ESnet
ENERGY SCIENCES NETWORK

Biological and Environmental Research Network Requirements Review Final Report

August 2022 – April 2023

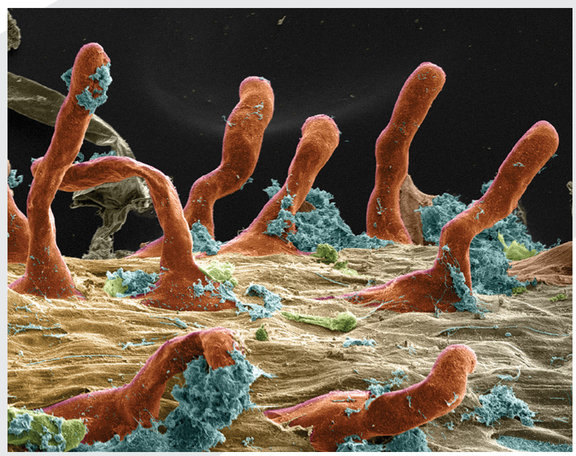

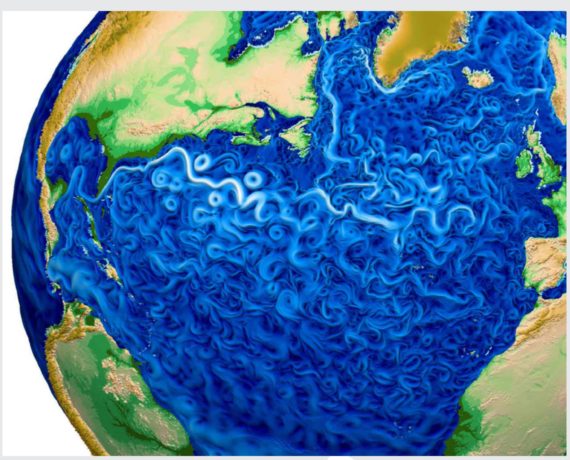

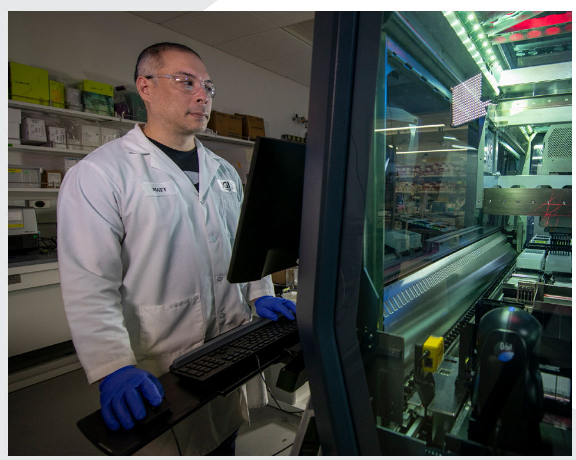

BERKELEY LAB

U.S. DEPARTMENT OF ENERGY
Office of Science

Biological and Environmental Research Network Requirements Review Final Report

August 2022 – April 2023

Office of Biological and Environmental Research, DOE Office of Science
Energy Sciences Network (ESnet)

ESnet is funded by the US DOE, 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 Biological and Environmental Research.

This is LBNL report number LBNL-2001542[1,2].

Cover Images:
(**Top left**) An Atmospheric Radiation Measurement (ARM) cloud radar in Brazil, courtesy of DOE ARM user facility
(**Top right**) Magnified view of a red pine (Pinus resinosa) root and associated microbiome, courtesy of DOE Environmental Molecular Sciences Laboratory
(**Bottom left**) Visulization from DOE's Energy Exascale Earth System Model, which simulates aspects of Earth's variability at weather-scale resolution and investigates decadal changes in climate, courtesy of DOE Energy Exascale Earth System Model (E3SM)
(**Bottom right**) Genomics sequencing seen in the Synthetic Biology Lab, at the DOE Joint Genome Institute, courtesy of Lawrence Berkeley National Laboratory, Joint Genomics Institute

1 https://doi.org/10.2172/1888810
2 https://escholarship.org/uc/item/3mz7h3mm

Participants and Contributors

Deb Agarwal, *Lawrence Berkeley National Laboratory*

Lewis Allen, *Pacific Northwest National Laboratory*

Sasha Ames, *Lawrence Livermore National Laboratory, ESGF, PCMDI*

Adam Arkin, *Lawrence Berkeley National Laboratory*

Dave Bader, *Lawrence Livermore National Laboratory, E3SM*

Vanessa Bailey, *Pacific Northwest National Laboratory*

Chris Beecroft, *Lawrence Berkeley National Laboratory*

Craig Bingman, *Great Lakes Bioenergy Research Center*

Shane Canon, *Lawrence Berkeley National Laboratory*

Gianni Carlo Contardo, *Lawrence Livermore National Laboratory*

Stephen Chan, *Lawrence Berkeley National Laboratory*

You-Wei Cheah, *Lawrence Berkeley National Laboratory*

Xingyuan Chen, *Pacific Northwest National Laboratory*

Sandy Chin, *Lawrence Berkeley National Laboratory*

Shreyas Cholia, *Lawrence Berkeley National Laboratory*

Housen Chu, *Lawrence Berkeley National Laboratory*

William Collins, *Lawrence Berkeley National Laboratory, University of California, Berkeley*

David Cowley, *Pacific Northwest National Laboratory*

Eli Dart, *Energy Sciences Network*

Michael Degan, *Pacific Northwest National Laboratory*

Tim Donohue, *Great Lakes Bioenergy Research Center*

Paul J. Durack, *Lawrence Livermore National Laboratory, PCMDI, ESGF*

Emiley Eloe-Fadrosh, *Lawrence Berkeley National Laboratory*

Kjiersten Fagnan, *Lawrence Berkeley National Laboratory*

Ian Foster, *Argonne National Laboratory, ESGF*

Brian Fox, *Great Lakes Bioenergy Research Center*

Hector Garcia-Martin, *Lawrence Berkeley National Laboratory*

Corey Hall, *Argonne National Laboratory*

Zach Harlan, *Energy Sciences Network*

Carol Hawk, *Department of Energy Office of Science*

Susan Heinz, *Oak Ridge National Laboratory*

Valerie Hendrix, *Lawrence Berkeley National Laboratory*

John Hess, *Energy Sciences Network*

Robert Hetland, *Pacific Northwest National Laboratory*

Nicki Hickman, *Argonne National Laboratory*

Susan Hicks, *Oak Ridge National Laboratory*

Justin J. Hnilo, *Department of Energy Office of Science*

Forrest Hoffman, *Oak Ridge National Laboratory, ESGF*

Matthew Hudson, *Center for Advanced Bioenergy and Bioproduct Innovations*

Colleen Iversen, *Oak Ridge National Laboratory*

Robert Jacob, *Argonne National Laboratory*

Dan Jacobson, *Oak Ridge National Laboratory*

Roy Kamimura, *Lawrence Berkeley National Laboratory*

Jay Keasling, *Lawrence Berkeley National Laboratory*

Stephen Klein, *Lawrence Livermore National Laboratory, PCMDI*

Dan Klos, *Argonne National Laboratory*

Charles Koven, *Lawrence Berkeley National Laboratory*

Mark Kulawik, *Lawrence Berkeley National Laboratory*

Krzysztof Lazarski, *Argonne National Laboratory*

John Macauley, *Energy Sciences Network*

Ramana Madupu, *Department of Energy Office of Science*

Kurt Maier, *Pacific Northwest National Laboratory*

Douglas Mans, *Pacific Northwest National Laboratory*

Stan Martin, *Oak Ridge National Laboratory*

Jim Mather, *Pacific Northwest National Laboratory*

Lee Ann McCue, *Pacific Northwest National Laboratory*

Gerald Meehl, *National Center for Atmospheric Research*

Dylan Merrigan, *Los Alamos National Laboratory*

Ken Miller, *Energy Sciences Network*

Nigel Mouncey, *Lawrence Berkeley National Laboratory*

Elizabeth Murphy, *Center for Advanced Bioenergy and Bioproduct Innovations*

Dirk Norman, *Great Lakes Bioenergy Research Center*

Travis O'Brien, *Indiana University*

Arthur Panganiban, *Lawrence Berkeley National Laboratory*

Gilberto Pastorello, *Lawrence Berkeley National Laboratory*

Stephanie Pennington, *Pacific Northwest National Laboratory*

Giri Prakash, *Oak Ridge National Laboratory*

Georg Rath, *Lawrence Berkeley National Laboratory*

Saubhagya Rathore, *Oak Ridge National Laboratory*

Roy Rich, *Smithsonian Environmental Research Center*

Jadwiga Richter, *National Center for Atmospheric Research*

Nan Rosenbloom, *National Center for Atmospheric Research*

Boris Sadkhin, *Argonne National Laboratory*

Henrik Scheller, *Lawrence Berkeley National Laboratory*

Corinne Scown, *Lawrence Berkeley National Laboratory*

Blake Simmons, *Lawrence Berkeley National Laboratory*

Richard Simon, *Lawrence Berkeley National Laboratory*

Surya Singh, *Pacific Northwest National Laboratory*

Leslie Stoecker, *Center for Advanced Bioenergy and Bioproduct Innovations*

Gary Strand, *National Center for Atmospheric Research*

Rune Stromsness, *Lawrence Berkeley National Laboratory*

Cory Stuart, *Oak Ridge National Laboratory*

Nick Thrower, *Great Lakes Bioenergy Research Center*

Margaret Torn, *Lawrence Berkeley National Laboratory, University of California, Berkeley*

Christopher Tracy, *Energy Sciences Network*

Gerald Tuskan, *Oak Ridge National Laboratory*

Charuleka Varadharajan, *Lawrence Berkeley National Laboratory*

Terri Velliquette, *Oak Ridge National Laboratory*

Jiali Wang, *Argonne National Laboratory*

Andrew Wiedlea, *Energy Sciences Network*

Linda Winkler, *Argonne National Laboratory*

Stan Wullschleger, *Oak Ridge National Laboratory*

Yiying Xiong, *Great Lakes Bioenergy Research Center*

Jill Zhang, *Lawrence Livermore National Laboratory, E3SM*

Jason Zurawski, *Energy Sciences Network*

Report Editors

Eli Dart, *Energy Sciences Network,* dart@es.net

Zach Harlan, *Energy Sciences Network,* zharlan@es.net

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

John Hess, *Energy Sciences Network,* jhess@es.net

Justin J. Hnilo, *Department of Energy Office of Science,* justin.hnilo@science.doe.gov

John Macauley, *Energy Sciences Network,* macauley@es.net

Ramana Madupu, *Department of Energy Office of Science,* ramana.madupu@science.doe.gov

Ken Miller, *Energy Sciences Network,* ken@es.net

Christopher Tracy, *Energy Sciences Network,* ctracy@es.net

Andrew Wiedlea, *Energy Sciences Network,* awiedlea@es.net

Jason Zurawski, *Energy Sciences Network,* zurawski@es.net

Table of Contents

1 Executive Summary

About ESnet

The Energy Sciences Network (ESnet) is the high-performance network user facility for the US Department of Energy (DOE) Office of Science (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 Advanced Scientific Computing Research (ASCR) program and managed and operated by the Scientific Networking Division at Lawrence Berkeley National Laboratory (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 they need 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.

Requirements Review Purpose and Process

ESnet and ASCR use requirements reviews 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. A requirements review regularly, and comprehensively, surveys major science stakeholders' plans and processes in order 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 intermediate updates scheduled every off year. 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.

This Review

Between August 2022 and April 2023, ESnet and the Office of Biological and Environmental Research (BER) of the DOE SC organized an ESnet requirements review of BER-supported activities. Preparation for these events included identification of key stakeholders: program and facility management, research groups, and technology providers. Each stakeholder group was asked to prepare formal case study documents about its relationship to the BER ESS program to build a complete understanding of the current, near-term, and long-term status, expectations, and processes that will support the science going forward. A series of pre-planning meetings better prepared case study authors for this task, along with guidance on how the review would proceed in a virtual fashion.

The BER program supports scientific research and facilities to achieve a predictive understanding of complex biological, Earth, and environmental systems with the aim of advancing the nation's energy and infrastructure security. The program seeks to discover the underlying biology of plants and microbes as they respond to and modify their environments. This knowledge enables the reengineering of microbes and plants for energy and other applications. BER research also advances understanding of the dynamic processes needed to model the Earth system, including atmospheric, land masses, ocean, sea ice, and subsurface processes.

Over the last three decades, BER has transformed biological and Earth system science. We helped map the human genome and lay the foundation for modern biotechnology. We pioneered the initial research on atmospheric and ocean circulation that eventually led to climate and Earth system models. In the last decade, BER research has made considerable advances in biology underpinning the production of biofuels and bioproducts from renewable biomass, spearheaded progress in genome sequencing and genomic science, strengthened the predictive capabilities of ecosystem and global scale models using the world's fastest computers, and evaluated model predictions in underrepresented biomes such as the Arctic and the tropics.

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

- The Joint Genome Institute (JGI).
- The DOE Systems Biology Knowledgebase (KBase).
- The National Microbiome Data Collaborative (NMDC).
- The Environmental Molecular Sciences Laboratory (EMSL).
- The Atmospheric Radiation Measurement Program (ARM).
- The Earth System Grid Federation (ESGF).
- The Environmental Systems Science Data Infrastructure for a Virtual Ecosystem (ESS-DIVE) and Self-Driving Field Laboratories Research.
- The Structural Biology Center (SBC) at the APS.
- The Center for Advanced Bioenergy and Bioproducts Innovation (CABBI).
- The Center for Bioenergy Innovation (CBI).
- The Great Lakes Bioenergy Research Center (GLBRC).
- The Joint BioEnergy Institute (JBEI).
- Energy Exascale Earth System Model (E3SM).
- AmeriFlux network.
- Next-Generation Ecosystem Experiments (NGEE), NGEE-Arctic and NGEE-Tropics.
- Coastal Efforts.

- Calibrated and Systematic Characterization, Attribution, and Detection of Extremes (CASCADE).
- The Cooperative Agreement To Analyze variabiLity, change and predictabilitY in the Earth SysTem (CATALYST).
- Multifacility Workflows in BER and Relationship to the IRI Initiative.

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 recommendations that will guide future interactions between BER, 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.
- **Recommendations**: potential strategic or tactical activities, investments, or opportunities that are recommended to be evaluated and potentially pursued to address the challenges laid out in the findings.

The review participants spanned the following roles:

- Subject-matter experts from the BER activities listed previously.
- ESnet Site Coordinators Committee (ESCC) members from BER activity host institutions, including the following DOE labs and facilities: Argonne National Laboratory (ANL), Los Alamos National Laboratory (LANL), LBNL, Lawrence Livermore National Laboratory (LLNL), Oak Ridge National Laboratory (ORNL), and Pacific Northwest National Laboratory (PNNL).
- Networking and/or science engagement leads from the ASCR high-performance computing (HPC) facilities.
- DOE SC staff spanning both ASCR and BER.
- ESnet staff.

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

- **Facility Management and Readiness:**
  - The challenges with enabling BER multifacility workflows are significant. Most are not bracketed by technology but are sociological, policy, or related to deeply scientific constraints. [Section 5.19, Focus Groups]
  - BER maintains and operates a number of multifacility workflows, with several more expected in the coming years [Section 5.19]:
    - The JGI has a number of strategic user facility partners that operate in a multifacility workflow paradigm.
    - KBase works with the JGI to facilitate seamless access to data to facilitate modeling and analysis.
    - The NMDC, EMSL, and JGI, through the JAWS effort, have partnered to use the services for NMDC analysis workflows.

    - The EMSL and JGI are partners in the FICUS program, which provides researchers with access to the world-class resources of multiple user facilities through a single proposal.

- **Scientific Data Management:**
  - In the next 2-5 years the BER community will see growing data volumes due to improvements in instrument resolution, model complexity, and increases in the rate of data generation due to laboratory automation: [Section 5.1, Section 5.4, Section 5.5, Section 5.6, Focus Groups]
    - Sequencing technologies continue to evolve to enable larger capacities at cheaper costs.
    - CMIP7 models will be four to five times larger than CMIP6, which places the evolution of data volumes on a path to exascale.
    - There will be significant data growth from both internal capacity growth and external data being imported into data portals.
  - BER Facilities have a mature set of tools for data management and providing remote data access. The main challenge in enabling multifacility workflows in the coming years will be improving how these systems connect to each other and exchange information. The major challenges will be in developing standards around metadata, data exchange, authentication and authorization, policies, and tooling to support all of these areas. [Section 5.19]
  - The "last foot" problem of data movement, e.g., where some tools are better run with laptops, still persists in a number of BER workflows. Adapting workflows to use DTNs to move data "closer" is the best mitigation, followed by streaming or bulk downloading to a laptop resource as a final step. [Focus Groups]

- **Scientific Workflow:**
  - Data reuse by the BER community is expected to increase as the value of the accumulated data improves from multimodal analyses [Section 5.4, Focus Groups]:
    - Instruments will increase in size and capability and could yield up to multiple terabytes per day.
    - Automation pipelines will be established, with expected data generation to increase to tens to hundreds of terabytes per day, and data-archive capacity is expected to grow to 500 PB.
    - Continued growth of the FICUS program will increase demand for transfer of data between facilities.
  - To meet the needs of field data from sensors and instruments, it will be critical to enable data acquisition, transformation, and analytics workflows that can integrate real-time data streams with predictive models, particularly for extreme events. The highly distributed nature of sensor networks, and the volumes of data collected, will necessitate a data "backbone" for AI/ML-based scientific workflows. [Section 5.7, Focus Groups]
  - It would be desirable to have the ability for moderate-scale HPC centers to remotely mount DOE HPC filesystems. This, combined with some sort of ability to preemptively cache data locally at the HPC sites before a job makes use of the data, would enable a distributed super-facility model to support BER workflows. [Section 5.17, Focus Groups]
  - Data transmission amounts between coupled BER facilities vary from megabytes and gigabytes when dealing with the sharing of raw sensor or observational data, to multiple

terabytes to petabyte scales when moving large models between HPC resources. Data movement tools are critical to managing the data mobility aspects, along with other methods such as portals, and the use of cloud sharing and storage. [Section 5.19]

  - New workflows to better stream, store, process, and share BER data will be needed in the future as longer paths between observation and result are created. [Focus Groups]

- **Computational and Storage Requirements:**
  - CMIP data volumes have grown, from ~1 GB in the early phases (~1989), to over 23 PB of replicated model data contributed by 49 institutions across 26 countries. Routinely, this project is transferring 100 TB of data between nodes weekly, with download and upload occurring concurrently. [Section 5.6]
  - The BER community is adopting dockerized microservices for portability (running on platforms that include NERSC's Spin). This has the potential to introduce network challenges, since all traffic goes through the container network layer. [Section 5.7]
  - Across the BER case studies, Globus use remains high, although not all facilities are users either due to security policy, licensing issues, or observations that performance is not measurably better. [Focus Groups]
  - Improvements to GPU hardware have made it such that data compression is much better, which has simplified processing (and network transfer) for some model and simulation activities. [Focus Groups]
  - The BER community takes advantage of interactive computing jobs (e.g., those that are operated via Jupyter) that can easily run on cloud-like computing platforms. Those that require access to midrange clusters, or HPC, often must use institutional resources, or allocations at the major HPC centers to accommodate their computing needs. DOE LCFs prioritize large-scale computation over moderate-scale analysis, thus some in the BER community must increasingly using medium-scale computational resources external to the LCFs when doing moderate-scale data analysis tasks. [Section 5.17, Focus Groups]

- **Remote Collaboration and Operational Requirements:**
  - Many BER user facilities are already designed around a remote model [Section 5.19]:
    - Users send physical samples to facilities for analyses and access results over data portals.
    - Multifacility models will expand on these by requiring coordination and information flow between facilities.
  - Federated authentication mechanisms among partner institutions would greatly ease data transfer and other collaboration efforts. [Section 5.16, Focus Groups]
  - For some collaborations with a strong international presence, high performance remains elusive. Data transfers that cross international boundaries can be incredibly slow, which makes collaboration challenging. There is a rich international ecosystem of connectivity (e.g., multiple 100 Gbps) across the Pacific, but local connectivity within a facility (or security policy) may be limiting performance. [Focus Groups]

- **Multifacility Computational Workflows:**
  - Workflows that require coupling capabilities across facilities are becoming more common within BER [Section 5.19, Focus Groups]:
    - The FICUS effort, with more multi-analytic couplings to increase over time.

    - Workflows that depend on data replication across sites to ensure that data is available are at the HPC facilities. This involves coordination of data movement and job management.
  - A core problem for addressing multifacility use cases remains a lack of a cohesive identify, authorization, and authentication strategy. It is still the case that each facility (and lab) has different ways of managing this, which complicates attempting to build multifacility workflows. [Focus Groups]
  - When a component of a coupled facility is unavailable (or has limited resources), the entire workflow suffers. If it were easier to move to other HPC facilities, or on temporal resources that may be located on ESnet, that could offer a layer of protection for some time-dependent use cases. [Focus Groups]
- **Domestic Networking for Local and Wide-Area Data Mobility:**
  - Transferring multiomics data across the wide-area network can take several days using traditional web protocols. The use of high-bandwidth data transfer mechanisms via Globus, is preferred at several BER community sites. [Section 5.3, Section 5.5, Section 5.7, Section 5.13, Focus Groups]
    - Most of the challenges for multifacility BER workflows are not at the networking layer, but the network could enable some models [Section 5.19]:
    - The ability to create overlay networks across ESnet between facilities could allow distributed resources to be coupled together more securely.
    - HTTP-based communications play an important role, especially in the context of pulling datasets from existing data repositories, reference databases, etc., and bringing them on to compute systems where workflows are run.
    - Allowing for high-performance data-movement capabilities natively over standard protocols like HTTP.
    - Hybrid commercial cloud and DOE HPC facility workflows, and the ability to enable a virtual network to facilitate orchestration of resources.
  - Most BER case studies have demonstrated effective use of networks, particularly as capacities have increased on backbones, regionals, and campuses. Several BER case studies have noted that data storage access and use remain a critical bottleneck [Focus Groups]:
    - Storage is a constant struggle to allocate, and it may not be located close to other parts of the workflow; this requires data movement several times.
    - Data-movement tools are much better, but still require a level of tuning on some institutional and laboratory resources that interact with large filesystems.
    - Not every BER case study has access to DTN resources by default (even those these are prevalent at most, if not all DOE-SC facilities), and may not know to ask how to get access.
  - There is interest in understanding some of ESnet's lessons learned using LEO satellites and 5G connectivity. Many remote sensor projects could benefit from this approach, as they build a standard model for pulling data from regions with poor connectivity. [Focus Groups]

- **Emerging Needs:**
  - The cybersecurity landscape at some of the institutions imposes many constraints upon scientific researchers. Each institution has independent cybersecurity policies, teams, and infrastructure, with different technical and policy requirements that must be traversed. [Section 5.10, Focus Groups]
  - It is expected that demands will increase for sensor networks with artificial intelligence (AI capabilities) as well as self-driving labs capabilities for BER use cases. To meet the needs of rapidly expanding field data from DOE Earth Systems sensors and instruments, it will be critical to enable data acquisition, transformation, and analytics workflows that can integrate real-time data streams with predictive models. [Section 5.19]
  - It is very likely that clouds over various forms will play a role in future BER multifacility workflows [Section 5.19]:
    - This could take the form of surging to commercial clouds to meet targeted turnaround times for analysis, hosting public datasets, or other web services.
    - This may also include cloud-like models provided by a future integrated research infrastructure.
    - Maintaining a replicated instance of data services on cloud resources for failover, redundancy, or backup should be considered to enable broader access, especially to resources/users outside the DOE complex.
    - Enabling analysis capabilities in the cloud for very large datasets already stored in the cloud (bringing compute to the data).
  - Lastly, ESnet will follow up with review participants on a number of high-level recommendations that were identified. These items are listed as guidance for future collaboration, and do not reflect formal project timelines. ESnet will review these with BER participations on a yearly basis, until the next requirements review process begins.
- **Facility Management and Readiness:**
  - ESnet will work with the BER community to assist in adapting workflows to use DTNs between facilities to facilitate data mobility. [Focus Groups]
  - ESnet will work with the BER community to address some of technological, sociological, and policy level constraints that come with the operation of BER multifacility workflows. [Section 5.19]
- **Scientific Data Management:**
  - ESnet will work with the CMIP7 model community in the coming years as the volume of the data increases toward the exabyte scale. To prepare for this, sites must be prepared in terms of network connectivity, storage, and data movement. [Section 5.6]
  - ESnet will assist the BER community on ways to best address data sharing through use of both service platforms like Globus, and commercial cloud infrastructure. [Section 5.16]
- **Scientific Workflow:**
  - The BER community, in conjunction with the ASCR HPC facilities, should explore new ways that HPC centers can improve data access for local and remote workflows and enable a distributed super-facility model to support future workflows. [Section 5.17] Suggestions include:
    - Creation and improvements of caching technology.

    - Data mirrors across the DOE HPC facilities.
    - Exploring behavior similar to a remote mounts of filesystems at ALCF, NERSC, and OLCF.
  - The BER community should explore workflow modifications to better stream, store, process, and share data. [Focus Groups]
- **Computational and Storage Requirements:**
  - DOE ASCR and BER, through efforts like IRI, will continue to address issues surrounding the availability of computation, storage, and data movement tools as they relate to both large-scale computation and moderate-scale analysis. [Section 5.17]
  - The BER community, ESnet, and R&E networking community projects, will work together to understand how containerized workflows may be applicable to current and future use cases. [Section 5.6]
  - The BER community and ESnet will continue to address the increases in data volumes, and ways that existing and future storage infrastructure must be understood and tuned to accept and serve this data efficiently. [Section 5.13, Section 5.16, Section 5.17, Section 5.18]
- **Remote Collaboration and Operational Requirements:**
  - The BER community, in cooperation with ASCR HPC facilities, must address the growing challenges of federated identify to streamline ways that users and institutions can access and transfer research data. [Section 5.16]
  - ESnet will continue to work with the ESGF to address performance problems to ESGF resources that are located around the world. [Section 5.6]
- **Multifacility Computational Workflows:**
  - The BER community, in conjunction with other facilities operated by ASCR, must address the lack of a cohesive identify, authorization, and authentication strategy in order to enable future multifacility use cases and workflows. [Focus Groups]
  - DOE SC, with stakeholders that include BER and ASCR facilities, will continue to address research workflow requirements as identified in the IRI effort, with a focus on ways to enable multifacility operation on a more regular basis. [Focus Groups]
- **Domestic Networking for Local and Wide-Area Data Mobility:**
  - The BER community and ESnet will continue to design and operate multifacility workflows with a particular focus on efforts that will adopt modern data transfer tools, portals, integration with cloud providers, and increased ESnet connectivity and services to core facilities. [Section 5.4, Section 5.5, Section 5.19, Focus Groups]
  - ESnet will work with the BER community to advise on possible ways that LEO satellites and 5G connectivity can be integrated into workflows that involve remote sensor environments. [Section 5.5, Section 5.7, Section 5.9, Section 5.14, Section 5.15, Section 5.16, Focus Groups]
  - The ESnet's Science Engagement Team will continue to assist with data mobility issues. [Section 5.6, Section 5.7, Section 5.10, Focus Groups]
- **Emerging Needs:**
  - The BER community and ESnet will continue to address the best common practices of sharing research data between platforms (commercial clouds and R&E based clusters) in the future to minimize the support overhead. [Section 5.1, Section 5.3, Focus Groups]

- The DOE SC community must work to together to address gaps in workforce, particularly as it relates to training and maintaining staff that assist with development and operations requirements: e.g., tasks that focus on maintenance, cybersecurity patching, and user support. [Section 5.2, Focus Groups]
- ESnet and the BER community will continue to discuss best common practices for interacting with cloud providers, both from the standpoint of peering and connectivity, but also ways to increase performance expectations. [Focus Groups]

2 Review Findings

The requirements review process helps to identify important facts and opportunities from the programs and user facilities that are profiled. The following sections outline a set of findings from the BER and ESnet requirements review. These points summarize important information gathered during the review discussions surrounding case studies and the BER-managed user program in general. These findings are organized by topic area for simplicity and by common themes:

- Facility Management and Readiness.
- Scientific Data Management.
- Scientific Workflow.
- Computational and Storage Requirements.
- Remote Collaboration and Operational Requirements.
- Multifacility Computational Workflows.
- Domestic Networking for Local and Wide-Area Data Mobility.
- Emerging Use Cases.

2.1 Facility Management and Readiness

- The JGI is a user facility hosted at LBNL. Numerous soft funded research projects operate within or in collaboration with the JGI [Section 5.1]:
  - The National Energy Research Scientific Computing Center (NERSC): development of tools and infrastructure for microbiome data science, optimized for high-performance computing environments.
  - The DOE System Biology Knowledgebase (KBase): a system to access and analyze data on scales ranging from single genes and individual genomes to metagenomes to systems-level modeling and understanding.
  - The NMDC: a community-driven national effort aimed to develop standards, processes, and infrastructure for an integrated microbiome data ecosystem.
  - The EMSL: a partner in the Facilities Integrating Collaborations for User Science (FICUS) program.
  - The HudsonAlpha Institute for Biotechnology[1]: a genome sequencing center that is integrated with SDM, providing genome sequencing for JGI analysis services.
- The DOE KBase is a software and data platform used to predict and design biological functions [Section 5.2]:
  - KBase integrates data, tools, and their associated interfaces into one environment, so users do not need to access them from numerous sources.
  - The open KBase platform enables external developers to integrate their analysis tools, facilitating distribution, comparative tool analysis, and access.
  - KBase users can perform large-scale analyses and combine multiple lines of evidence to create biological models.

1 https://www.hudsonalpha.org/

    - KBase is the first large-scale bioinformatics system that enables users to upload their own data, analyze it (along with collaborator and public data), build increasingly realistic models, and share and publish their workflows and conclusions.

- The NMDC leverages microbiome data science capabilities across three DOE national laboratories: LBNL, LANL, and PNNL, to deliver a science gateway for multiomics microbiome data. The NMDC Data Portal is a resource to discover data generated at the JGI and EMSL. [Section 5.3]
- The NMDC provides three defined products to support FAIR data [Section 5.3]:
    - The NMDC Submission Portal: a system designed to lower barriers to capture and adhere to community standards for sample contextual information (metadata) leveraging the Minimum Information about any (x) Sequence (MIxS) standards for environmental packages.
    - The NMDC Data Portal and public API: a system that facilitates seamless access for data discovery.
    - NMDC EDGE: a web application that provides community access to the NMDC data analysis workflows.
- EMSL is a BER user facility. EMSL [Section 5.4]:
    - Studies the role of molecular processes in controlling the function of biological and ecological systems across spatial and temporal scales.
    - Provides access to multimodal molecular science instruments, data analytics, production computing, and multiscale modeling.
- EMSL user research is conducted within the following three Science Areas [Section 5.4]:
    - Environmental Transformations and Interactions.
    - Functional and Systems Biology.
    - Computing, Analytics, and Modeling.
- The ARM user facility is a multi-laboratory effort and a key contributor to national and international climate research efforts [Section 5.5]:
    - The ARM facility is managed and operated by nine DOE laboratories (Argonne, Brookhaven, Lawrence Berkeley, Lawrence Livermore, Los Alamos, Oak Ridge, Pacific Northwest, National Renewable Energy, and Sandia) and actively serves approximately 1,000 science users each year.
    - ARM operates over 450 instruments to provide measurements of clouds, aerosols, precipitation, the surface energy balance, and the background atmospheric state.
    - ARM data is currently collected from three atmospheric observatories—Southern Great Plains (Oklahoma), North Slope of Alaska, and Eastern North Atlantic (the Azores).
    - There are three ARM mobile facilities, with plans to deploy to northwest Alabama in the Bankhead National Forest in mid-2023.
    - There are several ARM aerial capabilities: a Challenger 850 regional jet, mid-sized unmanned aerial system (UAS), and tethered balloon systems (TBS).
    - Data from these efforts is available through the ARM Data Center (ADC) at ORNL.
    - ARM personnel manage the collection, processing, quality assessment, archiving, and distribution of data from these instruments, which currently represent approximately 2,500 distinct data streams.

- ARM's vision is to [Section 5.5]:
  - Provide comprehensive and impactful field measurements to support scientific advancement of atmospheric process understanding.
  - Achieve the maximum scientific impact of ARM measurements through increased engagement with observational data by ARM staff, including the application of advanced data analytical techniques.
  - Enable advanced data analytics and community use of complex ARM datasets through the advancement of computing infrastructure and data analysis.
  - Accelerate and amplify the impact of ARM measurements on Earth system models by exploiting ARM and earth system model frameworks to facilitate the application of ARM data to model development.
- The ESGF is an international collaboration for the software that powers global climate change research, notably assessments by the Intergovernmental Panel on Climate Change (IPCC). ESGF manages a decentralized database for handling climate science data, with multiple petabytes of data at dozens of federated sites worldwide. It supports the Coupled Model Intercomparison Project (CMIP), whose protocols enable the periodic assessments carried out by the IPCC. [Section 5.6]
- ESGF is an interagency and international effort co-led by the DOE, and co-funded by the National Aeronautics and Space Administration (NASA), National Oceanic and Atmospheric Administration (NOAA), National Science Foundation (NSF), and international laboratories such as the Max Planck Institute for Meteorology German Climate Computing Centre (Deutsches Klimarechenzentrum, DKRZ), the Australian National University (ANU) National Computational Infrastructure (NCI), Institut Pierre-Simon Laplace (IPSL), and the Centre for Environmental Data Analysis (CEDA). [Section 5.6]
- ESS-DIVE is a data repository for Earth and environmental science data managed by LBNL. It stores and enhances access to critical information generated from research funded by or related to the BER ESS program. ESS-DIVE works to preserve, expand access to, and improve usability of critical data generated through research of terrestrial and subsurface ecosystems. [Section 5.7]
- The Surface Atmosphere Integrated Laboratory (SAIL) envirotranspiration application is deployed in western Colorado as a Self-Guiding Field Laboratory (SGFL), which has created an iterative two-way interaction between data gathering and model predictions, wherein data gathering is adapted based on model predictions utilizing a "decision engine": an algorithm that takes the gathered data as inputs and calculates a numerical decision matrix for assessing whether data gathering should be adjusted. NERSC HPC resources are leveraged for data analysis and scalable ML model training. [Section 5.7]
- SBC performs macromolecular crystallography at the APS. SBC focuses on structural biology and contributes to the expansion of existing programs and exploration of new opportunities in structural biology, proteomics, and genomics research with a major focus on medicine, bio-nanomachines, and biocatalysis highly relevant to energy resources, health, a clean environment, and national security. [Section 5.8]
- SBC is used for crystallographic experiments including [Section 5.8]:
  - Crystals of macromolecular assemblies with very large unit cells.
  - MAD/SAD phasing.
  - Crystals of membrane proteins.

- Small, weakly diffracting crystals.
- Ultra-high resolution crystallography.
- Cryo-crystallography.

- CABBI is a bioenergy research center at the University of Illinois with 20 partner institutions. CABBI will integrate recent advances in agronomics, genomics, and synthetic and computational biology to increase the value of energy crops using a "plants as factories" approach to grow fuels and chemicals in plant stems. CABBI will also use an automated foundry to convert biomass into valuable chemicals and ensure that its products are ecologically and economically sustainable. [Section 5.9]
- CABBI leverages key facilities at the University of Illinois, such as the Carl R. Woese Institute for Genomic Biology (IGB), Illinois Energy Farm, SoyFACE (free-air CO2 enrichment), the Integrative Bioprocessing Research Laboratory pilot plant, and the Illinois Biological Foundry for Advanced Biomanufacturing (iBioFAB), as well as facilities for plant transformation, field trials, computation, and diverse molecular, chemical, and physiological analyses at partner locations. [Section 5.9]
- CABBI is organized around three interconnected themes [Section 5.9]:
  - Provide an integrated economic and environmental framework for determining feedstock supply and its sustainability.
  - Provide a regionally adaptive yet national-scale platform for grass-based biorefining based on high-yielding feedstocks with improved environmental resilience.
  - Provide a broad set of platform microorganisms, and automated tools to engineer them, to produce value-added products from plant-produced feedstocks or substrates.
- CBI is a multi-institutional and interdisciplinary organization encompassing numerous institutions across the United States. The core DOE institutions involved are ORNL and the National Renewable Energy Laboratory (NREL). Major university partners include the University of California, University of Georgia, University of Tennessee, University of Maryland, University of Oregon, Washington State University, and Dartmouth College. CBI will accelerate domestication of bioenergy-relevant plants and microbes to enable fuels and coproduct development at multiple points in the bioenergy supply chain. [Section 5.10]
- The GLBRC is led by the University of Wisconsin–Madison with a mission of creating biofuels and bioproducts that are economically viable and environmentally sustainable. GLBRC focuses on three areas of research [Section 5.11]:
  - Sustainable cropping systems.
  - Efficient biomass conversion.
  - Field-to-product integration.
- The mission of the GLBRC has expanded to include collaborating with the other three BRCs on a variety of shared research objectives, putting an even larger emphasis on high-speed networking to allow for transferring of shared data over the next five years of research. [Section 5.11]
- The JBEI is a US DOE Bioenergy Research Center dedicated to developing advanced biofuels: liquid fuels derived from the solar energy stored in plant biomass that can replace gasoline, diesel, and jet fuels. JBEI partners include Brookhaven, Pacific Northwest and Sandia National Laboratories; Iowa State University; the University of California (UC) campuses at Berkeley, Davis, San Diego, and Santa Barbara; and TeselaGen Biotechnology, Inc. [Section 5.12]

- E3SM is a multi-laboratory Earth system modeling, simulation, and prediction project that optimizes how existing DOE laboratory resources are used to develop a leading-edge climate and Earth system model. The science questions the project aims to address are centered around water cycle changes and impacts, the evolution of the human-Earth systems and polar processes, and sea-level rise and coastal impacts. [Section 5.13]
- AmeriFlux is a network of more than 500 PI-managed sites measuring ecosystem CO2, water, and energy fluxes in North, Central, and South America [Section 5.14]:
  - The project was established to connect research on field sites representing major climate and ecological biomes, including tundra, grasslands, savanna, crops, and conifer, deciduous, and tropical forests.
  - Data produced by the network provides information and new knowledge central to the BER mission to understand complex biological and environmental systems, across spatial and temporal scales, by joining theory, observations, experiments, and models.
- The NGEEs, focusing on the Arctic and the tropics, are designed to advance the predictive power of Earth system models through understanding of the structure and function of Arctic terrestrial ecosystems, as well as the predictive understanding of tropical forest responses to global changes across scales. [Section 5.15]
- The NGEE Arctic project seeks to understand how surface and subsurface processes and properties are interconnected across permafrost-dominated tundra ecosystems through a series of collaborative investigations across a gradient of remote, permafrost landscapes in coordination with landowners from multiple Native corporations. [Section 5.15]
- NGEE–Tropics aims to fill the gaps in knowledge of tropical forest-climate system interactions, and develop a predictive understanding of how tropical forest carbon balance and climate system feedbacks will respond to changing environmental drivers. The overarching goal of NGEE–Tropics is to develop a greatly improved predictive understanding of tropical forests and Earth system feedback to changing environmental drivers. [Section 5.15]
- The Coastal Observations, Mechanisms, and Predictions Across Systems and Scales (COMPASS) program aims to improve fundamental scientific understanding, model representation, and predictive capacity of coastal systems. The project comprises two parts: a field study and a coastal modeling study [Section 5.16]:
  - COMPASS-FME (Field, Measurements, and Experiments) focuses on field studies and associated process and ecosystem modeling of two coastal interfaces.
  - COMPASS-GLM (Great Lakes Modeling) focuses on modeling and analysis of coastal systems in the Great Lakes Region.
- COMPASS-FME seeks to advance a scalable, predictive understanding of the fundamental biogeochemical processes, ecological structure, and ecosystem dynamics that distinguish coastal terrestrial–aquatic interfaces (TAIs) from the purely terrestrial or aquatic systems to which they are coupled. FME will focus on overarching long-term science questions [Section 5.16]:
  - What fundamental mechanisms control the structure, function, and evolution of coastal TAIs?
  - How do these fundamental mechanisms interact across spatial scales, and what interactions are most important to improving predictive models?
- COMPASS-FME is a multi-institutional project led by PNNL, and partners with ORNL, ANL, LANL, and LBNL, as well as the Smithsonian Environmental Research Center, the University of Toledo, and Heidelberg University. FME generates data across seven different sites in the

Chesapeake Bay and Western Lake Erie regions. This data provides broad spatial and temporal coverage of ecosystem processes. Data is remotely logged, collected through telemetry, and compiled for use by experimentalists and modelers to discern patterns of change and guide model improvements. [Section 5.16]

- COMPASS-GLM aims to enhance understanding of freshwater coastal systems, especially how they respond to climate warming, land use land cover change (LULCC), and other perturbations at watershed-to-regional scales. GLM will focus on overarching long-term science questions [Section 5.16]:
  - What multiscale mechanisms govern the structure, function, and dynamics of coastal systems at different spatial and temporal scales?
  - How do coastal systems respond to natural and anthropogenic influences?
  - Is it possible to generalize new process knowledge and predictive skill gained at a small number of sites or regions across the observed diversity of coastal systems?
- COMPASS-GLM modeling activities will center on two tasks. The first is focused on coupled, integrated atmosphere–land–lake regional modeling in the Great Lakes Region. The second will focus on high-resolution modeling of the Portage River watershed. GLM will share model output across tasks to drive the high-resolution watershed model with output from the atmospheric model to understand how the watershed responds to individual events and future climate conditions. [Section 5.16]
- The CASCADE project looks to identify and project climate extremes and how they are affected by environmental drivers. CASCADE will focus on four main research areas [Section 5.17]:
  - Description, detection, and designation of extremes in data and models.
  - Understanding and simulating the physical mechanisms that drive variability and change in the spatiotemporal characteristics of extreme events.
  - High-performance computing to detect and predict changes in weather extremes.
  - Experiments to explore how uncertainty in observations and models affects the understanding of extremes.
- The CASCADE project is collaborative work at LBNL, the University of California, Berkeley, and University of California, Davis campuses, Indiana University, and Iowa State University. CASCADE also coordinates with [Section 5.17]:
  - Earth system modeling efforts.
  - Land, ocean, and atmosphere diagnostics projects.
  - Stakeholder-driven science projects and scientific communities, including CMIP, SAMSI, and ARTIP.
- CASCADE is designed to address three major and interrelated scientific questions [Section 5.17]:
  - How have changes in the physical behavior of the coupled system altered the chances of encountering and the nature of extreme climate events?
  - To what degree are anthropogenic drivers responsible for altering the statistics, properties, and projections of extreme climate events?
  - What are the dominant uncertainties in detecting, attributing, and modeling extremes, how do these uncertainties affect the first two questions, and how can uncertainties be reduced?

- CATALYST is joint work between the DOE and the University Corporation for Atmospheric Research (UCAR) to perform research toward understanding modes of variability and change using models, observations, and process studies by using [Section 5.18]:
  - The DOE E3SM, the NCAR Community Earth System Model (CESM), and a hierarchy of simpler models.
  - Coupled Model Intercomparison Project (CMIP) multimodel datasets.
  - ML methods.
  - Numerous observational datasets.
- CATALYST features four interrelated research objectives (ROs) [Section 5.18]:
  - Address the limits of predictability of modes of variability using Earth system models (E3SM and the Community Earth System Model) and ML methods.
  - Use of a hierarchy of models to address relevant processes and feedbacks related to modes of variability and how they interact with each other.
  - Benchmark model representations, examine the role of changes, and address the likelihood and predictability of tipping points and irreversible changes.
  - Investigate the relationships between the modes of variability and high-impact events (e.g., flash droughts and precipitation extremes, atmospheric rivers, tropical cyclones).
- Some examples of coupled facility models within BER are as follows [Section 5.19]:
    - The JGI has a number of strategic user facility partners that operate in a multifacility workflow paradigm.
      - The JGI uses JAWS, Nextflow, Snakemake, and some additional systems for workflow management.
      - JAWS is the only resource able to submit work across distributed computing facilities.
      - Globus is leveraged for data movement.
      - NERSC is the primary repository for more than 14 PB of active and archival data.
    - KBase enables systems biology analyses, and works with the JGI to facilitate seamless access to data to facilitate modeling and analysis. KBase and the JGI partner through data-oriented APIs, with JGI services at LBNL and KBase production services at ANL.
    - The NMDC and JGI, through the JAWS effort, have partnered to use the services for NMDC analysis workflows. Initially, JAWS will handle NMDC workflows that operate at NERSC, but as the collaboration expands, other clusters with NMDC allocations will be integrated into JAWS. The NMDC links data and metadata from the same samples generated at the JGI and EMSL facilities.
    - The EMSL and JGI are partners in the FICUS program, which provides researchers with access to the world-class resources of multiple user facilities through a single proposal.
- The challenges with enabling BER multifacility workflows are significant. Most are not bracketed by technology but are sociological, policy, or related to deeply scientific constraints. [Section 5.19]

2.2 Scientific Data Management

- The JGI SDM services support analysis and publication pipelines for biological samples that are sent to the JGI for sequencing and analysis. The sequencing pipeline operates within the administrative boundaries of LBNL (using resources at the JGI, as well as some provided by LBNL IT and NERSC). [Section 5.1]
- Sequencing technologies continue to evolve to enable larger capacities at cheaper costs. The JGI has the two latest sequencers from Illumina and PacBio, which will increase the capacity by approximately threefold over current capacities. However, due to upstream (e.g., library creation) limitations, all this capacity will not be used. There will be significant data growth over the next five years from both internal capacity growth and external data being imported into JGI data portals. [Section 5.1]
- The JGI's data management system manages storage resources, and JAMO will always have a managed repository at NERSC, but this may change in the future. For egress to data centers other than NERSC, two models are under consideration [Section 5.1]:
  - Allowing JAWS to copy the files from NERSC to a remote site as needed.
  - Setting up managed repositories at remote data centers that JAMO can manage directly, and allowing transfer back-and-forth to NERSC.
- Underlying the KBase platform is a service-oriented architecture that runs across a distributed set of resources located at ANL and LBNL. [Section 5.2]
- KBase users primarily interact via a web-based platform for supporting data analysis and data discovery. Typically, this involves using public datasets that are already available in the platform, by uploading new data into KBase (or the JGI). The data (public or private) can be analyzed using a variety of applications that are exposed in the Jupyter-based Narrative interface. [Section 5.2]
- The NMDC does not anticipate any data-related resource constraints in the coming years. Risks are present in that datasets generated outside the DOE user facilities (the JGI and EMSL) may grow and outpace local storage. [Section 5.3]
- EMSL houses over 150 advanced and often one-of-a-kind instruments [Section 5.4]:
  - The raw data generated on EMSL instruments is uploaded to EMSL's data repository where it is stored in an active storage location during analysis, then moved to long-term storage in the EMSL data archive.
  - Data is typically processed prior to delivery to users. The data is often initially processed locally on the instrument control computers prior to transfer to a computer cluster or workstation for more intensive data analysis, and for delivery to users the fully processed data may take a variety of forms, from structured spreadsheets, to community-defined data formats, to image files.
  - Raw data and processed data are available to users when they authenticate through the EMSL user portal. Small data files (tens to hundreds of kilobytes) are also often emailed to users, and large data files (up to terabytes) are typically delivered using either https transfers or the Globus service.
- The next five years in EMSL will see growing data volumes due to improvements in instrument resolution, and large increases in the rate of data generation due to laboratory automation. Automation efforts will focus on mass spectrometry workflows. EMSL expects to provide variable autonomous operations in a multimillion-dollar Microbial Molecular Phenotyping (M2P) facility, for which DOE has approved the Mission Need (CD-0). [Section 5.4]

- Approximately 100 types of instruments are deployed at ARM sites. Major instruments include [Section 5.5]:
  - Profiling and scanning radars ranging in frequency from 915 MHz to 94 GHz to sample cloud and precipitation droplets.
  - A variety of lidars designed to measure profiles of temperature, humidity, water vapor, air motion, and aerosol properties.
  - In situ aerosol particle analyzers for studying physical and chemical properties of aerosols.
  - Wide-ranging in situ probes deployed on aircraft and tethered balloons to observe properties of clouds, aerosols, and the background atmospheric state.
  - ARM instruments that produce approximately 2 TB of raw data daily, and the resulting analysis can be more than 6 TB.
- ARM is exploring the application of new measurement technologies, such as phased array radar, differential absorption lidar, and an increasing number of miniaturized sensors for deployment on UAS and TBS. [Section 5.5]
- Many ARM users download a single dataset or a few datasets to study a particular phenomenon, evaluate other instruments (e.g., from a satellite), or evaluate or develop components of an Earth system model. These downloads can be gigabytes to terabytes in scale. Power users may download significantly more in bulk when performing larger studies. [Section 5.5]
- The ESGF mission is to [Section 5.6]:
  - Support current CMIP6 activities, and prepare for future assessments.
  - Develop data and metadata facilities for inclusion of observations and reanalysis products for CMIP6 use.
  - Enhance and improve current climate research infrastructure capabilities through involvement of the software development community and through adherence to sound software principles.
  - Foster collaboration across agency and political boundaries.
  - Integrate and interoperate with other software designed to meet the objectives of ESGF: e.g., software developed by NASA, NOAA, ESIP, and the European IS-INES.
  - Create software infrastructure and tools that facilitate scientific advancements.
- When considering the resolution ratio from a standard CMIP (~100 km) to a standard DYAMOND (~4 km) simulation, data volumes balloon by a factor of 625x. At this early stage, the most significant simulation planned is for 30 years with a 5 km coupled model, producing 40 TB per simulated month, and ~15 PB over the length of the proposed single simulation. [Section 5.6]
- The data growth over subsequent CMIP phases (CMIP3 -> 6) has been considerable, with an expansion of phase complexity, in addition to increasing grid resolutions that scale exponentially. For CMIP6, storage limitations as ESGF partner nodes have ensured that no one center has a complete copy of the CMIP6 archive, with high-value, and high-demand datasets replicated across the primary tier 1 nodes, but in an ad-hoc way that does not ensure data redundancy across the federation. [Section 5.6]
- It is likely that the output from the ultra-high resolution simulations being run in DYAMOND experiments will test network capabilities into the future, in the case that remote analysis is

even possible. If compute resources are not available next to published data, then both network and storage requirements will scale markedly to meet science needs. [Section 5.6]

- CMIP7 models will be four to five times larger than CMIP6, which places the evolution of data volumes on a path to exascale. Preparation to accommodate this, in terms of storage and data movement, will be required by many community repositories. [Section 5.6]
- ESS-DIVE collects, stores, manages, and shares data. The volume, complexity, and diversity of Earth and environmental science data make it challenging to capture, store, verify, analyze, and share information in a consistent manner. [Section 5.7]
- ESS-DIVE's long-term vision is to provide a repository that enables easy intuitive integration and management of data from models, experiments, and observations. In this vision, ESS-DIVE becomes a focus for enabling access to integrated data for visualizations, models, ML, and other analyses for knowledge discovery from these datasets. [Section 5.7]
- The data acquisition is performed using DECTRIS PILATUS 6M at 19ID, and Area Detector Systems Corporation (ADSC) Quantum 315 (Q315) at 19 BM. The maximum data rate for 19ID with PILATUS 6 M is 100 images per second (650 Mbytes per second), and for 19 BM – 1 full image per 2.7 seconds (30 Mbytes per second). [Section 5.8]
- CABBI's datasets include phenomic, genomic, transcriptomic, metabolomic, proteomic, and fluxomic data for both bioenergy feedstocks and the engineered microbes, in addition to bioeconomic, edaphic, and climatic data for sustainability determination. [Section 5.9]
- The CABBI data management plan focuses on each researcher or lab group being responsible for developing data collection and storage procedures that suit their work with the requirement that data is shared, well-documented, and backed up in multiple secure locations [Section 5.9]:
  - Researchers are required to describe their data (i.e., variables collected, temporal or spatial resolution), where it is stored, and who is responsible for it on CABBI's researcher website.
  - Once data has been analyzed and quality tested, researchers are then required to share datasets internally as manuscript development is undertaken.
  - Data-sharing locations include cloud-based locations, local servers, and web applications.
  - Researchers are required to publicly share supporting data after publication.
- The CABBI iBioFAB is an automated laboratory facility located in the IGB. The system supports rapid design, fabrication, validation/quality control, and analysis of genetic constructs and organisms [Section 5.9]:
  - A variety of data is generated for the various iBioFAB instrumentations. The number of files and the size are dependent on the analysis being performed. All data is retrieved for the computers attached to the instruments.
  - Data is manually aggregated for the various instruments and emailed to the requesting researchers for analysis.
  - There are general plans for the software team to develop standardized methods for storage and delivery.
  - While data collection and transfer are primarily performed at the facility and can make use of the University of Illinois network, at some future point, users of the facility may want to remotely operate instrumentation. This could potentially become a challenge for the current protocol.

