLS	Linac Coherent Light Source
LHC	Large Hadron Collider
LPA	laser-plasma accelerator
MAGPI	Mid-Atlantic GigaPop for Internet2
MAN	metropolitan area network
MB	megabyte
MEC	Matter in Extreme Conditions
ML	machine learning
MPEX	Material Plasma Exposure eXperiment
NERSC	National Energy Research Scientific Computing Center
OLCF	Oak Ridge Leadership Computing Facility
OMFIT	One Modeling Framework for Integrated Tasks
ORNL	Oak Ridge National Laboratory
PB	petabyte
PI	principal investigators
PPIC	Princeton Plasma Innovation Center

PPPL	Princeton Plasma Physics Laboratory
PSFC	Plasma Science and Fusion Center
R&E	research and education
SciDAC	Scientific Discovery through Advanced Computing
TB	terabyte
WDM	whole-device modeling
XFEL	X-ray free-electron laser
XGC	X-Point Included Gyrokinetic Code