- All data supporting CABBI publications is released either via Box or through publishing via a data repository. Subject-specific repositories like the National Center for Biotechnology Information (NCBI), AmeriFlux, DAAC, and Phytozome are preferred, but where those are unavailable, others, like university (Illinois Data Bank and Iowa State University DataShare) and public (Zenodo, Mendeley, and GitHub) repositories are used. [Section 5.9]
- CBI has identified research targets for a thriving bioeconomy [Section 5.10]:
  - Sustainable biomass feedstock crops using plant genomics and engineering.
  - Advanced processes to simultaneously break down and convert plants into specialty biofuels.
  - Valuable bioproducts, including chemical feedstocks, made from the lignin residue remaining after bioprocessing.
- CBI has identified four innovation targets that will enable the future bioeconomy [Section 5.10]:
  - Development of high-yielding, process-advantaged, sustainable plant feedstocks.
  - Creation of lower-cost consolidated bioprocessing to produce fermentation intermediates.
  - Valorization of lignin to atom-efficient biofuels and bioproducts.
  - Development of chemocatalytic conversion processes for biobased intermediates to produce a complete biomass-based sustainable aviation fuel (SAF).
- Hyperspectral instrumentation located both in the greenhouses and in the field is a major source of large data volumes for CBI. Each hyperspectral image is approximately 1 gigabyte of data, and the instruments can produce an image every three minutes. [Section 5.10]
- GLBRC produces heterogeneous data types, at a variety of time and spatial scales and at ever-increasing rates, with different management requirements, and has a variety of pre-existing data management practices. The key data management challenge is to leverage diverse resources and systems in support of the needs of individual research projects while optimizing information integration to promote collaboration and productive synergies across the center. [Section 5.11]
- GLBRC works with the following types of data [Section 5.11]:
  - Biomass compositional analysis.
  - DNA/RNA sequencing.
  - High-throughput (HT) chemical genomics.
  - Droplet-based microfluidics.
  - HT mass spectrometry (MS)-based protein and metabolite analysis.
  - Proteomic analysis.
  - Metabolite analysis.
  - Experimental fermentation lab (EFL).
  - Structural biology.
  - Cryo-EM and Cryo-ET (evapotranspiration) instrument access and training.
  - Imaging.
  - Separations.

- JBEI researchers, by virtue of the scope of the project, will generate data of many different types, which need to be preserved and shared with the broader scientific community [Section 5.12]:
  - JBEI data will be available to the scientific community beyond that published in peer-reviewed journals.
  - JBEI data will be preserved to the maximum degree possible through inclusion in public, long-lived databases.
  - Larger datasets will be deposited with recognized community databases.
  - The research community will be provided with other JBEI datasets through community resources such as KBase, the Experimental Data Depot (EDD), and the Inventory of Composable Elements (ICE) registry.
- JBEI research will generate multiple kinds of experimental data in a digital format that can be used by other researchers to validate research findings and potentially further analyzed [Section 5.12]:
  - Genomic data will be obtained through gene sequencing and will be created through construct design.
  - Multiple functional genomic data types (transcriptomics, proteomics, metabolomics, and fluxomics) will be generated in the characterization of engineered microbes, adapted microbes from adaptive laboratory evolution (ALE) experiments, and engineered plants.
  - Enzymatic data will be generated in the analysis of wild type and in engineering glycosyl hydrolases, lignases, and other biosynthetic enzymes.
  - NMR studies of plant biomass will generate multidimensional spectroscopic data.
  - Structural studies of these same proteins will generate crystallographic and cryo-electron microscopy data, and associated atomic models.
  - For DNA, RNA, and protein sequences, JBEI uses community data-exchange formats including FASTA, GenBank, and Synthetic Biology Open Language (SBOL).
  - For proteomic and metabolomic data, JBEI will use community standard data formats and processing methods for compound identification and validation.
  - For structural biology data, JBEI will make use of standard formats for internal workflows, including experimental data, molecular model, and geometric restraints formats.
- Earth system modeling is data intensive. Data transfer can be needed at every phase of the model development cycle for E3SM, including [Section 5.13]:
  - Running the model.
  - Archiving and publishing the simulation.
  - Post-processing and analyzing the simulation.
- AmeriFlux Core Sites, and the National Ecological Observation Network (NEON), typically upload datasets less than 100 GB in size for yearly data. Noncore sites have similar data volumes, but the cadence of upload (as well as mechanism) varies depending on access to data. [Section 5.14]
- AmeriFlux data products are publicly and freely accessible [Section 5.14]:
  - Most data users are based at universities, research institutions, or government agencies worldwide.

- Over the last seven years, about 5,000 users globally have downloaded AmeriFlux data, totaling around 28,000 unique downloads.
- Users can typically download data less than 100 GB in size, and in some cases may be interested only in a very small subset of data.

- Data transfers from measurement facilities to data repositories involve a preprocessing step to convert the data to a required format, and then upload it to a data processing facility via a web portal for QA/QC checks. Some measurement facilities transfer raw high-frequency data to NERSC through secure copy protocol (SCP). [Section 5.14]
- The NGEE projects, and their local teams, collaborate on data collection, sharing instrumentation and data protocols, doing data processing and data quality control, and generating data products, which are ultimately stored in data-archive facilities. It is anticipated that for long-term preservation, this data is also synchronized with the ESS-DIVE data repository. [Section 5.15]
- NGEE–Tropics field sites typically produce observations that can vary in size from a megabyte a few times a month, to a gigabyte over the course of a year. The frequency of sharing this depends heavily on data availability, network connectivity, and local staff availability to retrieve sensor data. Data is shared through a variety of tools, with cloud services being popular along with other more traditional methods like Globus and SCP. [Section 5.15]
- NGEE–Tropics uses data from collaborator repositories (e.g., NASA, NOAA, and NCAR), which can grow to TB scales and typically must use Globus for transfer. [Section 5.15]
- The data life cycle of field NGEE instruments is as follows [Section 5.15]:
  - Light computation happens at time of collection (in the field) to normalize results.
  - At the point of retrieval, data goes through additional processing to check for data integrity and completeness.
  - Data analysis typically happens at host institutions.
  - Many abacus variants are possible: dozens of KB-level files, millions of KB-level files, thousands of MB-level files, dozens of gigabyte-level files, all in many possible storage organization structures/directories.
- COMPASS-FME data from sensors amounts to datasets that are kilobytes to megabytes in size and are collected constantly. These are typically shared via cloud sharing services, and stored either at the facilities from the project partners or shared cloud resources. [Section 5.16]
- COMPASS-GLM model datasets will be shared between two multi-institutional task-level teams for atmosphere–land–lake regional modeling (Great Lakes Region) and high-resolution watershed modeling (Portage River Basin). The models are GB scale, and are stored on HPC resources at ANL, LBNL, and ORNL. Globus is used for data sharing, along with other tools such as web portals and SCP. [Section 5.16]
- COMPASS generates two broad categories of data, and makes heavy use of Jupyter Notebooks to document computational workflows including data sources, codes used, and all associated metadata. To that end, both efforts will generate, and work with [Section 5.16]:
  - Observational and experimental data generated by field and laboratory-based experiments.
  - Simulation data generated from multiscale modeling activities.

- COMPASS-GLM raw data output from WRF-FVCOM is approximately 200 MB over the course of a day. This is a combination of over 20 variables from the sensors. These are combined into multiday datasets that are approximately 10 GB in size. Raw data is shared via Globus, and used in modeling activities on HPC resources. The output files from the transient daily integrated hydrology simulations in ATS for the Portage basin (~1,000 km2) for 15 years of simulation are about 150 GB. [Section 5.16]
- Most major data transfer operations for COMPASS, both incoming and outgoing, make use of the Globus data mobility software. Smaller data transfers may use other protocols such as SFTP or https. [Section 5.16]
- The majority of data-related COMPASS-FME issues surround data sharing and tracking [Section 5.16]:
  - Cloud services, like Google Drive, do not have enough flexibility with permission options and leave the project vulnerable to data loss. For example, users must be given edit access to upload data, but this also gives users permission to delete or change data as well.
  - It is desirable to track dataset usage and version history to ensure that people are aware of how up-to-date a dataset is.
  - Sharing dataset availability and contact information associated with a dataset to the team is challenging. Often, people do not know whom to contact about locating a specific dataset, due to difficulty in searching.
- Climate model output is collected from several sources, each of which may result in TB datasets being retrieved using Globus primarily [Section 5.17]:
  - The Program in Climate Model Diagnostics and Intercomparison (PCMDI).
  - The Rutherford Appleton Laboratory in the UK.
  - The ESGF.
- CACADE leverages 14 PB of HPSS storage at NERSC, and routinely does analysis using containers that operate Jupyter [Section 5.17]:
  - 5 PB is devoted to retrospective simulations of the 20th century with and without effects of anthropogenic climate change.
  - 1.5 PB is devoted to Model Intercomparison Projects (MIPs) from the most recent Coupled Model Intercomparison Project (CMIP6).
- The global climate model output analysis community (currently thought to be in the hundreds) will retrieve model output from well-established community portals (e.g., CDG and ESGF). It is expected that these downloads could range from megabytes to terabytes in size, depending on the breadth of model that is retrieved. [Section 5.18]
- Facilities already have a mature set of tools for data management and providing remote data access. The main challenge in enabling multifacility workflows in the coming years will be improving how these systems connect to each other and exchange information. Data movement is in many ways the most easily addressed given the current capabilities of tools like Globus. The major challenges will be in developing standards around metadata, data exchange, authentication and authorization, policies, and tooling to support all of these areas. [Section 5.19]
- The "last foot" problem of data movement, e.g., where some tools are better run with laptops, still persists in a number of BER workflows. Data movement platforms such as Globus have simplified this from an application layer perspective, but there is nothing that can be done when

running over local wireless connectivity to personal computing devices. Adapting workflows to use DTNs to move data "closer" is the best mitigation, followed by streaming or bulk downloading to a laptop resource as a final step. [Focus Groups]

- Some larger BER facilities and experiments note they will see data rates increase in the coming years. ESGF (and all users of CMIP resources) expects to see an explosion as new models are released. These will require fast networks between the major HPC facilities, as well the locations of some of the ESGF nodes. It is expected that some simulations will be PB scale, approaching EB, and there will be hundreds of these in the coming years. [Focus Groups]

2.3 Scientific Workflow

- JGI Sequence data is uploaded to the National Institutes of Health's National Center for Biotechnology Information Sequence Read Archive (SRA), where it is freely accessible for research and education purposes. [Section 5.1]
- JAWS is designed to move the JGI away from bespoke and/or hand coded workflows, and to protect production analysis pipelines from outages and/or downtimes at partner HPC facilities (e.g., NERSC, LBNL IT's Lawrencium cluster, EMSL's [PNNL] Tahoma cluster, as well as private cloud clusters in AWS). [Section 5.1]
- The JGI is an important partner to KBase, and KBase users can use the JGI Portal to send datasets from the JGI directly into their workflows. This results in multiple terabytes of data transfer on a routine basis. [Section 5.2]
- KBase users can use Globus to transfer data from any data provider that makes data available via Globus. [Section 5.2]
- The NMDC can be viewed as a distributed data facility that combines JGI, EMSL, and NERSC resources. The NMDC data products are currently housed in multiple locations including EMSL and NERSC, based on the resource where the files were generated. The fully operational central metadata store maintains references to the locations of the data and access protocols, allowing the data to be retrieved when requested by users. [Section 5.3]
- EMSL expects to see increases in data transfers produced by other institutions for integration, analysis, and visualization in the next five years, and data products being copied to resources like KBase for analysis and to repositories like ESS-DIVE for storage. [Section 5.4]
- Data reuse by the community is expected to increase as the value of the accumulated data improves from multimodal analyses [Section 5.4]:
  - Instruments will increase in size and capability and could yield up to multiple terabytes per day.
  - Automation pipelines will be established, with expected data generation to increase to tens to hundreds of terabytes per day, and data-archive capacity is expected to grow to 500 PB.
  - Continued growth of the FICUS program will increase demand for transfer of data between facilities.
  - In the future, the M2P Facility will be in operation, increasing data generation to petabytes per day.
- The software used to manage data generated in EMSL has been custom designed to provide the staff who operate instruments with the tools needed to transfer instrument data and the associated user project metadata to EMSL's data archive. [Section 5.4]

- ARM is striving to develop a holistic and internally consistent description of the atmosphere. A goal in the next two to five years will be to develop more expanded datasets combining observations and model output to provide detailed descriptions of the atmosphere that are constrained by observations. Such integrated datasets will inevitably have large volumes, as they represent many variables in three-dimensional space and time, much like model output. Working with these comprehensive datasets will likely require more remote access, again similar to working with a large model or satellite data, in contrast with the more common mode of investigators downloading a few specific data streams and parameters today. [Section 5.5]
- All ARM measurements are remote; typically, few technical ARM staff are located at ARM observatories. This makes it challenging to modify the operation of instruments remotely, and to configure instruments to change their mode of operation in response to the environment. [Section 5.5]
- In the future, it will be highly desirable to operate UAS and/or TBS in an automated or semi-automated way such that ARM could obtain arrays of vertical profiles around an ARM observatory. [Section 5.5]
- The ADC archives high-resolution data bundles from LASSO model simulations over several ARM observatories [Section 5.5]:
  - These large eddy simulation (LES) model outputs are summarized and packaged within the ADC at ORNL to provide a concise representation of the LASSO simulations for sending to data users at other institutions.
  - These simulations represent a significant fraction of the data at the ADC and are expected to soon grow to several petabytes. When large volumes of LES output need to be transferred to the user's computing resources, ARM anticipated that many of these data transfers will use the ESnet infrastructure.
  - There is an opportunity to provide next-generation data connectivity for large data transfer between laboratory computing facilities to facilitate distributed data analysis.
  - Scientists using data analysis capabilities and visualization tools deployed in other leadership computing facilities would benefit by having seamless access to the full-resolution LASSO bundles archived at ORNL.
- ESGF P2P is a component architecture expressly designed to handle large-scale data management for worldwide distribution. The team of computer scientists and climate scientists has developed an operational system for serving climate data from multiple locations and sources. [Section 5.6]
- The DYnamics of the Atmospheric general circulation Modeled On Non-hydrostatic Domains (DYAMOND) project is an international consortium that has built on the momentum of CMIP, focused on ultra-high resolution global storm resolving models (GSRMs) with resolutions higher than 5 km globally. [Section 5.6]
- ESGF provides data discovery and access services to a broad global user community. Thousands of users (clients) discover data either via a web portal (graphical user interface [GUI]) or data search API. These services provide the endpoints so the users may download via HTTP, transfer via Globus, or access via the OPeNDAP protocol which enables limited data subsetting. [Section 5.6]
- In the future, more users will make use of "server-side" compute and remote analysis capabilities. By moving the compute close to the data, ESGF expects that the end products that need to be transferred back to the user will become trivially sized, e.g., MB-sized subsets or images. [Section 5.6]

- BER projects generate diverse observational, experimental, and modeling data: the aim is to become the repository of choice for ESS project data by working closely with the ESS community to build the repository based on community needs and priorities. ESS-DIVE accepts submissions of well-structured, standardized, high-quality data from data contributors and facilitates data processing, synthesis, and analysis capabilities accessible to data users around the world. Datasets include raw output from sensors, simulations, analyzed datasets, and user-contributed data. [Section 5.7]
- Initial ESS-DIVE work focused on developing and establishing a well-known repository for data. Ongoing and future work will expand the user base and functions through five key innovations [Section 5.7]:
  - Apply user-experience research methods to understand user needs.
  - Support early data archiving by projects.
  - Support scaling of the repository to a significantly broader portion of the ESS community.
  - Support search of extracted ESS-DIVE data with a fusion database.
  - Support federation with other repositories.
- To meet the needs of field data from sensors and instruments, it will be critical to enable data acquisition, transformation, and analytics workflows that can integrate real-time data streams with predictive models, particularly for extreme events. These can be used to train models using a hybrid of AI/ML and deterministic techniques, which can generate real-time predictions and feedback for sensors out in the field. The highly distributed nature of sensor networks, and the volumes of data collected, will necessitate a data "backbone" for AI/ML-based scientific workflows. [Section 5.7]
- SBC has implemented serial crystallography experiments, which enable the study of protein structures under physiological temperature and reduced radiation damage by collection of data from thousands of crystals. [Section 5.8]
- Synchrotron serial crystallography (SSX) has emerged as a valuable approach for low-dose room-temperature structural biology research that also allows for the study of dynamic processes in protein crystals, such as chemical transformations. For the real-time processing of the SSX data collection, SBC has developed a data analysis pipeline, Kanzus, to bridge the beamline with high-performance computing (HPC) and storage capabilities provided by Argonne Leadership Computing Facility (ALCF). [Section 5.8]
- SBC features a three-tier software architecture to handle experiments [Section 5.8]:
  - The lower level is Experimental Physics and Industrial Control System (EPICS) and beamline device controls.
  - The middle layer is SBCserver, which is an EPICS-based sequencer with the ability to be driven by socket commands.
  - The upper levels are GUIs, namely SBCcollect.
- SBC instrumentation consists of X-ray optical elements, beam position monitors, XRF measurements instruments, robotic sample mounting stages, visualization camera systems with auto-alignment software, and diffraction images recording detectors (ADSC Q315 and PILATUS3 6M). [Section 5.8]
- CABBI DNA sequencers are used extensively and produce large datasets [Section 5.9]:
  - The main data-producing technology is currently the Illumina system. The system's instruments generate up to 8 terabase pairs per day, which would translate to 2 to 20

terabytes depending on the data format, retention of quality values, nature of sequence, and the compression used.

- Current CABBI usage is around 30 terabases per year, which in raw data usually requires 50 terabytes or so in the formats usually employed.
- Data is retrieved from the DOE facilities the JGI and EMSL. Data is also sometimes shared between CABBI partner sites.

- CABBI uses a suite of models including three ecosystem models, an economic model, and a techno-economic analysis/life cycle assessment model. The researchers using these systems collaborate frequently, sharing both data and code [Section 5.9]:
  - Some modelers make use of BioCluster, IGB's fee-based high-performance computing resource. BioCluster contains 2,824 cores and over 27.7 TB of RAM with 1.3 PB of storage on a general parallel file system (GPFS).
  - Researchers archive their simulations, including input data, code, and results, in a repository stored on IGB's file archive server.
  - The uploaded files vary in number and size depending on the scope of the simulation and the model being used. Currently, archived simulations require 2 MB to 110 MB.
  - Researchers access this server via their individual networks using a secure file transfer protocol (SFTP) client.
- Data from various research components within CBI are shared, stored, and preserved by building upon existing integrations of resources. CBI datasets are TB to PB scale. The diverse data types in CBI can be categorized into basic data levels. The EarthData processing levels devised by NASA's Earth Observation System Data and Information System largely inspired CBI's data processing levels [Section 5.10]:
  - Level 0 data is raw, unprocessed data from original sources. Sources are either instruments or recorded observations.
  - Level 1 data is typically generated by a third-party provider on behalf of CBI.
  - Level 2 data is derived from data in Levels 0 and 1 and consists of data in standardized formats that are widely recognized within each field of study.
- CBI provides its users with a cloud-based Laboratory Information Management System (LIMS) repository based on the LabKey Server technology. The LabKey Server system houses Level 0 and Level 1 data. CBI investigators are encouraged to use the system to track and exchange their Level 0 and Level 1 data assets. The CBI LIMS can accommodate data dumps from both manual and automated sources. [Section 5.10]
- GLBRC provides researchers with access to an array of in-house analytical, computational, and business services. These include [Section 5.11]:
  - Experimental measurements and field observations.
  - Experimental designs and protocols.
  - Equipment descriptions, designs, and procedures for use.
  - Materials inventories and characterizations.
  - Algorithms and source code.
- GLBRC workflow typically involves [Section 5.11]:
  - Analysis of materials generated in field experiments.

- Different datasets being compiled for a number of core measurements (e.g., climate, soil, emissions, and other field sensor measurements).
- Composition and quality of materials determined through analysis at member institutions and national labs.
- Materials being transferred for processing through a variety of platforms generating additional rounds of data from DOE national labs and institutional core facilities.
- Additional analytic techniques possibly being applied at any stage of this process to assist with experimental goals.
- This data being transferred to, and stored in, the GLBRC data catalog.

- Optimizing the GLBRC workflow integration process to improve collaboration and productive synergies is a key challenge, and network capabilities play a critical role. [Section 5.11]
- JBEI has extensive, fully operational research facilities and resources at Emeryville Station East (ESE). Processed data is added to the JBEI Experimental Data Depot (EDD) to aid information transfer to individual researchers and data storage and communication with the biofuels community. Capabilities include [Section 5.12]:
  - Analytical laboratory: instruments to determine reaction outcomes.
  - Biomass pretreatment laboratory: to test the effects of mechanical and chemical processing of biomass.
  - Chemistry laboratory: to develop reactions, increase scale, study the effects of high temperature and pressure, and safely handle organic solvents and other hazardous materials.
  - Fermentation laboratory: to study the physiology of microbial strains engineered for production of biofuels or enzymes during fermentation.
  - Microfluidics laboratories: microfluidic tools for HT synthetic biology, using machine learning (ML).
  - Microscopy laboratory: provides researchers with a suite of spectral microscopes.
  - Plant growth facilities: provide precisely controlled indoor greenhouse facilities for growing genetically engineered feedstocks.
  - Robotics laboratory: commercial automation equipment to enable HT research.
  - Spectroscopy laboratory: specialized instruments to detect and characterize biofuel feedstocks.
  - Synthetic biology laboratories: multidisciplinary spaces that house specialized equipment for experimentation.
  - NextGen sequencing laboratory: next-generation sequencing research efforts such as RNA-seq analysis, genome resequencing applications, and general quality assurance and quality control (QA/QC) for synthetic biology and metabolic engineering.
  - X-ray crystallography facilities: protein purification equipment for the preparation of protein samples.
  - Computing server room: houses JBEI systems for data storage, data backup, bioinformatics analysis, Docker services, and web services.
  - Access to a computing cluster maintained by LBNL IT.

- Running E3SM generates a large amount of data [Section 5.13]:
  - Some data generated during the model development phase is one-off and not saved.
  - For production-level runs, the project requires a standard workflow to evaluate, archive, and publish the data. Storage allocation is available at NERSC and LLNL.
  - Moving data is important for post-processing and analyzing the simulations. For simulations from a group of ensemble runs, or for very long simulations, generated data are sometimes distributed across multiple machines, while the results need to be concatenated or evaluated against each other.
  - All data transfer has been done using Globus and via ESnet.
- AmeriFlux data's life cycle involves [Section 5.14]:
  - Data collection at each site, using a full suite of instruments.
  - Processing of high-frequency observations into half-hourly time series.
  - Preparation of the data in a standardized format.
  - Submission to the AmeriFlux website: a web data portal and a semi-automated data pipeline to accommodate data submissions, quality-checking, tracking, and releasing.
- Several NGEE initial field studies are underway at key sites, with measurements that address [Section 5.15]:
  - Impacts of thawing permafrost on greenhouse gas emissions.
  - Impacts of wildfire and thermokarst disturbance on tundra carbon balance.
  - Whether the Arctic is becoming wetter or drier.
  - Whether the Arctic is greening or browning.
  - Forest carbon cycle—hydrology interactions.
  - Nutrient limitations in tropical secondary forests.
  - Plant functional diversity response to climate change.
  - Regional variation in the causes of tree mortality.
- NGEE Arctic and NGEE–Tropics will each deliver an ecosystem model extending from bedrock to the top of the vegetative canopy/atmospheric interface that can be modeled at the scale of a high-resolution earth system model grid cell (~10 x 10 km2 grid size). [Section 5.15]
- The activities and field campaigns for NGEE Arctic at remote field sites in Arctic Alaska target uncertainties in model initialization, parameterization, and evaluation, and include the physical and remote monitoring of multiple surface and subsurface processes, ranging from abiotic variables such as temperature, moisture, and rain and snowfall, to biological variables from microbial activity to plant community composition and physiology. [Section 5.15]
- NGEE Arctic data workflow is described as follows [Section 5.15]:
  - Researchers are responsible for the collection and management of data until it is ready for submission to the project's archive and the data management team (DMT).
  - Researchers using the project archive create metadata using the online metadata editor (OME) or email to initiate submission.
  - Data (raw, processed, etc.) is submitted via OME upload, email, Globus, Dropbox, Google Drive, or other file-transfer system.

  - Researchers using other data-type specific archive facilities like AmeriFlux, JGL, or GitHub will follow their workflow submission protocols.
  - Data is shared with the public as CCby4, with some datasets reserved for team members until eventually released to the public with a manuscript publication.
  - Digital Object Identifiers (DOIs) are assigned using the DOE Office of Scientific and Technical Information (OSTI) application programming interface.
  - Data is currently stored on an NGEE Arctic server at ORNL located in Oak Ridge, Tennessee.
  - Data will initially be mirrored and eventually migrated to the DOE ESS-DIVE.
  - The NGEE Arctic project archive at ORNL currently has 233 datasets with 161 available to the public, 62 available to the team only, and 10 in the planned (forthcoming) stage.
- NGEE–Tropics workflows involve [Section 5.15]:
  - Receiving and processing field data from multiple sites.
  - Harmonizing this data in preparation for experimental analyses.
  - Synthesis analyses (multiple types of data combined).
  - Feeding analyzed data into simulations/model runs on HPC resources.
  - Synchronizing all publicly available datasets (currently 109) with ESS-DIVE, with automatic synchronization of new datasets.
- COMPASS-FME operates seven real-time environmental sensor clusters (e.g., synoptic nodes) at sites along Chesapeake Bay and Lake Erie. Each node contains data loggers and associated communication infrastructure, including cell phone and radio telemetry, as well as hardened solar power infrastructure. Each node produces about 5 Gb raw data per year, with 40 Gb of processed data annually. Current instrumentation includes [Section 5.16]:
  - Water quality measurements.
  - Meteorological parameters.
  - Soil measurements.
  - Tree health.
- COMPASS-FME aims to make data available through a collaborative network. The workflow will develop data distribution infrastructure and data products, and is designing a system that will facilitate interactivity with the data, as well as facilitating access. Starting with cloud services for data sharing and storage has revealed several areas of friction [Section 5.16]:
  - Performance when uploading large datasets is not predictable.
  - Permissions to release and share data with collaborators is not consistent, and often does not integrate with institutional identities.
  - Data changes are challenging to track over time.
- A common document-sharing platform supported/certified by all DOE laboratories will be helpful for collaborative work in COMPASS and potentially improve productivity. [Section 5.16]
- None of the core CASCADE partner facilities provide mechanisms to analyze the data “server side,” so it is necessary to use networks to transfer the data to local systems (e.g., Cori, etc.) where data can be operated on. [Section 5.17]

- It would be desirable to have the ability for moderate-scale HPC centers to remotely mount NERSC filesystems. This, combined with some sort of ability to preemptively cache data locally at the HPC sites before a job makes use of the data, would enable a distributed super-facility model to support the CASCADE workflow. [Section 5.17]
- Simulations are carried out on a number of supercomputers, including the NCAR/University of Wyoming, Cheyenne, and NERSC's Cori and Perlmutter, and others [Section 5.18]:
  - Outputs from these simulations are often transferred between the computing sites for analysis, depending on specific aspects of the simulations involved.
  - The preferred tool for data transfer is Globus.
- The typical process for the use by CATALYST of models for knowledge discovery involves the following steps [Section 5.18]:
  - Experimental design is created.
  - Resources (computing, storage, and so on) are determined, applied for, and allocated at a computing center.
  - Model is executed.
  - The output is analyzed and archived.
  - Output is made available via the CDG or ESGF as appropriate using data mobility tools and portal systems.
  - Results are published.
- The most significant change to current CATALYST practices will be reengineering of the CESM workflow, to enable the creation of single-field time series format data as the simulation is ongoing. This will enable the global user community to have easier and more efficient access to CESM results. [Section 5.18]
- CATALYST and CESM will continue to rely on Globus, the ESGF and CDG, and their follow-on projects to publish and deliver model output to the user community. [Section 5.18]
- Data transmission amounts between coupled BER facilities vary from megabytes and gigabytes when dealing with the sharing of raw sensor or observational data, to multiple terabytes to petabyte scales when moving large models between HPC resources. Tools like Globus are critical to managing the data mobility aspects, along with other methods such as SCP, HTTP portals, and the use of cloud sharing and storage. [Section 5.19]
- JAWS is a workflow service designed to move the JGI away from bespoke and/or hand coded workflows, and to protect the production analysis pipelines from outages and/or downtimes at partner HPC facilities. [Section 5.19]
  - This has enabled the JGI to move some of its analysis runs off NERSC, and into clusters such as LBNL IT's Lawrencium Cluster and EMSL's (PNNL) Tahoma cluster, as well as private cloud clusters in AWS.
  - It is expected that most of the JAWS workflows will operate within LBNL, split between NERSC and a JGI-dedicated cluster hosted within LBNL IT, with only a small fraction going off site to PNNL and AWS.
  - As a consequence, ESnet may not see much traffic between LBNL and other sites due to JAWS. The general trajectory of data growth will be tied to the growth of JAMO and NMDC data.
- New workflows to better stream, store, process, and share BER data will be needed in the future as longer paths between observation and result are created. [Focus Groups]

2.4 Computational and Storage Requirements

- The main computational systems at JGI with multisite footprints are part of the Advanced Analysis team, within the Data Science and Informatics department. Both systems are multisite distributed systems with significant internal backend data analysis and management processes that span multiple sites. Within Advanced Analysis, the two major teams are [Section 5.1]:
  - The Sequence Data Management (SDM) team, which supports the JGI Archive and Metadata Organizer (JAMO). JAMO manages over 14 PB of data storage and contains raw sequencing data and analysis artifacts for the JGI dating back to the late 1990s.
  - The JAWS team operates a distributed workflow service that allows analysis workflows to be distributed across multiple sites to provide resilience and additional capacity to JGI analysis. The sites involved are PNNL, LBNL (including NERSC, LBNL IT, and the JGI), and Amazon Web Services (AWS).
- The working set for JAMO is currently under 1 PB, but grows slowly and will likely be over 1 PB in the next couple of years. The archive in HPSS grows more rapidly, on the order of 1.5 PB per year currently. JAMO data is never deleted, so data usage grows monotonically. [Section 5.1]
- KBase is upgrading system capacity to add GPU nodes and additional storage (e.g., the goal is to be 90% of capacity). Recent storage additions now have the project at 58% of capacity, and KBase foresees the need to repeat the addition of storage regularly. In addition to the LBNL and ANL hardware, KBase has a NERSC allocation for HPC workloads and to use the Spin service. This is used to host KBase services, and it can also be used to leverage resources on the Perlmutter HPC resource. [Section 5.2]
- KBase does not have any immediate plans for using cloud services for routine computations; the ANL, LBNL, and NERSC resources are sufficient for the current demand. The cost of cloud computing is too great for the minimal “burst” use cases that may arise from time to time. [Section 5.2]
- The NMDC expects that the commercial cloud may be used for resilience for the front-end web service (search of metadata only), as well as to support "surges" in computing capacity. Since the NMDC does not have a "real-time" computing requirement, supporting the later use is not critical. The NMDC primarily utilizes resources at EMSL and NERSC to run its standardized workflows. These use an opensource workflow service and NMDC is planning to adopt JGI Analysis Workflow Service (JAWS) in the future. The NMDC is also exploring AWS infrastructure to run workflows. [Section 5.3]
- With automation will come the need for an order of magnitude or more increase in data storage, analytics, on-the-fly analysis, modeling, and computation. The EMSL data archive is expected to exceed 500 PB within the next five years to accommodate the increase in data generation. [Section 5.4]
- The ARM ADC is responsible for [Section 5.5]:
  - Archival storage for over 11,000 data products with a total repository holding of over 3.5 PB of data that dates to 1992. These include data from instruments, value-added products, model outputs, field campaigns, and PI-contributed data.
  - Data access and delivery options, which include THREDDS/OpenDAP, GlobusOnline, and near real-time data access API, automated data access via web services, advanced visualizations, and a big data analysis platform.
- ARM data is collected from instruments deployed in remote locations, and the data is sent to ORNL via a secured and encrypted transfer method. Data is deep archived at the ORNL mass

data storage system (HPSS), and an additional off-site data backup is done using the ANL's HPSS. [Section 5.5]

- Automated processes are used in continuously monitoring ARM data and metadata. These automated workflows assess the completeness of data and metadata. The remote instruments go through a detailed semi-automated quality analysis before users can access the data. If a data quality issue is identified, a data quality report (DQR) gets created by the reviewer. [Section 5.5]
- ARM has tiered storage and processing requirements for its data operations. The ADC uses a hybrid cyber infrastructure that includes a private cloud infrastructure operated within ORNL CADES for data flow operations and large-scale data analysis, as well as commercial cloud for processes and platform-as-a-service systems that do not require customized data resources. ARM leverages the ORNL Leadership Computing Facility (OLCF) to maintain and operate computing clusters for large-scale data processing, data analysis, and high-resolution modeling. [Section 5.5]
- The ARM ADC will continue to explore and adopt the use of commercial cloud services for nonroutine processes, operational support tools, platform-as-a-service applications, and cloud-optimized data services for cloud-based user-developed processes. [Section 5.5]
- CMIP data volumes have grown, from ~1 GB in the early phases (~1989), to over 23 PB of replicated model data contributed by 49 institutions across 26 countries. Routinely, this project is transferring 100 TB of data between nodes weekly, with download and upload occurring concurrently. [Section 5.6]
- The LLNL-ESGF node is depended upon as the primary CMIP6 node, with ~10 PB of storage allocated. It is the only ESGF node with CMIP3 data available, whereas a number of ESGF nodes continue to provide data through their portals. For this reason, it is considered the primary CMIP data source, and has been depended upon by the other tier 1 nodes to acquire data. [Section 5.6]
- The "modern" data analysis stack relies on Python-based open-source software, using JupyterHub and DASK gateways. In some instances, data is loaded into a local object store (S3). These services run on a Kubernetes cluster. Data loaded into object store is reformatted to an analysis-ready cloud-optimized (ARCO) format, e.g., Zarr. Data conversion to ARCO may require the download of remote data (10 GB to 2 TB) prior to conversion. [Section 5.6]
- For storage resources, the ESGF has obtained object storage grants, each from Google Cloud Storage (GCS) and AWS (Amazon), to host CMIP data in ARCO format (Zarr). The data conversion process would run on these cloud resources, but data would be transferred (to date typically HTTP download). [Section 5.6]
- Another difficulty in maintaining ESGF as a federation has been responsiveness and ease of operations with sites to upgrade their platforms. The hope is that containers will at least address an issue of version migration when there is no experience with a particular base version and limitations on what the development team can test for quality assurance purposes. [Section 5.6]
- ESS-DIVE data is hosted on resources at NERSC (under the Spin infrastructure and hosted on the community filesystem). Copies of the data and metadata are replicated to the DataONE federation. ESS-DIVE also has failover nodes on LBNL IT hardware (Science Virtual Machine [SVM] cluster) and at NCEAS (UC Santa Barbara). It is important to maintain high-bandwidth connections across these three facilities for data and service replication to support failover needs. [Section 5.7]

- For Globus Uploads, an external virtual machine (VM) on LBNL IT"s SVM infrastructure is used (advanced Globus features at NERSC are limited to users that are part of the NERSC allocation). Data is then transferred to NERSC. [Section 5.7]
- ESS-DIVE uses Dockerized microservices for portability (running on NERSC Spin Kubernetes). This has the potential to introduce network challenges, since all traffic goes through the container network layer. [Section 5.7]
- ESS-DIVE continually evaluates resource options such as cloud computing. The cloud computing model enables us to leverage a wide range of software, storage, and computing hardware resources under the "-as-a-service" model, where ESS-DIVE simply deploys and makes use of services in the cloud, rather than run something private. Both dev-ops and resilience strategies actively leverage resources on public clouds when those services meet the anticipated needs. [Section 5.7]
- The standard SBC data handling process is to collect data from sensors to local storage, while also copying this data simultaneously and providing it to the user's home institution or portable data storage. At the same time, the data is processed, usually several times on-site, and then after the experiment in the user's home institution. The design objective of the infrastructure is to give users flexibility in use of whatever protein crystallography data processing software they require. [Section 5.8]
- CABBI features centralized networking, data storage, and compute facilities at the University of Illinois are the main source of computer resources. These resources have connections to ESnet via the National Center for Supercomputing Applications (NCSA) on the Illinois campus. [Section 5.9]
- CABBI has access to Amazon, Google, and Microsoft cloud computing resources through the University of Illinois. To date, only AWS has been used to host a web tool developed by CABBI researchers. Additionally, cloud computation through AWS was explored in the past when a Windows operating system was required for specific software applications. It was found to be cost-prohibitive at that time, and local resources were used. For other computational needs, the available local resources are sufficient and cheaper to use. [Section 5.9]
- Cloud-based commercial file sharing and file management are used for files 150 GB or smaller. Files are shared both within CABBI with restrictive permissions, and publicly via a link from the research data website. [Section 5.9]
- The default max storage configuration for the ORNL CADES cloud virtual machines, where a lot of development is done, comes with 80 GB of space. This is insufficient for a single day's worth of data from the APPL facility. Storage requirements can be mitigated by mounting various devices to the CADES cloud, but there is a bottleneck with both the read and write speed of the devices themselves. This has limited some data transfer activities within the ORNL environment. [Section 5.10]
- GLBRC is a multi-institutional center with labs generating data from a variety of instruments. The output of the equipment varies in type and size, from spreadsheets (KB in size) to hundreds of gigabytes of sequencing data. It is important that this data is shared across teams and stored in a central repository, requiring high-speed network connectivity for the larger data types. [Section 5.11]
- The GLBRC uses a custom-built data catalog to store data and metadata along with custom-built applications that assist with collection and transfer of data from external data providers. [Section 5.11]

- Currently, the GLBRC uses cloud storage options for some smaller facility data transfers (Google Drive, OneDrive) but all compute and storage are located on campus. Over the next several years, GLBRC will be looking into using public cloud resources for applications (AWS, GCP) and to host the upcoming cross-BRC data-sharing portal, bioenergy.org. [Section 5.11]
- While transferring large amounts of data at high speed is critical to sharing and using compute resources outside of local resources, the storage of data is often an overlooked issue. GLBRC has sufficient storage space for the petabytes of research data that are expected, and has built this technology solution over a number of years of operation. A central data storage option provided by the DOE could be very beneficial. [Section 5.11]
- LBNL is a Google Workspace customer and has substantial discounts on services running in the Google Cloud Platform (GCP). JBEI is currently working on the instrumentation required to [Section 5.12]:
  - Archive data to GCS.
  - Provide access to Google's Compute Engine for projects that exceed the capabilities of local resources.
  - Deploy webapps to Google's App Engine.
- JBEI currently has 14 hardware systems dedicated to computational biology that are more than five years old. These will soon need to be replaced. Given the increasing focus on ML, JBEI should replace these with GPU-enabled systems. [Section 5.12]
- The next project phase will focus on supporting model development and conducting simulation campaigns on GPU-based machines as well as maintaining its excellent performance on traditional central processing unit (CPU) architectures. In addition to the existing pre-exascale GPU systems installed at the leadership computing facilities (LCFs) and at NERSC, E3SM is targeting two exascale systems: OLCF's Frontier and ALCF's Aurora. [Section 5.13]
- Major simulations are conducted on supported production-level machines (Chrysalis and Anvil at ANL, Cori-Haswell, Cori-KNL[Knight's Landing] and Perlmutter at NERSC, Summit at OLCF, Theta at ALCF, etc.). [Section 5.13]
- ACME1 has 10 PB designated to archive E3SM [Section 5.13]:
  - As of summer 2022, on ACME1, the E3SM volume reached more than 4 PB, comprising 2.9 PB of formal archive, and total ESGF publication volume of another 1.2 PB.
  - Estimated combined archive and publication growth is another 2.5–3.0 PB per year, having gained 1 PB in just the last five months.
- Every E3SM production simulation is required to be archived uniformly using Zstash, moved to NERSC HPSS for centralized deep storage, and then sent to the ACME1 server maintained at LLNL for data publication [Section 5.13]:
  - The massive data movements across platforms are done through Globus DTNs via ESnet.
  - The E3SM science campaign simulation results are being processed and published through the BER-supported ESGF, where all the datasets are cataloged and publicly discoverable through the ESGF search portal.
- The AmeriFlux data is stored in Berkeley Lab servers and in NERSC's CFS. An additional copy is stored in the high-performance storage system (HPSS) tape data storage, using metadata files that incorporate effective versioning, file naming, and backup strategies. [Section 5.14]

- In the past five years, the project has generated ~25 TB of data (primarily high frequency). In the next five years, AmeriFlux expects the project to generate on the order of 50–100 TB of data given the current workflow and project needs. [Section 5.14]
- Due to cost constraints, the use of cloud resources is somewhat limited. AmeriFlux has test instances on AWS but sees limited usage at this point as the costs are prohibitive for an instance that runs 24/7 instead of owning machines directly. This does impact redundancy to a certain extent. [Section 5.13]
- NGEE–Tropics uses Google Earth Engine for some aspects of the workflow, which results in MB to TB scales of transfer using web tools for transfer (e.g., Google Drive, etc.). [Section 5.15]
- For remote-sensing (RS) data use cases in NGEE (satellites or flown in aircraft/UAVs), heavy processing needs to happen close to where the data was collected [Section 5.15]:
  - Later analyses are conducted using GIS-type tools or numeric array data analysis tools.
  - RS data is large in nature, now easily going into PB and EB ranges, but almost always only parts of the data are needed for studies.
- NGEE simulation data is generated by running models like the Functionally Assembled Terrestrial Ecosystem Simulator (FATES) and E3SM for experiments that are more general (e.g., long-term climate predictions) or specialized (e.g., parameter sensibility analyses) [Section 5.15]:
  - Data generated by many of these runs becomes valuable for further analysis.
  - Some of these datasets can be deposited with data archives such as the project archive, but many are too large to be moved in their entirety, and must be accessed parts at a time.
- Local storage capacity is a recurring constraint when it comes to analyzing remote sensing or simulation output data for NGEE. Data corruption caused by unreliable connections is also a common issue for transferring data from inhospitable locations (not necessarily in a sense of "checksum corruption," but more of interrupted transfers that cannot be resumed later when connectivity is restored). [Section 5.15]
- COMPASS-GLM has access to HPC resources at facilities at the National Energy Research Scientific Center (NERSC) for modeling. All the data products are stored on and can be accessible from NERSC's HPSS. Future simulations will be performed on COMPASS-HPC, a midrange high-performance computing facility at PNNL. Data transfer and sharing between researchers across scales use Globus, which allows efficient, secure, and reliable data transfer across all supercomputers at DOE HPC facilities as well as personal computers. [Section 5.16]
- The COMPASS compute resource will have an on-premises Kubernetes deployment to facilitate containerized software deployment. It also hosts an Open OnDemand web UI to provide Jupyter hosting and other software-as-a-service products. Other DOE HPC user facilities may be engaged depending on the needs of researchers and the scale of their work. [Section 5.16]
- There is a risk that some COMPASS datasets may exceed the current per-dataset capacity at commonly used affiliated data archives, such as ESGF and ESS-DIVE [Section 5.16]:
  - In these cases, the archives are used to index the data, for which the researchers must still locate and secure long-term storage.
  - Similar constraints may appear for HPSS as large amounts of data from the regional climate modeling for long-term historic simulations and future projections are generated.

- There is also an extremely high demand for COMPASS computational resources for regional models for long-term simulations and projections. One of the reasons is that the input/output (I/O) reading and writing take a good portion (20 to 30%) of the entire calculation time. This cost can be significantly reduced by splitting the model output on each processor. However, it also takes wall clock time (instead of computing allocations) to stitch these split files back to a file that can be processed and analyzed for science problems. [Section 5.16]
- CASCADE creates and collects climate model simulations of the past, present, and future climate. Analysis of these simulations is typically focused on understanding how the statistics and properties of extremes could change in a warmer climate. Supercomputers at NERSC, OLCF, ALCF, IU (Big Red 200), and PNNL (Compy) are used to produce climate model simulations and analyze climate model outputs. Storage allocations at these facilities vary, but are typically terabytes to petabytes in size. CMIP7 will see a four to five times increase in data volumes, when related. [Section 5.17]
- DOE LCFs prioritize large-scale computation over moderate-scale analysis. Thus, the CASCADE team is increasingly using medium-scale computational resources external to the LCFs when doing moderate-scale data analysis tasks [Section 5.17]:
  - This results in repeated data transfers (e.g., to/from NERSC, or from external data repositories to external HPC systems like IU UITS described previously).
  - These transfers need to be done manually, and they often take O(1) days even when using transfer tools, like Globus, that nearly achieve 1 GBPs transfer speeds.
  - Given that large-scale computations—like simulations—will increasingly be done on LCFs, with the data stored/archived at those LCFs, this will likely continue to be a bottleneck.
- HPC facilities involved in the model creation have petabytes of storage available for the input to the model generation process, as well as to store output for later sharing [Section 5.18]:
  - Simulation output generated on Cheyenne can range in size from tens of megabytes to many terabytes, depending on the specific configuration of the model.
  - Large-scale simulations can generate data collections consisting of 100,000s of files and 1,000–2,000 PB of output.
  - Output is shared with well-established community portals (e.g., Climate Data Guide [CDG] and ESGF).
- CATALYST does acknowledge a perpetual "shortage" of disk space. This can be attributed to model output increases, but also needing a better strategy to manage available disk resources across project partners. [Section 5.18]
- Across the BER case studies, Globus use remains high, although not all facilities are users either due to security policy, licensing issues, or observations that performance is not measurably better. With respect to the latter, mounting storage remains a challenge since it requires lots of tuning to see high performance between facilities. [Focus Groups]
- Improvements to GPU hardware have made it such that data compression is much better, which has simplified processing (and network transfer) for some model and simulation activities. These new hardware packages will be available on new DOE-SC systems, and will revolutionize some of the ways that people interact with data. [Focus Groups]
- Computational needs are sufficient for most participants. Interactive jobs (e.g., those that are operated via Jupyter) are easily run on cloud-like computing platforms, and meet the needs of most users. Those that require access to midrange clusters, or HPC, can use institutional resources or allocations at the major HPC centers to accommodate their computing needs. [Focus Groups]

2.5 Remote Collaboration and Operational Requirements

- JGI's JAMO operates a data portal providing programmatic access to JAMO data, both a spinning disk "working set" as well as HPSS archives. Recent trends have been on the order of 1 to 2 PB per year and are accessed via the portal. [Section 5.1]
- JAWS supports compute clusters in AWS that scale out on-demand for specific workflows that require the reliability of AWS services. In the future, the JGI may move some of the core infrastructure into a cloud provider to maximize availability of services such as JAWS Central. [Section 5.1]
- KBase is converting Rancher Compose- and Docker Compose-based services to be deployed in Kubernetes, and in a reproducible configurable way that allows for the deployment on any Kubernetes environment. Once completed, this will enable more flexibility in deployment location across the DOE and commercial cloud world. [Section 5.2]
- As one of the distributed data sources for the NMDC, EMSL also routinely shares processed data and metadata (aggregate gigabytes), making that data findable and accessible through the NMDC data portal. [Section 5.4]
- EMSL does not currently use, or plan to use, academic or commercial cloud services for their primary data analysis or storage needs. However, EMSL users may use cloud resources if, or when, they choose to perform subsequent data analysis on the data generated in EMSL. [Section 5.4]
- Connectivity at the various ARM measurement facilities varies. Bandwidth and technologies will differ depending on deployment location and availability of connectivity [Section 5.5]:
  - Terrestrial connectivity provided by telecommunications companies (e.g., fiber, copper, digital subscriber line (DSL) service).
  - Cellular connectivity.
  - Point-to-point microwave radios.
  - Low Earth orbit (LEO) satellites.
- Data from measurement facilities is encrypted via a virtual private network (VPN) tunnel. Traffic between ANL and ORNL is unencrypted over ESnet. Data is served to users from here via file transfer protocol (FTP) and Globus. [Section 5.5]
- ESGF anticipates deployment of comparable replica nodes at ANL and ORNL by 2023. By that time, ESGF under DOE will operate with triple redundancy for CMIP data. Future data management remains uncertain. Most likely, the total unique data volume will exceed capacity for three labs. [Section 5.6]
- 25 to 30 collaborators provide data to CMIP by hosting data on an ESGF node at their site. Data is replicated to LLNL on an ongoing basis. Seven or eight of the sites have successfully deployed Globus as an option for ESGF data transfers. Most of the sites rely solely on HTTP services to make their data collections available. [Section 5.6]
- ESS-DIVE is being used to support a research use case supporting the development of Self-Driving Field Laboratories (SDFL). [Section 5.7]
- The DataONE network provides redundancy via replication to far-away network nodes and automatic self-healing capabilities. DataONE replication enables users to discover and download datasets if the main repository site is unavailable. [Section 5.7]
- ESS-DIVE acknowledges that interoperability across archives is a broader issue, involving

coordination across competing requirements, standards, and interfaces, which can only be sufficiently addressed by multi-institutional effort. [Section 5.7]

- CABBI has automated sensors for collecting environmental data at field trial locations in Illinois, Iowa, and Florida. The instrumentation includes [Section 5.9]:
  - Meteorological sensors (wind, temperature, precipitation, radiation).
  - Soil sensors (soil temperature and moisture).
  - Gas sensors (concentration, flux).
- CABBI field sites sometimes have permanent infrastructure (e.g., power and networking) to support automated collection of data and operation of instruments. Other times it may be necessary to manually retrieve data using removable media [Section 5.9]:
  - Field data is collected continuously all year. Each file holds 30 minutes of data collection, and is approximately 5 MB in size. At the end of the year, field data is compressed to 70 MB in size. The meteorological and soil data each year generates approximately 30 MB of data.
  - Remote sites can make use of commercial networking, which is expensive and less reliable than the facilities provided at a university.
  - Remote access is not feasible for many of the field measurements because they are done at the plant level. It would be expensive and unfeasible to connect sensors to specific plants that will need to be removed from the field at harvest.
- CBI has remote devices include field stations, edge instrumentation, and space-based assets. Not all of the partner institutions have high-speed networking capabilities, and this can be problematic regarding data transfer. Often the fastest transfer modality is shipping hard drives on FedEx. [Section 5.10]
- CBI is in the process of evaluating a Kubernetes environment hosted by OLCF. It is expected that this environment will become the main hosting environment for both data and compute. [Section 5.10]
- GLBRC routinely shares 100 GB datasets with a number of collaborators: JGI, EMSL, other BER-funded BRCs and national laboratories, and collaborators at the University of Wisconsin and Michigan State University. [Section 5.11]
- GLBRC collaboration with DOE light sources facilities requires low-latency, high-bandwidth network connections [Section 5.11]:
  - Diffraction experiments that generate crystallographic data are controlled and monitored in real time using secure virtual network computing (VNC) connections.
  - Existing network architecture is strained by the pace of data collection.
  - Datasets consisting of numerous 10 GB images are common, and many datasets are required for each solved structure. Acquisition can take minutes, and after planned upgrades dataset sizes will increase and acquisition times will decrease by orders of magnitude.
  - Even with rapid on-site data reduction, this pace of data collection will put new strain on network infrastructure and data storage for collaborations.
- JBEI will have access to scientific infrastructure and instrumentation at LBNL and partner institutions that are well suited to address the most challenging tasks [Section 5.12]:

  - At LBNL, these include dedicated biotechnology laboratories, specialized facilities for HT structural biology and genomics analysis, high-resolution imaging systems, extensive plant biology and greenhouse research areas, and a powerful supercomputing environment.
  - UC Berkeley (UCB) facilities include a biomolecular nanotechnology facility and nuclear magnetic resonance (NMR) facility.
  - JBEI has access to combustion research, supercomputing, and microsystems facilities at Sandia, UC Davis (UCD) feedstock growth facilities, and fungal research, proteomics, and bioinformatics facilities at PNNL.
  - In addition, JBEI researchers make use of national research facilities, often operated by DOE national laboratories.
- In addition to the facilities in JBEI's central location, approximately 10% of JBEI's researchers are housed in off-site locations at DOE facilities and laboratories, universities, and private industrial partners. [Section 5.12]
- The E3SM project is a collaborative effort across eight national laboratories (including LLNL, ANL, LANL, PNNL, SNL, ORNL, LBNL, and BNL) and more than 10 academic institutions. Data transfer is essential to share data with collaborators and the broader community. [Section 5.13]
- Members have access to E3SM data from at least one E3SM-supported facility. E3SM projects also make production runs available publicly through the DOE's ESGF for external projects. ESGF publication currently goes through the LLNL-ESGF node, with plans to also use the ANL ESGF node. [Section 5.13]
- Via the LBNL IT department, the E3SM project has set up a GCP cluster that is being used for running the E3SM nightly testing suite. [Section 5.13]
- AmeriFlux Core Sites are flux towers whose managing PIs and staff have agreed to deliver timely, high-quality, continuous data to the AmeriFlux database. These 49 sites represent nearly every major bioregion and vegetation type in the contiguous USA, span 15 states and multiple climate regions, and comprise contrasts of demography and disturbance. [Section 5.14]
- AmeriFlux is considering allowing a subset of sites to transfer (near real-time) raw high-frequency data to AMP. There is demand among data users for this transfer, as well as taking advantage of 5G capabilities for near real-time data transfer. AMP would process the high-frequency data to half-hourly data products using standardized protocols and codes. This will drastically increase the network demands, e.g., 1.5 to 5 GB per month per site, and also the computation demand, e.g., ~0.5 to 3 hours processing time for one site-month of data. [Section 5.14]
- AmeriFlux sites are at a mix of locations, some more urban with ample infrastructure and others in more remote areas where power is derived from solar panels and batteries. Similarly, networking ranges from those with wired connections to those where available options are cellular/satellite connections. Ideally, all data can be sent from the field sites to teams' institutions, where various data checks and processing occur. Practically, measurements are often recorded to a data logger or computer on-site. Then the site teams retrieve the data via physical visits or remote connection (e.g., cellular modem, radio transfer, Ethernet). Many site teams manage multiple field sites and, in rare cases, up to 40 to 50 field sites, e.g., the National Ecology Observatory Network (NEON). Usually, data transfer from field sites to designated institutions poses the most significant challenge in the data's life cycle. [Section 5.14]

- NGEE Arctic is led by ORNL with partnerships at LBNL, LANL, Brookhaven National Laboratory (BNL), and the University of Alaska Fairbanks, as well as multiple Alaskan Native communities. [Section 5.15]
- NGEE Arctic also collaborates with [Section 5.15]:
  - National Aeronautics and Space Administration (NASA) Arctic and Boreal Vulnerability Experiment (ABoVE) (a memorandum of understanding is in place between NASA ABoVE and NGEE Arctic).
  - Permafrost Carbon Network.
  - Tundra Trait Team.
  - Alfred Wegener Institute (AWI).
  - Study of Environmental Arctic Change (SEARCH).
  - International Land Model Benchmarking (ILAMB) project.
  - E3SM.
  - ARM.
  - Barrow Arctic Science Consortium.
  - Ukpeagvik Iñupiat Corporation (UIC) Science, and Sitnasuak, Mary's Igloo, Council, and Bering Strait Native corporations.
  - UAF Northwest Campus.
- NGEE–Tropics researchers collaborate with both field (e.g., research groups from Manaus, Brazil) and modeling groups (e.g., E3SM) on tropical forests research questions. [Section 5.15]
- The NGEE Arctic data ecosystem leverages several partnerships [Section 5.15]:
  - NGEE Arctic leverages the high-performance computing cluster and development area using the Compute and Data Environment for Science (CADES) infrastructure at the Oak Ridge Leadership Computing Facility (OLCF). This infrastructure hosts an E3SM land model container and tools to allow simulations in a common computing environment with datasets needed to run site- to regional-scale simulations.
  - NGEE Arctic uses ESS-DIVE for long-term data archival.
  - The NGEE Arctic project data archive at ORNL—an interim storage location before permanent data archival at ESS-DIVE—currently has 233 datasets with 161 available to the public, 62 available to the team only, and 10 in the planned (forthcoming) stage. The project archive currently has 3.5 TB of data.
  - The EMSL and JGI user facilities have been instrumental in improving the understanding of subsurface microbial communities across the Arctic tundra.
- Much of the NGEE project model development and use occurs at large computing centers such as OLCF and NERSC [Section 5.15]:
  - Model analysis is either on workstations or via Jupyter Notebooks running on large computing resources.
  - Datasets are stored in public repositories, and backed-up into tape regularly and automatically.
  - All publicly available datasets will be synchronized with ESS-DIVE.

- Federated authentication mechanisms among partner institutions would greatly ease data transfer and other collaboration efforts. This has been a particular problem in bringing the COMPASS-HPC computer online, as the registration and approval process was quite slow. [Section 5.16]
- Many BER user facilities are already designed around a remote model [Section 5.19]:
  - Users send physical samples to facilities for analyses and access results over data portals.
  - Multifacility models will expand on these by requiring coordination and information flow between facilities.
- For some collaborations with a strong international presence (e.g., ESGF), high performance remains elusive. At some ESGF nodes located in the Asia-Pacific region, the network performance can be measured in Kbps values, which makes downloading large data sets challenging. There is a rich international ecosystem of connectivity (e.g., multiple 100 Gbps) across the Pacific, but local connectivity within a facility (or security policy) may be limiting performance. ESnet will continue to work with communities to characterize and understand these observations. [Focus Groups]

2.6 Multifacility Computational Workflows

- Most of the JAWS workflows will operate within LBNL, split between NERSC and a JGI-dedicated cluster hosted within LBNL IT, with only a small fraction going off site to PNNL and AWS. [Section 5.1]
- The JGI manages petabytes of data on site, and regularly shares this data with EMSL and HudsonAlpha. Other entities that regularly share data with the JGI include KBase and the NMDC. [Section 5.1]
- The JGI does not foresee any need to expand into leadership HPC facilities: the workflows really require only midrange HPC resources. If workflows users begin to use more artificial intelligence/machine learning (AI/ML) tools, the JGI may eventually require access to clusters equipped with graphical processing units (GPUs). [Section 5.1]
- Expanding the KBase footprint at LBNL has been difficult due to facility constraints. As a result, KBase is migrating some deployments from LBNL to ANL, and will focus on using the LBNL resources to act as a mirror of the production site in the future. [Section 5.2]
- The NMDC relies on existing core partnerships and on building out new partnerships [Section 5.3]:
  - The NMDC provides access to JGI data sources. This partnership supports gigabytes of data exchanged between the platforms.
  - Through partnership with the NMDC, multiomics data generated in EMSL for microbiome research is findable and accessible through the NMDC Data Portal. This partnership supports gigabytes of data exchanged between the platforms.
  - In the next two years, the NMDC will leverage a number of partnerships, with each resulting in gigabytes of data exchanged between the respective platforms:
    - KBase, to enable integration of data across these platforms and perform advanced analyses.
    - ESS-DIVE, to facilitate integration of data generated by the research community and archived in the ESS-DIVE storage platform.
    - NEON, to enable integration of data generated from NEON field sites.

- The NMDC has established mechanisms for the JGI and EMSL to share metadata with the NMDC core metadata storage system about microbiome studies and samples [Section 5.3]:
  - Within the next two years, a core goal for the NMDC is to automate data exchange between these facilities, to handle the steady streams of data from the JGI and EMSL, as well as external agency partners.
  - In two to five years, the NMDC will be managing PB of data with variable access patterns; e.g., some users will download, others will access data through partner platforms for analysis.
- EMSL engages in multiple institutional collaborations, most notably with the FICUS program for collaborative access to DOE user facilities [Section 5.4]:
  - The FICUS program has supported joint user projects between EMSL and the JGI since 2014.
  - FICUS has expanded to include the National Ecological Observatory Network, the Bio-SANS beamline at the Center for Structural Molecular Biology, the Advanced Photon Source (APS), the Advanced Light Source (ALS), and the ARM user facility.
- Research universities participate in a climate modeling activity and may each produce around 100 GB to 10 TB of data. For this case, it is deemed impractical for those to run ESGF data nodes, and thus they will transfer the data to ESGF host centers prior to publication. [Section 5.6]
- Institutions that produce larger quantities of datasets (via modeling and/or observation gathering) e.g., 50+ TBs, should stand up an ESGF node and perform local publishing on site rather than transferring the data to another node to be published. This activity is presently done for CMIP6 data, and so the expectation is that it will continue for ongoing MIP phases. Some modeling centers have relationships with a node site and procedures to move data if the HPC running the model is not co-located with the data node. [Section 5.6]
- ESS-DIVE partners across the ESS community to develop data and metadata standards for the repository: these include ESS projects (e.g., SFAs and NGEEs); the broader digital library and standards communities (e.g., RDA, ESIP, DataONE, OGC, and DataCite); BER facilities and data systems (e.g., the NMDC, the ESGF, KBase, the JGI, the EMSL, and ARM); and other federal agencies (e.g., USGS, NASA, and NSF). ESS-DIVE will work to integrate with external data repositories (ESGF, ARM, EMSL, NMDC, and USGS) in the future. [Section 5.7]
- ESS-DIVE serves a broad range of national lab and university partners that contribute data to the system. This includes data coming from projects with a footprint at major DOE HPC facilities such as NERSC, ORNL, ANL, etc. (e.g. for large-model datasets). Additionally, ESS-DIVE serves partners from the BER facilities like the JGI and EMSL for environmental sample data. A new area of emphasis is centered around sensor networks in the field that are collecting large volumes of data. Dataset sizes range from megabytes to multiple terabytes. [Section 5.7]
- ESS-DIVE maintains two standby copies of ESS-DIVE: one copy in the LBNL Information Technology Division (LBNL IT) data center and the other copy at NCEAS in Santa Barbara. Each failover node can run a copy of the stack in either read-only or read/write modes. Either node can take over as primary in the event of a failover. [Section 5.7]
- ESS-DIVE sees the importance of streaming data, increased use of cloud resources, and federated data to facilitate these use cases. Additionally, integration and interoperability with other repositories will also be important. To enable a future-facing data repository, ESS-DIVE must create pathways to bridge the scientific data and metadata stored in other partner

repositories with the data in ESS-DIVE. Given that the data ecosystem is fundamentally a distributed enterprise, scientific discovery ultimately rests on being able to deliver insights that can fuse information from different data silos. [Section 5.7]

- As the rate of genomic data increases, increased network connectivity from CABBI to facilities like the JGI and EMSL will be necessary to continue these partnerships. [Section 5.9]
- CBI uses the Advanced Plant Phenotyping Laboratory (APPL) facility at ORNL, the PY-MBMS instrumentation at NREL, and a custom-designed facility for low-cost fermentation of biomass with co-treatment located at Dartmouth College. In addition to these specialized facilities, CBI uses standard molecular biology and field experimentation techniques at its various locations. CBI currently maintains active field sites located in Davis, California; Corvallis, Oregon; Clatskanie, Oregon; Knoxville, Tennessee; and Athens, Georgia that are equipped with field-based instrumentation. [Section 5.10]
- An Azure instance is being evaluated for CBI. The purpose of the evaluation is to ascertain whether Azure can be used in a hybrid on-prem/ off-prem environment to access the data. The main use case for the Azure instance will be to surface virtual machines that are staged with specific Microsoft Windows-only applications that need to run against the main data store. [Section 5.10]
- The two primary institutions within the GLBRC, the University of Wisconsin-Madison and Michigan State University (MSU), have multi-hundred gigabit backbone connections from campus and high-speed intra-campus connections that range from 1 to 100 Gb between buildings. [Section 5.11]
  - The primary data center for GLBRC, located within the Wisconsin Energy Institute, has redundant 40 Gb connections to the campus backbone and 10 Gb to all equipment internally.
  - High-performance data transfer is available between the primary institutions using Globus endpoints.
  - New network architectures with improved transfer speeds (Science DMZ) are under investigation by the University of Wisconsin's campus networking team for potential use in the next five years of research.
- JBEI routinely shares datasets that are .5 TB in size with other DOE facilities: the JGI, EMSL, ALS, CBI, CABBI, GLBRC, and the NMDC. Data transfer to some facilities could be improved, and is attributed to the tools used, as well as some of the networking abilities at remote sites. [Section 5.12]
- The following facilities play a more important role in model development and in conducting production simulations and evaluation [Section 5.13]:
  - Chrysilas and Anvil at ANL owned by E3SM.
  - Compy at PNNL (Shared by E3SM/RGMA/ESMD).
  - NERSC (Cori and Perlmutter).
  - ACME1: Archival system at LLNL with ESGF connection.
  - ALFC: Theta, Cooley, Polaris, Aurora (expected in 2023).
  - OLCF: CADES, Summit, Frontier, Andes.
  - ANL: the Joint Laboratory for System Evaluation (JLSE).
  - LANL: Darwin, Badger, Chicoma/Rome, Chicoma/GPU.

- LBNL: Lawrencium clusters, GCP via LBNL.
- LLNL: Livermore Computing HPC systems.
- PNNL: Research Computing.
- Sandia National Laboratories (SNL): Center for Computing Research and HPC systems.

- COMPASS-FME uses data across existing repositories: ARM, ESGF, ESS-DIVE, MSD-LIVE, NASA DAAC, USGS, NOAA, Great Lakes Observing System, NERR, NEON, AmeriFlux, LTERs, MarineGEO, the Smithsonian's Coastal Carbon Research Coordination Network (CCRCN) for coastal wetland carbon stocks, NOAA's National Estuarine Research Reserve System-Wide Monitoring Program Database, and others. [Section 5.16]
- Server-side analysis is desirable for CASCADE computation, but moderate-scale analyses (e.g., using O(10)-O(100) nodes) have historically had very long queue waiting times at NERSC, when compared with small jobs (O(1) node) or large jobs (O(1,000) nodes) with similar wall clock time requirements. These wait times can be upwards of several days, which is less than the network transfer time of the datasets. [Section 5.17]
- Workflows that require coupling capabilities across facilities are becoming more common within BER [Section 5.19]:
  - The FICUS effort is one notable example, but it is expected that the frequency and depth of this type of multi-analytic couplings will increase over time.
  - Several workflows that depend on data replication across sites to ensure that data is available are at the HPC facilities. This involves coordination of data movement and job management. The JGI and NMDC are using tools like JAWS to handle data movement across sites to facilitate these needs for data replication workflows.
- It is anticipated that multifacility efforts will be expanded to include DOE BES Light-source facilities and other non-BER user facilities [Section 5.19]:
  - In addition, there are already many examples of BER facilities coupling with computational capabilities provided by ASCR. This connectivity should only increase over the coming years driven by demands from the scientific community.
  - Efforts are underway to make information flow across the facilities more seamless, such as those to standardize sample metadata and methods of exchange. It is anticipated that BER will place even more emphasis on this in the coming years.
  - It is envisioned that an ecosystem will exist, beyond the five-year time frame, where common interfaces are defined across facilities that will enable scientists to more easily take advantage of the instruments and computational capabilities across the DOE-SC complex in a seamless manner.
- Multifacility BER workflows can take several forms [Section 5.19]:
  - In one model, a researcher has a common physical sample that must be sub-sampled, and sent for measurements at different facilities using different capabilities (e.g., sequencing, MS, imaging). These results are then synthesized by the researcher to gain a better understanding of structure, function, dynamics, etc.
  - Similarly, a researcher might use the capabilities at multiple facilities to understand mechanistic behavior across space and time, at a larger study level.
  - In another model, staff at the facilities interact directly to improve the process of science for the community, for example by the transfer of data between facilities (e.g., provide

genomics data to use in the analysis of proteomics data), or by providing surge capacity with computing resources. In each of these cases, computational and data capabilities are critical.

- A core problem for addressing multifacility use cases remains a lack of a cohesive identify, authorization, and authentication strategy. It is still the case that each facility (and lab) has different ways of managing this, which complicates attempting to build multifacility workflows. [Focus Groups]
- The use of multifacility workflows will remain high in the future. Coupling computing to an instrument facility, or combining multimodal data, is still of high research interest. Tools like Jupyter have simplified the language of workflows, and remain popular. Globus and other data movement platforms are widely adopted for mobility of research information. Efforts such as IRI must acknowledge this, and develop intelligent APIs to join resources that span facility boundaries. [Focus Groups]
- When a component of a coupled facility is unavailable (or has limited resources), the entire workflow suffers. As an example, if a system that provides computation is down for maintenance, or has a long wait time, it may be impossible to migrate the work to another resource unless the steps to make the workflow portable were taken. If it were easier to move to other HPC facilities, or on temporal resources that may be located on ESnet, that could offer a layer of protection for some time-dependent use cases. [Focus Groups]
- Some facilities want to upgrade their software and workflow mechanisms, but lack development resources to do so. There is a desire to have more portability and also address long-standing issues with storage performance and interactive feel for real-time applications. All of this is compounded by the fact that analysis routines are regularly designed to consume more data (e.g., volume and quantity), and both the hardware and software must scale to accommodate this. Some workflows must be converted to rely on data streams (versus bulk data movement) to better handle the larger data volumes in these analysis workflows. [Focus Groups]

2.7 Domestic Networking for Local and Wide-Area Data Mobility

- Neither JGI's JAMO nor JAWS intrinsically has special networking requirements for data acquisition or distribution. Both JAWS and JAMO make use of replicating data from the central facility to the remote sites using tools such as Globus, as well as rsync. [Section 5.1]
- The JGI does not foresee any data-related resource constraints that would be addressed by ESnet. JAMO data grows at a relatively high rate (currently 1.5 PB per year), but most data remain within the facility, with the exception of the aforementioned sharing and backup patterns. [Section 5.1]
- Transferring multiomics data across the wide-area network can take several days using traditional web protocols. The NMDC will therefore enable high-bandwidth data transfer mechanisms via Globus, a secured high-performance data transfer service allowing users to transfer large amounts of data between systems within and across organizations. [Section 5.3]
- EMSL and the JGI actively collaborate, sharing resources for data storage and providing resiliency. This collaboration currently requires routine transfer of up to terabytes of data for analysis, and episodic transfer of up to a PB of data for long-term archival. [Section 5.4]
- EMSL operates four data transfer nodes (DTNs) for Globus data transfers. The DTNs are located on a Science DMZ per ESnet best practices and provide access to data on EMSL's computing systems. The DTNs currently each provide 100 Gbit per second bandwidth and collectively transfer 250–350 TB per month. With growth in the FICUS program and additional collaborations with DOE facilities (described previously), EMSL expects data transfers to increase over the next five years to 500+ TB per month. EMSL plans an upgrade to 400 Gbit per second technology for the DTNs when ESnet6 becomes available. [Section 5.4]
- Anticipated network or data constraints are currently internal to EMSL and PNNL. Specifically, the Science DMZ connection at 100 Gbit per second is the fastest part of the entire EMSL network. This interface with ESnet will be a bottleneck only in the sense that the current DTN servers do not individually have fast enough buses for full 400 Gbit per second when that bandwidth is available. That said, the aggregate bandwidth of the four existing DTNs is 400 Gbit per second, and it is within EMSL's ability and discretion to upgrade these DTNs to a PCIe gen 5-based server when those enter general availability. [Section 5.4]
- The ADC's local network configuration consists of 40 Gbps core storage switches connecting back-end storage to hypervisors and processing nodes and 10 Gbps core switches for node-to-node communications. Those 10 Gbps switches also possess 40 Gbps connectivity to ORNL's backbone. [Section 5.5]
- ARM is looking to streamline communications from the measurement facilities to the ADC. This would entail migrating the ADC into the Science DMZ, then updating measurement facility routing to go to the ADC, among other tasks. This will ensure that all traffic is encrypted end to end and all data is routed as efficiently as possible, and maximize flexibility with regard to access control. [Section 5.5]
- During peak periods, the download rates from 10 reporting sites for HTTP-based downloads peak at ~1.15 PB per month total across all sites and 500 GB on "lighter usage" months. At LLNL, it has been observed that a peak rate of 300 MB per second for bursts of requests to data node servers running the HTTP-based download service. [Section 5.6]
- It is important for ESS-DIVE to support high-bandwidth, low-latency connections to users submitting data to ESS-DIVE. Data collected by ESS-DIVE is generally from other ESS projects, and is most likely generated on remote resources. ESS-DIVE is essentially the

repository where ESS projects store their data for long-term archival. This includes national labs, universities, and individual computing resources. [Section 5.7]

- ESS-DIVE expects to be able to store large datasets on ESS-DIVE (~1 TB datasets) to support model data archiving needs in the future. This will require the use of high-performance mechanisms like Globus and high-bandwidth networks. [Section 5.7]
- The IGB, where CABBI is headquartered and housed on the Illinois campus, works within the University of Illinois's infrastructure to provide network access [Section 5.9]:
  - For internal data sharing, the IGB provides CABBI 10 TB of shared storage. Storage is accessible only using SFTP clients.
  - Globus is integrated into the biotransfer system of IGB's high-performance computing resource, BioCluster.
  - A perfSONAR node will be made available for validating end-to-end performance.
  - These networking services are expected to stay the same for the next three to four years.
- CBI has some data-related resource constraints on the network to address. Performance to remote sites, such as NREL, can sometimes be below 1 Gbps. For large datasets (e.g., PB scale), this can take weeks to transfer. This level of friction for data that is already publicly available is a serious hindrance to the utility of ESnet. [Section 5.10]
- Local area network (LAN) and wide-area network (WAN) services for building 978 where JBEI is housed at LBNL are provided by LBNLnet. The LAN consists of two switches connected by a 10 G link with 10 G uplinks to the WAN. The primary WAN link is a 100 G link through ESnet. JBEI has a secondary 10 G link through the Corporation for Education Network Initiatives in California (CENIC). [Section 5.12]
- The massive data movements across platforms are done through Globus data transfer via ESnet. This software and network connection is the central piece that the project relies on. [Section 5.13]
- ESnet has adequately addressed the present need of data transfer in E3SM. A period over the last year when the Globus transfer performance from ANL to NERSC HPSS dropped significantly was noticed. The simulation team had to get a workaround to first transfer to NERSC's Community File System (CFS) then to HPSS. The problem was resolved after several weeks/months, which suggested diagnosing and fixing this type of issue currently is time consuming. [Section 5.13]
- As data is processed to a half-hour resolution at each site team's institution before submitting to AmeriFlux, the data transfer is less constrained by the network, except for some sites that choose to transfer raw high-frequency data. If future scientific needs change and require more sites to transfer high-frequency data to AMP, network infrastructure can be an issue. [Section 5.14]
- Moving data from field sites to analysis and archival locations, as well as accessing large, remote datasets, will continue to be a crucial part of projects like NGEE. New tooling to facilitate data transfers from edge locations with unreliable networking without creating additional problems like file duplication or file corruption caused at transfer time is already a pressing need. Easier access to large remote datasets, e.g., remotely accessing portions of interest from a large dataset as if it was a local file in an HPC system, or from within a Jupyter Notebook, will likely become common requirements for conducting data processing and analysis. [Section 5.15]
- Data movement uses FTP, SCP, Dropbox/GDrive, Globus.org, and interfaces between these (e.g., creating a GDrive endpoint in Globus to connect a system that cannot run GDrive to a

system that cannot run Globus). Simple tools like SCP will continue to be critical for conditions of limited compute power, like field locations equipped with only dataloggers. [Section 5.15]

- Most of the challenges for multifacility BER workflows are not at the networking layer, but the network could enable some models [Section 5.19]:
  - For example, the ability to create overlay networks across ESnet between facilities could allow distributed resources to be coupled together more securely.
  - HTTP-based communications play an important role, especially in the context of pulling datasets from existing data repositories, reference databases, etc., and bringing them on to compute systems where workflows are run.
  - Allowing for high-performance data-movement capabilities natively over standard protocols like HTTP, instead of through proprietary tools and protocols. Tools like Globus are very useful but ultimately require a fair amount of custom configuration for each site involved, and advanced capabilities require site licenses.
  - Hybrid commercial cloud and DOE HPC facility workflows, and the ability to enable a virtual network to facilitate orchestration of resources.
- Most BER case studies have demonstrated effective use of networks, particularly as capacities have increased on backbones, regionals, and campuses. Several participants noted that when moving information from institution to institution, observed speeds are high and ease of use has improved incredibly over the past several years. Several BER case studies have noted that data storage access and use remain a critical bottleneck [Focus Groups]:
  - Storage is a constant struggle to allocate, and it may not be located close to other parts of the workflow; this requires data movement several times.
  - Data-movement tools are much better, but still require a level of tuning on some institutional and laboratory resources that interact with large filesystems.
  - Not every BER case study has access to DTN resources by default (even those these are prevalent at most, if not all DOE-SC facilities), and may not know to ask how to get access.
- As real-time sensor data becomes more widely available from remote locations using tools like LEO satellites, it will be necessary to link analysis and storage to accept and process this data in a more real-time fashion. It is expected that this use case could challenge some aspects of the multifacility model not in terms of data volume, but more quantity of resources that will have to be accessed. [Focus Groups]
- A number of participants have taken advantage of ESnet's Science Engagement Team to assist with data mobility issues. These range from helping to improve international connectivity to ESGF nodes, to helping to tune DTNs, to advising on the best way to deal with remote resources that have minimal connectivity options. Participants have noted that there can be a simplified "how to ask for help" resource that is posted to the ESnet website and Fasterdata, along with continued outreach to user groups. [Focus Groups]
- There is interest in understanding some of ESnet's lessons learned using LEO satellites and 5G connectivity. Many remote sensor projects could benefit from this approach, as they build a standard model for pulling data from regions with poor connectivity. [Focus Groups]

2.8 Emerging Needs

- Collaborators who provide important cluster resources may require changes in file transport mechanisms. The JGI currently uses rsync or ssh and S3 for communications between backend cluster resources (S3 for AWS clusters). The JGI expects to expand to Google Cloud, which uses an S3 compatible protocol to manage object storage. The JGI does not have any plans to add Aspera support, and would prefer to avoid doing so to minimize the support overhead. [Section 5.1]
- KBase is challenged by a lack of staff to assist with development and operations requirements: e.g., tasks that focus on maintenance, cybersecurity patching, and user support. Future DOE support platforms that simplify infrastructure support would result in more time to focus on project-affiliated work. [Section 5.2]
- ARM is looking to replace slower Internet service provider (ISP) connections with services like LEO satellites, as the technology becomes more widely available, to ensure data is transferred to the ADC as efficiently as possible. It will be beneficial to increase throughput at all locations to ensure data is flowing optimally and provide better near real-time data flow options and effective data delivery to users. [Section 5.5]
- ESS-DIVE is considering maintaining a replicated instance of the repository on cloud resources for failover for the next phase. ESS-DIVE will be evaluating the options based on overall cost, services, and capabilities. This includes enabling cloud-based data ingest mechanisms, cloud-based data analysis tools, and cloud storage for replicating data. [Section 5.7]
- In the future, as CABBI moves towards more remote measurements in both the laboratory and the field, network security will become more critical. Ensuring outside actors are unable to take control of equipment will be necessary not only for the protection of the equipment and data, but also for public safety. Determining the appropriate security restrictions without compromising the commitment to ensuring CABBI research is usable and reproducible by others will be challenging. [Section 5.9]
- The cybersecurity landscape at some of the institutions imposes many constraints upon scientific researchers. Each institution has independent cybersecurity policies, teams, and infrastructure, with different technical and policy requirements that must be traversed [Section 5.10]:
  - At ORNL, the CADES infrastructure is logically separated from ORNL/OLCF. Hence, to move data between the two organizations, one must traverse two firewalls—each run by separate cybersecurity groups—and run through a series of DTNs and jump servers before the data can be ingested into the application database for broad access.
  - This creates layers of friction that impede seamless data transfer between indigent organizations within a single entity such as ORNL, as well as between partner institutions.
- Currently, there are no major outstanding data-related resource constraints. However, AmeriFlux anticipates that as real-time data needs come to the forefront, both networking and storage space demands will scale up significantly. There is a discussion within the community that AmeriFlux may slowly shift towards more real-time data processing and usage, where APIs will see much heavier usage in the next five years. In addition, tools like real-time data exploration with visualization will drive up both networking and data needs. [Section 5.14]
- NGEE Arctic notes that a data team must become an integral, well-funded, component of large projects, with the appropriate toolboxes to facilitate interaction among data and models. This requires skills in data handling, preservation, and access by empiricists and modelers, all of which must be trained and integrated. [Section 5.15]

- COMPASS desires a secure method of interacting with commercial cloud storage providers (Dropbox, Google Drive, etc.), or a similar in-house service, to help ease transition into using more sustainable utilities when working with collaborators. Key features of this storage would be the ability for all collaborators—at national labs, research institutes, or universities—to access project-related files such as documents, presentations, and small (say, 1 GB or less) datasets. [Section 5.16]
- The CASCADE SFA is unlikely to make use of cloud services any time soon, largely because of costs required to host the magnitudes of data that can be operated on.
- CATALYST is not currently using cloud services for model output storage and distribution. However, the CESM project has engaged in early and preliminary studies of the utility of AWS [Section 5.18]:
  - Some older datasets (CESM1 Large Ensemble, for example) have been "Zarr-ified" (e.g., compressed using the Zarr tool) and placed on AWS as tests.
  - Analysis of the usability of AWS-hosted data for the user community is ongoing.
  - The cloud may become an important resource in the future, but CATALYST will continue to rely on CDG and ESGF for the distribution of model output.
- It is expected that demands will increase for sensor networks with artificial intelligence (AI capabilities) as well as self-driving labs capabilities for BER use cases. To meet the needs of rapidly expanding field data from DOE Earth Systems sensors and instruments, it will be critical to enable data acquisition, transformation, and analytics workflows that can integrate real-time data streams with predictive models. Live data streams from sensors can then be used to train models using a hybrid of AI/ML and deterministic techniques, which can, in turn, generate real-time predictions and feedback for sensors out in the field. [Section 5.19]
- It is very likely that clouds over various forms will play a role in future BER multifacility workflows [Section 5.19]:
  - This could take the form of surging to commercial clouds to meet targeted turnaround times for analysis, hosting public datasets, or other web services.
  - This may also include cloud-like models provided by a future integrated research infrastructure.
  - Maintaining a replicated instance of data services on cloud resources for failover, redundancy, or backup should be considered to enable broader access, especially to resources/users outside the DOE complex.
  - Enabling analysis capabilities in the cloud for very large datasets already stored in the cloud (bringing compute to the data).
- Most participants do some commercial cloud computing, but cost remains the biggest obstacle to use. ESnet's peering with cloud providers has simplified some of the data workflow, but the use of DOE-SC resources for computing remains popular and necessary. Participants acknowledge that having more information on the ways that ESnet peers with cloud providers would be useful information when negotiating contracts. [Focus Groups]

3. Review Recommendations

ESnet recorded a set of high-level recommendations from the BER user facilities and ESnet requirements review that extend ESnet's ongoing support of BER-funded collaborations. Based on the key findings, the review identified several recommendations for BER, ASCR, ESnet, and DOE HPC facilities to jointly pursue. These items are listed as guidance for future collaboration, and do not reflect formal project timelines. ESnet will review these with BER participations on a yearly basis, until the next requirements review process begins.

These are also organized by topic area for simplicity and follow common themes:

- Facility Management and Readiness.
- Scientific Data Management.
- Scientific Workflow.
- Computational and Storage Requirements.
- Remote Collaboration and Operational Requirements.
- Multifacility Computational Workflows.
- Domestic Networking for Local and Wide-Area Data Mobility.
- Emerging Use Cases.

3.1 Facility Management and Readiness

- ESnet will work with the BER community to adapt workflows to use DTNs between facilities to facilitate data mobility, instead of personal computing devices that often run using local wireless connectivity. This "last foot" problem of data movement can be solved by adopting platforms such as Globus at the application layer, and coupling this with highly capable data movement devices on fast networks, such as those located at HPC centers. [Focus Groups]
- ESnet will work with the BER community to address some of technological, sociological, and policy level constraints that come with the operation of BER multifacility workflows. [Section 5.19]

3.2 Scientific Data Management

- ESnet will work with the CMIP7 model community in the coming years as the volume of the data increases to the exabyte scale. To prepare for this, sites must be prepared in terms of network connectivity, storage, and data movement. [Section 5.6]
- ESnet will assist the BER community on ways to best address data sharing through use of both service platforms like Globus, and commercial cloud infrastructure. [Section 5.16]

3.3 Scientific Workflow

- The BER community, in conjunction with the ASCR HPC facilities, should explore new ways that HPC centers can improve data access for local and remote workflows and enable a distributed super-facility model to support future workflows. [Section 5.17] Suggestions include:
  - Creation and improvements of caching technology.

- Data mirrors across the DOE HPC facilities.
- Exploring behavior similar to a remote mounts of filesystems at ALCF, NERSC, and OLCF.

- The BER community should explore workflows to better stream, store, process, and share data. Efforts like this will be needed in the future as longer paths between observation and result are created. [Focus Groups]
- The BER community should investigate the use of a common document-sharing platform that can be supported and certified by all DOE laboratories to assist with collaborative work and potentially improve productivity. [Section 5.16]

3.4 Computational and Storage Requirements

- ESGF will work with ESnet, and R&E networking community projects, to understand how containerized workflows may be applicable to future deployments. An overall goal of this effort will be to ease the burden of upgrades, and improve operational metrics. [Section 5.6]
- CBI will work with ESnet to learn more about Kubernetes deployment. [Section 5.10]
- E3SM and NERSC staff will continue to work on ways to ease the hardware bottleneck to efficient use of the HPSS system. Doing so will improve the overall operation of E3SM's data-intensive workflows. [Section 5.13]
- COMPASS, ESGF, and ESS-DIVE will continue to address the growing volumes of datasets, and ways they can be stored and indexed. For the data to be useful, it must be searchable and downloadable for the research community. [Section 5.16]
- DOE ASCR and BER, through efforts like IRI, will continue to address issues surrounding the availability of computation, storage, and data movement tools as they relate to both large-scale computation and moderate-scale analysis. [Section 5.17]
- The BER community will continue to address the perpetual "shortage" of disk space, and how it relates to operations in the coming years. This can be attributed to model output increases, but also needing a better strategy to manage available disk resources across project partners. [Section 5.18]
- Globus will continue to work with the BER community on matters of tool adoption, security policy requirements, and ways to improve user experience through integration with storage. [Focus Groups]

3.5 Remote Collaboration and Operational Requirements

- ESnet will continue to work with the ESGF to address performance problems to ESGF resources that are located around the world. [Section 5.6]
- ESS-DIVE will continue to work with LBNL to address limitations with the Cloudflare service, and how it impacts data uploads and downloads. [Section 5.7]
- The BER community, in cooperation with ASCR HPC facilities, must address the growing challenges of federated identify to streamline ways that users and institutions can access and transfer research data. [Section 5.16]

3.6 Multifacility Computational Workflows

- ESnet will work with the GLBRC to understand and improve data transfer. [Section 5.11]
- CASCADE will continue to work with NERSC, and other DOE HPC facilities, to address the needs of moderate-scale computing requirements, and how they can be accessed on shorter time scales. This observation relates to the ongoing work of the IRI activity, in that portable workflows may be able to utilize other DOE resources in the future, instead of being tied to a single multifacility partner. [Section 5.17]
- The BER community, in conjunction with other facilities operated by ASCR, must address the lack of a cohesive identify, authorization, and authentication strategy in order to enable future multifacility use cases and workflows. [Focus Groups]
- DOE SC, with stakeholders that include BER and ASCR facilities, will continue to address research workflow requirements as identified in the IRI effort, with a focus on ways to enable multifacility operation on a more regular basis. The BER community in particular desires portability between HPC facilities, to allow for opportunities to store and compute where resources are available. Additional needs include help in modernizing workflows to take advantage of new technologies in analyses and data movement, as well as optimizing software and hardware for increased data volumes. [Focus Groups]

3.7 Domestic Networking for Local and Wide-Area Data Mobility

- ESnet will work with BER facilities and experiments to address connectivity concerns, as data rates increase in the coming years. Handling these increased data requirements will require fast networks between the major HPC facilities, as well collaboration locations. It is expected that some simulations will be PB scale, approaching EB, and there will be hundreds of these in the coming years. [Focus Groups]
- It is recommended that EMSL experiment with the Globus Modern Research Data Portal (MRDP) for future portal work. [Section 5.4]
- ESnet and PNNL will collaborate to improve data movement performance in support of EMSL. PNNL will be upgrading network connections in future years to ESnet, and will also work to upgrade campus connectivity and remove several lower-bandwidth bottlenecks. [Section 5.4]
- It is recommended that ARM experiment with the Globus MRDP instance for future portal work. [Section 5.5]
- It is recommended that ESGF continue to work with ESnet to help characterize and improve performance to remote resources. This can apply to domestic and international deployments. Slow performance must be mitigated as the CMIP7 era starts, and data sizes increase. [Section 5.6]
- ESnet will advise CABBI on ways they can improve network access using wireless and satellite options to data transfer from remote environmental instrumentation. [Section 5.9]
- CBI will work with ORNL, NREL, and ESnet to address slow data-transfer problems. This will involve migration of data to ORNL resources that have access to the institutional Science DMZ and performance testing with Globus. [Section 5.10]
- ESnet will work with GLBRC to understand performance limitations through the MSU campus, and potentially design a Science DMZ infrastructure. [Section 5.11]

- The BRCs will collaborate on inter-BRC shared research objectives, many of which will increase the need for efficient data transfer over the next five years. [Section 5.12]
- It is recommended that ESnet and LANL collaborate on ways to effectively share data with other NGEE collaborators. [Section 5.15]
- The BER community and ESnet will continue to design and operate multifacility workflows with a particular focus on efforts that will adopt modern data transfer tools, portals, integration with cloud providers, and increased ESnet connectivity and services to core facilities. [Section 5.19]
- The BER community and ESnet will address the growing data volumes that its experiments and facilities require in the coming years by evaluating network connectivity, adopting high-speed data movement software and hardware approaches, and working to eliminate the bottleneck to storage infrastructure. [Focus Groups]
- ESnet will work with the BER community to advise on possible ways that LEO satellites and 5G connectivity can be integrated into workflows that involve remote sensor environments. [Focus Groups]
- The ESnet's Science Engagement Team will continue to assist with data mobility issues. These range from helping to improve international connectivity, helping to tune DTNs, and advising on the best way to deal with remote resources that have minimal connectivity options. [Focus Groups]

3.8 Emerging Needs

- The BER community and ESnet will continue to address the best common practices of sharing research data between platforms (commercial clouds and R&E based clusters) in the future to minimize the support overhead. As collaborations and facilities migrate toward cloud computing resources, collaborators who provide cluster resources may require changes in file transport mechanisms. Each provider may have different mechanisms for transport and storage, and these may not translate between providers. [Section 5.1, Section 5.3, Focus Groups]
- The DOE SC community must work to together to address gaps in workforce, particularly as it relates to training and maintaining staff that assist with development and operations requirements: e.g., tasks that focus on maintenance, cybersecurity patching, and user support. Future DOE support platforms that simplify infrastructure support would result in more time to focus on project-affiliated work, along with efforts to grow staff through training in affiliated cyberinfrastructure technologies. [Section 5.2, Focus Groups]
- The BER community and ESnet will discuss options that enable fast and secure methods for interacting with commercial cloud storage providers (Dropbox, Google Drive, etc.), to help ease transition into using more sustainable utilities when working with collaborators. [Section 5.16]
- ESnet and the BER community will continue to discuss best common practices for interacting with cloud providers, both from the standpoint of peering and connectivity, but also ways to increase performance expectations. [Focus Groups]
- The BER community will work to address the growing need for artificial intelligence (AI capabilities), as well as self-driving labs capabilities, across BER use cases. It will be critical to enable data acquisition, transformation, and analytics workflows that can integrate real-time data streams with predictive models. Live data streams from sensors can then be used to train models using a hybrid of AI/ML and deterministic techniques, which can, in turn, generate real-time predictions and feedback for sensors out in the field. [Section 5.19]

4 Requirements Review Structure

The requirements review is designed to be an in-person event; however, the COVID-19 pandemic has changed the process to operate virtually and asynchronously for several aspects. 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 recommendations that can be implemented in the coming years.

4.1 Background

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

- 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.

4.2 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 study template has the following sections:

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 (LAN, MAN, and WAN) capabilities that connect the science experiment/facility/data source to external resources and collaborators.

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

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

5 BER Case Studies

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 PIs funded by the Office of BER of the DOE SC. These case studies were discussed virtually between August 2022 and April 2023.

The case studies were presented, and are organized in this report, in a deliberate format to present an overview based on individual experiments, larger facilities, and in some cases the encompassing laboratory environments that provide critical resources for operation. The case studies profiled include:

- JGI
- DOE KBase
- NMDC
- EMSL
- ARM
- ESGF
- ESS-DIVE and Self-Driving Field Laboratories Research
- SBC at the APS
- CABBI
- CBI
- GLBRC
- JBEI
- E3SM
- AmeriFlux network
- NGEE Arctic and NGEE-Tropics
- Coastal Efforts
- CASCADE
- CATALYST
- Multifacility Workflows in BER and Relationship to the IRI Initiative

Each of these documents contains a complete set of answers to the questions posed by the organizers:

- 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?

A summary of each will be presented prior to the case study document, along with a "Discussion Summary" that highlights key areas of conversation from authors and attendees. These brief write-ups are not meant to replace a full review of the case study, but will provide a snapshot of the discussion and focus during the in-person review.

5.1 JGI

The US DOE JGI, a DOE SC user facility, is managed by DOE's BER. JGI provides HT DNA sequencing, synthesis and analysis services to the user community in support of BER's bioenergy and environmental missions.

5.1.1 Discussion Summary

- The main computational systems at JGI with multisite footprints are part of the Advanced Analysis team, within the Data Science and Informatics department. Both systems are multisite distributed systems with significant internal backend data analysis and management processes that span multiple sites. Within Advanced Analysis, the two major teams are:
  - The Sequence Data Management (SDM) team, which supports the JGI Archive and Metadata Organizer (JAMO). JAMO manages over 14 PB of data storage and contains raw sequencing data and analysis artifacts for the JGI dating back to the late 1990s.
  - The JAWS team operates a distributed workflow service that allows analysis workflows to be distributed across multiple sites to provide resilience and additional capacity to JGI analysis. The sites involved are PNNL, LBNL (including NERSC, LBNL IT, and the JGI), and Amazon Web Services (AWS).
- The JGI SDM services support analysis and publication pipelines for biological samples that are sent to the JGI for sequencing and analysis. The sequencing pipeline operates within the administrative boundaries of LBNL (using resources at the JGI, as well as some provided by LBNL IT and NERSC).
- JGI Sequence data is uploaded to the National Institutes of Health's National Center for Biotechnology Information Sequence Read Archive (SRA), where it is freely accessible for research and education purposes.
- The JGI is a user facility hosted at LBNL. Numerous soft funded research projects operate within or in collaboration with the JGI:
  - The National Energy Research Scientific Computing Center (NERSC): development of tools and infrastructure for microbiome data science, optimized for high-performance computing environments.
  - The DOE System Biology Knowledgebase (KBase): a system to access and analyze data on scales ranging from single genes and individual genomes to metagenomes to systems-level modeling and understanding.
  - The NMDC: a community-driven national effort aimed to develop standards, processes, and infrastructure for an integrated microbiome data ecosystem.
  - The EMSL: a partner in the Facilities Integrating Collaborations for User Science (FICUS) program.
  - The HudsonAlpha Institute for Biotechnology[2]: a genome sequencing center that is integrated with SDM, providing genome sequencing for JGI analysis services.
- The JGI manages petabytes of data on site, and regularly shares this data with EMSL and HudsonAlpha. Other entities that regularly share data with the JGI include KBase and the NMDC.

2 https://www.hudsonalpha.org/

- The working set for JAMO is currently under 1 PB, but grows slowly and will likely be over 1 PB in the next couple of years. The archive in HPSS grows more rapidly, on the order of 1.5 PB per year currently. JAMO data is never deleted, so data usage grows monotonically.
- JGI's JAMO operates a data portal providing programmatic access to JAMO data, both a spinning disk "working set" as well as HPSS archives. Recent trends have been on the order of 1 to 2 PB per year and are accessed via the portal.
- Sequencing technologies continue to evolve to enable larger capacities at cheaper costs. The JGI has the two latest sequencers from Illumina and PacBio, which will increase the capacity by approximately threefold over current capacities. However, due to upstream (e.g., library creation) limitations, all this capacity will not be used. There will be significant data growth over the next five years from both internal capacity growth and external data being imported into JGI data portals.
- The JGI's data management system manages storage resources, and JAMO will always have a managed repository at NERSC, but this may change in the future. For egress to data centers other than NERSC, two models are under consideration:
  - Allowing JAWS to copy the files from NERSC to a remote site as needed.
  - Setting up managed repositories at remote data centers that JAMO can manage directly, and allowing transfer back-and-forth to NERSC.
- JAWS is designed to move the JGI away from bespoke and/or hand coded workflows, and to protect production analysis pipelines from outages and/or downtimes at partner HPC facilities (e.g., NERSC, LBNL IT's Lawrencium cluster, EMSL's [PNNL] Tahoma cluster, as well as private cloud clusters in AWS).
- Most of the JAWS workflows will operate within LBNL, split between NERSC and a JGI-dedicated cluster hosted within LBNL IT, with only a small fraction going off site to PNNL and AWS.
- Neither JGI's JAMO nor JAWS intrinsically has special networking requirements for data acquisition or distribution. Both JAWS and JAMO make use of replicating data from the central facility to the remote sites using tools such as Globus, as well as rsync.
- Collaborators who provide important cluster resources may require changes in file transport mechanisms. The JGI currently uses rsync or ssh and S3 for communications between backend cluster resources (S3 for AWS clusters). The JGI expects to expand to Google Cloud, which uses an S3 compatible protocol to manage object storage. The JGI does not have any plans to add Aspera support, and would prefer to avoid doing so to minimize the support overhead.
- JAWS supports compute clusters in AWS that scale out on-demand for specific workflows that require the reliability of AWS services. In the future, the JGI may move some of the core infrastructure into a cloud provider to maximize availability of services such as JAWS Central.
- The JGI does not foresee any data-related resource constraints that would be addressed by ESnet. JAMO data grows at a relatively high rate (currently 1.5 PB per year), but most data remain within the facility, with the exception of the aforementioned sharing and backup patterns.
- The JGI does not foresee any need to expand into leadership HPC facilities: the workflows really require only midrange HPC resources. If workflows users begin to use more artificial intelligence/machine learning (AI/ML) tools, the JGI may eventually require access to clusters equipped with graphical processing units (GPUs).

5.1.2 JGI Facility Profile

BER's bioenergy and environmental missions mirror DOE's and national priorities to develop renewable and sustainable sources of biofuels and bioproducts from plant biomass, to understand the biological processes controlling greenhouse gas accumulation in the atmosphere (especially carbon dioxide and methane, key factors in global climate change), and to gain insights into biogeochemical processes controlling the cycling of key nutrients in environments for sustainable bioenergy production or the mobility of heavy metals and radionuclides at contaminated sites for which DOE has stewardship responsibilities. Projects with direct relevance in these areas will have the best chance for selection. Projects focused on organisms for comparative purposes, on model systems for microbe-microbe or plant-microbe interactions, or on development of improved sequencing-based technologies are also welcomed but the applicant should clearly outline how the proposed work will advance BER-mission relevant science. Projects primarily focused on human health, food/animal agriculture, wastewater treatment, or bioremediation of organics will not be considered. Projects targeting marine systems must clearly demonstrate relevance and translatability to freshwater, coastal or terrestrial systems. JGI has a staff of roughly 280 employees who working on a diverse portfolio of projects[3].

5.1.2.1 Science Background

JGI's main computational systems with multisite footprints are part of the Advanced Analysis team with the Data Science and Informatics department. Within Advanced Analysis, the two major teams are the SDM team which supports JAMO and the JAWS team. Both systems are multisite distributed systems with significant internal backend data analysis and management that span multiple sites. In addition, the JAMO service is accessible via data portals that provide programmatic and interactive browser access.

The SDM services support analysis and publication pipelines for biological samples that are sent to JGI for sequencing and analysis. The core service is the JAMO, which manages a 14.5+ PB data store (14.5+ PB in HPSS, roughly 1 PB on spinning disk at NERSC) containing raw sequence data and analysis artifacts for JGI dating back to the late 1990's. Most of the sequencing pipeline operates within the boundaries of LBNL (JGI, LBNL IT and NERSC) however there is a remote sequencing service which uploads sequence data to the central services at LBNL. JAMO's raw sequence data held in HPSS is currently being backed up to remote sites for disaster recovery purposes via external processes, but in the future the cross-site replication will be integrated into JAMO. In addition, when sequence data is made publicly available, it is uploaded to the National Institute of Health's National Center for Biotechnology Information SRA archive, where it is freely accessible for research and education purposes.

The JAWS team operates a distributed workflow service that allows analysis workflows to be distributed across multiple sites, in order to provide resilience and additional capacity to JGI analysis. It uses a collection of services which are the de facto standard for workflow execution and high-speed data movement. The sites involved are PNNL, LBNL (NERSC, LBNL IT and JGI) and AWS.

5.1.2.2 Collaborators

JGI is a user facility hosted at LBNL, there are numerous soft funded research projects that operate within or in collaboration with the JGI. As strategic user facility partners for JGI, the data partners/collaborators are:

NERSC: Through its partnership with NERSC and the DOE ESnet, the JGI has unique access to both expertise and infrastructure to facilitate these efforts. Particular areas of emphasis will be the development of tools and infrastructure for microbiome data science, and the development of new informatics tools specifically optimized for high-performance computing environments. NERSC supercomputer facilities have historically served as the major computing resource for running JGI analysis workflows (JAMO and JAWS), and NERSC Data services such as the cluster filesystem and HPSS archives are used to store JAMO data.

3 https://jgi.doe.gov/our-science/

KBase: KBase is a critical partner in enabling systems biology analyses, such as metabolic modeling, that go well beyond genome annotation and analysis. The JGI and KBase seek to enable users to seamlessly access and analyze data on scales ranging from single genes and individual genomes to metagenomes to systems-level modeling and understanding. JGI and KBase are designated DOE SC Public Reusable Research (PuRe) Data Resources are authoritative providers of data or capabilities in their respective subject area. KBase and JGI partner through data oriented APIs, with JGI services at LBNL and KBase production services at ANL. Data volume for these APIs are currently several orders of magnitude smaller than the backend services for JAMO and JAWS, and are not reported as a major integration with JGI in this document.

NMDC: In 2017, a stakeholder workshop involving JGI leadership and hosted by the American Society for Microbiology brought together representatives from academia, industry, government, and philanthropic funding agencies to conceptualize the NMDC. These efforts serve as the foundation of a community-driven national effort aimed to develop standards, processes, and infrastructure for an integrated microbiome data ecosystem. NMDC and the JAWS team have partnered to use JAWS services for NMDC analysis workflows, initially JAWS will handle NMDC workflows that operate at NERSC. In the future, as the collaboration expands, other clusters with NMDC allocations will be integrated into JAWS.

EMSL: EMSL and JGI are partners in the FICUS program created in 2014 to accelerate ambitious user research projects. FICUS represents a unique opportunity for researchers to combine the power of genomics and molecular characterization in one proposed research project that uses a broad range of the capabilities of each facility, and generate datasets beyond what each of these facilities could generate by itself. The program provides researchers with access to the world-class resources of multiple user facilities through a single proposal.

HudsonAlpha Institute for Biotechnology: HudsonAlpha is a major collaborator with partnerships on many genomics research projects. The HudsonAlpha Genome Sequencing Center is integrated with JGI's Sequence Data Management services, providing genome sequencing for JGI analysis services.

As future collaborators are identified and integrated, it is expected that JGI will use JAWS for workflow management, and Globus for data movement, and use NERSC as the primary repository. The long-term strategy for storage and data management will depend on where JGI is able to establish resources/capabilities in the future. LBNL IT is capacity limited, and NERSC can become oversubscribed. One possibility is to leverage PNNL/EMSL, or the commercial cloud, for additional computing capabilities.

JGI's data management system currently requires the ability to manage storage, but there are plans to remove this requirement. JAMO will always have a managed repository at NERSC, or at least there are no plans on changing this. There should be a centralized way of managing the life cycle of files on disk. For egress to data centers other than NERSC there are two models under consideration.

- Let JAWS copy the files from NERSC to the remote site. Remote data will not be managed nor recognized by JAMO and will need to be lifecycled by JAWS or some other mechanism.
- Set up managed repositories at remote data centers that JAMO will manage the life cycle of these repositories and manage the transfer back-and-forth to NERSC.

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
NERSC - LBNL	Primary	Filesystem, HPSS	14.5PB	Daily updates (Constant)	This is source	None
KBase - ANL+LBNL	Consumer of primary data	API based gateway	Not available	Infrequent	N	None
NMDC - LBNL (eventually ANL, LANL, EMSL)	Consumer of scheduler services	API	No historical data yet	No historical data yet	Yes	Not yet
EMSL	Secondary	Globus	1.5PB	Daily updates	DR backup	None
Hudson Alpha	Primary-ish	Globus & rsync+ssh	400TB	Daily updates	See Section 3	None

Table 5.1.1: JGI Collaboration Space

5.1.2.3 Instruments and Facilities

JAMO/Sequence Data Management

The Integrated Genomics Building at Berkeley Lab houses significant wet-lab facilities as well as PacBio and Illumina sequencers. In addition, samples are also sequenced at the HudsonAlpha facility, and the sequence data is then copied to JAMO.

As data is brought into JAMO, it is available for analysis on the NERSC Cluster filesystem and also stored in the NERSC HPSS system for archival storage. The working set for JAMO is currently under 1PB, but grows slowly and will likely be over 1PB in the next couple of years. The archive in HPSS grows more rapidly, on the order of 1.5PB/year currently. JAMO data is never deleted, so data usage grows monotonically.

After data has been loaded into HPSS by JAMO, out of band disaster recovery backups copy the data from the NERSC HPSS into the HPSS systems at ORNL and PNNL, using Globus. Depending whether one of the Globus endpoints is an HPSS system (EMSL has a DTN endpoint in HPSS), the transfer rate is either aggressively tuned for maximum disk I/O or else is bottlenecked by tape archive I/O rates. The backup procedures do not have a requirement for high data rates. During 2022, the copies to ORNL and PNNL were performed manually using scripts for the initial synchronization. Moving forward, the plan is to have HPSS automate the cross site incremental archives.

When data is ready for publication, it is uploaded to NIH's NCBI using the Aspera protocol, where it becomes available to the general public via the NCBI SRA portal.

Estimated JAMO/SDM Data Transfer in TB					
	FY 2023	FY 2024	FY 2025	FY 2026	FY 2027
ORNL (Globus)	407	250	290	340	390
EMSL (Globus)	950	250	290	340	390
NCBI SRA (Aspera)	50	100	175	200	250
HudsonAlpha (rsync+ssh)	100	60	70	80	90
HudsonAlpha - (Globus)	300	232	270	310	360

Table 5.1.2: Estimated JAMO Backend data transfer

JAMO operates a data portal providing programmatic access to JAMO data, both a spinning disk "working set" as well as HPSS archives. Data portal activity has been uneven over the last several years, making it difficult to extrapolate future traffic, however the last several years of data usage are presented, to provide an idea of the scale of utilization. Aside from the outlier in 2018, recent trends have been on the order of 1000-2000 TB/year. As a general estimate, it may be reasonable to estimate an upper bound for data transfer in the range of 2000 TB/year.

Fiscal Year	Data Transfer in TB
2014	600
2015	1790
2016	600
2017	2260
2018	6450[4]
2019	660
2020	1060
2021	1110

Table 5.1.3: Historical Yearly Data Portal Data Transfers in TB

4 2018 data includes 5.3 TB from single user repeatedly downloading same data 818 times

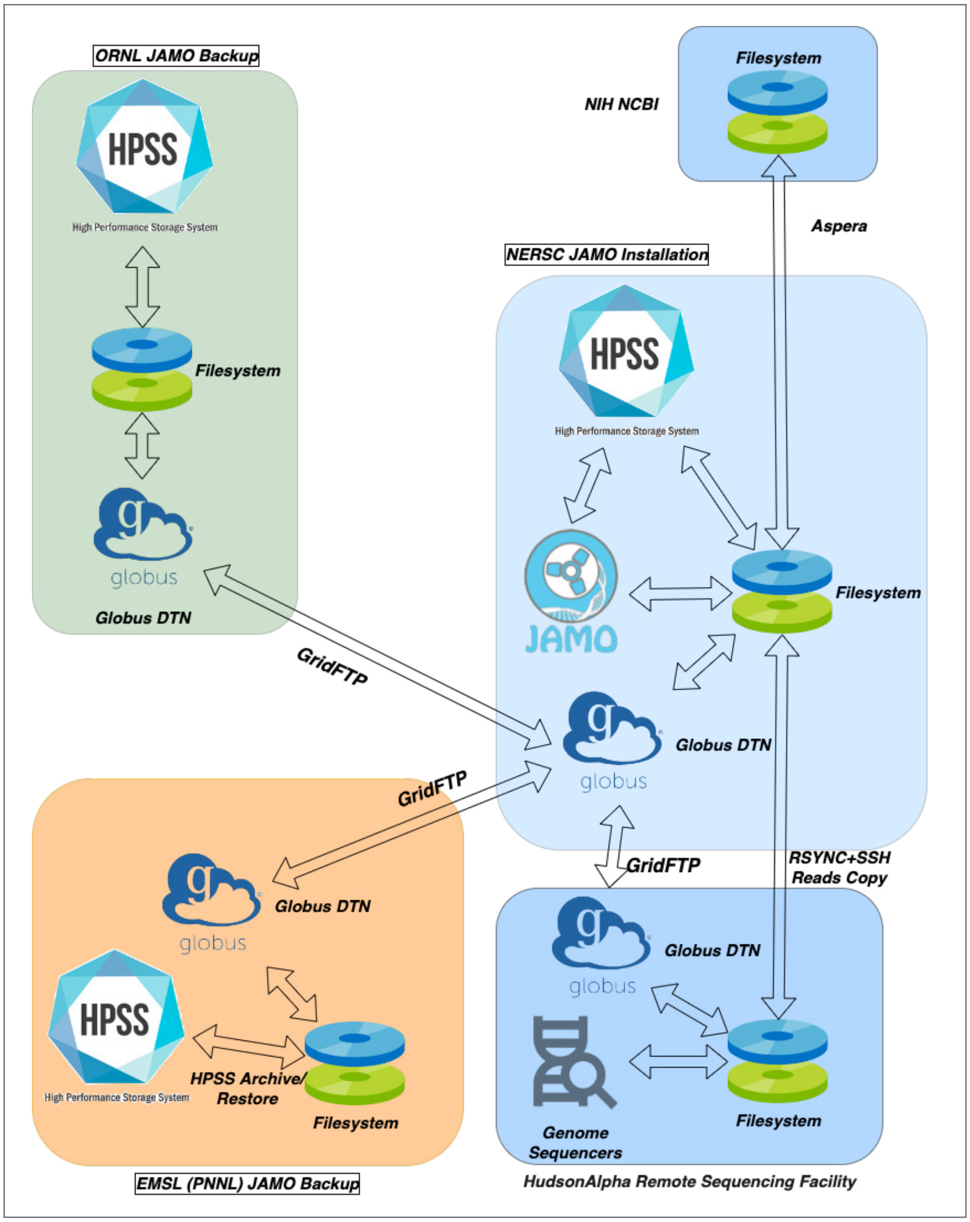

Figure 5.1.1: JAMO Main Services, Sites and Protocols

Sequencing technologies continue to evolve to enable larger capacities at cheaper costs. JGI has the two latest sequencers from Illumina and PacBio, which will increase the capacity by approximately 3-fold over current capacities. However, due to upstream (e.g., library creation) limitations, all this capacity will not be used. There will be significant data growth over the next five years from both internal capacity, but also from external data being imported into JGI data portals.

5.1.2.4 Generalized Process of Science

JAWS is a workflow service designed to move JGI away from bespoke and/or hand coded workflows, and to protect the production analysis pipelines from outages and/or downtimes at partner HPC facilities. By standardizing on the Workflow Description Language and the widely adopted Cromwell workflow engine from the Broad Institute, JGI can avoid some of the pitfalls of siloed one-off, rigid workflows that are rigidly hardcoded in disparate languages. The requirements of writing site agnostic Workflow Description Language (WDL) workflows also advances the goal of improved maintainability for JGI workflows.

We have started creating a community of WDL users within JGI who can support each other, as well as share workflows in a modular fashion. JAWS adds a layer above Cromwell that monitors the availability of different compute clusters, and distributes the necessary input data to the appropriate site, launches the workflows and then collects the output data, in the form of analysis artifacts, and performance metadata. This output data returns to the central service where users can collect the results of the analysis and analyze performance characteristics. This has enabled JGI to move some of its analysis runs off NERSC, and into clusters such as LBNL IT's Lawrencium Cluster, EMSL's (PNNL) Tahoma cluster, as well as private cloud clusters in AWS. The strategy is to use free allocations of cluster resources in partner institutions whenever possible, and if certain analyses prefer to run in AWS (and if budget available to support it) they can be routed to AWS.

Depending on the services supported, the central dispatcher (JAWS Central) and the service that resides with the local batch scheduler (JAWS Site) use either Globus (Gridftp) or AWS S3 for transporting source and output data, as well as reference data. Least recently used caching on source data is used to avoid excess network traffic, and Cromwell's call caching is used to avoid re-running analysis tasks which have already been run with the same inputs.

For remote procedure call communication between JAWS Central and Site, depending on the security preferences of each site, either REST requests over HTTP or JavaScript Object Notation (JSON) messages over RabbitMQ (AMQP) are used for communication. For many DOE sites, AMQP has proven to be the preferred transport, while REST has been easiest to use with AWS.

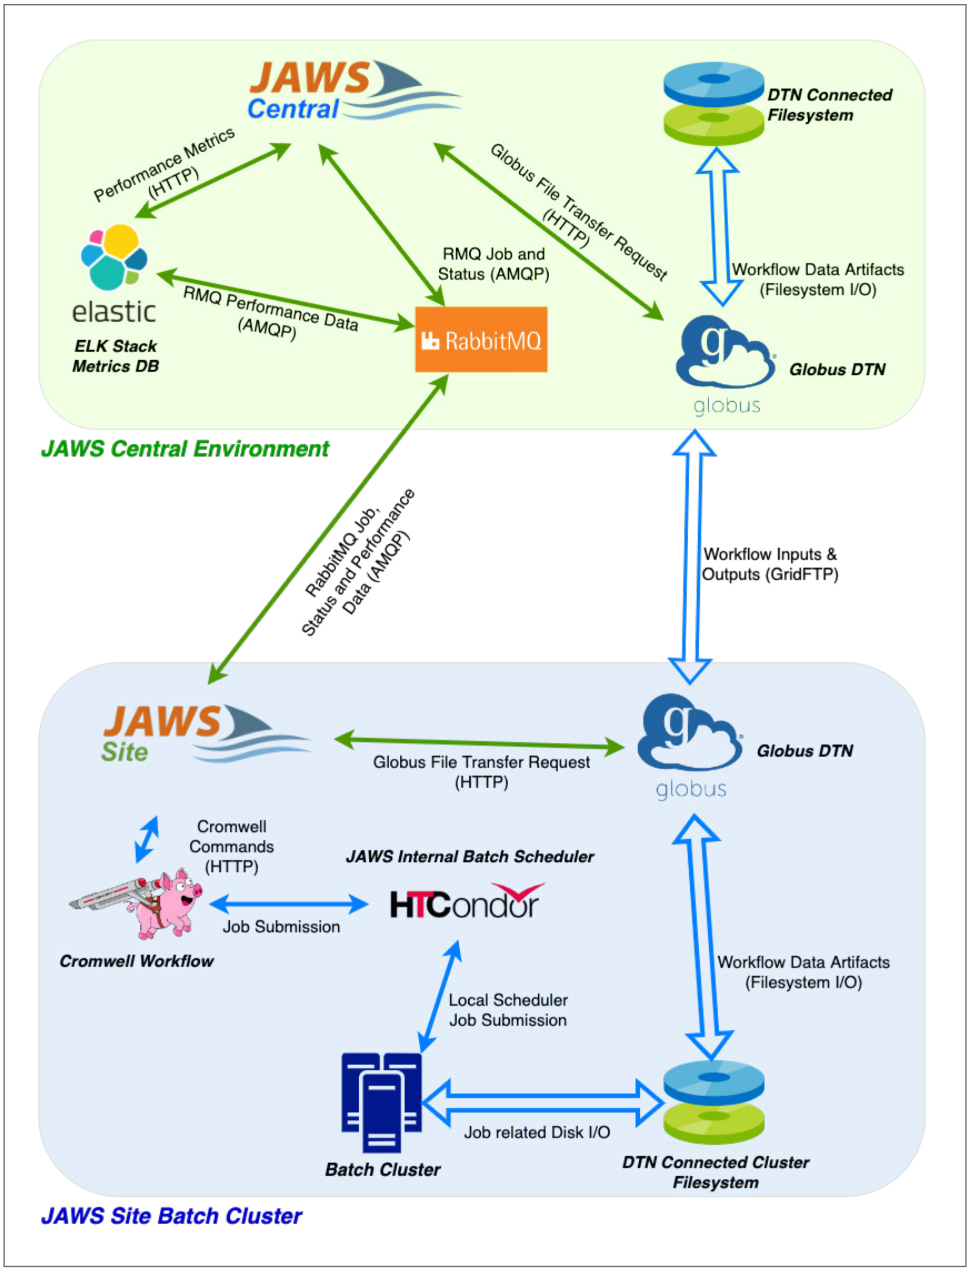

Figure 5.1.2: Example of JAWS Central and JAWS Site services and protocols

Because JAWS has recently gone into production, there is not historical data from which to extrapolate network traffic usage. Many of these workflows have historically operated within NERSC's clusters and network usage data is opaque to the JGI team. The addition of new workflow customers such as NMDC who do not have historical production workflow metrics also complicates the process of estimation.

It is expected that most of the JAWS workflows will operate within LBNL, split between NERSC and a JGI-dedicated cluster hosted within LBNL IT, with only a small fraction going off site to PNNL and AWS. As a consequence, ESnet may not see much traffic between LBNL and other sites due to JAWS. The general trajectory of data growth will be tied to the growth of JAMO data + NMDC Data.

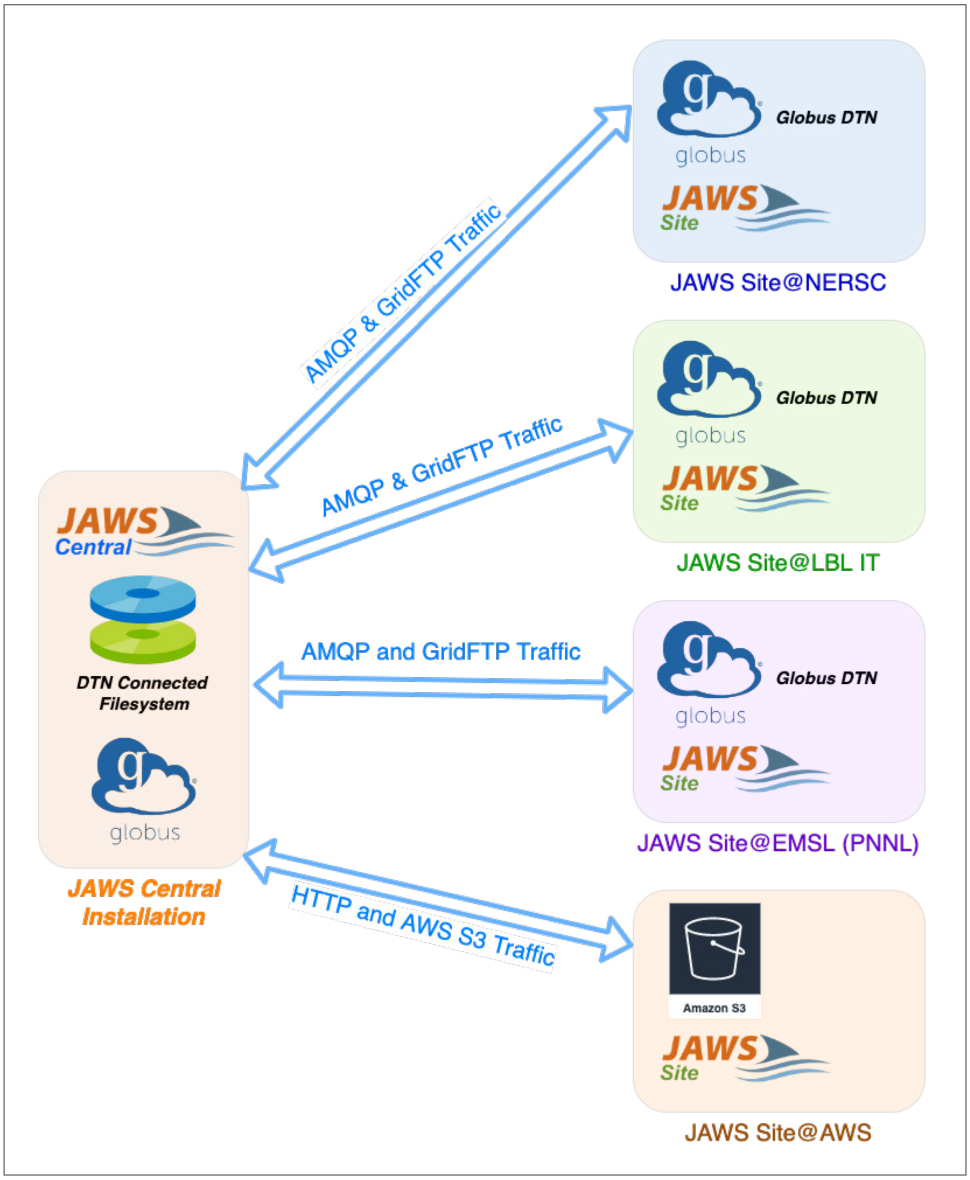

Figure 5.1.3: Current JAWS Site and Protocol Interactions

5.1.2.5 Remote Science Activities

JGI is a complex enterprise with multiple simultaneous scientific processes. A visit to the main website[5] provides an overview of the internal scientific research projects, the tools and data sources JGI operates as well as the user programs that are offered to provide sequencing and analysis of biological samples. JGI cannot reasonably summarize them in this document, however it is possible to present the relationships between JAMO and the other data sources and analysis groups. Each one of the groups and databases shown have their own scientific processes, with JAMO sitting at the core as a data and metadata repository.

5 https://jgi.doe.gov/

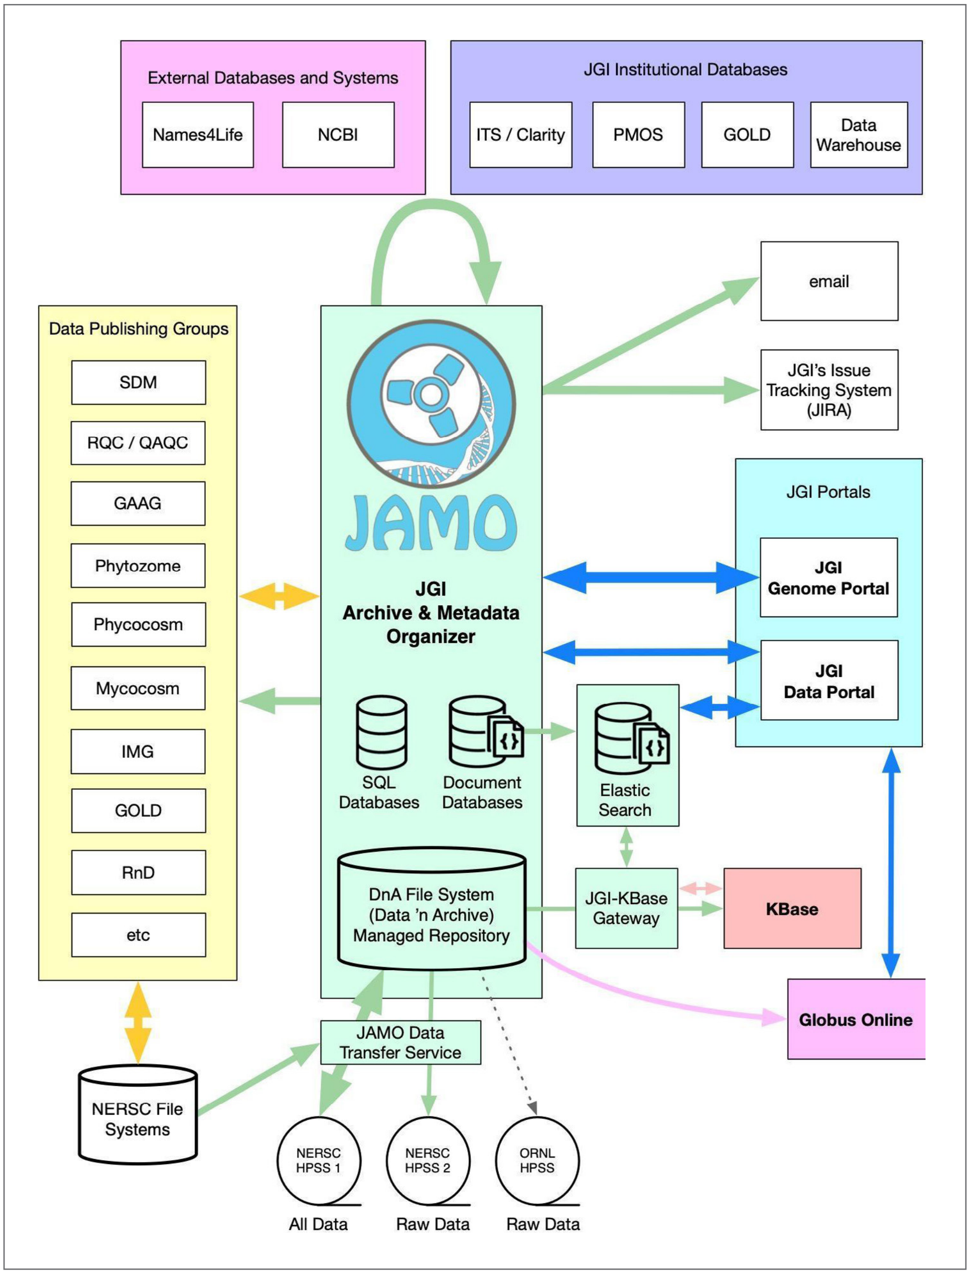

Figure 5.1.4: JAMO relationships to other JGI scientific activities

As a common platform for running analysis workflows, JAWS serves multiple groups within JGI, such as the "Data Publishing Groups" on the left side of Figure 5.1.4.

5.1.2.6 Software Infrastructure

JAWS has two categories of collaborators:

1. Collaborators who run workflows
2. Collaborators that provide cluster resources

We expect collaborators who run workflows to be JGI users, or groups closely affiliated with JGI/LBNL such as NMDC. JGI users will submit their workflow requests from the sites where JGI users have shell accounts, such as NERSC, LBNL IT or within hosts operated by JGI's internal HPC staff. NMDC users are distributed across EMSL(PNNL), LBNL, LLNL and LANL. These users will submit workflow requests from NERSC or EMSL. These users use WDL files that specify input and output files using either filesystem paths or HTTP/S3 URLs. When users specify their inputs using filesystem paths, JAWS transparently moves files from the user's home system to remote compute clusters using Globus. This should not change because WDL files only really support these types of locations.

Collaborators who provide important cluster resources may require changes in file transport mechanisms. JGI currently uses rsync/ssh and S3 for communications between backend cluster resources (S3 for AWS clusters). JGI expects to expand to Google Cloud, which uses an S3 compatible protocol for accessing their object storage. JGI does not have any plans to add Aspera support, and would prefer to avoid doing so in order to minimize the support overhead.

JGI does not foresee any need to expand into leadership HPC facilities - the workflows only really require midrange HPC resources. If workflows users begin to use more AI/ML tools, JGI may eventually require access to GPU-equipped clusters.

5.1.2.7 Network and Data Architecture

The data architecture has been summarized at the high level in preceding diagrams. Neither JAMO nor JAWS intrinsically have special networking requirements for data acquisition or distribution, however the use of HPC clusters, NERSC's HPSS and Cluster Filesystem and Globus takes advantage of the high-performance interconnects and well-tuned HT LAN/WAN provided by partner HPC facilities.

Both JAWS and JAMO make use of replicating data from the central facility to the remote sites. If there were a widely supported, low(ish) latency mechanism for filesystem replication, JGI would be able to replace the rsync and Globus based file copies.

5.1.2.8 Cloud Services

Currently JAWS supports compute clusters in AWS that scale out on-demand for specific workflows that require the reliability of AWS services. In the future, JGI may move some of the core infrastructure into a cloud provider to maximize availability of services such as JAWS Central. In addition, GCP is a major partner of LBNL, and it may make economic sense to migrate cloud-based compute clusters to GCP (or other cloud vendor). A cloud service that has high-speed peering with ESnet, and lower egress fees would be a compelling option for JAWS (and JGI in general).

There are three broad ways in which cloud-based services will integrate with JAWS in the future:

1. **Using AWS and GCP as compute resources for JAWS workflows.**
   JAWS already has a compute pool in AWS for running genome assembly workflows, and JGI can expand to cover other workflows. There is a plan related to the testing of on-demand compute pools in GCP in spring 2023, to compare costs and ease of maintenance. A drawback JGI encountered was the difficulty of programmatic access to the NERSC Globus endpoint from within AWS - as a consequence JGI have built S3 support into JAWS to support file movements between JAWS Central and AWS-hosted clusters.
2. **Moving core JAWS infrastructure into a more resilient AWS or GCP environment.**
   Core JAWS infrastructure is now currently to sitewide outages at LBNL/NERSC - this undermines the basic premise of JAWS as providing workflow services independent of any site's outage. Pre-emptive power shutdowns have not been an issue in recent years, but were a regular feature

from about 2017-2020; given the trend of climate change, JGI needs to be prepared for a return to wildfires induced power outages. Designing the core JAWS infrastructure to run as resilient cloud applications in AWS and GCP would avoid JAWS downtime when LBNL is shutdown.

3. **Architecture designs based Cloud native API practices**
We currently support WDL workflows that identify inputs and outputs as filesystem paths on the local clusters. This has resulted in portability issues not only regarding local file access, but also support of full Portable Operating System Interface (POSIX) semantics on cluster filesystems. Using S3 api instead of a POSIX based filesystem for transferring workflow inputs and outputs would avoid these issues. JGI is also dependent on executing binaries for batch scheduler interactions. Ideally API endpoints can be used for:
   - Accessing data used as inputs and outputs for workflows
   - Interfacing with batch schedulers such as Slurm (instead of executing scheduler command-line tools)
   - metrics collection such as a Prometheus or Open Telemetry instead of the custom psutils based metrics JAWS currently depends on

   Ideally all these services would have an integrated Identity and Access Management service for controlling access to APIs, instead of a hodgepodge of site and service specific credentials and JAWS credentials.

   We hope that as part of the nascent integrated research infrastructure (IRI) activities, more laboratory partners can migrate to offering API based access to their services, based on widely supported (official or de facto) standards.

5.1.2.9 Data-Related Resource Constraints

We do not foresee any data-related resource constraints that would be addressed by ESnet. JAMO data grows at a relatively high rate (currently 1.5 PB/year) which is well within the capacity of the existing tape silos for HPSS, with the main issues being administrative approval for additional quota.

5.1.2.10 Outstanding Issues

None to report at this time.

5.1.2.11 Facility Profile Contributors

JGI Representation

Stephen Chan, *JGI Advanced Analysis Manager,* sychan@lbl.gov

Chris Beecroft, *JAMO/SDM team lead,* kmfagnan@lbl.gov

Kjiersten Fagnan, *JGI CIO,* kmfagnan@lbl.gov

Nigel Mouncey, *JGI,* nmouncey@lbl.gov

Georg Rath, *JGI,* gbrath@lbl.gov

ESCC Representation

Rune Stromsness, *LBNL,* rstrom@lbl.gov

Richard Simon, *LBNL,* rsimon@lbl.gov

5.2 DOE KBase

The DOE KBase is an open software and data platform that aims to enable researchers to predict and ultimately design biological function. KBase's unified data model allows users to perform integrated analyses across plants, microbes, and their communities with a wide range of tools that interoperate across the tree of life, and to publish their data, methods, results, and thoughts in persistent, citable, executable, and reusable electronic narratives that allow scientists to build on the work of others.

5.2.1 Discussion Summary

- The DOE KBase is a software and data platform used to predict and design biological functions:
  - KBase integrates data, tools, and their associated interfaces into one environment, so users do not need to access them from numerous sources.
  - The open KBase platform enables external developers to integrate their analysis tools, facilitating distribution, comparative tool analysis, and access.
  - KBase users can perform large-scale analyses and combine multiple lines of evidence to create biological models.
  - KBase is the first large-scale bioinformatics system that enables users to upload their own data, analyze it (along with collaborator and public data), build increasingly realistic models, and share and publish their workflows and conclusions.
- Underlying the KBase platform is a service-oriented architecture that runs across a distributed set of resources located at ANL and LBNL.
- The JGI is an important partner to KBase, and KBase users can use the JGI Portal to send datasets from the JGI directly into their workflows. This results in multiple terabytes of data transfer on a routine basis.
- KBase users can use Globus to transfer data from any data provider that makes data available via Globus.
- KBase users primarily interact via a web-based platform for supporting data analysis and data discovery. Typically, this involves using public datasets that are already available in the platform, by uploading new data into KBase (or the JGI). The data (public or private) can be analyzed using a variety of applications that are exposed in the Jupyter-based Narrative interface.
- KBase is upgrading system capacity to add GPU nodes and additional storage (e.g., the goal is to be 90% of capacity). Recent storage additions now have the project at 58% of capacity, and KBase foresees the need to repeat the addition of storage regularly. In addition to the LBNL and ANL hardware, KBase has a NERSC allocation for HPC workloads and to use the Spin service. This is used to host KBase services, and it can also be used to leverage resources on the Perlmutter HPC resource.
- KBase does not have any immediate plans for using cloud services for routine computations; the ANL, LBNL, and NERSC resources are sufficient for the current demand. The cost of cloud computing is too great for the minimal "burst" use cases that may arise from time to time.
- KBase is converting Rancher Compose- and Docker Compose-based services to be deployed in Kubernetes, and in a reproducible configurable way that allows for the deployment on any Kubernetes environment. Once completed, this will enable more flexibility in deployment location across the DOE and commercial cloud world.

- Expanding the KBase footprint at LBNL has been difficult due to facility constraints. As a result, KBase is migrating some deployments from LBNL to ANL, and will focus on using the LBNL resources to act as a mirror of the production site in the future.
- KBase is challenged by a lack of staff to assist with development and operations requirements: e.g., tasks that focus on maintenance, cybersecurity patching, and user support. Future DOE support platforms that simplify infrastructure support would result in more time to focus on project-affiliated work.

5.2.2 KBase Facility Profile

The DOE KBase is a software and data platform designed to meet the grand challenge of systems biology: predicting and designing biological function. KBase integrates data, tools, and their associated interfaces into one unified, scalable environment, so users do not need to access them from numerous sources or learn multiple systems in order to perform sophisticated systems biology analyses. KBase's open platform enables external developers to integrate their analysis tools, facilitating distribution, comparative tool analysis, and access to enterprise-class computing.

5.2.2.1 Science Background

Users can perform large-scale analyses and combine multiple lines of evidence to model plant and microbial physiology and community dynamics. KBase is the first large-scale bioinformatics system that enables users to upload their own data, analyze it (along with collaborator and public data), build increasingly realistic models, and share and publish their workflows and conclusions. KBase aims to provide a knowledgebase: an integrated environment where knowledge and insights are created and multiplied.

Underlying the KBase platform is a service-oriented architecture that runs across a distributed set of resources located at ANL and LBNL. Components include:

- **narrative.kbase.us** is hosted at ANL. This is a production system, and a small subset of production workers run on Cori at NERSC in addition to the majority of the resources that run at ANL
- **appdev.kbase.us** is hosted at LBNL but uses compute resources and authentication provided by ANL. Future incarnations of this environment will operate completely at ANL.
- **next.kbase.us** is a development environment at LBNL to test components before they become operational. This resource will be deprecated in favor of staging.kbase.us at ANL
- **ci.kbase.us** is also a development environment located at LBNL.

The KBase project has members from national laboratories and a number of universities. Collaboration exists with many universities. Users of the system span the entire globe and upload data via the KBase narrative UI. An important partner facility is the JGI, and the user can use the JGI Portal to send datasets from JGI into KBase. Users also have the ability to use Globus to transfer data from any data provider that makes their data available via Globus.

5.2.2.2 Collaborators

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
Argonne Laboratory	Primary	rsync	500 TB	Ad hoc	No	No
Berkeley Laboratory	Primary	rsync	50 TB	Ad hoc	No	No
JGI	Primary	JGI Portal Transfer	Many TB	Ad hoc	No	No
The public	Secondary	Web interface GitHub Private FTP Google Drive	N/A	Ad hoc	Yes, after analysis, sent back in web UI	Yes, slow network
Globus	Primary	Globus Transfer	N/A	Ad hoc	Yes, via Globus	No
NERSC	Primary	Compute workers at Cori to run specific HPC jobs transferred via HTTP	Varies GB - TB	Ad hoc	Yes, HTTP	No
DOE BER SFAs	Secondary	Web Interface	Varies GB-TB	Ad hoc	Yes, via download or Globus	No
NMDC	Secondary	Globus + KBase Upload Jobs	TB	Ad hoc	No	No

Table 5.2.1: KBase Collaboration Space

5.2.2.3 Instruments and Facilities

KBase provides:

CPUs	Memory	Temporary Storage
3360	47T	287TB

Table 5.2.2: KBase Worker node Profiles

Profile of Kernel-based Virtual Machine (KVM) nodes which run all VMs that host critical services

CPUs	Memory	Temporary Storage
1120	7T	33TB

Table 5.2.3: KBase KVM node Profiles

Profile of storage nodes (ElasticSearch, Mongo, Arango, Reference Data, Private Docker Registry)

CPUs	Memory	Temporary Storage
192	1.1T	798TB

Table 5.2.4: KBase Storage node Profiles

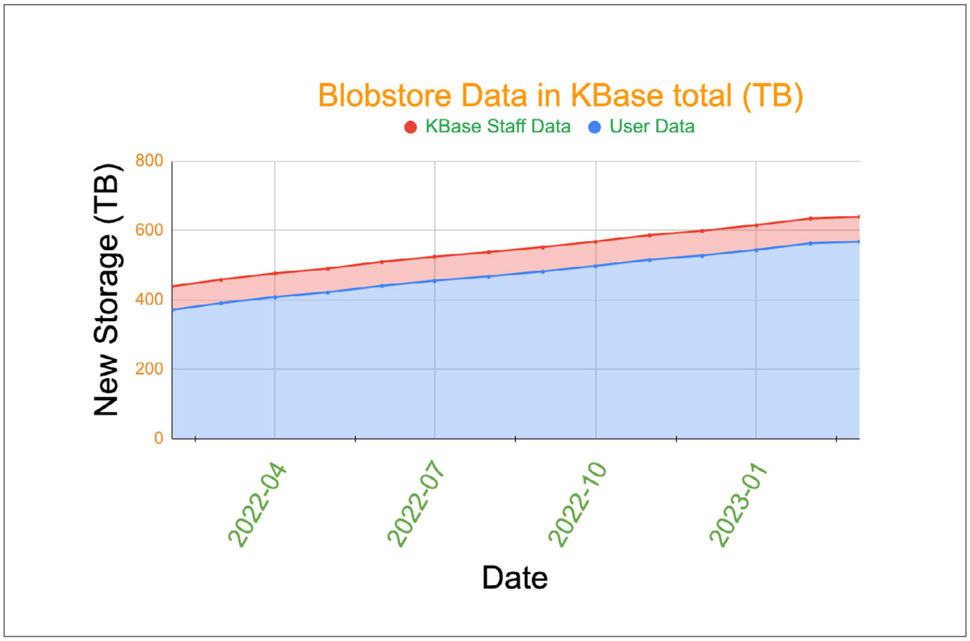

Figure 5.2.1: KBase storage use over time

Profile of object storage (Minio)

CPUs	Memory	Temporary Storage
192	1.1T	798TB

Table 5.2.4: KBase Storage node Profiles

In the current horizon, KBase plans to add GPU nodes and additional storage. KBase is currently in the works for planning future capacity needs based on usage trends and upcoming collaborations, such as with Genome Resolved Open Watershed (GROW), which is one of the largest data depositors. This is accomplished by monitoring usage metrics for job queues and by monitoring storage used across the systems. In the future, KBase would like to be able to do off-site replication of the largest storage repository (our object store)

In addition to LBNL and ANL hardware, KBase has NERSC allocation to use their Spin service to host services and can run “HPC” jobs on their Cori machines.

5.2.2.4 Generalized Process of Science

KBase is a web-based platform for supporting data analysis and data discovery for biological science. Users can access public datasets that are already available in the platform or they can upload their own data. The data, public or private, can be analyzed using a variety of applications that are exposed in the Jupyter-based Narrative interface. Users can share this data with their collaborators and can publish the results. This enables their data to be discoverable by other KBase users.

We expect this general pattern to continue for the next two to five years and beyond. In addition, KBase is working on a variety of novel capabilities geared towards new methods of discovery and predictions. These tools will enable users to discover relevant data in the platform to their own data or data of interests. New predictive tools will use a large range of public datasets in the platform to generate integrative models.

5.2.2.5 Remote Science Activities

KBase enables multifaceted remote science activities for its end users and third-party developers.

For developers, KBase provides a secondary environment for app development which runs at LBNL but also leverages compute resources and the app catalog from the main KBase platform. Developers must make their code available on GitHub, and make any required datasets open. When developers are ready to publish their new KBase App, they provide the KBase App Catalog with a link to their GitHub Repository and then their application Docker image is built by the App Catalog and then saved into the KBase Docker Registry. In addition, their publicly available reference data is downloaded via FTP or HTTP and saved into the KBase Reference Data repository, in which the data is then made available as a versioned resource for KBase Apps to use. Developers can then bump the version number and provide new reference data as it is made available, so the end users can use the most up-to-date data for their analysis. The reference data and application Docker images are kept as artifacts at KBase to be reused throughout the system.

Every user is provided with a data staging directory (Staging Area) which resides on the KBase Data Transfer Node. Users can leverage the JGI Data Portal to send biological datasets from the JGI data store to incorporate into their analysis. In addition, the Staging Area runs a Globus endpoint, and users can transfer data from any other Globus endpoint, such as their home institution, public repositories such as Phytozome, or even their own laptop. Once the data has been transferred to the staging area, it can then be imported into the KBase workspace and analyzed in a narrative using the applications and reference data that have been registered into the App Catalog. Once the user is satisfied with their analysis, they can share their narrative for private review by other collaborators or even make it available publicly for publication and sharing with the community.

5.2.2.6 Software Infrastructure

KBase is a form of data management and analysis platform in and of itself. The platform is built on various common third-party tools, as well as custom software. KBase provides Narratives, built on Jupyter Notebooks, where users can run KBase Apps to perform reproducible analyses. A complete list of apps provided by the KBase App Catalog[6]. These apps run on in-house and HPC partner compute resources using the HTCondor Job Scheduler.

KBase apps are created using the KBase SDK, a tool that allows in-house and third-party developers, such as Science Focus Area collaborators, to wrap any free, open-source biological analysis tools and bring them into the system. In addition to apps, KBase provides a range of "core" microservices written by KBase developers and a catalog of "dynamic" microservices written by both KBase and community developers. These services are used to programmatically access, import, manage, and analyze data on the system. Core services Docker images are built using GitHub Actions and stored in the GHCR.io Docker image repository. Dynamic services Docker images are built using the KBase Catalog and stored in a private Docker registry. The web services are used by the KBase website UI and by compute jobs.

These microservices interact with databases to store research data and metadata, which include MongoDB (NoSQL database), ArangoDB (graph-based multimodal database), ElasticSearch (search database), and MinIO (our s3-compatible object store). Additionally, messages are sent to message queues, which use Apache Kafka, to trigger events within the platform. Data imports are run through a KBase Staging and Upload system, which includes both web services and collections of apps used for importing and exporting data into and out of the KBase workspace and ObjectStore via the Staging Area and Globus.

All of the above infrastructure is managed by the KBase DevOps team using custom scripts, ansible, rancher, rancher2, and xcat.

6 https://kbase.us/applist/

Critical Tools for Internal Functional/Operation Processes

For internal tools and collaboration within the KBase project, KBase relies heavily on Slack, Zoom, and Google Docs. Miro Boards are also used for collaboration, brainstorming, and product designs. For Scrum processes, Retrium is used for retrospectives and Jira for Kanban boards. KBase heavily relies on GitHub/Gitlab for collaboration within and with 3rd-party developers. Lastly, a variety of custom in-house tools for data analysis are used that wrap a wide range of public, open-source bioinformatics applications and utilities.

We do not anticipate this to change significantly over the two to five years or beyond.

5.2.2.7 Network and Data Architecture

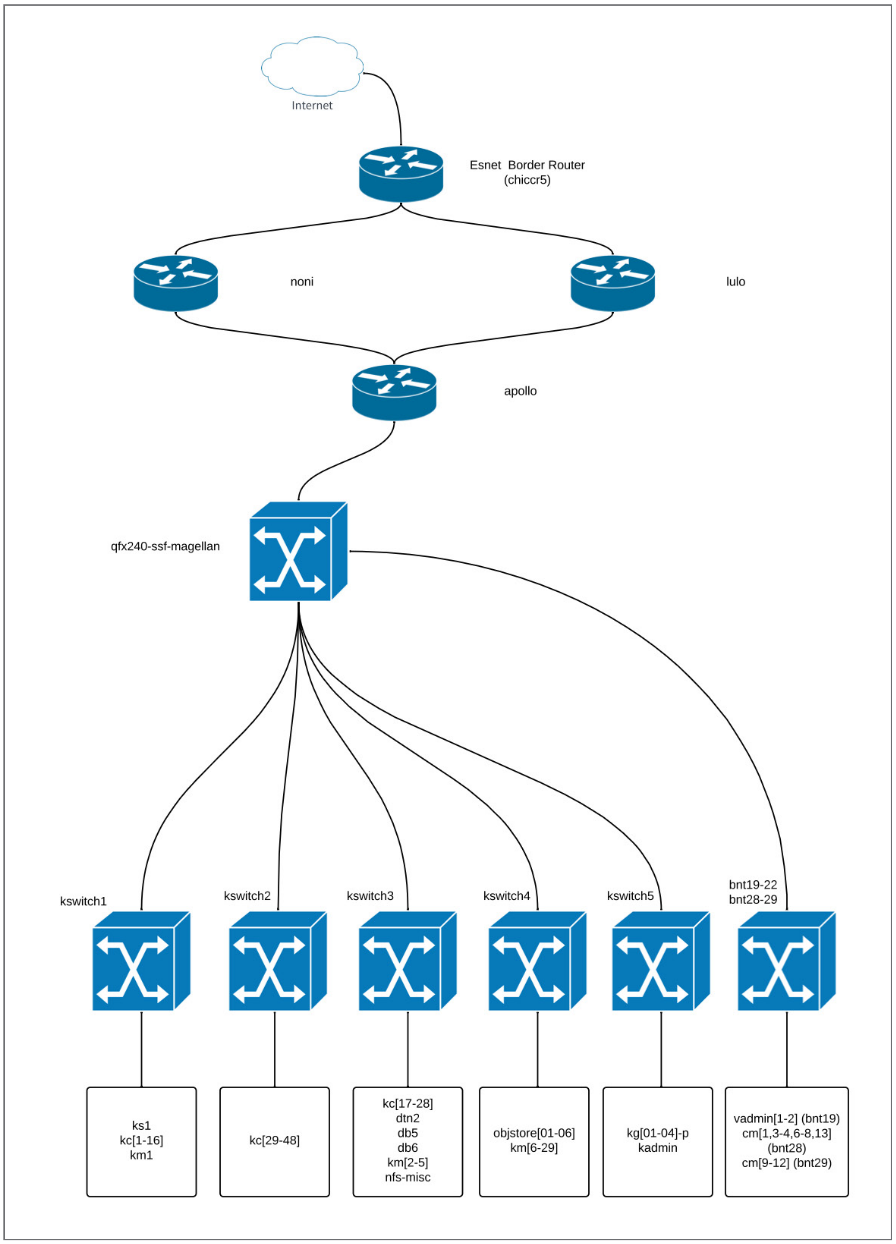

Figure 5.2.2: ANL Network Diagram

The ANL Network diagram gives an overview of the network infrastructure. The core compute, storage, and database services run on kswitches, which are Mellanox switches. Some of the admin and older compute nodes run on Blade Network Technologies (BNT) switches, which are rj45 copper switches.

Our Globus connection resides on the Data Transfer Node (DTN2) (kswitch3) which then sends data to the KBase workspace, on kswitch2, which then sends data to the minio object store, behind kswitch4.

Finally, each of these switches are behind a Juniper qfx5120-48y, which in turn is connected to the border routers.

There is no plan at this time to use the Science DMZ or perfSONAR. KBase uses Iperf3 to diagnose network issues.

We do not anticipate this to change significantly over the two to five years or beyond.

For LBNL:

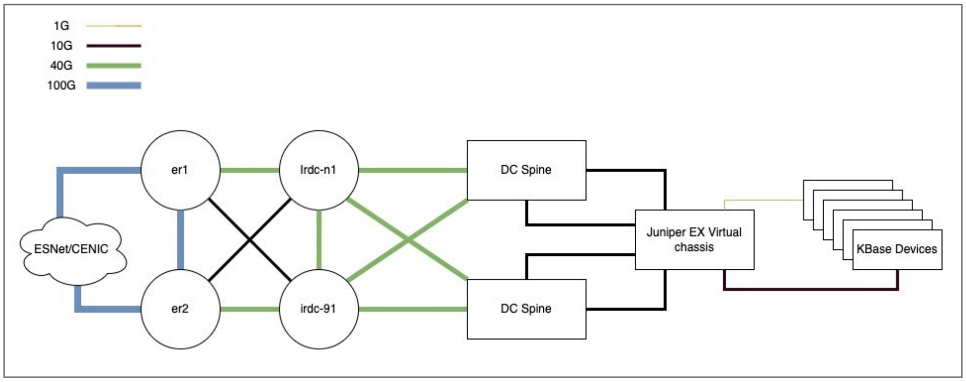

Figure 5.2.3: LBNL Network Diagram

Devices are attached to a set of Juniper switches through a combination of 1G and 10G connections. From there the switches have 40G total bandwidth up to the next level. From there is 160G total bandwidth to the routing layer and then 100G of total bandwidth to the border. From the border there are 2 x 100G connections to ESnet and 1x 100G + 1x10G connections to CENIC, all of which can transit to the Internet. The Primary uplink switch is a Mellanox.

5.2.2.8 Cloud Services

Provider	Current Usage and Plans
AWS	Domain name system (DNS) Hosting
GCE	API Keys Backups of data for certain databases PLAN: DNS Hosting , no ETA
CLOUDFLARE	SSL Provider + Security
Spin	Hosted environment to test deployment on environments. A target to achieve portability for the KBase system in unprivileged environments. MVP for subset of KBase Services is already hosted here

Table 5.2.6: KBase Cloud Usage

We currently do not have any immediate plans for using cloud services for routine computations, on-premises computing and a NERSC allocation exist that are sufficient for the current demand. In the past, KBase has made use of Cloud to handle brief bursts in demands, and this could occur in the future. In general, cloud usage has been minimized due to its higher cost.

Currently, KBase is working on converting rancher and Docker-compose based services to be deployed in Kubernetes, and in a reproducible configurable way that allows us to deploy on any Kubernetes environment, anywhere, whether on premises or in the cloud. In addition, playbooks are being created on how to migrate data stores when deploying new environments.

We do not anticipate this to change significantly over the next two to five years or beyond. However, KBase has been engaged in the ongoing IRI efforts and this is seen as a promising possibility for addressing the needs of KBase.

5.2.2.9 Data-Related Resource Constraints

Due to facility constraints at LBNL, it has been difficult to expand the footprint at LBNL. As a result, KBase is migrating some of the deployments from LBNL to ANL, so that it can focus on using the LBNL resources to act as a mirror of the production site in the future. In addition, KBase has been in the process of standing up a deployment of KBase in the Spin system at NERSC.

In addition, disaster recovery backups are being made to the public cloud but these cannot currently allow us to continue operations.

Many of these challenges are exacerbated by limited DevOps staffing which must focus on maintenance, cybersecurity patching, and user support, so capacity for project-based work is reduced. If a future IRI platform provided the infrastructure and capabilities to maintain distributed replicated data storage and services, this would greatly enhance KBase operations.

Due to increasing storage demands, KBase is refreshing storage in an effort to keep the used capacity at below 90%. KBase has just finished adding storage and is at 58% of capacity, and foresee the need to repeat the addition of storage as the need arises.

5.2.2.10 Outstanding Issues

Our limited DevOps staffing is primarily focused on maintenance, cybersecurity patching, and user support, so capacity for project-based work is reduced. KBase believes a future IRI platform could help address some of these challenges.

5.2.2.11 Facility Profile Contributors

KBase Representation

Boris Sadkhin, *ANL,* bsadkhin@anl.gov

Dan Klos, *ANL,* dklos@anl.gov

Adam Arkin, *LBNL,* aparkin@lbl.gov

Shane Canon, *LBNL,* SCanon@lbl.gov

Roy Kamimura, *LBNL,* royk@lbl.gov

ESCC Representation

Rune Stromsness, *LBNL,* rstrom@lbl.gov

Richard Simon, *LBNL,* rsimon@lbl.gov

Linda Winkler, *ANL,* winkler@mcs.anl.gov

Corey Hall, *ANL,* chall@anl.gov

5.3 NMDC

The NMDC[7] is a science gateway for Findable, Accessible, Interoperable and Reusable (FAIR) multiomics microbiome data that leverages existing resources, capabilities, and expertise across three DOE National Laboratory partners: LBNL, LANL, and PNNL. The vision of the NMDC is to drive a microbiome data-sharing network connecting data, people, and ideas to advance microbiome innovation and discovery.

5.3.1 Discussion Summary

- The NMDC leverages microbiome data science capabilities across three DOE national laboratories: LBNL, LANL, and PNNL, to deliver a science gateway for multiomics microbiome data. The NMDC Data Portal is a resource to discover data generated at the JGI and EMSL.
- The NMDC provides three defined products to support FAIR data:
  - The NMDC Submission Portal: a system designed to lower barriers to capture and adhere to community standards for sample contextual information (metadata) leveraging the Minimum Information about any (x) Sequence (MIxS) standards for environmental packages.
  - The NMDC Data Portal and public API: a system that facilitates seamless access for data discovery.
  - NMDC EDGE: a web application that provides community access to the NMDC data analysis workflows.
- The NMDC relies on existing core partnerships and on building out new partnerships:
  - The NMDC provides access to JGI data sources. This partnership supports gigabytes of data exchanged between the platforms.
  - Through partnership with the NMDC, multiomics data generated in EMSL for microbiome research is findable and accessible through the NMDC Data Portal. This partnership supports gigabytes of data exchanged between the platforms.
  - In the next two years, the NMDC will leverage a number of partnerships, with each resulting in gigabytes of data exchanged between the respective platforms:
    - KBase, to enable integration of data across these platforms and perform advanced analyses.
    - ESS-DIVE, to facilitate integration of data generated by the research community and archived in the ESS-DIVE storage platform.
    - NEON, to enable integration of data generated from NEON field sites.
- The NMDC has established mechanisms for the JGI and EMSL to share metadata with the NMDC core metadata storage system about microbiome studies and samples:
  - Within the next two years, a core goal for the NMDC is to automate data exchange between these facilities, to handle the steady streams of data from the JGI and EMSL, as well as external agency partners.
  - In two to five years, the NMDC will be managing PB of data with variable access patterns; e.g., some users will download, others will access data through partner platforms for analysis.

7 https://microbiomedata.org/

- Transferring multiomics data across the wide-area network can take several days using traditional web protocols. The NMDC will therefore enable high-bandwidth data transfer mechanisms via Globus, a secured high-performance data transfer service allowing users to transfer large amounts of data between systems within and across organizations.
- The NMDC can be viewed as a distributed data facility that combines JGI, EMSL, and NERSC resources. The NMDC data products are currently housed in multiple locations including EMSL and NERSC, based on the resource where the files were generated. The fully operational central metadata store maintains references to the locations of the data and access protocols, allowing the data to be retrieved when requested by users.
- The NMDC expects that the commercial cloud may be used for resilience for the front-end web service (search of metadata only), as well as to support "surges" in computing capacity. Since the NMDC does not have a "real-time" computing requirement, supporting the later use is not critical. The NMDC primarily utilizes resources at EMSL and NERSC to run its standardized workflows. These use an opensource workflow service and NMDC is planning to adopt JGI Analysis Workflow Service (JAWS) in the future. The NMDC is also exploring AWS infrastructure to run workflows.
- The NMDC does not anticipate any data-related resource constraints in the coming years. Risks are present in that datasets generated outside the DOE user facilities (the JGI and EMSL) may grow and outpace local storage.

5.3.2 NMDC Facility Profile

The NMDC was launched in July 2019 to work towards developing a science gateway for FAIR multiomics microbiome data through leveraging existing resources, capabilities, and expertise across DOE National Laboratory partners in collaboration with the microbiome research community. In March 2021, the team launched the NMDC Data Portal[8], a resource to find and access multiomics microbiome data generated at the flagship DOE user facilities, the JGI at LBNL and the EMSL at PNNL. Underpinning the data available to the research community within the NMDC Data Portal is a schema that captures the relationships between biological and environmental samples and their metadata, and the multiomics data objects generated by the production-quality data processing workflows developed at the DOE user facilities. The NMDC has also created an enabling environment for FAIR multiomics microbiome data by building partnerships across a variety of DOE projects, including working with DOE Science Focus Area (SFA) projects and complementary DOE resources such as ESS-DIVE and KBase (Figure 5.3.1). With this collaborative focus, the NMDC Pilot consolidated core capabilities in metadata standards for sample and environmental descriptors; standardized data analysis workflows; interfaces for sample metadata collection and microbiome data search and access; and a robust strategic engagement plan.

8 https://data.microbiomedata.org

5.3.2.1 Science Background

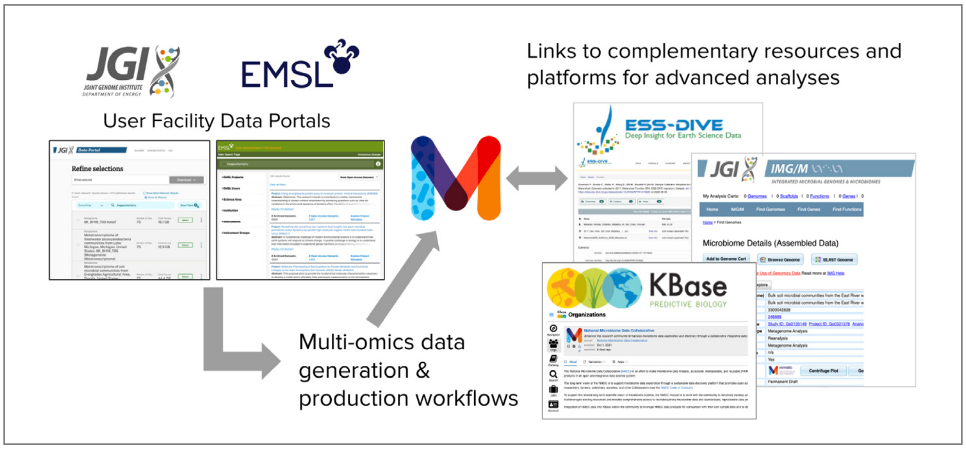

Figure 5.2.2: LBNL Network Diagram

The NMDC currently supports FAIR[9] data through three defined Product Initiatives:

- **The NMDC Submission Portal:** a system designed to lower barriers to capture and adhere to community standards for sample contextual information leveraging the MIxS standards developed by the Genomics Standards Consortium[10]. The flexible framework addresses the critical gap of collecting the necessary metadata describing each stage of the data cycle: a research study, biosamples, experimental procedures for data generation and computational processing of the data.
- **The NMDC Data Portal and public API:** a system that facilitates seamless access for data discovery. The Data Portal infrastructure relies on a distributed framework and sets the foundation for future development to address an immediate and pressing need for access to integrated multiomics studies.
- **NMDC EDGE:** a web application developed to provide broad community access to the NMDC data analysis workflows. The lightweight architecture enables installation and configuration in different computing environments, thereby addressing challenges in bioinformatics software portability.

9 Data Management Plan: https://dmphub.cdlib.org/dmps/doi:10.48321/D1NK60
10 https://github.com/GenomicsStandardsConsortium/mixs

5.3.2.2 Collaborators

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
JGI	Primary	Portals, Globus	O(10GB)	ad hoc	No	No
EMSL	Primary	Portals, Globus	O(1-10GB)	ad hoc	No	No
KBase	Secondary	API	O(1-10GB)	ad hoc	No	No
ESS-DIVE	N/A	other - link to NMDC data		ad hoc	No	No
NEON (NSF)	Primary	API	O(10-100GB)	ad hoc	No	Not sure
GeneLab (NASA)	Primary	API		ad hoc	No	Not sure
NCBI / EBI	Primary	Data transfer		ad hoc	No	No

Table 5.3.1: NMDC Collaboration Space

As a DOE data resource, NMDC has built robust collaborative partnerships and is actively expanding on those partnerships. Partners coordinate their activities related to data, infrastructure, and engagement with the research community to improve the user experience to find, access, and reuse microbiome data and complementary environmental metadata. These collaborations currently require transfer of datasets of O(1-10 GB) on an ad hoc basis (Table 5.3.1); data transfers are expected to grow in frequency over the next two to five years with data transfers approaching weekly occurrence.

5.3.2.2.1 JGI

JGI is a BER User Facility dedicated to energy and the environment, providing the research community with access to the most advanced genomic science capabilities in support of DOE's research mission. The facility provides integrated and state-of-the-art genomic technologies including HT sequencing, DNA synthesis, metabolomics and computational analysis that enable systems-based scientific approaches to these challenges. The program[11] at JGI generates metagenome and metatranscriptome data that are included in the NMDC database.

5.3.2.2.2 EMSL

EMSL is a BER User Facility that provides the research community with access to premier multimodal molecular science instruments, data analytics, production computing, and multiscale modeling to study biotic and abiotic processes in a biological or environmental systems context. The MS instruments[12] and facilities in EMSL generate metaproteomics, metabolomics and organic matter data that are included in the NMDC database.

11 https://jgi.doe.gov/our-science/science-programs/metagenomics/
12 https://www.emsl.pnnl.gov/science/instruments-resources#Omics_&_Mass_Spectrometry

5.3.2.2.3 KBase

KBase[13] is a software and data science platform that enables users to upload their data, analyze it alongside collaborator and public data, build models, and share and publish reproducible workflows and conclusions. The NMDC–KBase collaboration will streamline data flow between these facilities through the NMDC public APIs, client tools, and high-performance data transfer mechanisms, as well as existing interfaces provided by KBase.

5.3.2.2.4 ESS-DIVE

The Environmental System Science Data Infrastructure for a Virtual Ecosystem (ESS-DIVE[14]) is a data repository designed to collect, store, manage, and share Earth and environmental systems data created through research sponsored by DOE. The NMDC–ESS-DIVE collaboration will facilitate data integration by aligning metadata standards used across these systems, cross-referencing sample identifiers and long-term persistent identifiers (e.g., DOIs), and streamlining data sharing through APIs.

5.3.2.2.5 NEON

The National Ecological Observatory Network (NEON[15]) is a continental-scale observation facility operated by Battelle and designed to collect long-term open access ecological data to better understand how US ecosystems are changing. The NMDC–NEON collaboration will incorporate microbiome data generated across the NEON network into the NMDC infrastructure, to make it readily findable and accessible by the research community for integration with data generated at the DOE user facilities.

5.3.2.3 Instruments and Facilities

NMDC is a distributed data resource that makes microbiome data findable and accessible on the NMDC Data Portal including data generated through the laboratory resources of the user programs at JGI and EMSL, data generated by independent investigators and stored at KBase and ESS-DIVE, and data generated by additional national resources such as NEON. The computational infrastructure for NMDC is supported by, and distributed across, NERSC and EMSL.

13 https://www.kbase.us/
14 https://ess-dive.lbl.gov/
15 https://www.neonscience.org/

5.3.2.4 Generalized Process of Science

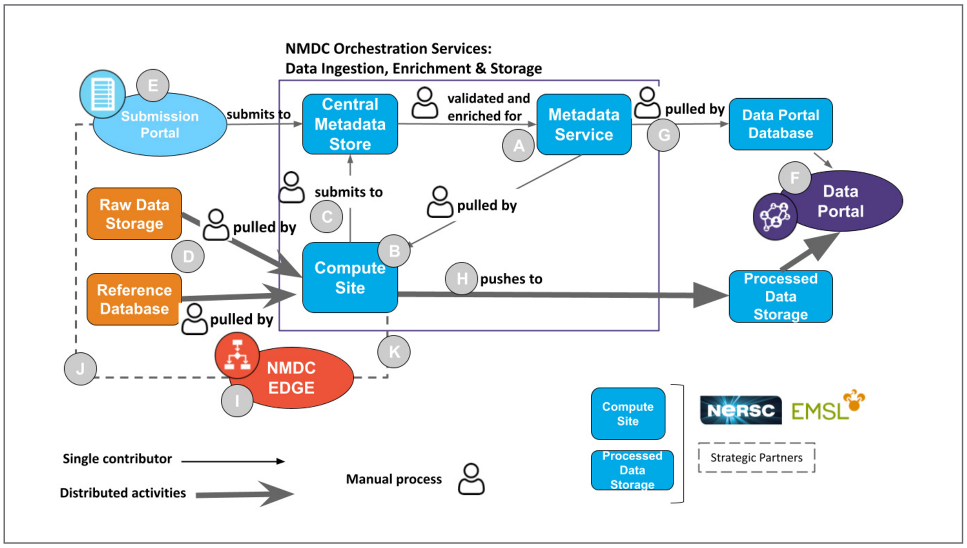

Figure 5.3.2: Simplified diagram of the NMDC infrastructure components.

Elements of the NMDC infrastructure, the communication between the elements and the data-sharing processes are illustrated in Figure 5.3.2. (A) The metadata service provides access to the central metadata store. (B) An NMDC compute site runs an instance of the NMDC services for workflow execution and metadata registration; current sites include EMSL and NERSC. (C) Workflows generate metadata that is submitted through a service and validated prior to storage. (D) Individuals, partners, or facilities provide the necessary raw data and reference databases for workflow execution; in future, NMDC will design an upload service to accommodate non-user facility data. (E) The Submission Portal serves as a main entry point for the external community to submit standardized metadata to the NMDC. (F) The Data Portal is the main access point for exploring NMDC data and metadata graphically. (G) The database that supports the Data Portal consumes data from the central metadata store instance to maintain a responsive Data Portal. (H) The data products generated by the NMDC workflows are housed at the compute sites (e.g., EMSL and NERSC); Globus is used to support high-performance data transfers where needed. (I) NMDC EDGE provides standalone access to the NMDC workflows for individuals and research teams. (J) In future, NMDC EDGE will, in concert with the Submission Portal, become an entry point for individuals to contribute both metadata and processed multiomics data. (K) In future, NMDC EDGE workflow execution will be integrated with the compute sites.

In the next two to five years, the focus will be on optimizing and automating several processes described above. Specific relevant activities include:

- **Automation and coordination with JGI & EMSL User Facilities:** The JGI and EMSL are key partners for the NMDC because they support the environmental research community, have established processes for collecting metadata from researchers, and execute many of the NMDC workflows as part of their data-generation process. The NMDC is collaborating closely with the user facility staff to incorporate the multiomics data generated from user proposals supported by the FICUS program. NMDC is enhancing the Data Submission Portal and associated APIs to facilitate sharing metadata from the respective JGI and EMSL data management systems, and automate sharing of the processed multiomics data.
- **Internal automation to scale pulling user facility processed omics data:** For JGI and EMSL users who want their data available in the NMDC Data Portal, NMDC plans to automate the processes to identify the required metadata in the NMDC metadata store, then query and pull

the associated processed data through the user facility's APIs. Reprocessing of raw data will only be triggered when non-standard workflows not typically run by JGI or EMSL are needed, such as the metagenome reads-based taxonomy analysis workflow. The combination of APIs and workflow triggers will enable seamless sharing of high-quality data from the JGI and EMSL.

- **High-performance Data Transfer:** Transferring multiple terabytes of omics data across the WAN can take several days using traditional web protocols. NMDC will therefore adopt Globus for secured high-performance data transfers. It can tune performance parameters, maintain security, monitor progress, validate correctness and resume transfers if there is an interruption. For NMDC, Globus will enable high-performance data transfer across DOE facilities, as well as between NMDC EDGE, other NMDC resources and external systems such as KBase. Efficient and automated data transfer will play an important role in streamlining pre- and post-workflow processing steps across facilities.

NMDC anticipates steady streams of data from JGI and EMSL, as well as external agency partners. In two to five years NMDC will be managing O(PB) of data with variable access patterns; e.g., some users will download, others will access data through partner platforms for analysis.

Globus will be the primary data transfer protocol between DOE sites, while HTTP-based tools will facilitate user-data downloads through the browser and the API. NMDC does not plan to support the rsync, Aspera, or FTP protocols.

5.3.2.5 Remote Science Activities

5.3.2.5.1 NMDC EDGE

NMDC EDGE is an online resource designed to provide easy access to NMDC bioinformatic workflows for microbiome researchers who have data, but without the computational resources to generate NMDC-compliant results. NMDC EDGE is currently running on DOE and NSF resources. Because NMDC maintains only downstream (processed) data products to populate the Data Portal, the primary data movement will be metadata and data products from those compute resources to the main NMDC repository.

5.3.2.5.2 JGI/EMSL

NMDC is currently a distributed data facility that combines JGI, EMSL and NERSC resources. The NMDC data products are currently housed in multiple locations including EMSL and NERSC, based on the resource where the files were generated. The fully operational central metadata store maintains references to the locations of the data and access protocols, allowing the data to be retrieved when requested by users. In order to scale to an ever-increasing number of studies and datasets, NMDC will further develop this central data store.

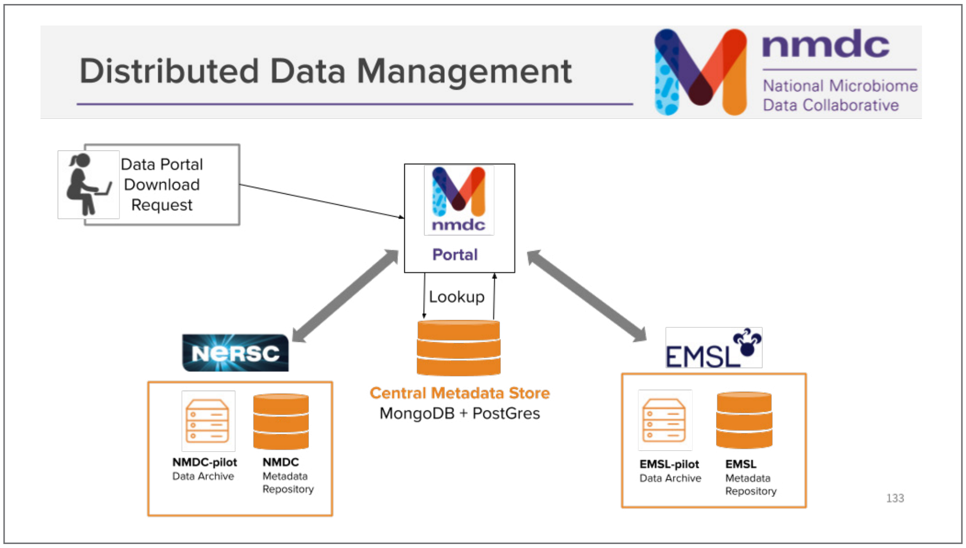

Figure 5.3.3: Simplified diagram of the NMDC infrastructure components.

5.3.2.5.3 External Partners:

Because of the intimate ties with the DOE user facilities, the NMDC has focused on obtaining and integrating cross-institute sample metadata acquisition for the NMDC Data Portal. However, the majority of microbiome data, including pertinent environmental microbiome data from international researchers, is not generated by DOE user facilities, and making this data findable and accessible through NMDC is key to the NMDC becoming an integral part of the microbiome research community as a whole. These external environmental microbiome data generators will need streamlined processes to provide sample and environment metadata, along with processed multiomics data through the NMDC workflows at their respective sites. NMDC expects to enable external researchers and research teams to use the Submission Portal for metadata submission and to establish strategic partners to serve as external computing sites (e.g., NSF's ACCESS [Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support]) supporting NMDC EDGE for multiomics data processing. The NMDC expects to automate data staging, submission, and all post-processing workflow outputs, along with high-performance data transfer mechanisms to facilitate data movement between external resources and the central NMDC infrastructure.

- Present–two years: A publicly accessible NMDC EDGE instance is hosted at the San Diego Supercomputing Center (SDSC) through NSF's Extreme Science and Engineering Discovery Environment (XSEDE) program; a mechanism for transfer of data products from NMDC EDGE to the NMDC data store is under development. Transfer of O(10GB) of data from NERSC (JGI data storage) and EMSL to the NMDC data store occurs on approximately a quarterly basis.
- Next two to five years: NMDC anticipates expanding the instances of NMDC EDGE to provide additional access points for external partner scientists to process their data and share data products with the NMDC data store. Transfer of O(10GB) of data from NERSC (JGI data storage) and EMSL to the NMDC data store will continue, at an expected increased frequency (at least monthly).
- Beyond five years: NMDC will use on-demand scalable computing services provided by either a federal government entity (NSF, DOE, etc.) or commercial entities such as AWS to support NMDC EDGE services and transfer of data products to the NMDC data store will be automated. Transfer of O(100GB) of data from NERSC (JGI data storage), EMSL, and external partners (NEON) to the NMDC data store will occur on an expected monthly basis.

5.3.2.6 Software Infrastructure

Sustainable infrastructure is a critical part of NMDC's long-term strategy. Automation and integration will ensure that NMDC can scale to enable growth in data, users, and science. Modularity will allow rapid innovation and adaptation as technology evolves.

NMDC has three primary software products, a workflow orchestration component, and a set of database and metadata services for the data management pipeline. This includes the following components:

- NMDC Data Portal to enable broad FAIR data access to the microbiome community
- NMDC EDGE to enable NMDC workflows for external users
- Submission Portal for standardized metadata submission management
- Workflow tools for data processing
- Internal "orchestration services" APIs for metadata management and database interactions.

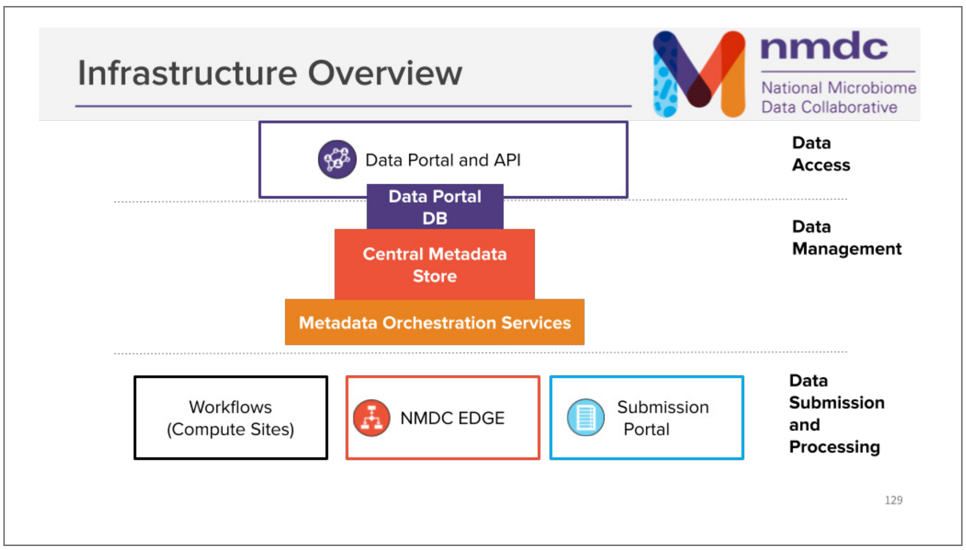

Figure 5.3.4: NMDC Infrastructure Overview

Our software is built on the following technologies:

- Data Portal, Submission Portal: Nginx, Vue.js, FastAPI, PostGres, Kubernetes
- NMDC Edge: Docker containers running on VMs (on SDSC/XSEDE)
- Workflows: Cromwell, Condor, JAWS, Docker
- Orchestration Services: MongoDB, FastAPI, Dagster, Kubernetes

In the next two to five years, NMDC will focus on automation of workflows, and integration of data processes across partner FICUS facilities (EMSL and JGI). This will enable data from workflows running at EMSL and JGI to automatically be integrated into NMDC. There are plans to use the JAWS framework to coordinate management of workflows and data movement across JGI - EMSL - NMDC (NERSC) facilities. JAWS uses Globus for data movement and Cromwell / HTC Condor for job management.

NMDC plans to support automated data submission and integration for DOE facilities, individual researchers, and external partners. This includes programmatic access through the public NMDC API both for data submissions and for search and discovery of data.

NMDC is also investing in improvements to database infrastructure for scalable ingest and search. On the Data Portal there are plans to implement advanced search, visualization, and analysis features.

5.3.2.7 Network and Data Architecture

NMDC uses the NERSC Spin platform to host NMDC web and database services, including the Data and Submission Portals, Databases (MongoDB and PostgreSQL), and API Services. NERSC Spin is a Kubernetes service for managed deployment of Docker containers.

The distributed data infrastructure for NMDC is supported by NERSC and EMSL, and relies on the ESnet high-performance network to support data sharing among the elements of NMDC's infrastructure (illustrated in Figure 5.3.2). Data products are stored on the NERSC Community FileSystem and EMSL Data Storage platforms (see Figure 5.3.3), and can be pulled into central NERSC resources from those systems. The NMDC EDGE platform runs on cluster resources provided by NSF (ACCESS[16], formerly XSEDE) at the SDSC, and in the near future (two to five years), NMDC will be able to pull data products from NMDC EDGE processing into the central NERSC resources. NMDC will use Globus to move multiomics data across these sites and will also support HTTP interfaces to pull the reference databases (from JGI and EMSL) used by the workflows into NMDC. Key data transfer paths will include:

1. EMSL <-> NERSC/JGI
2. NMDC EDGE (Xsede SDSC) <-> NERSC

In the long term (beyond five years), NMDC will pull data from a wider array of data sources (e.g, NCBI, EBI, NASA GeneLab, NFDI4Microbiota) and will consider deploying NMDC workflows on additional remote compute facilities (e.g., AWS, DOE Leadership Class Compute Centers, NSF ACCESS, USDA SCINet). This would include real-time data submission from remote facilities, and federated data across external sites and cloud resources.

5.3.2.8 Cloud Services

NMDC is exploring the commercial cloud for resilience for the front-end web service (metadata), as well as to support "surges" in computing capacity. NMDC does not have a "real-time" computing requirement, which does not make the later use critical to support. There are plans to look into NMDC data in the cloud, but cost considerations will play a factor here.

NMDC is currently exploring AWS infrastructure to run NMDC workflows via JAWS. It is noted that NMDC EDGE runs on external SDSC resources - NMDC effectively treats / manages these as cloud-based VMs and has the ability to migrate or move to other cloud resources as needed. Thus, to support NMDC EDGE services in the longer term (beyond five years), NMDC will explore AWS and/or ACCESS (the successor to XSEDE) to create a cloud-computing-based dynamic cluster that is similar to AWS ParallelCluster.

5.3.2.9 Data-Related Resource Constraints

NMDC uses data storage on the NERSC community filesystem and EMSL data systems, which are sufficient capacity for current data needs. Availability of data is contingent on those systems being available, thus additional data storage options (e.g., Cloud resources) may be considered if higher Service Level Agreements (SLAs) for data availability are desired.

In the next two to five years there is a small chance that data storage could become a challenge if more external datasets are ingested along with their sources. NMDC may need to consider federating with other facilities or cloud resources under this model.

NMDC's plans beyond five years will be heavily dependent on adoption outside of NMDC outside of DOE-SC. NMDC would plan to expand the federated data model to support this.

16 https://access-ci.org/

5.3.2.10 Outstanding Issues

None to report at this time.

5.3.2.11 Facility Profile Contributors

NMDC Representation

Emiley *Eloe-Fadrosh, LBNL,* eaeloefadrosh@lbl.gov

Kjiersten Fagnan, *LBNL,* kmfagnan@lbl.gov

Shreyas Cholia, *LBNL,* scholia@lbl.gov

Lee Ann McCue, *PNNL,* leeann.mccue@pnnl.gov

ESCC Representation

Rune Stromsness, *LBNL,* rstrom@lbl.gov

Richard Simon, *LBNL,* rsimon@lbl.gov

5.4 EMSL

The EMSL is a DOE SC, user facility sponsored by the Biological and Environmental Research (BER) program. At EMSL, scientists focus on fundamental biological and environmental research. EMSL collaborates on projects with researchers from academia, other government laboratories, and industry. Scientists around the world can partner with us to use world-class laboratory space, expertise, and equipment.

5.4.1 Discussion Summary

- EMSL is a BER user facility. EMSL:
  - Studies the role of molecular processes in controlling the function of biological and ecological systems across spatial and temporal scales.
  - Provides access to multimodal molecular science instruments, data analytics, production computing, and multiscale modeling.
- EMSL user research is conducted within the following three Science Areas:
  - Environmental Transformations and Interactions.
  - Functional and Systems Biology.
  - Computing, Analytics, and Modeling.
- EMSL engages in multiple institutional collaborations, most notably with the FICUS program for collaborative access to DOE user facilities:
  - The FICUS program has supported joint user projects between EMSL and the JGI since 2014.
  - FICUS has expanded to include the National Ecological Observatory Network, the Bio-SANS beamline at the Center for Structural Molecular Biology, the Advanced Photon Source (APS), the Advanced Light Source (ALS), and the ARM user facility.
- EMSL and the JGI actively collaborate, sharing resources for data storage and providing resiliency. This collaboration currently requires routine transfer of up to terabytes of data for analysis, and episodic transfer of up to a PB of data for long-term archival.
- As one of the distributed data sources for the NMDC, EMSL also routinely shares processed data and metadata (aggregate gigabytes), making that data findable and accessible through the NMDC data portal.
- EMSL does not currently use, or plan to use, academic or commercial cloud services for their primary data analysis or storage needs. However, EMSL users may use cloud resources if, or when, they choose to perform subsequent data analysis on the data generated in EMSL.
- EMSL expects to see increases in data transfers produced by other institutions for integration, analysis, and visualization in the next five years, and data products being copied to resources like KBase for analysis and to repositories like ESS-DIVE for storage.
- EMSL houses over 150 advanced and often one-of-a-kind instruments:
  - The raw data generated on EMSL instruments is uploaded to EMSL's data repository where it is stored in an active storage location during analysis, then moved to long-term storage in the EMSL data archive.
  - Data is typically processed prior to delivery to users. The data is often initially processed locally on the instrument control computers prior to transfer to a computer cluster or workstation for more intensive data analysis, and for delivery to users the fully processed

data may take a variety of forms, from structured spreadsheets, to community-defined data formats, to image files.

  - Raw data and processed data are available to users when they authenticate through the EMSL user portal. Small data files (tens to hundreds of kilobytes) are also often emailed to users, and large data files (up to terabytes) are typically delivered using either https transfers or the Globus service.

- The next five years in EMSL will see growing data volumes due to improvements in instrument resolution, and large increases in the rate of data generation due to laboratory automation. Automation efforts will focus on mass spectrometry workflows. EMSL expects to provide variable autonomous operations in a multimillion-dollar Microbial Molecular Phenotyping (M2P) facility, for which DOE has approved the Mission Need (CD-0).
  - With automation will come the need for an order of magnitude or more increase in data storage, analytics, on-the-fly analysis, modeling, and computation. The EMSL data archive is expected to exceed 500 PB within the next five years to accommodate the increase in data generation.
- Data reuse by the community is expected to increase as the value of the accumulated data improves from multimodal analyses:
  - Instruments will increase in size and capability and could yield up to multiple terabytes per day.
  - Automation pipelines will be established, with expected data generation to increase to tens to hundreds of terabytes per day, and data-archive capacity is expected to grow to 500 PB.
  - Continued growth of the FICUS program will increase demand for transfer of data between facilities.
  - In the future, the M2P Facility will be in operation, increasing data generation to petabytes per day.
- The software used to manage data generated in EMSL has been custom designed to provide the staff who operate instruments with the tools needed to transfer instrument data and the associated user project metadata to EMSL's data archive.
- EMSL operates four data transfer nodes (DTNs) for Globus data transfers. The DTNs are located on a Science DMZ per ESnet best practices and provide access to data on EMSL's computing systems. The DTNs currently each provide 100 Gbit per second bandwidth and collectively transfer 250–350 TB per month. With growth in the FICUS program and additional collaborations with DOE facilities (described previously), EMSL expects data transfers to increase over the next five years to 500+ TB per month. EMSL plans an upgrade to 400 Gbit per second technology for the DTNs when ESnet6 becomes available.
- It is recommended that EMSL experiment with the Globus Modern Research Data Portal (MRDP) for future portal work.
- Anticipated network or data constraints are currently internal to EMSL and PNNL. Specifically, the Science DMZ connection at 100 Gbit per second is the fastest part of the entire EMSL network. This interface with ESnet will be a bottleneck only in the sense that the current DTN servers do not individually have fast enough buses for full 400 Gbit per second when that bandwidth is available. That said, the aggregate bandwidth of the four existing DTNs is 400 Gbit per second, and it is within EMSL's ability and discretion to upgrade these DTNs to a PCIe gen 5-based server when those enter general availability.

5.4.2 EMSL Facility Profile

To create a secure bioeconomy and a predictive understanding of the Earth system, EMSL's vision is for a research community empowered to study the role of molecular processes in controlling the function of biological and ecological systems across spatial and temporal scales and to enable a predictive understanding of the living Earth system. As a BER User Facility, EMSL contributes to this future state through its mission to provide access to premier multimodal molecular science instruments, data analytics, production computing, and multiscale modeling to enable researchers to study biotic and abiotic processes and understand their function in a systems context for energy and environmental security and infrastructure resilience. Engaging and empowering the user community is a critical element of EMSL's strategy to deliver on the mission and vision.

5.4.2.1 Science Background

User research is conducted within the following three **Science Areas**:

- **Environmental Transformations and Interactions** focuses on the mechanistic and predictive understanding of environmental (physiochemical, hydrological, biogeochemical), microbial, plant and ecological processes in above and belowground ecosystems, the atmosphere, and their interfaces. EMSL provides the experimental, computational, and simulation expertise to investigate and model cycling, transformation, and transport of critical nutrients, elements, contaminants, and atmospheric aerosols.
- **Functional and Systems Biology** focuses on elucidating and harnessing the biochemical pathways that connect gene functions to complex phenotypic responses through a deep understanding of interactions within cells, among cells in communities, and between cellular membrane surfaces and their environment for microbes (archaea, bacteria, algae), fungi, and plants.
- **Computing, Analytics, and Modeling** focuses on using state-of-the-art experimental data to develop a predictive understanding of biological and environmental systems through advanced data analytics, visualization, and computational modeling and simulation.

Further information on EMSL's Science Areas, User Program, current and past user research projects, instruments and resources, Data Management Policy and Strategic Plan can be found on EMSL's website[17].

5.4.2.2 Collaborators

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
PIs & staff at US academic laboratories (300-400 users annually)	EMSL is primary source of generation and storage of data; users may store secondary copies	email, EMSL user portal, SFTP, Globus	2 MB – 2 GB	Ad hoc	N	N/A
PIs & staff at DOE national laboratories (70-80 users annually)	EMSL is primary source of generation and storage of data; users may store secondary copies	email, EMSL user portal, SFTP, Globus	2 MB – 2 GB	Ad hoc	N	N/A

17 https://www.emsl.pnnl.gov/

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe “other”)	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
PIs and staff at Foreign academic laboratories (40-60 users annually)	EMSL is primary source of generation and storage of data; users may store secondary copies	Email, EMSL user portal, SFTP, Globus	2 MB – 2 GB	Ad hoc	N	N/A
JGI; Berkeley, CA	FICUS collaboration: Each institute stores primary copies of their data	Globus	1 GB – 1 TB	Ad hoc	N	N/A
NEON; multiple field sites	FICUS collaboration: Each institute stores primary copies of their data	Emerging collaboration; https and/or SFTP likely method	1 – 10 GB	Ad hoc	N	N/A
Bio-SANS at the Center for Structural Molecular Biology; Oak Ridge, TN	FICUS collaboration: Each institute stores primary copies of their data	Emerging collaboration; https and/or SFTP likely method	1 – 10 GB	Ad hoc	N	N/A
Advanced Photon Source; Lemont, IL	FICUS collaboration: Each institute stores primary copies of their data	Emerging collaboration; https and/or SFTP likely method	1 – 10 GB	Ad hoc	N	N/A
Advanced Light Source; Berkeley, CA	FICUS collaboration: Each institute stores primary copies of their data	Emerging collaboration; https and/or SFTP likely method	1 – 10 GB	Ad hoc	N	N/A
ARM; Oak Ridge, TN	FICUS collaboration: Each institute stores primary copies of their data	Emerging collaboration; https and/or SFTP likely method	1 – 10GB	Ad hoc	N	N/A
NMDC; Berkeley, CA	EMSL stores primary copy of data; partner stores secondary copy	Https data transfer	0.1 MB – 2 MB	Ad hoc	N	N/A
ESS-DIVE; Berkeley, CA	EMSL stores primary copy of data; partner stores secondary copy	Https data transfer	2 MB – 2 GB	Ad hoc	N	N/A
KBase; Berkeley, CA and Lemont, IL	EMSL stores primary copy of data; partner stores secondary copy	Emerging collaboration; Globus, https, SFTP likely methods	1 – 10 GB	Ad hoc	N	N/A
Pacific Northwest Cryo-EM Center	Oregon Health & Science University is primary source of data generation and the EMSL archive provides data processing and storage; users may store secondary copies	Globus	5 – 75TB	Daily	N	N/A

Table 5.4.1: EMSL Collaboration Space

For the broad scientific community (researchers at US academic laboratories, DOE national laboratories and Foreign academic laboratories), EMSL user facility access is determined by a competitive peer-review proposal process. Upon proposal acceptance, the laboratory instruments and computational resources in EMSL are available to users without regard to location or institutional affiliation. Individual principal investigators (PIs) and research teams who use EMSL resources are expected to publish research results in the open scientific literature.

EMSL also engages in multiple institutional collaborations, most notably with the growing FICUS program for collaborative access to DOE user facilities. The FICUS program[18] has supported joint user projects between EMSL and the JGI[19] since 2014, and has since expanded to include NEON[20] , the Bio-SANS[21] beamline at the Center for Structural Molecular Biology[22], the APS[23], the ALS[24], and the ARM[25] user facility.

Computational staff at EMSL and JGI also actively collaborate, sharing resources for data storage and providing resiliency for JGI's high-performance computing needs. This collaboration currently requires routine transfer of up to terabytes of data for analysis on EMSL's Tahoma[26] resource, and episodic transfer of up to a PB of data for long-term archival. As one of the distributed data sources for the NMDC[27], EMSL also routinely shares processed data and metadata (aggregate gigabytes), making those data findable and accessible through the NMDC data portal[28].

We expect increases in the next five years of data produced by other institutions being moved to EMSL for integration, analysis, and visualization, and data products being copied to resources like KBase[29] for analysis and to repositories like ESS-DIVE[30] for storage.

5.4.2.3 Instruments and Facilities

EMSL houses over 150 advanced and often one-of-a-kind instruments operated and maintained by expert EMSL staff. Many instruments are housed in highly specialized laboratory spaces, including: a state-of-the-art computing space for Tahoma, EMSL's 0.93 petaFLOPS hybrid architecture high-performance computer (HPC); the Quiet Wing, which features eight acoustically, vibrationally, and electromagnetically shielded bays housing advanced electron microscopes (e.g., a Krios cryo-EM system and a helium ion microscope); a plant sciences lab with controlled growth chambers, phytotrons, and various root and rhizosphere imaging capabilities; and a highly modified lab space for the 21 tesla Fourier transform ion cyclotron resonance mass spectrometer (21T FTICR-MS), one of only two in the world.

EMSL's diverse instruments and resources are grouped into the following categories:

- Aerosol Characterization
- Analytical
- Biological Sample and Cell Separations
- Chemical Imaging
- Flow & Transport

---

18 https://www.emsl.pnnl.gov/basic/ficus-program-information/1872
19 https://jgi.doe.gov/
20 https://www.neonscience.org/about
21 https://www.ornl.gov/content/bio-sans-capabilities
22 https://www.ornl.gov/facility/csmb
23 https://www.aps.anl.gov/
24 https://als.lbl.gov/
25 https://www.arm.gov/
26 https://www.emsl.pnnl.gov/MSC/UserGuide/tahoma/tahoma_overview.html
27 https://microbiomedata.org/
28 https://data.microbiomedata.org/
29 https://www.kbase.us/
30 https://ess-dive.lbl.gov/

- Plant Growth & Soil Incubation
- Omics & MS
- NMR & EPR
- Optical Microscopes
- Sequencers
- Structural Tomography & Topography
- High-Performance Computing (HPC) and Visualization
- Synthetic Surfaces

Information about the specific instruments in these groups is found on the EMSL web page[31].

Users work with EMSL staff members to identify the instruments and resources appropriate to their research needs, and to determine the best access mechanism: (1) the user physically on-site working directly with EMSL staff, (2) the user working at their home institution and delivering samples to EMSL for analysis, (3) the user working remotely using EMSL computational resources directly, or (4) the user accessing open, public data for reuse.

The raw data generated on EMSL instruments is uploaded to EMSL's data repository (NEXUS) where it is stored in an active storage location during analysis, then moved to long-term storage in the EMSL data archive (Aurora). Data are typically processed prior to delivery to users and may take a variety of forms, from structured spreadsheets to community-defined data formats (e.g., mzML) to image files (e.g., tiff). Raw and processed data are available to the users when they authenticate through the EMSL user portal[32]. Small data files (10's to 100's of KB) are also often emailed to users, and large data files (up to terabytes) are typically delivered using either https transfers or the Globus[33] service. In accordance with EMSL's data management policy[34], data access by the mechanisms described above is restricted to the users' research team for a brief period, after which data is made open to the broad scientific community through EMSL's Data Portal[35].

The next five years in EMSL will see growing data volumes due to improvements in instrument resolution, and large increases in the rate of data generation due to laboratory automation. EMSL is pursuing both traditional fixed (HT) and variable autonomous operations, to support and amplify the complexity of science performed for BER researchers. Automation efforts within the next 2 years will focus on MS workflows for organic matter and soils analyses, to greatly enhance the throughput for highly subscribed but largely standardized analyses with fixed automation. Within the next five years, EMSL expects to provide variable autonomous operations in a multimillion dollar Microbial Molecular Phenotyping (M2P) Facility, for which DOE has approved the Mission Need (CD-0).

With these plans for automation will come the need for an order of magnitude or more increase in data storage, analytics, on-the-fly analysis, modeling, and computation. The Aurora data archive currently houses over 33 PB of data (of 75+ PB capacity), and is expected to exceed 500 PB within the next five years to accommodate the increase in data generation. Planned EMSL infrastructure modifications will modernize power and cooling to the computing facilities, provide 100 Gbit per second internal bandwidth to support data movement from instruments to storage locations, and upgrade the ESnet connection when ESnet6 becomes available (400 Gbit per second) to deliver data to users.

31 https://www.emsl.pnnl.gov/science/instruments-resources
32 https://nexus.emsl.pnnl.gov/Portal/
33 https://www.globus.org/
34 https://www.emsl.pnnl.gov/basic/data-management-policy/1243
35 https://search.emsl.pnnl.gov/

EMSL's Data Portal is also being updated and redesigned to accommodate the expected increases in data generation and to improve data access. Data reuse by the community is expected to increase as the value of the accumulated data improves from multimodal analyses, FAIR compliance, and expanding breadth that includes molecular, cellular, soil, sediment, and aerosol samples, across ecosystem, regional, and continental scales.

- Present: High-resolution mass spectrometers and advanced imaging platforms are driving the most data growth in EMSL, yielding up to multiple terabytes per day. During the next 24 months, this rate of data generation is expected to continue, and the Aurora data-archive capacity is expected to grow from 75+ PB to 100 PB.
- Next two to five years: EMSL's fixed automation pipelines will be established, with expected data generation to increase to 10s to 100s of terabytes per day and the Aurora data-archive capacity is expected to grow to 500 PB. Continued growth of the FICUS program (see Section 5.4.2.2. Collaborators) will increase demand for transfer of data between facilities. Established metadata requirements will align EMSL data with FAIR principles and create demand for accessibility and ease of data transfer for data reuse.
- Beyond five years: The M2P Facility will be in operation, increasing data generation to petabytes per day.

5.4.2.4 Generalized Process of Science

Researchers in the community with an approved EMSL user project typically work with project managers and scientific staff in EMSL to determine which instruments and measurements are needed to address their research questions. They then deliver physical samples to the facility, where EMSL staff members prepare the samples, make the measurements, process the data and return results and data to the research team.

The data generated by instruments in EMSL are often initially processed locally on the instrument control computers prior to transfer to a computer cluster or workstation for more intensive data analysis. The software and methods used vary widely for the 150+ major instruments in EMSL. The data delivered to the user are typically the fully processed and analyzed results in the form of flat text files, structured spreadsheets or images.

5.4.2.5 Remote Science Activities

While EMSL does not remotely operate instruments or resources outside of the facility, users may obtain remote access to resources within the facility (e.g., the high-performance computer Tahoma) to perform user research.

In addition, EMSL's engagement in the user facility collaboration program, FICUS[36], is growing (described in Section 5.4.2.2), and will require network infrastructure to transfer data between facilities.

- Present: Approximately quarterly transfer of up to 100 TB of genomic data from JGI to incorporate into the analysis of proteomic data in EMSL.
- Next two to five years: With expansion of the FICUS program to engage additional facilities, capability is developed to transfer terabytes of data from multiple facilities to EMSL.
- Beyond five years: With a robust FICUS program, transfer of up to 100 TB of data from those facilities to EMSL becomes routine, to support processing, data integration and modeling using EMSL computing resources

36 https://www.emsl.pnnl.gov/basic/ficus-program-information/1872

5.4.2.6 Software Infrastructure

The software used to manage data generated in EMSL has been custom designed to provide the staff who operate instruments with the tools needed to transfer instrument data and the associated user project metadata to EMSL's data archive. The software also supports the search and retrieval of instrument data for analysis and reuse by enabling the selection of data to create a "cart" and subsequent download by https transfer.

5.4.2.7 Network and Data Architecture

EMSL's internal network core is built around a fiber infrastructure that provides connectivity to standard 8-port network switches in laboratory spaces to attach instruments and computational resources directly to the EMSL core network. Isolated instrumentation networks use the building fiber to interconnect lab spaces throughout the EMSL building, and in some cases extending into other buildings. The EMSL core network has multiple connections into the PNNL core network through redundant fiber paths. The existing fiber plant has the potential to run 10 Gbit per second with updated switches and optics, though little of it has been so upgraded. EMSL's HPC lab has extensive 10 Gbit per second networking with a few 100 Gbit per second links.

EMSL operates four DTNs for Globus data transfers. The DTNs are located on a Science DMZ per ESnet best practices and provide access to data on EMSL's computing systems. The DTNs currently each provide 100 Gbit per second bandwidth and collectively transfer 250 – 350 terabytes per month. With growth in the FICUS program and additional collaborations with DOE facilities (described above), EMSL expects data transfers to increase over the next five years to 500+ terabytes per month. EMSL plans an upgrade to 400 Gbit per second technology for the DTNs when ESnet6 becomes available.

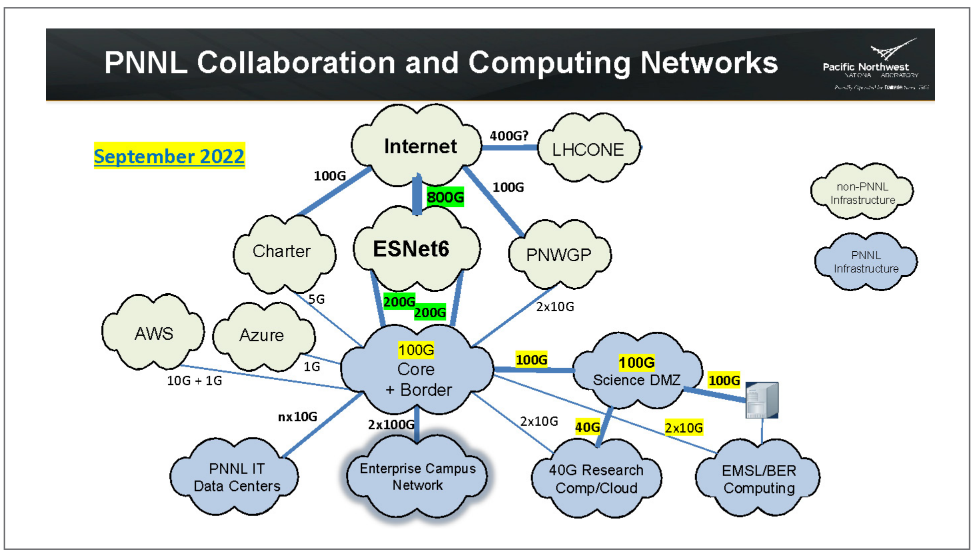

Figure 5.4.1: EMSL compute resources are located in the "EMSL/BER Computing" network at lower right in the diagram. The Globus DTNs connect at 100 gigabit per second each to the Science DMZ network.

5.4.2.8 Cloud Services

EMSL does not currently use, or plan to use, academic or commercial cloud services for their primary data analysis or storage needs. However, EMSL users may use cloud resources if, or when, they choose to perform subsequent data analysis on the data generated in EMSL.

5.4.2.9 Data-Related Resource Constraints

Anticipated network or data constraints are currently internal to EMSL and PNNL. Specifically, the Science DMZ connection (Figure 5.4.1) at 100 Gbit per second is the fastest part of the entire EMSL network. This interface with ESnet will be a bottleneck only in the sense that the current DTN servers do not individually have fast enough buses for full 400 Gbit per second when that bandwidth is available. That said, the aggregate bandwidth of the four existing DTNs is 400 Gbit per second and it is within EMSL's ability and discretion to upgrade these DTNs to a PCIe gen 5-based server when those enter general availability.

5.4.2.10 Outstanding Issues

None to report at this time.

5.4.2.11 Facility Profile Contributors

EMSL Representation

Lee Ann McCue, *PNNL,* leeann.mccue@pnnl.gov

David Cowley, *PNNL,* david.cowley@pnnl.gov

ESCC Representation

Surya Singh, *PNNL,* surya.singh@pnnl.gov

Lewis Allen, *PNNL,* lewis.allen@pnnl.gov

5.5 ARM

The ARM user facility is a multi-laboratory, US DOE scientific user facility, and a key contributor to national and international climate research efforts.

5.5.1 Discussion Summary

- The ARM user facility is a multi-laboratory effort and a key contributor to national and international climate research efforts:
  - The ARM facility is managed and operated by nine DOE laboratories (Argonne, Brookhaven, Lawrence Berkeley, Lawrence Livermore, Los Alamos, Oak Ridge, Pacific Northwest, National Renewable Energy, and Sandia) and actively serves approximately 1,000 science users each year.
  - ARM operates over 450 instruments to provide measurements of clouds, aerosols, precipitation, the surface energy balance, and the background atmospheric state.
  - ARM data is currently collected from three atmospheric observatories—Southern Great Plains (Oklahoma), North Slope of Alaska, and Eastern North Atlantic (the Azores).
  - There are three ARM mobile facilities. Once is currently deployed in La Jolla, CA, another will be deployed in northwest Alabama in the Bankhead National Forest in mid-2023, and another will be deployed in Cape Grim, Australia in early 2024.
  - There are several ARM aerial capabilities: a Challenger 850 regional jet, mid-sized unmanned aerial system (UAS), and tethered balloon systems (TBS).
  - Data from these efforts is available through the ARM Data Center (ADC) at ORNL.
  - ARM personnel manage the collection, processing, quality assessment, archiving, and distribution of data from these instruments, which currently represent approximately 2,500 distinct data streams.
- ARM's vision is to:
  - Provide comprehensive and impactful field measurements to support scientific advancement of atmospheric process understanding.
  - Achieve the maximum scientific impact of ARM measurements through increased engagement with observational data by ARM staff, including the application of advanced data analytical techniques.
  - Enable advanced data analytics and community use of complex ARM datasets through the advancement of computing infrastructure and data analysis.
  - Accelerate and amplify the impact of ARM measurements on Earth system models by exploiting ARM and earth system model frameworks to facilitate the application of ARM data to model development.
- Approximately 100 types of instruments are deployed at ARM sites. Major instruments include:
  - Profiling and scanning radars ranging in frequency from 915 MHz to 94 GHz to sample cloud and precipitation droplets.
  - A variety of lidars designed to measure profiles of temperature, humidity, water vapor, air motion, and aerosol properties.
  - In situ aerosol particle analyzers for studying physical and chemical properties of aerosols.
  - Wide-ranging in situ probes deployed on aircraft and tethered balloons to observe properties of clouds, aerosols, and the background atmospheric state.

  - ARM instruments that produce approximately 2 TB of raw data daily, and the resulting analysis can be more than 6 TB.
- ARM is exploring the application of new measurement technologies, such as phased array radar, differential absorption lidar, and an increasing number of miniaturized sensors for deployment on UAS and TBS.
- Many ARM users download a single dataset or a few datasets to study a particular phenomenon, evaluate other instruments (e.g., from a satellite), or evaluate or develop components of an Earth system model. These downloads can be gigabytes to terabytes in scale. Power users may download significantly more in bulk when performing larger studies.
- The ARM ADC is responsible for:
  - Archival storage for over 11,000 data products with a total repository holding of over 3.5 PB of data that dates to 1992. These include data from instruments, value-added products, model outputs, field campaigns, and PI-contributed data.
  - Data access and delivery options, which include THREDDS/OpenDAP, GlobusOnline, near real-time data access API, automated data access via web services, advanced visualizations, and a big data analysis platform.
- ARM data is collected from instruments deployed in remote locations, and the data is sent to ORNL via a secured and encrypted transfer method. Data is deep archived at the ORNL mass data storage system (HPSS), and an additional off-site data backup is done using the ANL's HPSS.
- Automated processes are used in continuously monitoring ARM data and metadata. These automated workflows assess the completeness of data and metadata. The remote instruments go through a detailed semi-automated quality analysis before users can access the data. If a data quality issue is identified, a data quality report (DQR) gets created by the reviewer.
- ARM has tiered storage and processing requirements for its data operations. The ADC uses a hybrid cyber infrastructure that includes a private cloud infrastructure operated within ORNL for data flow operations and large-scale data analysis, as well as commercial cloud for processes and platform-as-a-service systems that do not require customized data resources. ARM leverages the ORNL Leadership Computing Facility (OLCF) to maintain and operate computing clusters for large-scale data processing, data analysis, and high-resolution modeling.
- The ARM ADC will continue to explore and adopt the use of commercial cloud services for nonroutine processes, operational support tools, platform-as-a-service applications, and cloud-optimized data services for cloud-based user-developed processes.
- ARM is striving to develop a holistic and internally consistent description of the atmosphere. A goal in the next two to five years will be to develop more expanded datasets combining observations and model output to provide detailed descriptions of the atmosphere that are constrained by observations. Such integrated datasets will inevitably have large volumes, as they represent many variables in three-dimensional space and time, much like model output. Working with these comprehensive datasets will likely require more remote access, again similar to working with a large model or satellite data, in contrast with the more common mode of investigators downloading a few specific data streams and parameters today.
- All ARM measurements are remote; typically, few technical ARM staff are located at ARM observatories. This makes it challenging to modify the operation of instruments remotely, and to configure instruments to change their mode of operation in response to the environment.

- In the future, it will be highly desirable to operate UAS and/or TBS in an automated or semi-automated way such that ARM could obtain arrays of vertical profiles around an ARM observatory.
- The ADC archives high-resolution data bundles from LASSO model simulations over several ARM observatories:
  - These large eddy simulation (LES) model outputs are summarized and packaged within the ADC at ORNL to provide a concise representation of the LASSO simulations for sending to data users at other institutions.
  - These simulations represent a significant fraction of the data at the ADC and are expected to soon grow to several petabytes. When large volumes of LES output need to be transferred to the user's computing resources, ARM anticipates that many of these data transfers will use the ESnet infrastructure.
  - There is an opportunity to provide next-generation data connectivity for large data transfers between laboratory computing facilities to facilitate distributed data analysis.
  - Scientists using data analysis capabilities and visualization tools deployed in other leadership computing facilities would benefit by having seamless access to the full-resolution LASSO bundles archived at ORNL
- Connectivity at the various ARM measurement facilities varies. Bandwidth and technologies will differ depending on deployment location and availability of connectivity:
  - Terrestrial connectivity provided by telecommunications companies (e.g., fiber, copper, digital subscriber line (DSL) service).
  - Cellular connectivity.
  - Point-to-point microwave radios.
  - Low Earth orbit (LEO) satellites.
- Data from measurement facilities is encrypted via a virtual private network (VPN) tunnel. Traffic between ANL and ORNL is unencrypted over ESnet. Data is served to users from here via file transfer protocol (FTP), Globus, and web services (APIs).
- The ADC's local network configuration consists of 40 Gbps core storage switches connecting back-end storage to hypervisors and processing nodes and 10 Gbps core switches for node-to-node communications. Those 10 Gbps switches also possess 40 Gbps connectivity to ORNL's backbone.
- ARM is looking to streamline communications from the measurement facilities to the ADC. This would entail migrating the ADC into the Science DMZ, then updating measurement facility routing to go to the ADC, among other tasks. This will ensure that all traffic is encrypted end to end and all data is routed as efficiently as possible, and maximize flexibility with regard to access control.
- ARM is looking to replace slower Internet service provider (ISP) connections with services like LEO satellites, as the technology becomes more widely available, to ensure data is transferred to the ADC as efficiently as possible. It will be beneficial to increase throughput at all locations to ensure data is flowing optimally and provide better near real-time data flow options and effective data delivery to users.
- It is recommended that ARM experiment with the Globus MRDP instance for future portal work.

5.5.2 ARM Facility Profile

ARM data is currently collected from three atmospheric observatories—Southern Great Plains, North Slope of Alaska, and Eastern North Atlantic—that represent the broad range of climate conditions around the world, as well as from the three ARM mobile facilities and ARM aerial facilities. Data from these atmospheric observatories, as well as from past research campaigns and the former Tropical Western Pacific observatory, are available at no charge through the ADC.

Nine DOE national laboratories share the responsibility of managing and operating ARM. Along with these laboratories, several constituent groups help provide scientific guidance and develop ARM priorities. ARM also collaborates with many national and international partners.

5.5.2.1 Science Background

ARM operates over 450 instruments to provide measurements of clouds, aerosols, precipitation, the surface energy balance and the background atmospheric state. These measurements are made from a network of six continuously operating ground-based atmospheric observatories and aerial platforms that are deployed on an episodic basis. These measurements are used by researchers to study the interactions among atmospheric constituents (e.g., the life cycle of aerosols and their effects on cloud properties) and the effects of these processes on climate.

ARM's vision is to provide the research community with the best array of field observations and supporting state-of-the-art data analytics to significantly improve the representation of challenging atmospheric processes in Earth system models.

The ARM facility is managed and operated by nine DOE laboratories (Argonne, Brookhaven, Lawrence Berkeley, Lawrence Livermore, Los Alamos, Oak Ridge, Pacific Northwest, National Renewable Energy, and Sandia) and actively serves approximately 1,000 science users each year. ARM science users are based in national laboratories, universities and industry in countries around the world.

The ARM data life cycle is explained in the following figure.

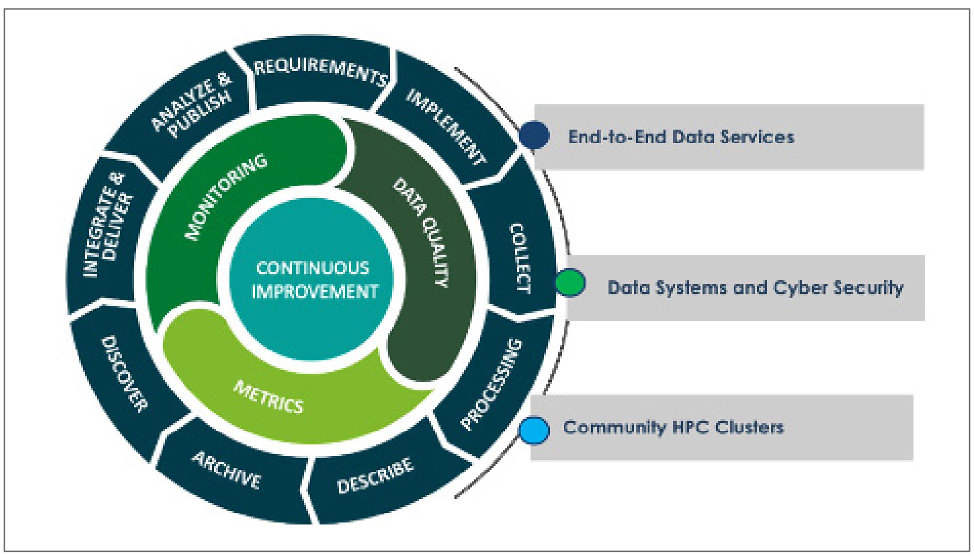

Figure 5.5.1: ARM data life cycle

The ADC provides a robust integrated data and computing ecosystem to advance understanding of atmospheric radiation. The ADC currently holds more than 3.5 Petabytes of data from over 11,000 data products, including data collected from field campaigns, instruments, value-added products, and principal investigator–contributed data. The ADC achieves FAIR compliance by providing effective data discovery, distribution, and interoperability with its newly developed data discovery and delivery tools that use modern and scalable architecture with powerful metadata search features. In addition, the ADC offers computing infrastructure to support next-generation atmospheric model simulations, petascale data storage, big data analytics, and ML for atmospheric and climate science research.

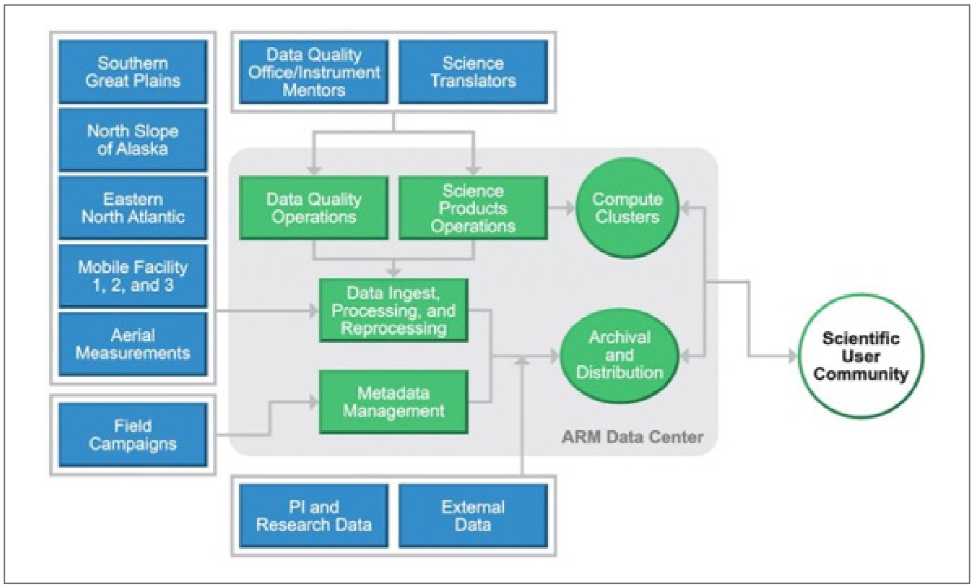

Figure 5.5.2: ARM Data Services manages end-to-end data flow

The metadata is captured and shared using a variety of community-developed standards such as Federal Geographic Data Committee (FGDC) and ISO 19115 to achieve maximum reusability and interoperability. The ADC currently archives over 11,000 data products with a total repository holding of over 3.5 petabytes of data that date back to 1992. These include data from instruments, value-added products, model outputs, field campaigns, and PI-contributed data.

5.5.2.2 Collaborators

By its very nature, ARM is a collaborative entity. As a national scientific user facility, ARM was designed to provide scientists with atmospheric observations needed to conduct their research. While ARM works closely with the Atmospheric System Research (ASR) program to meet the objectives of DOE's Earth and Environmental Systems Sciences Division, ARM also supports research by scientists with diverse programmatic and institutional affiliations around the United States—and the world. In addition, ARM frequently leverages measurements obtained by other organizations to provide a more complete description of the environment around its observatories and regularly coordinates with other agencies on field campaigns. Though each agency has its own goals and priorities, the coordination of observational activities produces more comprehensive data and leads to broader scientific outcomes. These collaborations are key to ARM's success.

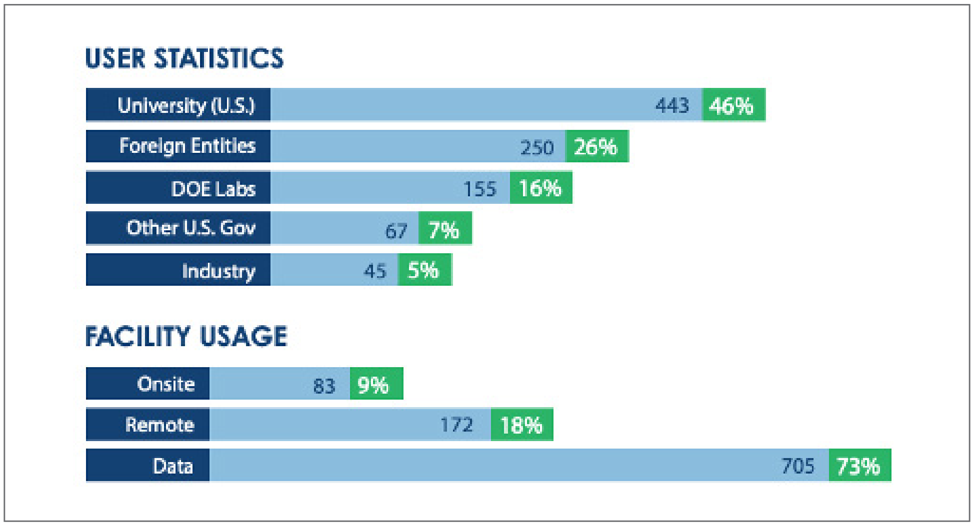

Figure 5.5.3: ARM Users data and facility usage statistics

Comprehensive user and facility metrics are available in the ARM Annual Report[37]

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
AERONET[38]	Secondary	Data transfer	Not available	Annually	N	N
BSRN[39]	Secondary	Data transfer	Not available	Not available	N	N
MOSAiC	Secondary	Data API	Not available	As needed	N	N
ISMN[40]	Secondary	Data transfer	Not available	Not available	N	N
ANL HPSS	Secondary	Data transfer	Varies	Daily	N	Off-site backup

Table 5.5.1: ARM Collaboration Space

5.5.2.3 Instruments and Facilities

The ARM facility provides measurements from the continuous operation of approximately 450 instruments distributed across six ground-based atmospheric observatories in diverse locations ranging from the Arctic to the tropics to Antarctica. These observatories include three fixed location observatories in Oklahoma, Alaska, and the Azores and mobile facilities that are deployed on a proposal-driven basis. These ground-based measurements are augmented by instruments on aerial platforms that are deployed on an episodic basis, also in response to proposals by ARM users. ARM personnel also manage the collection, processing, quality assessment, archiving, and distribution of data from these instruments, which currently represent approximately 2,500 distinct datastreams.

37 https://arm.gov/publications/annual-reports/docs/doe-sc-arm-21-023.pdf
38 https://aeronet.gsfc.nasa.gov
39 https://bsrn.awi.de
40 https://ismn.geo.tuwien.ac.at/en

ARM instruments provide comprehensive measurements of the atmosphere in support of climate research. There are approximately one hundred types of instruments deployed at ARM sites. Major instruments include:

- Profiling and scanning radars ranging in frequency from 915 MHz to 94 GHz to sample cloud and precipitation droplets
- A variety of lidars designed to measure profiles of temperature, humidity, water vapor, air motion, and aerosol properties.
- In situ aerosol particle analyzers for studying physical and chemical properties of aerosols
- Wide-ranging in situ probes are deployed on aircraft and tethered balloons to observe properties of clouds, aerosols, and the background atmospheric state

For approximately a decade, ARM operated a Gulfstream 159 (G-1) turboprop aircraft to conduct most of its in situ sampling. The G-1 was retired at the end of 2018. ARM is currently in the process of modifying a Challenger 850 regional Jet to serve as its primary aerial platform. ARM also operates a mid-sized UAS with an approximately 100-pound payload. ARM also operates several Tethered Balloon Systems (TBS), also with a payload capacity of approximately 100 pounds.

ARM routinely re-deploys one or both of its "mobile" facilities each year to sample new regimes each year. In 2023, ARM will undertake an intermediate length deployment of at least five years in northwest Alabama in the Bankhead National Forest. This observatory will be remarkable because it will include a tower that rises above the tree canopy in the national forest to study profiles of temperature, humidity, radiation, carbon, aerosols, and other parameters. The observatory will also include several ancillary sites to provide information about spatial variability of selected parameters.

In addition to continued deployments of the ARM mobile facilities and the development of the new aircraft, possible developments further in the future include:

- Application of new measurement technologies such as phased array radar, differential absorption lidar, and an increasing number of miniaturized sensors for deployment on UAS and TBS.
- Development of a cloud chamber for studying the interactions of clouds, aerosols, and the environment in a controlled setting.

5.5.2.4 Generalized Process of Science

ARM's vision is to:

- Provide comprehensive and impactful field measurements to support scientific advancement of atmospheric process understanding.
- Achieve the maximum scientific impact of ARM measurements through increased engagement with observational data by ARM staff, including the application of advanced data analytical techniques.
- Enable advanced data analytics and community use of complex ARM datasets through the advancement of computing infrastructure and data analysis.
- Accelerate and amplify the impact of ARM measurements on Earth system models by exploiting ARM and earth system model frameworks to facilitate the application of ARM data to model development.

With the above vision, the ARM facility continues to deploy measurement and data capabilities according to the science community's needs. A detailed description of both near-term and long-term vision is published in the ARM's 2021 Decadal Vision[41].

Many ARM users download a single dataset or a few datasets to study a particular phenomenon, evaluate other instruments (e.g., from a satellite), or evaluate or develop components of an Earth system model. This type of work is expected to continue.

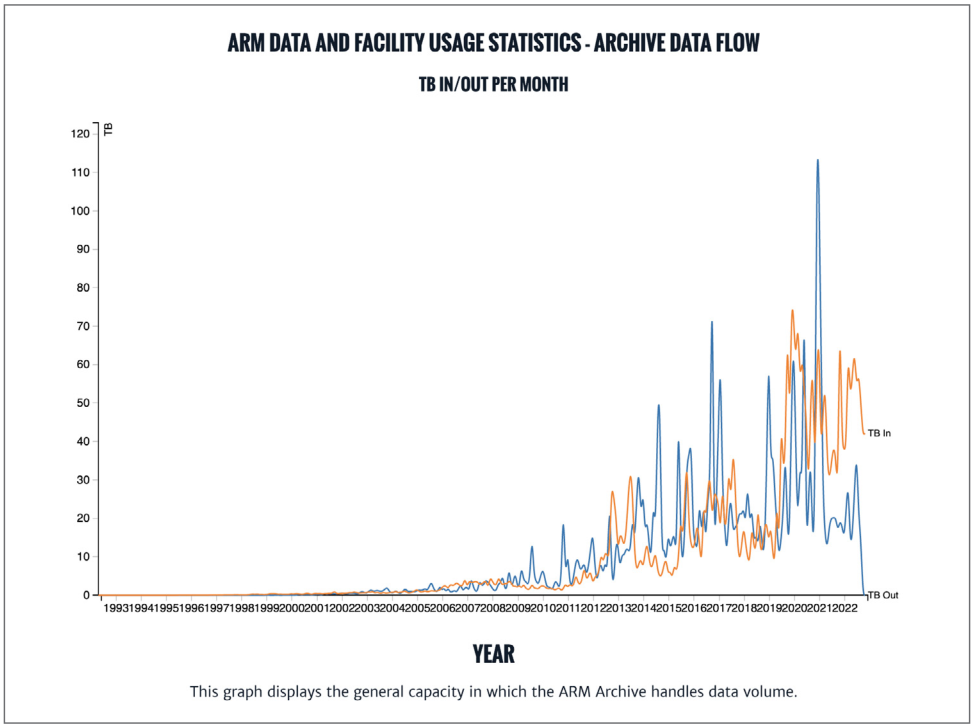

Figure 5.5.4: ARM Usage Statistic[42]

Increasingly, ARM is striving to develop a holistic and internally consistent description of the atmosphere. This work is aided by the application of the high-resolution modeling LASSO[43] framework, which has been applied to date over the ARM Southern Great Plains site studying shallow convection and over a mobile facility deployment in Argentina to study deep convection. This link between ARM observations and LASSO will continue. A goal in the next two to five years will be to develop more expanded datasets combining observations and model output to provide detailed descriptions of the atmosphere that are constrained by observations. Such integrated datasets will inevitably have large volumes as they represent many variables in three-dimensional space and time - much like model output. Working with these comprehensive datasets will likely require more remote access, again similar to working with a large model or satellite data, in contrast with the more common mode of investigators downloading a few specific data streams and parameters today.

41 https://www.arm.gov/publications/programdocs/doe-sc-arm-20-014.pdf
42 https://www.archive.arm.gov/archive-usage/tbinandout.html
43 https://www.arm.gov/capabilities/modeling/lasso

5.5.2.5 Remote Science Activities

To a significant degree, all ARM measurements are remote. There typically are a few technical ARM staff located at ARM observatories, but scientists do not typically spend much if any time at these observatories and even most ARM staff work remotely.

What is changing is the ability to modify the operation of instruments remotely and to configure instruments to change their mode of operation in response to the environment, thereby maximizing the extraction of information from the environment.

In the future, it would be highly desirable to operate UAS and/or TBS in an automated or semi-automated way such that ARM could obtain arrays of vertical profiles around an ARM observatory. Currently such profiles are constrained to remote-sensing (RS) instruments or staffed UAS or TBS activities.

5.5.2.6 Software Infrastructure

ARM Data Services establish end-to-end data security and integrity using a robust data pipeline. After the data is collected from instruments deployed in remote locations, it is securely transferred to ORNL via encrypted data transfer protocols. Different versions of data are deep archived (raw data from instruments, processed data, value-added products etc.) at the ORNL mass data storage system (HPSS), and an additional off-site data backup is done using the ANL's HPSS. Data Provence, versioning details, and change histories are stored in various metadata databases. Robust automated processes are used in continuously monitoring the data and metadata. These automated workflows assess the completeness of data (using automated MD5 checks before and after transferring) and metadata (using the checks applied in the metadata management tool and standards compliance engine).

The instruments go through a detailed semi-automated quality analysis before users can access the data. If a data quality issue is identified, a DQR gets created by the reviewer (instrument specialist, data developer or the ARM data quality office).

5.5.2.7 Network and Data Architecture

Connectivity at the various ARM measurement facilities varies. Bandwidth and technologies used are as follows:

- Southern Great Plains (SGP), Oklahoma
  - Main Site - 1 Gbps fiber from OneNet
  - 17 x Extended Facilities (EFs) - A mix of cellular (AT&T and T-Mobile) and local DSL service
  - 7 x Intermediate Facilities (IFs) - Point-to-Point microwave radios back to the main site. Due to terrain, some have to route through other IFs.
- Northern Slope Alaska (NSA) - Utqiagvik, AK
  - Main Site - 10 Mbps link from ASTAC
    - 3 Starlink Satelite Receivers (25–200 Mbps)
    - Future plans for ASTAC fiber connectivity (100 Mbps down, 30 Mbps up)
  - 1 x Extended Facility - Shipping point-to-point wireless radios to the site for installation.
- Eastern North Atlantic (ENA), Graciosa Island, Azores
  - 100 Mbps link from a local ISP

        - This link loses bandwidth when it crosses from Europe to the US. The highest speed observed is 60 Mbps. Near line speed is observed within Europe.
- Mobile Facilities
    - Links vary depending on the deployment. Current:
        - AMF1 - 1 Gbps in San Diego (Eastern Pacific Cloud Aerosol Precipitation Experiment (EPCAPE) deployment)
        - AMF2
            - LEO satellite service in Cape Grim, Australia (upcoming Cloud and Precipitation Experiment at Kennaook (CAPE-K) deployment)
        - AMF3 - 1 Gbps provided by a local ISP (upcoming southeast US deployment)
- ANL (ANL), Lemont, IL
    - All network traffic to/from the measurement facilities flows through ANL. This includes data flow, remote connectivity, etc.
    - Traffic between ANL and the measurement facilities are encrypted via a VPN tunnel. Traffic between ANL and ORNL is unencrypted over ESnet.
- ORNL, Oak Ridge, TN – ADC
    - All data is sent here, including data from external sources
    - Data is served to users from here via FTP and Globus

The measurement facilities' and ANL's local network configurations are fairly consistent. 10 Gbps core switches connect back-end storage with local hypervisors. Connectivity to all other devices is 1 Gbps.

The ADC's local network configuration consists of 40 Gbps core storage switches connecting back-end storage to hypervisors and processing nodes and 10 Gbps core switches for node-to-node communications. Those 10 Gbps switches also possess 40 Gbps connectivity to ORNL's backbone.

In the next two years, ARM is looking at streamlining communications from the measurement facilities to the ADC. This would entail migrating the ADC into the Science DMZ, then updating measurement facility routing to go to the ADC, among other tasks. This will ensure that all traffic is encrypted end-to-end, all data is routed as efficiently as possible, and maximize flexibility with regard to access control. Also, it will be beneficial to replace slower ISP connections with services like Starlink to ensure data is transferred to the ADC as efficiently as possible.

In the next five years, it will be beneficial to increase throughput at all locations to ensure data is flowing optimally, provide better near real-time data flow options, and effective data delivery to users.

5.5.2.8 Cloud Services

ARM has complex tiered storage and processing requirements for its data operations. The ADC uses a hybrid cyber infrastructure that includes a private cloud infrastructure operated within the ORNL Computing facility for data flow operations and large-scale data analysis, as well as commercial cloud for processes and platform-as-a-service systems that do not require customized data resources. In addition, ARM leverages OLCF to maintain and operate computing clusters for large-scale data processing, data analysis, and high-resolution modeling. ARM uses a shared parallel storage cluster offered by OLCF for HPC projects. All ARM data is archived within the ORNL mass storage system (HPSS) and allows users to seamlessly access the data from various computing resources and scale the analysis from small to large clusters without moving the data. In addition, ORNL Cyber Security provides automated and manual moderated protection for all ADC operations.

Currently, ARM uses cloud services for select system/application backups as well as a few web services that are being updated to run at the ADC.

As a future strategy, the ADC will continue to explore and adopt the use of commercial cloud services for nonroutine processes, operational support tools, platform-as-a-service applications, and cloud-optimized data services for cloud-based user-developed processes.

5.5.2.9 Data-Related Resource Constraints

There is an opportunity to provide next-generation data connectivity for large data transfer between laboratory computing facilities to facilitate distributed data analysis. For example, the ADC archives high-resolution data bundles from LASSO model simulations over several ARM observatories. These LES model outputs are summarized and packaged within the ADC at ORNL to provide a concise representation of the LASSO simulations for sending to data users at other institutions. In some instances, scientists using data analysis capabilities and visualization tools deployed in other Leadership Computing Facilities would benefit by having seamless access to the full-resolution LASSO bundles archived at ORNL. These simulations already represent a significant fraction of the data at the ADC and are expected to soon grow to several petabytes. When large volumes of LES output need to be transferred to the user's computing resources, ARM anticipates that many of these data transfers will use the ESnet infrastructure.

5.5.2.10 Outstanding Issues

None to report at this time.

5.5.2.11 Facility Profile Contributors

ARM Representation

Giri Prakash, *ORNL,* palanisamyg@ornl.gov

Cory Stuart, *ORNL,* stuartcg@ornl.gov

Jim Mather, *PNNL,* Jim.Mather@pnnl.gov

Nicki Hickman, *ANL,* nhickmon@anl.gov

ESCC Representation

Susan Hicks, *ORNL,* hicksse@ornl.gov

Linda Winkler, *ANL,* winkler@mcs.anl.gov

Corey Hall, *ANL,* chall@anl.gov

Surya Singh, *PNNL,* surya.singh@pnnl.gov

Lewis Allen, *PNNL,* lewis.allen@pnnl.gov

5.6 The Earth System Grid Federation (ESGF)

The ESGF is an international collaboration supporting the software that powers most global climate change research. Most notably, assessments by the Intergovernmental Panel on Climate Change (IPCC) rely on ESGF tools and data.

ESGF manages the first-ever internationally federated and decentralized database for handling climate science data, in excess of 30 (32.6+ PB currently) petabytes of data at dozens of federated sites. It is recognized as the leading infrastructure for the management and access of large distributed data volumes for climate change research. It supports the Coupled Model Intercomparison Project (CMIP) and other associated MIP projects, whose protocols enable the periodic assessments carried out by the IPCC.

5.6.1 Discussion Summary

- The ESGF is an international collaboration for the software that powers global climate change research, notably assessments by the IPCC. ESGF manages a decentralized database for handling climate science data, with multiple petabytes of data at dozens of federated sites worldwide. It supports the CMIP, whose protocols enable the periodic assessments carried out by the IPCC.
- ESGF is an interagency and international effort co-led by the DOE, and co-funded by the National Aeronautics and Space Administration (NASA), National Oceanic and Atmospheric Administration (NOAA), National Science Foundation (NSF), and the IS-ENES consortium (e-infrastructure of the European Network for Earth System Modelling (ENES)), including partner institutions Deutsches Klimarechenzentrum (DKRZ), Centre for Environmental Data Analysis (CEDA), Institut Pierre-Simon Laplace (IPSL), in addition to the the Australian National University (ANU) and National Computational Infrastructure (NCI).
- The The ESGF mission is to:
  - Support current CMIP6, and previous CMIP5 and CMIP3 activities, plus other supporting activities (e.g. input4MIPs, obs4MIPs), and prepare for future assessments.
  - Develop data and metadata facilities for inclusion of observations and reanalysis products for CMIP6 use.
  - Enhance and improve current climate research infrastructure capabilities through involvement of the software development community and through adherence to sound software principles.
  - Foster collaboration across agency, political and funding boundaries.
  - Integrate and interoperate with other software designed to meet the objectives of ESGF: e.g., software developed by NASA, NOAA, ESIP, and the European IS-ENES, along with the ENES-RI follow-on.
  - Create software infrastructure and tools that facilitate scientific advancements.
- When considering the resolution ratio from a standard CMIP (~100 km) to a standard DYAMOND (~4 km) simulation, data volumes balloon by a factor of 625x. At this early stage, the most significant simulation planned is for 30 years with a 5 km coupled model, producing 40 TB per simulated month, and ~15 PB over the length of the proposed single simulation.
- The data growth over consecutive CMIP phases (CMIP3 -> 6) has been considerable, with an expansion of phase complexity, in addition to increasing grid resolutions that scale exponentially. For CMIP6, storage limitations as ESGF partner nodes have ensured that no one center has a complete copy of the CMIP6 archive, with high-value, and high-demand datasets replicated

across the primary tier 1 nodes (US-DOE: LLNL, ORNL, DKRZ, CEDA), but in an ad-hoc way that does not ensure data redundancy across the federation.

- It is likely that the output from the ultra-high resolution simulations being run in DYAMOND experiments will test network capabilities into the future, in the case that remote analysis is even possible. If compute resources are not available next to published data, then both network and storage requirements will scale markedly to meet science needs.
- CMIP7 models may be four to five times larger than CMIP6, which places the evolution of data volumes on a path to exascale. Preparation to accommodate this, in terms of storage and data movement and network/hardware throughput, will be required by many community repositories.
- ESGF P2P is a component architecture expressly designed to handle large-scale data management for worldwide distribution. The team of computer scientists and climate scientists has developed an operational system for serving climate data from multiple locations and sources.
- The DYnamics of the Atmospheric general circulation Modeled On Non-hydrostatic Domains (DYAMOND) project is an international consortium that has built on the momentum of CMIP, focused on ultra-high resolution global storm resolving models (GSRMs) with resolutions higher than 5 km globally.
- ESGF provides data discovery and access services to a broad global user community. Thousands of users (clients) discover data either via a web portal (graphical user interface [GUI]) or data search API. These services provide the endpoints so the users may download via HTTP, transfer via Globus, or access via the OPeNDAP protocol which enables limited data subsetting.
- In the future, more users will make use of "server-side" compute and remote analysis capabilities. By moving the compute close to the data, ESGF expects that the end products that need to be transferred back to the user will become trivially sized, e.g., MB-sized subsets or images.
- CMIP data volumes have grown, from ~1 GB in the early phases (~1989), to over 25.8 PB currently/today of replicated model data contributed by 49 institutions across 26 countries. Routinely, this
project is transferring 100s TB of data between nodes weekly, with download and upload occurring concurrently.
- The LLNL-ESGF node is depended upon as the primary CMIP6 node, with ~10 PB of storage allocated. It is the only ESGF node with CMIP3, CMIP5 and CMIP6 data available, whereas a number of the internationally federated ESGF nodes also house a subset of the CMIP5 archive, and continue to provide data through their portals. For this reason, it is considered the primary CMIP data source, and has been depended upon by the other tier 1 nodes to acquire data.
- The "modern" data analysis stack relies on Python-based open-source software, using JupyterHub and DASK gateways. In some instances, data is loaded into a local object store (S3). These services run on a Kubernetes cluster. Data loaded into object store is reformatted to an analysis-ready cloud-optimized (ARCO) format, e.g., Zarr. Data conversion to ARCO may require the download of remote data (10 GB to 2 TB) prior to conversion.
- For storage resources,some ESGF collaborators have obtained time-limited object storage grants, each from Google Cloud Storage (GCS) and AWS (Amazon), to host CMIP data in ARCO format (Zarr). The data conversion process would run on these cloud resources, but data would be transferred (to date typically HTTP download).

- Another difficulty in maintaining ESGF as a federation has been responsiveness and ease of operations with sites to upgrade their platforms. The hope is that containers will at least address an issue of version migration when there is no experience with a particular base version and limitations on what the development team can test for quality assurance purposes.
- ESGF anticipates deployment of comparable replica nodes at ANL and ORNL by 2023. By that time, ESGF under DOE will operate with triple redundancy for CMIP data. Future data management remains uncertain. Most likely, the total unique data volume will exceed capacity for three labs.
- 25 to 30 collaborators provide data to CMIP by hosting data on an ESGF node at their site. Data is replicated to LLNL on an ongoing basis. Seven or eight of the sites have successfully deployed Globus as an option for ESGF data transfers. Most of the sites rely solely on HTTP services to make their data collections available.
- Research universities participate in a climate modeling activity and may each produce around 100 GB to 10 TB of data. For this case, it is deemed impractical for those to run ESGF data nodes, and thus they will transfer the data to ESGF host centers prior to publication.
- Institutions that produce larger quantities of datasets (via modeling and/or observation gathering) e.g., 50+ TBs, should stand up an ESGF node and perform local publishing on site rather than transferring the data to another node to be published. This activity is presently done for CMIP6 data, and so the expectation is that it will continue for ongoing MIP phases. Some modeling centers have relationships with a node site and procedures to move data if the HPC running the model is not co-located with the data node.
- During peak periods, the download rates from 10 reporting sites for HTTP-based downloads peak at ~1.15 PB per month total across all sites and 500 GB on "lighter usage" months. At LLNL, it has been observed that a peak rate of 300 MB per second for bursts of requests to data node servers running the HTTP-based download service.
- Some ESGF nodes, particularly those hosted in Asia, have had poor performance when accessing datasets: dropped connections, poor speed, and poor responsiveness to resolve issues.
- It is recommended that ESGF continue to work with ESnet to help characterize and improve performance to remote resources. This can apply to domestic and international deployments. Slow performance must be mitigated as the CMIP7 era starts, and data sizes increase.

5.6.2 ESGF Facility Profile

The ESGF Peer-to-Peer (P2P) enterprise system is a collaboration that develops, deploys and maintains software infrastructure for the management, dissemination, and analysis of model output and observational data. ESGF's primary goal is to facilitate advancements in Earth System Science. It is an interagency and international effort led by the DOE, and co-funded by NASA, NOAA, NSF, and international laboratories such as the IS-ENES consortium (ENES, DKRZ, CEDA, IPSL, ANU and NCI). The ESGF mission is to:

- Support current CMIP6 activities, and prepare for future assessments
- Develop data and metadata facilities for inclusion of observations and reanalysis products for CMIP6 use
- Enhance and improve current climate research infrastructure capabilities through involvement of the software development community and through adherence to sound software principles
- Foster collaboration across agency, political and funding boundaries
- Integrate and interoperate with other software designed to meet the objectives of ESGF:

e.g., software developed by NASA, NOAA, ESIP, and the European IS-ENES and ENES-RI follow-on

- Create software infrastructure and tools that facilitate scientific advancements

ESGF P2P is a component architecture expressly designed to handle large-scale data management for worldwide distribution. The team of computer scientists and climate scientists has developed an operational system for serving climate data from multiple locations and sources. Model simulations, satellite observations, and reanalysis products are all being served from the ESGF P2P distributed data archive.

5.6.2.1 Science Background

This profile covers the use case of two distinct projects that operate out of Lawrence Livermore National Laboratory. The Coupled Model Intercomparison Project phases 3, 5, 6(CMIP6)[44], and the DOE contribution to the international DYAMOND[45] project focused on storm-resolved ultra high-resolution atmospheric model simulations.

The origins of the CMIP project extend back to 1989, with the founding of the Program for Climate Model Diagnosis and Intercomparison (PCMDI) at LLNL. In the initial phases, LLNL hardware (then NERSC) was used to run atmospheric-only models contributing to the Atmospheric Model Intercomparison Project phase 1 (AMIP1, ~1990) and phase 2 (AMIP2 ~1993), with the data from these simulations solely stored on LLNL-NERSC hardware. Most recently, the LLNL-ESGF node has been actively supporting the subsequent phases CMIP3 (~2003-2008), CMIP5 (~2008-2013) and the active CMIP6 project (2018-) which continues today, and have been collectively identified as MIPs. During this time data volumes have grown, from ~1 GB in the early phases, to 14 unique and more than 25 PB replicated (non-unique, copies) of model data contributed by 49 institutions across 26 countries. Routinely, this project is transferring ~100 TB/week of data between nodes, with download (monitored) and upload (unmonitored; to other federated nodes) occurring concurrently. The LLNL-ESGF node is depended upon as the primary CMIP6 node, with ~10 PB of storage allocated, with the next largest tier 1 nodes having a subset of this (DKRZ 6 PB, CEDA 4 PB), depending on LLNL-ESGF for the bulk of publications. The LLNL-ESGF node is the only ESGF node with CMIP3 data available, whereas a number of ESGF nodes continue to provide data through their portals. For this reason, it is considered the primary CMIP data source, and has been depended upon by the other tier 1 nodes to acquire data from China and other difficult to connect locations. Once data is housed on the LLNL-ESGF hardware, data custodianship is assumed by the PCMDI and ESGF projects. A similar situation occurs with the input4MIPs and obs4MIPs projects, with far lower data volumes (storage and transfer). The next phase, CMIP7 is currently in an early planning phase, in addition to a follow-on project CMIP6Plus that aims to continue to leverage the ESGF global infrastructure to facilitate ongoing collaborative MIP activities.

44 https://esgf-node.llnl.gov/projects/cmip6/
45 https://www.esiwace.eu/the-project/past-phases/dyamond-initiative

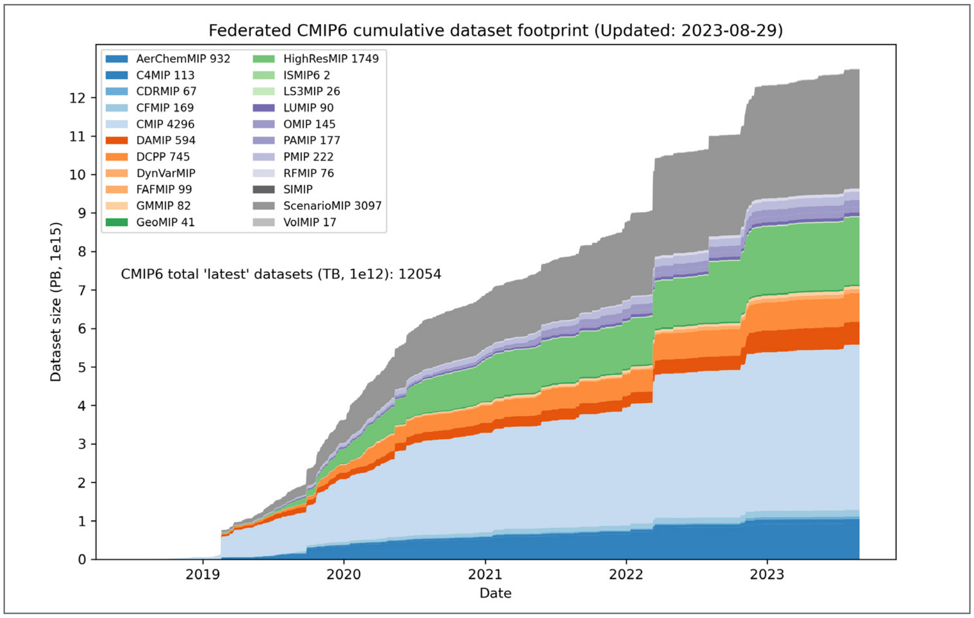

Figure 5.6.1: CMIP6 Dataset Size

ESGF anticipates deployment of comparable replica nodes at ANL and ORNL by 2023. By that time, ESGF under DOE will operate with triple redundancy for CMIP data. Future data management remains uncertain. Most likely, the total unique data volume will exceed capacity for three labs. Likely the top 5-20% of most popular data would be replicated across all three sites on a permanent basis. The remaining 80-95% of data would be partitioned for permanent storage. For local accesses to support analysis work at the site, such data must be transferred as needed to a local cache at the lab.

The DYAMOND project is an international consortium that has built on the momentum of CMIP, focused on ultra-high resolution GSRMs with resolutions higher than 5 km globally. The DYAMOND1 phase began in 2016 and was focused on a 40-day simulation in an atmosphere-only configuration. DYAMOND2 began in 2020 broadened to include coupled atmosphere-ocean configurations, again targeting 40-day simulations. The DOE SCREAM atmospheric model has contributed a use case in the E3SM profile, and so rather ESGF shall focus on the MIP aspect of the project. When considering the resolution ratio from a standard CMIP (~100 km) to a standard DYAMOND (~4 km) simulation, data volumes balloon by a factor of 625x. At this early stage, the most significant simulation planned is for 30-years with a 5 km coupled model, producing 40 TB per simulated month, and ~15 PB over the length of the proposed single simulation. For data exchange, it is likely that contributors will make available lower resolution versions of their model output, along with subsets of the full-resolution data. Analysis on the full-resolution output will most likely have to occur on hardware physically co-located with the data; this will require dedicated network connectivity to enable cross-institutional collaborations. Further details regarding the DOE SCREAM contribution to DYAMOND can be found in the E3SM Profile document.

5.6.2.2 Collaborators

The table below lists the major international replication centers for CMIP data and North American host sites for CMIP. The replica sites provide subsets of the entire CMIP data collection. In addition, 25-30 collaborators provide data to CMIP by hosting data on an ESGF node at their site. Data is replicated to LLNL using the Synda application over HTTP on an ongoing basis. 7-8 sites have successfully deployed Globus as an option for ESGF data transfers. Most of the sites rely solely on HTTP services to make their data collections available.

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
REPLICA SITES						
CEDA / Appleton, UK	Both	http/Globus	~2 GB	daily/ongoing	N	N/A
DKRZ / Hamburg, Germany	Both	http/Globus	2 GB	daily/ongoing	N	N/A
NCI / Canberra, Australia	Both	http/Globus	2 GB	daily/ongoing	N	N/A
Addictional Publishing SITES (No Replicas)						
CCCma / Victoria, Canada	Primary	http	2 GB	ad hoc	N	N/A
NCAR Boulder, CO	Primary	http/Globus	2 GB	ad hoc	N	N/A
GFDL Princeton, NJ	Primary	http	2 GB	ad hoc	N	N/A
NASA-GSFC/ NCCS Greenbelt, MD	Primary	http	2 GB	ad hoc	N	N/A

Table 5.6.1: ESGF Collaboration Space

Dataset sizes are based on current CMIP6 metrics 6 million "dataset" entries (single simulation variable output) ~13 PB distinct (un-replicated) volume reported[46] as at October 2022.

5.6.2.3 Instruments and Facilities

Research universities participate in a climate modeling activity and may each produce around 100 GB-10 TB of data. For this case, it is deemed impractical for those to run ESGF data nodes, and thus they will transfer the data to ESGF host centers prior to publication. Consider two example sites. Site A has a Globus endpoint so they schedule a Globus transfer to move the data to the target ESGF site to be published. Site B does not and the staff researcher cannot get up to speed on the use of Globus, so will opt to use an FTP or HTTP-based transfer for publishing. In the next two years these transfers are not integrated into the ESGF publishing process, so must be handled offline with humans in the loop. Within the two-to five-year timeframe ESGF anticipates that publication procedures will be in place to automate the process (Globus or HTTP transfer) coordinated with data publication into a single operation.

At current publication occurs at an ESGF "data node" site that hosts services for data, so the publication software sees the same data locally as what is served out to the community via the data services.

Other institutions may produce datasets by collecting and curating observations from instruments in the field. One such related activity expected to produce such data for ESGF is obs4MIPs, the project goal is to bring observational data inline with the standards applied to model output. The current volume is 10s of gigabytes, with an expectation toward terabytes as higher temporal frequency, and spatial resolution data is identified. The expectation is for these sites to make use of the updated ESGF publication procedure.

46 http://esgf-ui.cmcc.it/esgf-dashboard-ui/federated-view.html

Institutions that produce larger quantities of datasets (via modeling and/or observation gathering) e.g., 50+ TB should stand up an ESGF node and perform local publishing on site rather than transferring the data to another node to be published. This activity is presently done for CMIP6 data and so the expectation is that it will continue into the future for ongoing MIP phases. Some modeling centers have relationships with a node site and procedures to move data if the HPC running the model is not co-located with the data node.

5.6.2.4 Generalized Process of Science

ESGF provides data discovery and access services to a broad global user community. 1000s of users (clients) discover either via a web portal (GUI) or data search API. These services provide the endpoints so the users may download via http, transfer via Globus or access via the OPeNDAP protocol which enables limited data subsetting.

During peak periods, the download rates from 10 reporting sites for HTTP-based downloads (many additional sites do not report download metrics) peak at ~1.15 PB per month total across all sites and 500 GB on "lighter usage" months. At LLNL a peak rate of 300 MB per second was observed for bursts of requests to a data node server running the HTTP-based download service.

We expect that over time more users will make use of "server-side" compute and remote analysis capabilities. By moving the compute close to the data, ESGF expects that the end products that need to be transferred back to the user will become trivially sized, e.g., MB-sized subsets or images.

5.6.2.5 Remote Science Activities

The CMIP6 phase included contributions from 24 sub-project MIPs (activities), with each activity administered through a MIP panel. These panels are regionally distributed, and often in the early stages, data associated with each MIP is not replicated across the federation. In these instances, is it common for initial analyses to leverage the OPeNDAP connectivity provided alongside the ESGF local instance, with data being read across global networks, often unrelated to the ESGF federated nodes. Such uses, at least in examples to date, are most often an order of magnitude fewer than direct data downloads, or local data analysis using adjacent compute resources.

5.6.2.6 Software Infrastructure

The "modern" data analysis stack is Python-based open-source, using JupyterHub and DASK gateway. In some instances, data is loaded into a local object store (S3). These services run on a Kubernetes cluster. Data loaded into object store is reformatted to an analysis-ready cloud-optimized (ARCO) format, e.g., Zarr. Data conversion to ARCO may require the download of remote data (10 GB-2 TB) prior to conversion. DASK Gateway allows for users to parallelize their computations, typically using locally stored data, but could be on remote object stores.

5.6.2.7 Network and Data Architecture

ORNL connects to ESnet via redundant border routers, each of which currently connects to a diverse ESnet router at 100G. The expectation is that these connections will soon be upgraded to 400G connections. The ORNL border routers connect the ORNL Enterprise network and OLCF to ESnet. This connectivity is depicted in the diagram below.

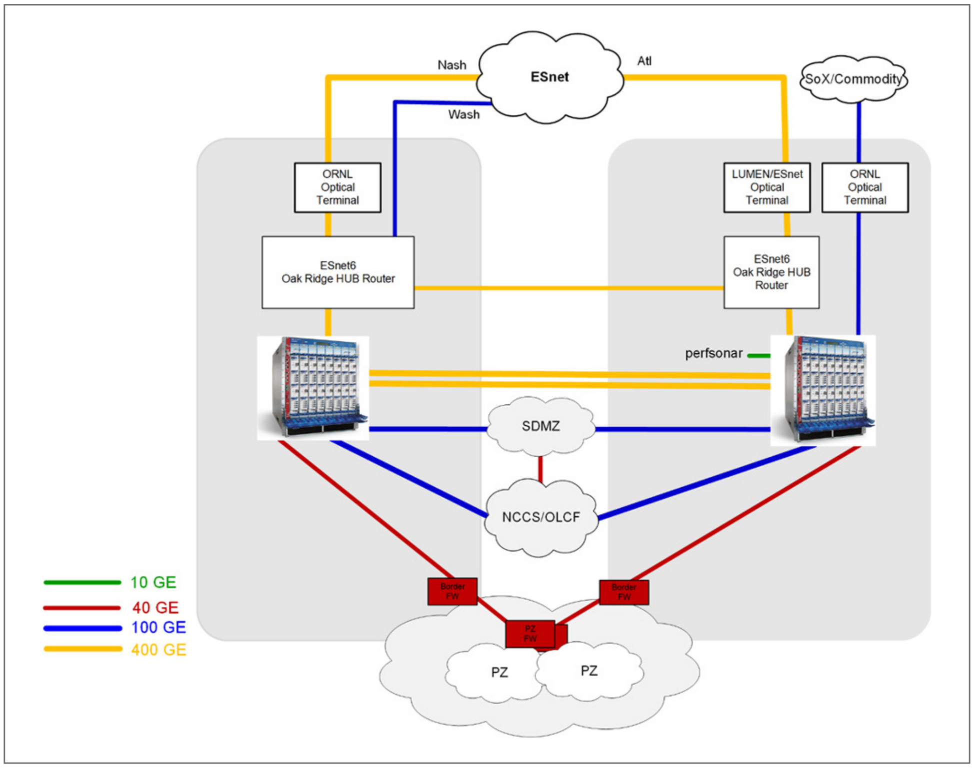

Figure 5.6.2: ORNL Network Diagram

The Enterprise firewalls are connected to the border routers at 40G currently. ORNL does use a Science DMZ architecture for high-performance data transfer. This environment connects to the border routers with 10/40/100G DTN connections available. Globus is the approved transfer method. A border perfSONAR node is connected to the border router and participates in the ESnet grid.

5.6.2.8 Cloud Services

Clarification is needed regarding how DOE plans to leverage cloud services for use by ESGF. The community, though, has experience in this regard from the Pangeo effort to enable cloud-based analysis using the software stack described in Section 6. For storage resources the project has obtained object storage grants, each from GCS (Google) and AWS (Amazon) to host CMIP data in ARCO format (Zarr). The data conversion process would run on these cloud resources but data would be transferred (to date typically HTTP download). This is a scalable process so concurrent transfers are possible depending on the compute resource allotment for the conversion process. Additional conversion is done on an ad-hoc based on a particular scientific need to populate object store and facilitated by the "Pangeo-Forge" capability that runs using cloud-hosted (GitHub) workflows.

Amazon has made additional storage available, and if community resources present themselves to undertake a large-scale conversion process for CMIP6, then the community would benefit from additional cloud-hosted data. The consideration is that compute resources that can leverage such storage and will be paid for by a sponsor will be limited.

5.6.2.9 Data-Related Resource Constraints

The data growth over subsequent CMIP phases (CMIP3 -> 6) has been considerable, with an expansion of phase complexity, in addition to increasing grid resolutions that scale exponentially. For CMIP6, storage limitations as

ESGF partner nodes have ensured that no one center has a complete copy of the CMIP6 archive, with high-value, and high-demand datasets replicated across the primary tier 1 nodes (LLNL, DKRZ, CEDA, NCI), but in an ad-hoc way that does not ensure data redundancy (copies at more than a single node) across the federation.

In addition to storage limitations, it is likely that the output from the ultra-high resolution simulations being run in DYAMOND experiments will test network capabilities into the future, in the case that remote analysis is even possible. If compute resources are not available next to published data, then both network and storage requirements will scale markedly to meet science needs.

5.6.2.10 Outstanding Issues

An ongoing concern has been that connectivity with some ESGF nodes, particularly those hosted in Asia, has been poor, for example: dropped connections, poor speed, and poor responsiveness to resolve issues. Another difficulty in maintaining ESGF as a federation has been responsiveness and ease of operations with sites to upgrade their platforms. The hope is that containers will at least address an issue of version migration when there is no experience with a particular base version and limitations on what the development team can test for quality assurance purposes.

To give a better sense of how data transfer rates have been observed during a long period of intensive transfers, ESGF can point to the activity of replication of the entire CMIP replica collection from LLNL to ANL and ORNL over the period of Feb to May 2022. The chart below shows the peak transfer rates for concurrent Globus transfers. LLNL is limited to three DTN servers comprising the Globus endpoint.

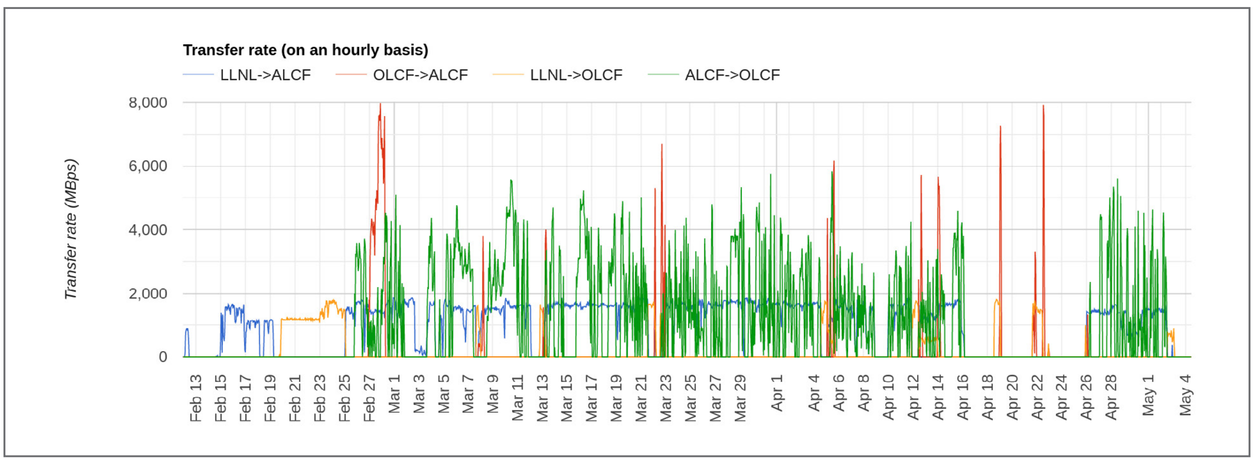

Figure 5.6.3: ESGF Replication from LLNL to ORNL and ANL

5.6.2.11 Facility Profile Contributors

ESGF Representation

Forrest Hoffman, *ORNL*, hoffmanfm@ornl.gov

Ian Foster, *ANL*, foster@anl.gov

Program for Climate Model Diagnosis and Intercomparison (PCMDI)

Paul J. Durack, *LLNL*, durack1@llnl.gov

Sasha K. Ames, *LLNL*, ames4@llnl.gov

Stephen Klein, *LLNL*, klein21@llnl.gov

ESCC Representation

Susan Hicks, *ORNL*, hicksse@ornl.gov

Linda Winkler, *ANL* ,winkler@mcs.anl.gov

Corey Hall, *ANL*, chall@anl.gov

Gianni Carlo Contardo, *LLNL*, gianni@llnl.gov

5.7 The Environmental Systems Science Data Infrastructure for a Virtual Ecosystem (ESS-DIVE) and Self-Driving Field Laboratories Research

The U.S. DOE's ESS-DIVE is a new data repository for Earth and environmental science data funded by the Data Management program within the Climate and Environmental Science Division at LBNL. It stores and enhances access to critical information generated from research funded by or related to the DOE's SC BER under its Subsurface Biogeochemical Research (SBR) and Terrestrial Ecosystem Science (TES) Programs in the Environmental Systems Science (ESS) activity.

5.7.1 Discussion Summary

- ESS-DIVE is a data repository for Earth and environmental science data managed by LBNL. It stores and enhances access to critical information generated from research funded by or related to the BER ESS program. ESS-DIVE works to preserve, expand access to, and improve usability of critical data generated through research of terrestrial and subsurface ecosystems.
- ESS-DIVE collects, stores, manages, and shares data. The volume, complexity, and diversity of Earth and environmental science data make it challenging to capture, store, verify, analyze, and share information in a consistent manner.
- ESS-DIVE's long-term vision is to provide a repository that enables easy intuitive integration and management of data from models, experiments, and observations. In this vision, ESS-DIVE becomes a focus for enabling access to integrated data for visualizations, models, ML, and other analyses for knowledge discovery from these datasets.
- BER projects generate diverse observational, experimental, and modeling data: the aim is to become the repository of choice for ESS project data by working closely with the ESS community to build the repository based on community needs and priorities. ESS-DIVE accepts submissions of well-structured, standardized, high-quality data from data contributors and facilitates data processing, synthesis, and analysis capabilities accessible to data users around the world. Datasets include raw output from sensors, simulations, analyzed datasets, and user-contributed data.
- Initial ESS-DIVE work focused on developing and establishing a well-known repository for data. Ongoing and future work will expand the user base and functions through five key innovations:
  - Apply user-experience research methods to understand user needs.
  - Support early data archiving by projects.
  - Support scaling of the repository to a significantly broader portion of the ESS community.
  - Support search of extracted ESS-DIVE data with a fusion database.
  - Support federation with other repositories.
- To meet the needs of field data from sensors and instruments, it will be critical to enable data acquisition, transformation, and analytics workflows that can integrate real-time data streams with predictive models, particularly for extreme events. These can be used to train models using a hybrid of AI/ML and deterministic techniques, which can generate real-time predictions and feedback for sensors out in the field. The highly distributed nature of sensor networks, and the volumes of data collected, will necessitate a data "backbone" for AI/ML-based scientific workflows.
- ESS-DIVE is being used to support a research use case supporting the development of Self-Driving Field Laboratories (SDFL).

- The DataONE network provides redundancy via replication to far-away network nodes and automatic self-healing capabilities. DataONE replication enables users to discover and download datasets if the main repository site is unavailable.
- ESS-DIVE acknowledges that interoperability across archives is a broader issue, involving coordination across competing requirements, standards, and interfaces, which can only be sufficiently addressed by multi-institutional effort.
- ESS-DIVE partners across the ESS community to develop data and metadata standards for the repository: these include ESS projects (e.g., SFAs and NGEEs); the broader digital library and standards communities (e.g., RDA, ESIP, DataONE, OGC, and DataCite); BER facilities and data systems (e.g., the NMDC, the ESGF, KBase, the JGI, the EMSL, and ARM); and other federal agencies (e.g., USGS, NASA, and NSF). ESS-DIVE will work to integrate with external data repositories (ESGF, ARM, EMSL, NMDC, and USGS) in the future.
- ESS-DIVE serves a broad range of national lab and university partners that contribute data to the system. This includes data coming from projects with a footprint at major DOE HPC facilities such as NERSC, ORNL, ANL, etc. (e.g. for large-model datasets). Additionally, ESS-DIVE serves partners from the BER facilities like the JGI and EMSL for environmental sample data. A new area of emphasis is centered around sensor networks in the field that are collecting large volumes of data. Dataset sizes range from megabytes to multiple terabytes.
- ESS-DIVE maintains two standby copies of ESS-DIVE: one copy in the LBNL Information Technology Division (LBNL IT) data center and the other copy at NCEAS in Santa Barbara. Each failover node can run a copy of the stack in either read-only or read/write modes. Either node can take over as primary in the event of a failover.
- ESS-DIVE sees the importance of streaming data, increased use of cloud resources, and federated data to facilitate these use cases. Additionally, integration and interoperability with other repositories will also be important. To enable a future-facing data repository, ESS-DIVE must create pathways to bridge the scientific data and metadata stored in other partner repositories with the data in ESS-DIVE. Given that the data ecosystem is fundamentally a distributed enterprise, scientific discovery ultimately rests on being able to deliver insights that can fuse information from different data silos.
- ESS-DIVE data is hosted on resources at NERSC (under the Spin infrastructure and hosted on the community filesystem). Copies of the data and metadata are replicated to the DataONE federation. ESS-DIVE also has failover nodes on LBNL IT hardware (Science Virtual Machine [SVM] cluster) and at NCEAS (UC Santa Barbara). It is important to maintain high-bandwidth connections across these three facilities for data and service replication to support failover needs.
- For Globus Uploads, an external virtual machine (VM) on LBNL IT"s SVM infrastructure is used (advanced Globus features at NERSC are limited to users that are part of the NERSC allocation). Data is then transferred to NERSC.
- ESS-DIVE uses Dockerized microservices for portability (running on NERSC Spin Kubernetes). This has the potential to introduce network challenges, since all traffic goes through the container network layer.
- ESS-DIVE continually evaluates resource options such as cloud computing. The cloud computing model enables us to leverage a wide range of software, storage, and computing hardware resources under the "-as-a-service" model, where ESS-DIVE simply deploys and makes use of services in the cloud, rather than run something private. Both dev-ops and resilience strategies actively leverage resources on public clouds when those services meet the anticipated needs.

- It is important for ESS-DIVE to support high-bandwidth, low-latency connections to users submitting data to ESS-DIVE. Data collected by ESS-DIVE is generally from other ESS projects, and is most likely generated on remote resources. ESS-DIVE is essentially the repository where ESS projects store their data for long-term archival. This includes national labs, universities, and individual computing resources.
- ESS-DIVE expects to be able to store large datasets on ESS-DIVE (~1 TB datasets) to support model data archiving needs in the future. This will require the use of high-performance mechanisms like Globus and high-bandwidth networks.
- The Surface Atmosphere Integrated Laboratory (SAIL) envirotranspiration application is deployed in western Colorado as a Self-Guiding Field Laboratory (SGFL), which has created an iterative two-way interaction between data gathering and model predictions, wherein data gathering is adapted based on model predictions utilizing a "decision engine": an algorithm that takes the gathered data as inputs and calculates a numerical decision matrix for assessing whether data gathering should be adjusted. NERSC HPC resources are leveraged for data analysis and scalable ML model training.
- ESS-DIVE is considering maintaining a replicated instance of the repository on cloud resources for failover for the next phase. ESS-DIVE will be evaluating the options based on overall cost, services, and capabilities. This includes enabling cloud-based data ingest mechanisms, cloud-based data analysis tools, and cloud storage for replicating data.
- Currently LBNL resources behind Cloudflare have a very small upload limit of 500 MB, which makes uploading large datasets nearly impossible. ESS-DIVE was able to work with LBNL IT to increase this limit to 100 GB, but this is considered a hard limit. ESS-DIVE must bypass Cloudflare due to limits on data uploads larger than 100 GB (for API and data portal services).

5.7.2 ESS-DIVE Facility Profile

The mission of ESS-DIVE is to preserve, expand access to, and improve usability of critical data generated through DOE-sponsored research of terrestrial and subsurface ecosystems in support of the DOE's efforts to address some of society's most pressing energy and environmental challenges. ESS-DIVE is a member of the DataONE Federation, and the Federation of Earth System Information Partners (ESIP).

Scientists with expertise in Earth and Environmental Sciences and Computing Sciences from Berkeley Lab are developing ESS-DIVE in collaboration with the ESS community. The repository's infrastructure is being developed in collaboration with NERSC, the National Center for Ecological Analysis and Synthesis (NCEAS).

5.7.2.1 Science Background

ESS-DIVE is a data repository for Earth and environmental sciences. ESS-DIVE collects, stores, manages, and shares data. The volume, complexity, and diversity of Earth and environmental science data make it challenging to capture, store, verify, analyze, and share information in a consistent manner.

DOE-BER projects generate diverse observational, experimental, and modeling data. To meet DOE ESS goals, the aim is to establish ESS-DIVE as the repository of choice for ESS project data. To achieve this vision, ESS-DIVE works closely with the ESS community to build the repository based on community needs and priorities. The ESS-DIVE repository meets the DOE's Earth and Environmental Systems Sciences Division (EESSD) strategic goals by promoting submissions of well-structured, standardized, high-quality data from data contributors; and by facilitating data processing, synthesis, and analysis capabilities accessible to data users around the world.

In the FY17-FY20 phase, ESS-DIVE focused on developing and establishing ESS-DIVE as a well-known repository where ESS projects could store and publish their data. In the FY21-24 phase, ESS-DIVE is expanding the user-base and function of ESS-DIVE through five key innovations:

1. apply user-experience research methods to understand user needs.
2. support early data archiving by projects.
3. support scaling of the repository to a significantly broader portion of the ESS community.
4. support search of extracted ESS-DIVE data with a fusion database.
5. support federation with other repositories.

ESS-DIVE's innovative strategies, strong alliances with ongoing BER data projects, and inclusive community-based approach help optimize the value of BER data. Key priorities relevant to large data and networking needs are highlighted (3. and 5.)

ESS-DIVE is also being used to support a research use case supporting the development of SDFLs. This emerging use case is included in this report in Section 5.7.2.4.2, since ESS-DIVE is envisioned as a key component of ongoing work supporting SDFL in the next few years, and the main location for landing data from this research.

5.7.2.2 Collaborators

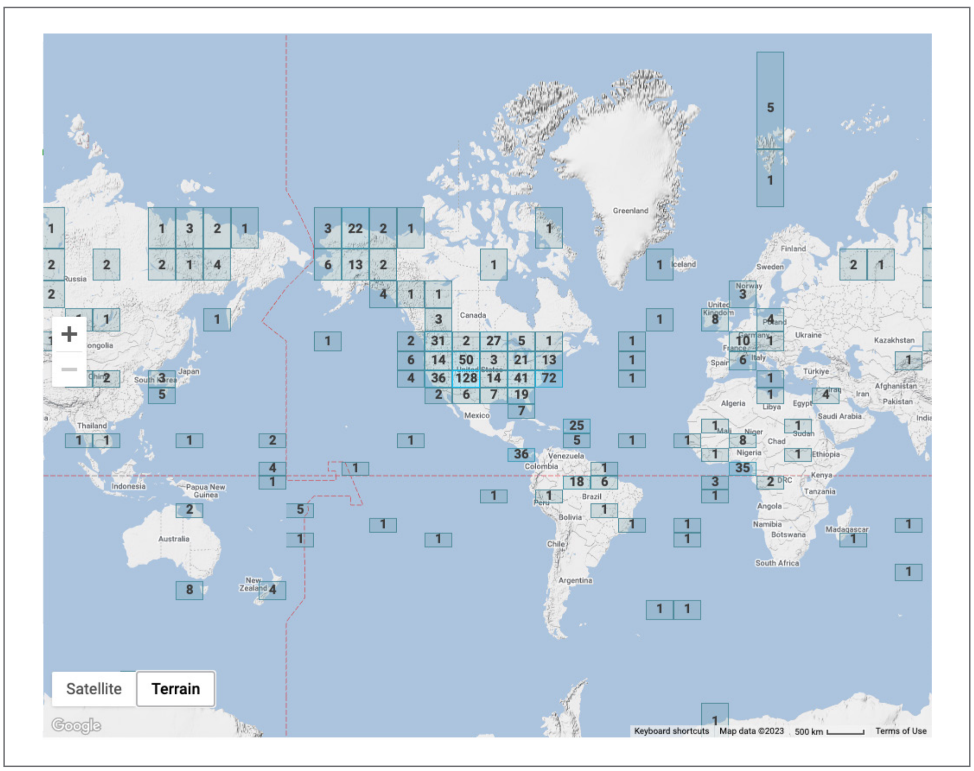

Figure 5.6.2: The ESS-DIVE repository has been providing users data since September 2017 and accepting data since April 2018.

ESS-DIVE is partnering with several projects across the ESS community to develop data and metadata standards for the repository through use of ESS-DIVE community funds. In phase I, the first three years of the project (FY18 - FY20), ESS-DIVE focused on building the underlying repository infrastructure and developing partnerships with ESS projects (e.g., SFAs and NGEEs); the broader digital library and standards communities (e.g., RDA, ESIP, DataONE, OGC, and DataCite); BER facilities and data systems (e.g., NMDC, ESGF, KBase, JGI, EMSL, and ARM); and other federal agencies (e.g., USGS, NASA, NSF).

During phase II, over the next three years of the project (FY21 - FY24), ESS-DIVE seeks to expand the ESS-DIVE user community and prepare the repository for the next ten years of operation and growth. ESS-DIVE focuses on providing a scalable, robust repository and long-term curation of ESS data that adhere to FAIR principles, with the goal of increasing the ease and capacity of storing data in the repository, and data usability. ESS-DIVE will use user-experience research methods to understand user needs and priorities, and will continue innovative community partnerships to standardize data and enable searches of the data in the repository, not just the metadata, through a "fusion database." Finally, ESS-DIVE isinitiating work towards integrating ESS-DIVE with other external data repositories (ESGF, ARM, EMSL, NMDC, USGS). The vision for phase II is to transform the data-archival user experience and advance the ability of scientists to discover and use the data in analysis and modeling. Listed below are:

1. Collaborators (DOE BER data systems with whom ESS-DIVE is involved with in data and metadata sharing and integration activities)
2. Community Partners (Groups of ESS-DIVE data contributors that are co-developing data and metadata standards with ESS-DIVE)
3. ESS-DIVE project software and infrastructure development partners
4. Projects submitting data (Users that are Data Contributors for ESS-DIVE)

5.7.2.2.1 BER Collaborators

- NMDC – API integration, Data linking, Sample Data standardization
- KBase – Sample Data standardization
- EMSL – Data linking, Sample Data standardization
- ESGF – Large Model Data

5.7.2.2.2 Community Partners:

- Ranjeet Devarakonda, Susan Heinz, and Terri Velliquette, ORNL
- Ben Bond-Lamberty and Stephanie Pennington, PNNL
- Amy Goldman and Huiying Ren, PNNL
- Alistair Rogers and Kim Ely, BNL
- Kristin Boye, SLAC
- Pamela Weisenhorn, ANL

5.7.2.2.3 ESS-DIVE Core Development Partners

- National Center for Ecological Analysis and Synthesis (NCEAS) - Software development partner for Metacat software and location for backup instance of ESS-DIVE
- DataONE - Federated network for replication of ESS-DIVE Metadata and Data

5.7.2.2.4 Projects Submitting Data to ESS-DIVE:

ESS-DIVE has over 400 projects approved to store data. The list below are the projects that have stored data in ESS-DIVE:

- A Comprehensive Framework for Modeling Emissions from Tropical Soils and Wetlands
- A general mechanistic framework for cross-scale understanding of hot spots and hot moments in carbon and water fluxes
- A Global, High-Resolution River Network Model for Improved Flood Risk Prediction
- A Multiscale Approach to Modeling Carbon and Nitrogen Cycling within a High Elevation Watershed
- Accounting for hydrological and microbial processes on greenhouse gas budgets from river systems
- Advanced Global Atmospheric Gases Experiment (AGAGE)
- Advancing a Watershed Hydro-Biogeochemical Theory: Linking Water Travel Time and Reaction Rates Under Changing Climate
- AmeriFlux Management Project
- AmeriFlux US-IB2 Fermi National Accelerator Laboratory- Batavia (Prairie site), AmeriFlux US-Bo1: Bondville, Water foraging with dynamic roots in E3SM; The role of roots in terrestrial ecosystem memory on intermediate timescales (PI Berkelhammer)
- Arctic Shrub Expansion, Plant Functional Trait Variation, and Effects on Belowground Carbon Cycling
- Belowground Biogeochemistry Scientific Focus Area
- Can microbial functional traits predict the response and resilience of decomposition to global change? / Steven Allison
- Carbon Dioxide Information Analysis Center at AppState (CDIAC-FF)
- Climate Change-Terrestrial Ecosystem Science SFA
- Coastal Wetland Carbon Cycling Processes in a Warmer Climate
- COMPASS-FME
- Consequences of Plant Nutrient Uptake for Soil Carbon Stabilization
- Constraining Physical Understanding of Aerosol Loading, Biogeochemistry, and Snowmelt Hydrology from Hillslope to Watershed Scale in the East River Scientific Focus Area
- Defense Coastal/Estuarine Research Program (DCERP)
- Development of a molecularly informed biogeochemical framework for reactive transport modeling of subsurface carbon inventories, transformations and fluxes
- Dissolved Oxygen Probe System For Real-Time, In Situ Subsurface Monitoring
- Duke Forest FACE - Forest-Atmosphere Carbon Transfer and Storage (FACTS-I)
- EAGER SitS: Can remotely imaged vegetation characteristics provide a window into soil nutrient cycles?
- Early Career Research Program: Watershed Perturbation-Response Traits Derived Through Ecological Theory

- Early Career Research Program: Watershed Perturbation-Response Traits Derived Through Ecological Theory - Worldwide Hydrobiogeochemistry Observation Network for Dynamic River Systems (WHONDRS)
- Ecohydrological controls on root and microbial respiration in the East River watershed of Colorado
- Effects of Hurricane Disturbance and Increased Temperature on Carbon Cycling and Storage of a Puerto Rican Forest: A Mechanistic Investigation of Above- and Belowground Processes
- Effects of Warming on Tropical Forest Carbon Cycling: Investigating Temperature Regulation of Key Tropical Tree and Soil Processes
- Environmental System Science Data Infrastructure for a Virtual Ecosystem; Carbon Dioxide Information Analysis Center (CDIAC), ORNL, Oak Ridge, TN
- ESS-DIVE
- ExaSheds
- Forecasting Carbon Storage as Eastern Forests Age: Joining Experimental and Modeling Approaches at the University of Michigan Biological Station AmeriFlux Site
- Free Air CO2 Enrichment Model Data Synthesis (FACE-MDS)
- Free-Air CO2 Enrichment (FACE) Experiment Synthesis Activities
- Functional-type modeling approach and data-driven parameterization of methane emissions in wetlands
- Groundwater Quality SFA
- High latitude application and testing of Earth system models
- High-resolution Estimation of Carbon Stock and Changes in Tropical Forests of Sabah for Capturing Climate Change Funds
- HiLAT-RASM project
- IDEAS Applications: Reactive-Transport and Integrated Hydrology Model Development (IDEAS = Interoperable Design of Extreme Scale Application Software)
- iNAIADS
- Incorporating the Hydrological Controls on Carbon Cycling in Floodplain Ecosystems into earth system models
- Interdisciplinary Research for Arctic Coastal Environments (InteRFACE)
- Linking Nutrient Reactivity and Transport in Subsurface Flowpaths Along a Terrestrial-Estuarine Continuum
- Marine Unexploded Ordnance (UXO) Characterization Based on Autonomous Underwater Vehicle (AUV) Technology
- Mechanistic and Predictive Understanding of Needle Litter Decay in Semi-arid Montane Ecosystems Experiencing Unprecedented Vegetation Mortality
- Mercury Critical Interfaces SFA
- Microbial environmental feedbacks and the evolution of soil organic matter
- Modelling Microbes to Predict Post-fire Carbon Cycling in the Boreal Forest across Burn Severities

- Net Methylation Potential of Mercury Sulfides for Different Groups of the Methylating Microbial Community
- Next-Generation Ecosystem Experiments (NGEE) Arctic
- Next-Generation Ecosystem Experiments (NGEE) Tropics
- North Carolina AmeriFlux Core Site Cluster
- ORNLTESSFA (Oak Ridge National Lab's Terrestrial Ecosystem Science Scientific Focus Area (ORNL TES SFA))
- Particulate organic matter (POM) transport and transformation at the terrestrial–aquatic interface
- Quantifying Distributed Exchanges of Groundwater with River Corridors
- Quantifying Subsurface Biogeochemical Variability in a High Altitude Watershed During Winter Isolation
- Quantifying Variability and Controls of Riverine Dissolved Organic Carbon Exported to Arctic Coastal Margins of North America
- Quantitative, trait-based microbial ecology to accurately model the impacts of nitrogen deposition on soil carbon cycling in the Anthropocene
- Real-time Dissolved Oxygen Monitoring for Understanding Biogeochemical Processes at the Columbia River Hanford Reach and the East River Floodplain
- Reducing Uncertainties in Biogeochemical Interactions through Synthesis and Computation (RUBISCO)
- Regional Vulnerability of Permafrost Carbon to Climate Change: A Multifactor Experiment and Model Network
- River Corridor and Watershed Biogeochemistry SFA
- River Corridor and Watershed Biogeochemistry SFA Worldwide Hydrobiogeochemistry Observation Network for Dynamic River Systems (WHONDRS)
- Scaling from Flux Towers to Ecosystem Models: Regional Constraints on Carbon Cycle Processes from Atmospheric Carbonyl Sulfide
- Science Area 1: Standard Award: Model-Data Fusion to Examine Multiscale Dynamical Controls on Snow Cover and Critical Zone Moisture Inputs
- Seasonal controls on dynamic hyporheic zone redox biogeochemistry
- Soil Carbon Biogeochemistry
- Soil Carbon Response to Environmental Change
- Sonoma Water Wildfires
- Space and Time Dynamics of Transpiration in the East River Watershed: Biotic and Abiotic Controls (DE-SC0019210)
- Subalpine and Alpine Species Range Shifts with Climate Change: Temperature and Soil Moisture Manipulations to Test Species and Population Responses
- Subalpine and Alpine Species Range Shifts with Climate Change: Temperature and Soil Moisture Manipulations to Test Species and Population Responses (Alpine Treeline Warming Experiment)

- Subsurface Biogeochemistry of Actinides SFA
- Terrestrial Ecosystem Science at Berkeley Lab
- Testing mechanisms of how mycorrhizal associations affect forest soil carbon and nitrogen cycling
- The Advanced Global Atmospheric Gases Experiment (AGAGE)
- The Carbon-Nutrient Economy of the Rhizosphere: Improving Biogeochemical Prediction and Scaling Feedbacks From Ecosystem to Regional Scales
- The Role of Pyrophilous Microbes in the Breakdown and Sequestration of Pyrogenic Organic Matter
- Trace Metal Dynamics and Limitations on Biogeochemical Cycling in Wetland Soils and Hyporheic Zones
- Trace Metal Dynamics and Limitations on Biogeochemical Cycling in Wetland Soils and Hyporheic Zones, PI Jeffrey G. Catalano
- Tropical Forest Response to a Drier Future: Turnover Times of Soil Organic Matter, Roots, Respired CO2, and CH4 Across Moisture Gradients in Time and Space
- Understanding and Modelling Current and Future Coastal Wetland Methane
- Understanding Soil Microbial Sources of Nitrous Acid and their Effect on Carbon-Nitrogen Cycle Interactions
- Using Root and Soil Traits to Forecast Woody Encroachment Dynamics in Mesic Grassland
- Vegetation Survival-Mortality (SUMO)
- Water foraging with dynamic roots in E3SM; The role of roots in terrestrial ecosystem memory on intermediate timescales
- Watershed Function SFA
- WHONDRS

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
ESS-DIVE Data Contributors (science projects spread across national labs, universities etc.) - see list 5.7.2.2.4 for details	Both	Upload: Web Portal, REST API, Globus Download: Web Portal, REST API, Globus	Mean file size: 123 MB range: 0-273 GB file count: 13,093 (datasets can vary and consist of multiple files)	Ad hoc - largely driven by external project deadlines	N	
NCEAS + DataOne for data replication of all ESS-DIVE data and metadata. (University of Chicago, UC Santa Barbara at Goleta and Downtown Santa Barbara). NCEAS is also a software development partner on ESS-DIVE. (See 5.7.2.2.3)	DataOne Stores Secondary Copy of ESS-DIVE data	HTTP-based APIs	Total - 2.2 TB	Daily	N	Depends on network connection between NERSC and DataOne nodes
BER Partner facilities (See 5.7.2.2.1) at NERSC, EMSL, ORNL, LBNL - this is an effort still in its early stages to enable cross data repository sharing.	Both	HTTP (web portal and APIs)	TBD (will cover a subset of data to be shared across facilities)	TBD (will cover a subset of data to be shared across facilities)	Likely yes. mechanism TBD	Data Transfers can be slow when done outside the DTN framework

Table 5.7.1: ESS-DIVE Collaboration Space

5.7.2.3 Instruments and Facilities

ESS-DIVE is the data repository with data holdings from multiple DOE ESS projects (and some other related efforts that are relevant to the DOE ESS mission). ESS-DIVE does not generate its own data - rather it works with projects and researchers that are collecting data from the field, lab, and models. ESS-DIVE data stores its primary copy of the data on NERSC infrastructure (on the community filesystem). The data repository is hosted as a web service on the NERSC Spin Kubernetes platform, which serves as the primary mode of data ingest and egress. Backup instances of the services and data are also hosted on LBNL IT infrastructure and at NCEAS (UC Santa Barbara).

We will also present some initial research into Self-Driving Field Laboratories, a new initiative that is expected to create the need for real-time streaming distributed data services that will generate data that may be archived on ESS-DIVE.

5.7.2.3.1 ESS-DIVE Core

ESS-DIVE serves a broad range of national lab and university partners that contribute data to the system. This includes data coming from projects with a footprint at major DOE HPC facilities such as NERSC, ORNL, ANL etc. (e.g. for large model datasets). Additionally, ESS-DIVE serves partners from the BER facilities like JGI and EMSL for environmental sample data. A new area of emphasis is centered around sensor networks in the field that are collecting large volumes of data.

In the next two to five years ESS-DIVE foresees increasing growth in the following types of data:

1. Increasing data volumes from models (typically generated at large DOE facilities). Towards this end ESS-DIVE has generated a set of guidelines for archiving this type of model data - "Guidelines for Publicly Archiving Terrestrial Model Data to Enhance Usability, Intercomparison, and Synthesis."[47]
2. Collaborations with DOE FICUS facilities (EMSL, JGI) will drive increased sharing of sample metadata. These data products will be increasingly cross linked, and in some cases replicated across multiple user facing systems such as KBase and NMDC.

5.7.2.3.2 Self-Driving Field Laboratories

Beyond five years, Self-Driving Field Laboratories will be an increasingly important use case. See Section 5.7.2.4.2 for more information.

To meet the needs of rapidly expanding field data from DOE Earth Systems sensors and instruments, it will be critical to enable data acquisition, transformation, and analytics workflows that can integrate real-time data streams with predictive models, particularly for extreme events. Live data streams from sensors can then be used to train models using a hybrid of AI/ML and deterministic techniques, which can, in turn, generate real-time predictions and feedback for sensors out in the field. The two-way connectivity between models and data can also enable adaptive sampling, tailored to measuring the perturbation events of interest, and collecting the most optimal data streams when and where it is necessary.

The highly distributed nature of sensor networks, and the volumes of data collected, will necessitate a data "backbone" for AI/ML-based scientific workflows. Such a platform would support: 1. data ingestion from multiple sources across multiple domains; 2. data analysis and creation of hybrid AI/ML models, through synthesis and fusion of datasets; 3. data feedback loop for sensors enabling real-time decision-making for optimal data acquisition. The integrated platform should provide core capabilities allowing users to store, process, analyze, integrate, curate, search, and disseminate data. A toolbox of data analysis algorithms and data processing pipeline software should be maintained and available to the data producers. The ultimate aim would be to enable AI/ML driven workflows where researchers across domains can perform advanced analyses across multiple datasets and generate predictive models that can then be used directly by field sensors to target optimal, real-time data acquisition. AI/ML also provides the opportunity to perform quality assurance and control (QA/QC) against real-time streams on-the-fly, including metadata extraction, data validation, and sensor error checking.

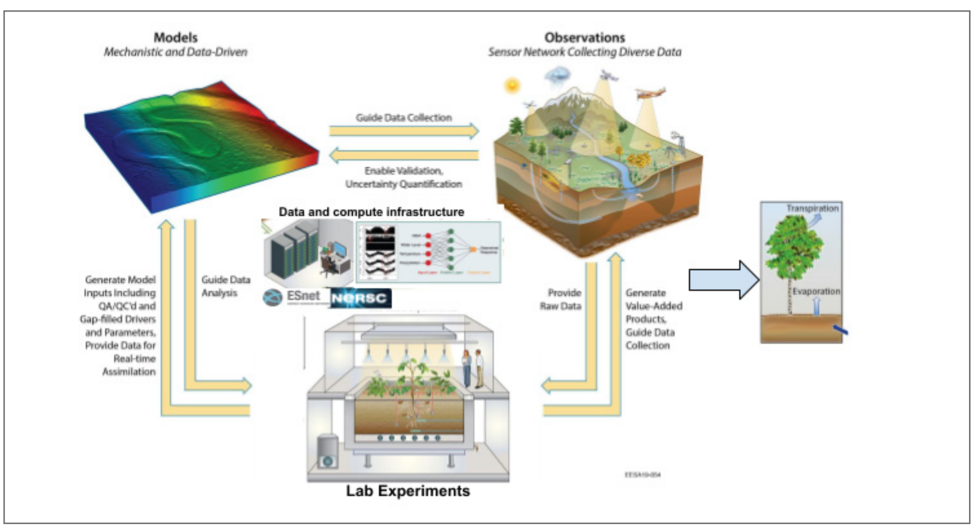

Figure 5.7.2: Self-Driving Field Laboratories

47 https://datascience.codata.org/article/10.5334/dsj-2022-003/

5.7.2.4 Generalized Process of Science

In developing and running the ESS-DIVE repository, several key principles provided guidance:

- Prioritize community needs and service;
- Build a scalable, reliable repository that can evolve for >30 years, without sacrificing user experience;
- Adhere to digital library best practices and FAIR data principles, except when they directly conflict with community needs;
- Incorporate software best practices into Standard Operating Procedures for infrastructure development and maintenance; and
- Store data for the benefit of the data contributor.

Understanding and prioritizing ESS community needs is central to the success of ESS-DIVE. A key goal is to establish a presence within the scientific and digital library communities. Community engagement activities are organized into three areas - outreach, standards development, and user support. ESS-DIVE community engagement is involved in all stages of the data publication and intake process, including development of data standards and the entire publication life cycle.

5.7.2.4.1 ESS-Dive Submission & Publication

The process for submitting and publishing data packages allows users to iteratively supplement, package, and update their data and metadata. This data package publication process involves several steps (See Figure 5.7.3). First, an authorized contributor logs into ESS-DIVE, uploads data files, and enters metadata. Next, the user submits a data package. In the third step, the user requests to publish the data package. An ESS-DIVE administrator reviews the data package, works with the user to address metadata quality issues, and assigns a DOI through OSTI, if needed. The data package then becomes publicly available on the ESS-DIVE portal and federated catalogs. And finally, the user can make revisions to the package.

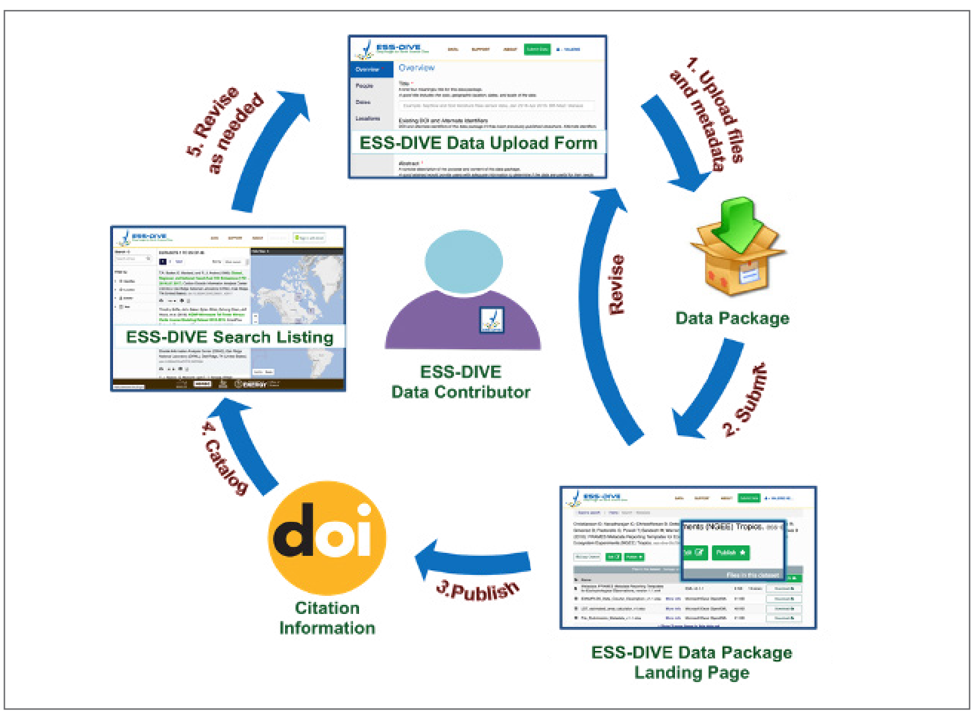

Figure 5.7.3: ESS-DIVE publication process.

We provide users three ways of submitting data packages: Graphical web interface, REST API service, or Globus. The web interface consists of interactive web forms that guide the user through the process of entering metadata for a data package and uploading files, it provides tool tips and documentation to aid users. The API service is a web service enabling users to upload metadata and data through a programmatic interface. The REST interface allows users to use scripts to create or change metadata, replace data files and add new data files up to 500GB/upload, and has already proven effective for bulk data submissions. Beyond these limits, it is necessary to use an out of band mechanism to enable users to send us data via Globus (hosted on a virtual machine [VM] at LBNL IT). ESS-DIVE support staff internally publishes this data to the system via the API by copying data into NERSC from the LBNL system.

Preservation services ensure that data is reliably stored, retained, and retrieved. It provides infrastructure for storing data packages and metadata using standards that are compatible with the DataONE federation and API. This allows us to replicate data packages on DataONE nodes.

ESS-DIVE's long-term vision is to provide a repository that enables easy intuitive integration and management of data from models, experiments, and observations. In this vision, ESS-DIVE becomes a focus for enabling access to integrated data for visualizations, models, ML, and other analyses for knowledge discovery from these datasets.

Long-term ESS-DIVE goals include the following:

- Store and distribute increasingly large and heterogeneous data with fast access mechanisms and open data licenses;
- Build data integration tools that connect and synthesize distributed datasets across systems (e.g., ESS-DIVE, ESGF, ARM, AmeriFlux, NASA, USGS, EPA, KBase, NMDC) and enable users to easily discover, access, and integrate diverse datasets;
- Provide multiscale data assimilation tools to enable real-time integration of observation data with simulation codes;
- Enable data analytics and computational capabilities for data mining and deep learning, advanced statistical and information theory algorithms for time-series and spatial analyses, and core libraries for data preprocessing such as QA/QC, subsetting, gridding;
- Support a computational framework that enables community development of scripts, gateways, and app-based tools with analytics engines to enable users to discover, query, subset, process, analyze, and store data (similar to or based on cloud infrastructure such as GCP, Google Earth Engine, Amazon Web Services Cloud, Microsoft Azure Cloud).
- Provide interactive visualizations and narrative interfaces built using recent advances in web-based tools to enable data exploration and knowledge discovery; and
- Enable an open-source software repository for sharing programs developed by the community (e.g., QA/QC and data processing scripts).

5.7.2.4.2 Self-Driving Field Labs

Despite significant advances in instrumentation, automation, and AI/ML, the complete cycle of field science research — the siting, collecting, moving, and analyzing observations, and integration of collected data with models and lab experiments, is predominantly a manual, time-intensive process often taking many months. Currently, instrument deployments and the integration of data with models requires significant human intervention at every step of the process, which significantly affects the amount, quality, and resolution of the data being collected, and their integration with models. Numerous studies such as ecosystem responses to climate change and disturbance, energy resources, and resilient infrastructure need to collect multistream data from heterogeneous field sites and feed this data into models to infer scientific insights or make decisions.

With long delays between data production and knowledge generation, it is impossible to achieve the real-time system understanding necessary to guide adaptive measurement or support timely decision-making. In addition, large heterogeneity across field sites challenges the selection of representative sites for data collection, and instrumentation deployment is often conducted in an ad-hoc fashion with questionable representativeness of the entirety of the site. This significantly limits scaling of the observations to greater spatiotemporal extents, and the ability to adapt instrument deployment to optimally collect data when and where they are needed the most.

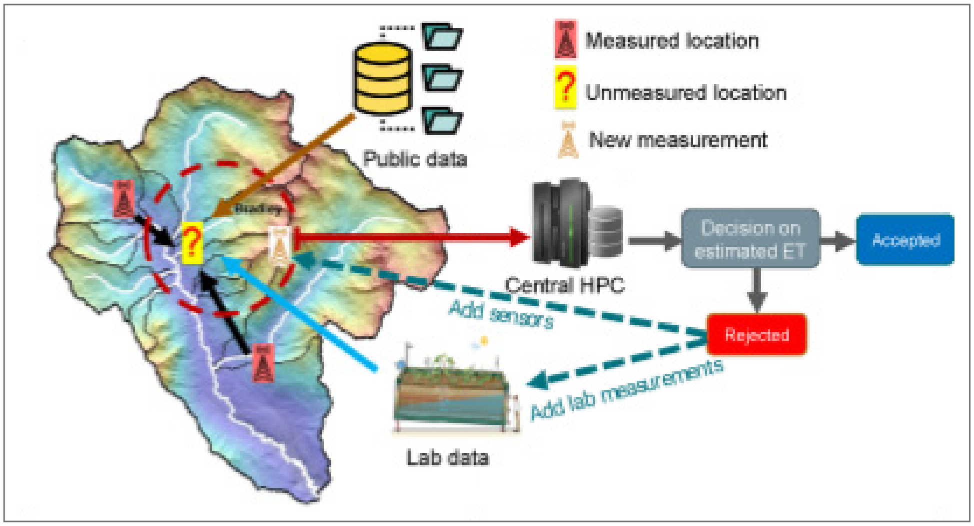

Figure 5.7.4: Workflow for the proprosed self-guiding field laboratory used for the ET case

5.7.2.4.2.1 Surface Atmosphere Integrated Laboratory (SAIL) Envirotranspiration application in Colorado.

Research underway supporting SGFL involves the development of an iterative two-way interaction between data gathering and model predictions, wherein data gathering is adapted based on model predictions utilizing a "decision engine" – an algorithm that takes the gathered data as inputs and calculates a numerical decision matrix for assessing whether data gathering should be adjusted. Such a decision matrix is tailored toward the specified scientific goals. In the case of evapotranspiration (ET), the main scientific goal is to reduce ET uncertainty across a large heterogeneous field site in Colorado, which makes the uncertainty calculation that is based on a joint-ML approach, the focus of the decision engine. Multiple ET models and data assimilation approaches are integrated to dynamically automate sampling and uncertainty reduction. The NERSC high-performance computing (HPC) user facility will be leveraged by the SGFL to provide a centralized location for data assimilation, HT uncertainty and sensitivity analysis, and scalable ML model training. This proposed SGFL will be developed via use of an exemplar Earth science use case important for water resource management.

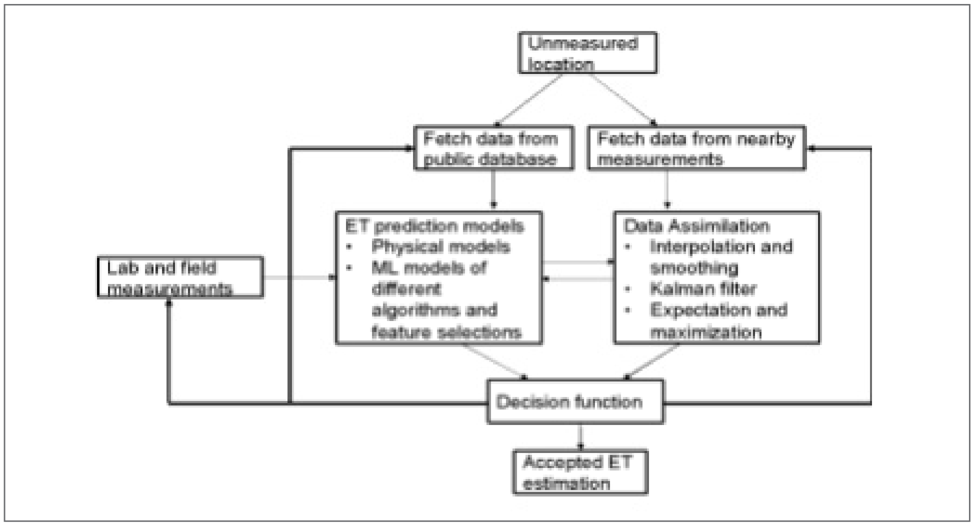

Figure 5.7.5: Detailed modeling and data assimilation workflow for the proposed SGFL

Workflow summary: The proposed ET focused SGFL development aims to integrate multimodal (time-series, ground-based, geospatial, and airborne/RS) datasets with process-based and ML models into a decision engine to guide adaptive strategies for scaling local measurements to watershed scales to reduce uncertainties in ET quantification. This SGFL capability will integrate measurement and model benchmarking and calibration capabilities of the SMARTSoils testbed to adjust instrument configuration, model parameterization and calibrations.

5.7.2.4.2.2 FY23/24 Results

1. SGFL capability for adaptive data acquisition and assimilation integrating field, lab, and model components;
2. an ET use case for accurate ET quantification across large heterogeneous field to support water resource management;
3. the establishment of a well-integrated, collaborative team and capabilities across EESA and the Computing Sciences Area (CSA) to advance the next-generation self-guiding field laboratories to support better field scale research and experimentation.

5.7.2.4.2.3 Data and Infrastructure & Networking

The exact amount of data that will be collected as part of SGFL work in the next several years is not yet known. These will probably be on the order of 10s of gigabytes over the course of the next few years, which can be handled using existing ESS-DIVE and NERSC resources which have been allocated.

In future years, however, the same trends that are causing nuclear physics, light sources, and science measurement campaigns to greatly grow data collection and management capabilities will also begin to drive similar trends for field science applications, and in particular for the SGFL. In addition, ESS-DIVE anticipates greater use of bi-directional data flows between the program “back end” at NERSC & ESS-DIVE to the field, as sensors are retasked semi-autonomously based on predictions from analytic tools and simulations driven by prior measurements. Recursive loops driving dynamic field sampling will need to operate in conjunction with control systems which may be used for mobile sensor systems, as there will be a need to coordinate and synchronize edge compute resources (and simplified codes being executed at the mobile edge) with HPC simulations being executed at NERSC or other DOE HPC centers.

Wireless edge capabilities will also play a large role in the success of these efforts, as communication links will need to be established beyond the physical span of optical connectivity to include 5G, satellite, and other wireless connections as needed. These radio frequency (RF) links will be operated in environments that are typically remote, have limited access to power, and will need to maintain connectivity for data movement to/from ESS-DIVE despite weather and other field conditions and limitations.

As designs and requirements are identified, ESS-DIVE will update information on SGFL during BER mid-cycle updates, etc.

5.7.2.5 Remote Science Activities

The ESS-DIVE primary ESS-DIVE data is hosted on resources at NERSC (under the Spin infrastructure and hosted on the community filesystem). Copies of the data and metadata are replicated to the DataONE federation. ESS-DIVE also has failover nodes on LBNL IT hardware (SVM cluster) and at NCEAS (UC Santa Barbara) - it is important to maintain high-bandwidth connections across these 3 facilities for data and service replication to support failover needs.

The DataONE network provides redundancy via replication to far-away network nodes and automatic self-healing capabilities. DataONE replication enables users to discover and download datasets if the main repository site is unavailable. ESS-DIVE maintains two standby copies of ESS-DIVE: one copy in LBNL IT data center and the other copy at NCEAS in Santa Barbara. Each failover node is able to run a copy of the stack in either read-only or read/write modes. Either node can take over as primary in the event of a failover event. The primary takes hourly snapshots of the search index and metadata database. These backups along with the data on each failover node is kept in sync with the primary on an hourly basis.

Data collected by ESS-DIVE is generally from other ESS projects, and is most likely generated on remote resources. ESS-DIVE is essentially the repository where ESS projects store their data for long-term archival. This includes national labs, universities, and individual computing resources. As such it is important for ESS-DIVE to support high-bandwidth, low-latency connections to users submitting data to ESS-DIVE.

Currently ESS-DIVE runs on compute resources hosted at NERSC, with failover nodes at NCEAS and LBNL IT. Data is submitted through HTTP interfaces - either through the web UI or a programmatic HTTP API. Currently it is important to maximize these capabilities. Moving forward in the next two to five Years, ESS-DIVE will need to increase support for large datasets stored in the repository. This will mean ensuring high-bandwidth interfaces like Globus to support large user uploads/downloads. ESS-DIVE also needs to replicate this data to other locations for redundancy and failover. Moving forward ESS-DIVE would like to add the ability to perform real-time replication for failover nodes. ESS-DIVE will also limit DataONE replication to 50 GB due to storage limitations. ESS-DIVE would like this to grow to larger datasets, to cover the entire repository holdings. It is expected that these will grow from the current 2.2 TB to 500 – 1,000 TB in the next two to five years.

The ESS-DIVE archive and catalog will be federated on external repositories using existing interfaces. This will enable users to search for and access data seamlessly across available archives, including those within ESS, EESSD, and existing data systems (e.g., NMDC, ESGF, EMSL, ARM, USGS, Sciencebase). The ESS-DIVE front-end will need the capacity to support searches and references to data on other archives. ESS-DIVE will publish ESS-DIVE metadata and make it available for consumption via other repositories and data search engines (such as the Google Dataset Search service). ESS-DIVE will begin this work with repositories that have a common metadata and API layer as part of DataONE. ESS-DIVE notes that interoperability across archives is a broader issue, involving coordination across competing requirements, standards, and interfaces, which can only be sufficiently addressed by multi-institutional effort. The approach should not be considered a replacement for these types of efforts since it is only within scope for us to tackle this through specific engagements with partner repositories. For example, the WHONDRS[48] datasets will be good potential candidates to test integration between ESS-DIVE, NMDC, and EMSL. The conversations with these repository teams and plans for initial proofs of concept are in development.

We continue to engage in the broader data standards and infrastructure discussions taking place around common APIs and platforms for accessing data. ESS-DIVE has centered development efforts on common, open standards for tools, data standards, and infrastructure that have the potential for wide adoption by the community (OpenAPI, IGSN, EML, JSON-LD etc.). This allows us to interoperate with other data providers using well-established, common tools. ESS-DIVE approaches this integration from the standpoint of being able to:

- Link to and reference data stored by external data providers directly from the data
- Serve up common metadata for search across repositories to allow some degree of metadata translation and exchange; and
- Communicate with API services exposed by other data providers through development of new tools.

In the longer horizon (beyond five years), ESS-DIVE sees the importance of streaming data, increased use of cloud resources, and federated data to facilitate these use cases. Additionally, integration and interoperability with other repositories will also be important. To enable a future-facing data repository, ESS-DIVE must create pathways to bridge the scientific data and metadata stored in other partner repositories with the data in ESS-DIVE. Given that the data ecosystem is fundamentally a distributed enterprise, scientific discovery ultimately rests on being able to deliver insights that can fuse information from different data silos.

5.7.2.6 Software Infrastructure

ESS-DIVE partners with the NCEAS/DataONE network and leverages the NCEAS MetaCat and MetaCat UI software for the repository platform. The ESS-DIVE infrastructure components (Figure 5.7.4) are based on established, well maintained, and open-source MetaCat software[49]. The ESS-DIVE web interface accessed at data.ess-dive.lbl.gov supports key user features including:

- File upload and metadata editing to create data packages for submission;
- Single sign-on based authentication using ORCID14;
- Automatic metadata generation from form fields;
- Full text and faceted search of metadata; and
- Basic user permissions management.

The ESS-DIVE user interface (UI) is implemented by leveraging DataONE's Metacat UI, which has a rich set of features including support for relevant metadata standards, metadata editing, data upload, data package search, repository statistics, data package usage information, and user authentication using ORCID[50]. The ESS-DIVE UI customizations include:

- Styling to harmonize with ESS-DIVE branding and theme;
- User registration workflow to allow authorization of new data contributors;
- Support for the ESS-DIVE review and data publication process;

48 https://data.ess-dive.lbl.gov/portals/WHONDRS
49 https://github.com/nceas/metacat
50 Hendrix, V., C. Snavely, and S. Cholia. 2019. ESS-DIVE: A Scalable Community Repository for Managing Earth and Environmental Science Data. https://doi.org/10.17605/OSF.IO/8C3FA

- Configuration via files to enable seamless UI deployment for different environments; and
- Customization of metadata (JSON-LD) embedded in the data package landing pages to enable Google dataset searches[51]

In addition to MetaCat and MetaCatUI, ESS-DIVE has developed support services for managing users and controlled lists of metadata providing programmatic REST API access to datasets (api.ess-dive.lbl.gov) and handling large volumes of data (Globus). These services have been integrated with the main ESS-DIVE archive and are integral to the mission of long-term stewardship of DOE ESS data.

Our current software stack includes:

- Metacat (Tomcat based service) for data management - preserve data for the long term with persistent identifiers.
- Metacat UI front end - for interacting with archive
- Postgres backend - stores dataset metadata
- Solr - index for data search
- API Services - Uses FastAPI and Django web framework technologies to provide public metadata, user and dataset management interfaces
- Globus for handling large data volumes

We are in the process of deploying a secondary "tier-2" storage layer, where metadata is searchable and available up in the Metacat system, but it links out to files directly on the filesystem, so that ESS-DIVE can support very large datasets and hierarchical data. These will be directly accessed from the filesystem via Globus or HTTP bypassing the ESS-DIVE Metacat System, and will serve the growing data (500 TB-1 PB) needs in the next two to five years.

Looking beyond that horizon ESS-DIVE may rely on different storage backends including use of object storage technologies to manage data, and federated access to data including use of multiple facilities and cloud resources.

5.7.2.7 Network and Data Architecture

The ESS-DIVE infrastructure components (Figure 5.7.4) are deployed on a modern microservices architecture. The main archive (data.ess-dive.lbl.gov) consists of three microservices: a web front-end, an application server, and a database. The number of web front-end microservices can be scaled up to meet user demand and allows ESS-DIVE system copies to be run alongside each other or independently. The archive is supported by additional microservices (e.g., api.ess-dive.lbl.gov), which support additional features for users and administrators.

51 https://datasetsearch.research.google.com/

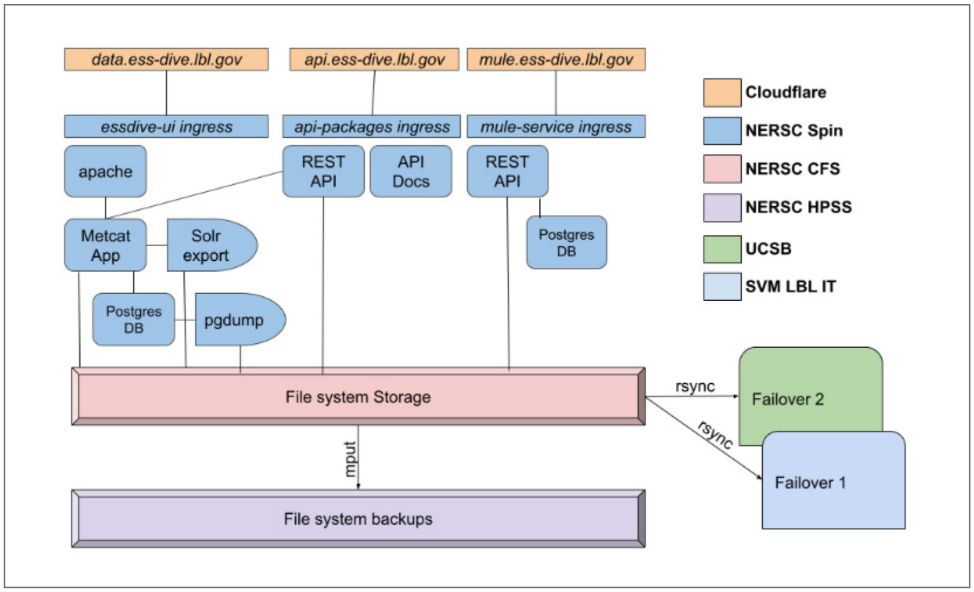

Figure 5.7.6: ESS-DIVE Data Architecture

Main Archive Infrastructure

ESS-DIVE runs at NERSC on Spin (medium blue boxes), the NERSC platform based on Docker containers, a microservices technology[52]. This deployment stores scientific data, metadata and log files in the NERSC Community File System (NERSC CFS) as seen in the pink box above, a global file system available on all NERSC computational systems that is a large, permanent, medium-performance file system[53]. All data that is stored in NERSC CFS is backed up daily to NERSC's HPSS[54] as seen in the purple box above. HPSS is a modern, flexible, performance-oriented mass storage system intended for long-term storage of data that is not frequently accessed.

Failover and Redundancy

On an hourly basis, all ESS-DIVE data is synchronized to two separate failover systems outside of NERSC. One system is at UC Santa Barbara (green box) provided by NCEAS and the other is in a different building at Berkeley Lab (light blue box) provided by Berkeley Lab Information Technology (LBNL IT). Both systems are identically configured and can bring up the entire ESS-DIVE deployment as seen above in the medium blue, pink and purple boxes (Figure 5.7.4). Additionally, all public datasets are replicated automatically via the MetaCat software over the DataONE infrastructure. Each dataset has at least two replicas outside of ESS-DIVE. These replicas are verified minimally every 90 days by calculating and comparing checksums across all replicas.

Other Data Services

For Globus Uploads ESS-DIVE uses an external VM on LBNL ITs SVM infrastructure (Globus Sharing Upload features at NERSC are limited to users that are part of the NERSC allocation). Data is then transferred to NERSC

ESS-DIVE is also deploying a 500 TB Ceph based storage platform attached to the LBNL IT Lawrencium cluster and VM systems, to provide dedicated scalable storage for the project (outside shared NERSC filesystem resources). This will serve as a location for data synthesis and analysis for ESS-DIVE and related datasets.

52 Snavely, C., G. Alvarez, V. Hendrix, S. Cholia. and S. Lasiewski. 2018. Spin: A Docker-based Platform for Deploying Science Gateways at NERSC. https://doi.org/10.6084/m9.figshare.7071770.v3
53 https://docs.nersc.gov/filesystems/community/
54 https://docs.nersc.gov/filesystems/archive/

Cloudflare for Public Access

Public access to the major services (data.ess-dive.lbl.gov, api.ess-dive.lbl.gov) is managed via Cloudflare which gives an integrated set of L3-L7 network services from a single dashboard[55]. ESS-DIVE uses Cloudflare to manage DNS minimally as ESS-DIVE cannot take full advantage of Cloudflare's DNS proxied services to optimize, cache, and protect all requests coming into ESS-DIVE. The reason for this is that Cloudflare's proxied service does not allow incoming requests to be larger than 500MB. Since ESS-DIVE is a data repository which allows upload sizes up to 500 GB this did not fit the requirements.

Future

In the long term, ESS-DIVE may consider a more federated production infrastructure taking advantage of cloud resources and partner facilities with transparent failover and high availability across multiple locations. ESS-DIVE will also investigate an infrastructure that supports streaming data from sensor networks.

5.7.2.8 Cloud Services

ESS-DIVE uses various cloud-based software-as-a-service offerings, for support services.

Our overall strategy is to continually evaluate resource options such as cloud computing to support ESS-DIVE. The cloud computing model enables us to leverage a wide range of software, storage, and computing hardware resources under the "-as-a-service" model, where ESS-DIVE can deploy and leverage the scaling and availability of cloud-based services. Both dev-ops, and resilience strategies, actively leverage resources on public clouds when those services meet project needs. These cloud software services include:

- GitHub
- CircleCI
- Dockerhub
- Zenhub
- Gitbook
- Jira

Moving forward cloud storage and computation may play a part in the repository deployments. Maintaining a replicated instance of the repository on cloud resources for failover is under consideration for the next phase of ESS-DIVE. ESS-DIVE will be evaluating the options based on overall cost, services, and capabilities. This includes enabling cloud-based data ingest mechanisms , cloud-based data analysis tools, and cloud storage for replicating data.

In the long run, ESS-DIVE sees a use case for having replicated copies of ESS-DIVE data on public cloud resources. Object storage technologies for data with S3 compatible interfaces are also being considered. This would give us direct and transparent access to storage objects, while paving the way for a hybrid cloud / on-premises data model, where data could live in either location and is accessed through a common interface.

Additionally, ESS-DIVE is also considering a model where the data indexes and web services may live in the cloud while raw storage for large datasets would still live on ESS-DIVE managed hardware due to cost considerations.

55 https://www.cloudflare.com/

5.7.2.9 Data-Related Resource Constraints

ESS-DIVE is currently rolling out capabilities to store, upload, and download increasingly larger datasets. In the next two to five years it is expected to be able to store large datasets on ESS-DIVE (datasets in the 1–10 TB range) to support model data archiving needs[56]. There will be an increasing need for high-performance mechanisms like Globus and high-bandwidth networks to support these use cases. Additionally, it is expected that the repository will grow to the petabyte range, and ESS-DIVE will need to work with DataONE partners to enable data replication of these larger volumes. Data Replication to DataONE currently limited to files < 50 GB.

In terms of the current infrastructure, as a primarily HTTP web-based service ESS-DIVE faces the Cloudflare data limitations and issues. Currently LBNL resources behind Cloudflare have a very small upload limit of 500 MB which makes it unusable for uploading large files. ESS-DIVE were able to work with LBNL IT to increase this limit to 100 GB, but this is considered a hard limit. ESS-DIVE is forced to bypass Cloudflare due to limits on data uploads larger than 100 GB (for API and data portal services), which results in a misalignment with LBNL security policies requiring Cloudflare.

Additionally, use of Cloudflare may force data transfers down a sub-optimal network route where data is moving in and out of Cloudflare rather than through ESnet. E.g., data transfers between NERSC and LBNL hosts, being effectively routed externally and back through Cloudflare.

5.7.2.10 Outstanding Issues

SDFL research is at an early stage, but ESS-DIVE anticipates that it will grow in the future. Current work is being funded via a combination of LBNL LDRD and DOE resources, via cross-disciplinary collaboration between the Computing Sciences (CS) and EESA areas.

As resources like ESS-DIVE decrease the time and burdens upon researchers seeking to share data sources, SDFL is an affiliated set of efforts to increase the velocity of data collection from field sensor systems to repositories like ESS-DIVE, and to improve the ability for sensor systems to dynamically adapt data collection sampling to better support field science efficiency.

5.7.2.11 Facility Profile Contributors

ESS-DIVE Representation

Deb Agarwal, *LBNL,* daagarwal@lbl.gov

Shreyas Cholia, *LBNL,* scholia@lbl.gov

Charuleka Varadharajan, *LBNL,* cvaradharajan@lbl.gov

Valerie Hendrix, *LBNL,* vchendrix@lbl.gov

ESCC Representation

Rune Stromsness, *LBNL,* rstrom@lbl.gov

Richard Simon, *LBNL,* rsimon@lbl.gov

56 https://datascience.codata.org/article/10.5334/dsj-2022-003/

5.8 The Structural Biology Center (SBC) at the APS

The SBC is a national user facility for macromolecular crystallography at the APS. The facility uses advanced instrumentation, state-of-the-art software, and methods and HT technologies. Located at Sector 19 of the APS, ANL, SBC provides the scientific community with two experimental beamlines: an insertion device, 19-ID, and a bending-magnet, 19-BM. The beamlines are well suited to a wide range of crystallographic experiments including:

- Crystals of macromolecular assemblies with very large unit cells
- MAD/SAD phasing
- Crystals of membrane proteins
- Small, weakly diffracting crystals
- Ultra-high resolution crystallography
- Cryo-crystallography

5.8.1 Discussion Summary

- SBC performs macromolecular crystallography at the APS. SBC focuses on structural biology and contributes to the expansion of existing programs and exploration of new opportunities in structural biology, proteomics, and genomics research with a major focus on medicine, bio-nanomachines, and biocatalysis highly relevant to energy resources, health, a clean environment, and national security.
- SBC is used for crystallographic experiments including:
  - Crystals of macromolecular assemblies with very large unit cells.
  - MAD/SAD phasing.
  - Crystals of membrane proteins.
  - Small, weakly diffracting crystals.
  - Ultra-high resolution crystallography.
  - Cryo-crystallography.
- SBC has implemented serial crystallography experiments, which enable the study of protein structures under physiological temperature and reduced radiation damage by collection of data from thousands of crystals.
- The data acquisition is performed using DECTRIS PILATUS 6M at 19ID, and Area Detector Systems Corporation (ADSC) Quantum 315 (Q315) at 19 BM. The maximum data rate for 19ID with PILATUS 6 M is 100 images per second (650 Mbytes per second), and for 19 BM – 1 full image per 2.7 seconds (30 Mbytes per second).
- Synchrotron serial crystallography (SSX) has emerged as a valuable approach for low-dose room-temperature structural biology research that also allows for the study of dynamic processes in protein crystals, such as chemical transformations. For the real-time processing of the SSX data collection, SBC has developed a data analysis pipeline, Kanzus, to bridge the beamline with high-performance computing (HPC) and storage capabilities provided by Argonne Leadership Computing Facility (ALCF).

- SBC features a three-tier software architecture to handle experiments:
  - The lower level is Experimental Physics and Industrial Control System (EPICS) and beamline device controls.
  - The middle layer is SBCserver, which is an EPICS-based sequencer with the ability to be driven by socket commands.
  - The upper levels are GUIs, namely SBCcollect.
- SBC instrumentation consists of X-ray optical elements, beam position monitors, XRF measurements instruments, robotic sample mounting stages, visualization camera systems with auto-alignment software, and diffraction images recording detectors (ADSC Q315 and PILATUS3 6M).
- The standard SBC data handling process is to collect data from sensors to local storage, while also copying this data simultaneously and providing it to the user's home institution or portable data storage. At the same time, the data is processed, usually several times on-site, and then after the experiment in the user's home institution. The design objective of the infrastructure is to give users flexibility in use of whatever protein crystallography data processing software they require.

5.8.2 SBC Facility Profile

The SBC mission is to advance and promote scientific and technological innovation in support of the DOE mission by providing world-class scientific research and advancing scientific knowledge:

- SBC provides advanced data collection facilities at the APS to the national user community and maintains high-profile user programs
- SBC exploits major advances in macromolecular X-ray crystallography and addresses the most challenging structural biology problems to advance scientific knowledge
- SBC is an important component of integrated biosciences and contributes to the expansion of existing programs and exploration of new opportunities in structural biology, proteomics and genomics research with a major focus on medicine, bio-nanomachines, and biocatalysis that are highly relevant to energy resources, health, a clean environment and national security

5.8.2.1 Science Background

The SBC (Rosenbaum & Westbrook, 1997a) at ANL was established by the DOE as a national user facility to serve the expanding and highly demanding macromolecular crystallography community. The SBC was intended to provide highly efficient and flexible X-ray beamlines capable of not only dealing with the very large volume of crystallographic projects but also of satisfying the requirements of the most challenging and diverse crystallographic experiments, like, for example, as a result of major initiatives such as the Protein Structure Initiative (structural genomics) and related initiatives in the wake of the Human Genome Project. Examples of the latter are crystals with very large unit cells (viruses, macromolecular assemblies), very small crystals (microcrystals) and very weakly diffracting crystals (ribosome and other macromolecular assemblies). In addition, the beamline design and instrumentation was intended to meet the demanding criteria laid down by crystallographic techniques such as multi-and single-wavelength anomalous diffraction (MAD/SAD) (Hendrickson, 1985, 1991) and direct methods applied to macromolecules (Sheldrick, 1998; Weeks & Miller, 1999). Finally, the design called for open access to the sample area to improve crystal mounting, positioning and visualization and for the appropriate space to prepare samples for cryofreezing. The SBC collaborative access team operates Sector 19 of the APS (Moncton et al., 1989). The sector consists of two beamlines: bending magnet (19BM) and undulator insertion device (19ID).

5.8.2.2 Collaborators

The SBC user base is very large. Every year SBC has above 200 user's groups visiting personally or remotely using beamlines and performing multiple experiments. The standard data handling process is to collect data from sensors to local storage, while also copying this data simultaneously and providing it to the user's home institution or portable data storage. At the same time the data is processed, usually several times, on-site and then after their experiments in their home institutions. The design objective of the infrastructure is to give users flexibility in use of whatever protein crystallography data processing software they require.

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
Beamline Users	Primary	Globus	GB to TB	Ad hoc	N	N
ALCF	Secondary	Globus	GB to TB	Daily	Y	N

Table 5.8.1: SBC Collaboration Space

5.8.2.3 Instruments and Facilities

SBC instrumentation consists of X-ray optical elements, beam position monitors, XRF measurements instruments, robotic sample mounting stages, visualization camera systems with auto-alignment software, diffraction images recording detectors (ADSC Q315 and PILATUS3 6M).

X-ray optical elements consists of:

- primary aperture
- The monochromator is a double-crystal design with constant height and direction of the monochromatic output beam and a sagittally bent second crystal for horizontally focusing.
- vertically focusing mirror
- horizontal and vertical collimation guard slits
- CRL (compative refractive lenses) at 19ID

All motors are controlled by the Delta Tau PMAC Turbo Ultralight motor controllers, using MACRO fiber-ring (for each beamline, around 64 motors). The software interface consists of Delta Tau device driver integrated into SNL sequencers (State Notation Language), and presented to the EPICS by multiple input/output controllers (IOCs) on different platforms (Windows, Linux, Solaris).

XRF instrumentation consists of Vortex detector, Mercury and Saturn MCA (Multi Channel Analyzer) electronics and EPICS IOCs.

Beam Position Monitors are quad pin diodes assemblies placed at several points at each beamline for positional and quantitative feedback from optical elements.

At 19ID there is Rigaku 6-axis robotic arm sample handling robot, with the endefector adapted to cryo conditions.

At 19BM there is an ALS-type pneumatic motors, cryo-sample mounting robot, designed and built by SBC.

Visualization camera systems are based on high-resolution visible light, CCD camera and optics including on-axis mirror.

5.8.2.4 Generalized Process of Science

Recently, SBC has implemented serial crystallography experiments, a fixed-target approach with a new 3D-printed mesh-holder optimized for sample handling. Serial synchrotron crystallography enables the study of protein structures under physiological temperature and reduced radiation damage by collection of data from thousands of crystals.

Synchrotron serial crystallography (SSX) has emerged as a valuable approach for low-dose room-temperature structural biology research that also allows for the study of dynamic processes in protein crystals, such as chemical transformations.

The data acquisition is performed using DECTRIS PILATUS 6M at 19ID, and ADSC Quantum 315 (Q315) at 19BM. The maximum data rate for 19ID with PILATUS 6M is 100 images per second (650 Mbytes per seccond), and for 19BM – 1 full image per 2.7 seconds (30 Mbytes per second).

There are two major types of experiments at SBC: traditional protein crystallography, and the goal over the years was to give the users ability to process data during the data collection. Traditional protein crystallography experiments use several pipelines developed over the years, using multiple software approaches, dependent on the user's choice. SSX experiments require maximum data collection rates. For the real-time processing of the SSX data collection, SBC developed a data analysis pipeline, Kanzus, to bridge the beamline with high-performance computing (HPC) and storage capabilities provided by Argonne Leadership Computing Facility (ALCF).

5.8.2.5 Remote Science Activities

None to report at this time.

5.8.2.6 Software Infrastructure

SBC has developed a 3-tier software architecture to handle experiments. The lower level is EPICS and beamline device controls, the middle layer is SBCserver, which is EPICS-based sequencer (written in SNL language) multithreaded and multiprocess software, with ability to be driven by socket commands with different "personalities" (different command syntaxes), and the upper level are GUIs, with the major one – SBCcollect, written in C/C++/Java/GTK+/Python, and other experiment interfaces, mostly written in Python and MEDM.

5.8.2.7 Network and Data Architecture

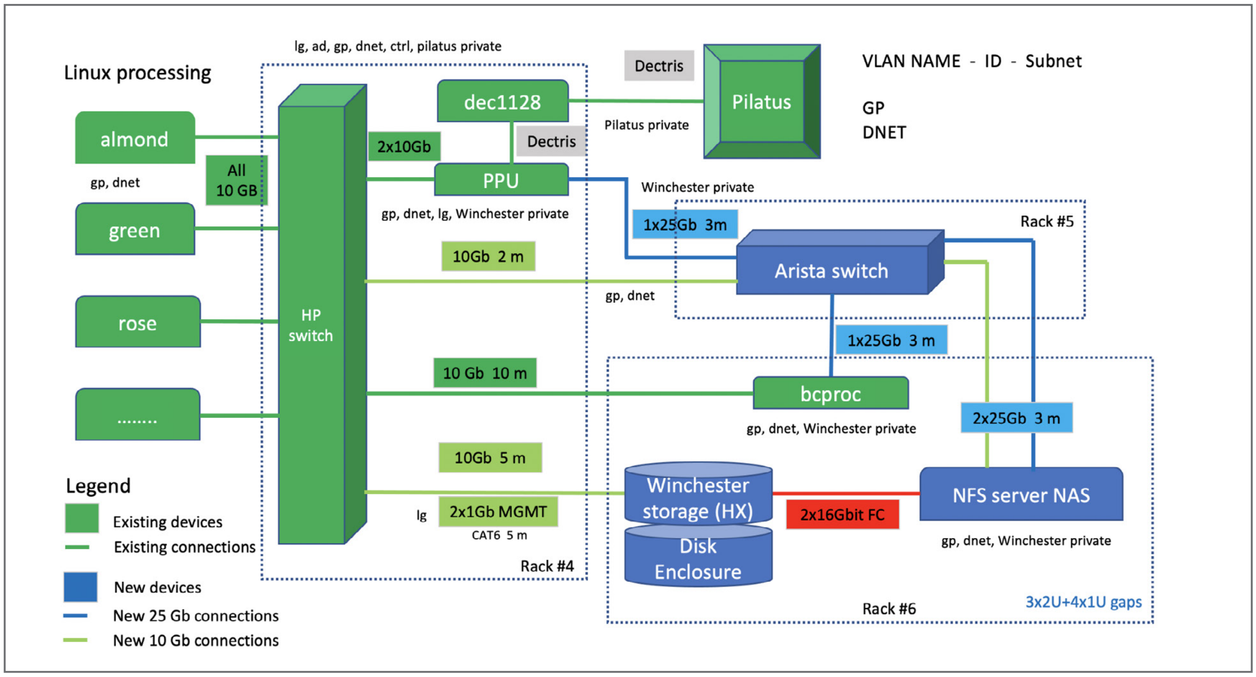

Figure 5.8.1: SBC ID Figure Network Diagram

5.8.2.8 Cloud Services

None to report at this time.

5.8.2.9 Data-Related Resource Constraints

None to report at this time.

5.8.2.10 Outstanding Issues

None to report at this time.

5.8.2.11 Facility Profile Contributors

SBC Representation

Krzysztof Lazarski, *ANL,* lazarski@anl.gov

ESCC Representation

Linda Winkler, *ANL,* winkler@mcs.anl.gov

Corey Hall, *ANL,* chall@anl.gov

5.9 CABBI

The CABBI is a bioenergy research center comprising over 300 researchers at the University of Illinois and 20 partner institutions. CABBI leverages key facilities at the University of Illinois such as the Carl R. Woese Institute for Genomic Biology (IGB), Illinois Energy Farm, SoyFACE: the Integrative Bioprocessing Research Laboratory pilot plant, and the Illinois Biological Foundry for Advanced Biomanufacturing (iBioFAB), as well as facilities for plant transformation, field trials, computation, and diverse molecular, chemical, and physiological analyses at partner locations.

5.9.1 Discussion Summary

- CABBI is a bioenergy research center at the University of Illinois with 20 partner institutions. CABBI will integrate recent advances in agronomics, genomics, and synthetic and computational biology to increase the value of energy crops using a "plants as factories" approach to grow fuels and chemicals in plant stems. CABBI will also use an automated foundry to convert biomass into valuable chemicals and ensure that its products are ecologically and economically sustainable.
- CABBI leverages key facilities at the University of Illinois, such as the Carl R. Woese Institute for Genomic Biology (IGB), Illinois Energy Farm, SoyFACE (free-air CO2 enrichment), the Integrative Bioprocessing Research Laboratory pilot plant, and the Illinois Biological Foundry for Advanced Biomanufacturing (iBioFAB), as well as facilities for plant transformation, field trials, computation, and diverse molecular, chemical, and physiological analyses at partner locations.
- CABBI is organized around three interconnected themes:
  - Provide an integrated economic and environmental framework for determining feedstock supply and its sustainability.
  - Provide a regionally adaptive yet national-scale platform for grass-based biorefining based on high-yielding feedstocks with improved environmental resilience.
  - Provide a broad set of platform microorganisms, and automated tools to engineer them, to produce value-added products from plant-produced feedstocks or substrates.
- CABBI's datasets include phenomic, genomic, transcriptomic, metabolomic, proteomic, and fluxomic data for both bioenergy feedstocks and the engineered microbes, in addition to bioeconomic, edaphic, and climatic data for sustainability determination.
- The CABBI data management plan focuses on each researcher or lab group being responsible for developing data collection and storage procedures that suit their work with the requirement that data is shared, well-documented, and backed up in multiple secure locations:
  - Researchers are required to describe their data (i.e., variables collected, temporal or spatial resolution), where it is stored, and who is responsible for it on CABBI's researcher website.
  - Once data has been analyzed and quality tested, researchers are then required to share datasets internally as manuscript development is undertaken.
  - Data-sharing locations include cloud-based locations, local servers, and web applications.
  - Researchers are required to publicly share supporting data after publication.
- The CABBI iBioFAB is an automated laboratory facility located in the IGB. The system supports rapid design, fabrication, validation/quality control, and analysis of genetic constructs and organisms:

- A variety of data is generated for the various iBioFAB instrumentations. The number of files and the size are dependent on the analysis being performed. All data is retrieved for the computers attached to the instruments.
  - Data is manually aggregated for the various instruments and emailed to the requesting researchers for analysis.
  - There are general plans for the software team to develop standardized methods for storage and delivery.
  - While data collection and transfer are primarily performed at the facility and can make use of the University of Illinois network, at some future point, users of the facility may want to remotely operate instrumentation. This could potentially become a challenge for the current protocol.
- All data supporting CABBI publications is released either via Box or through publishing via a data repository. Subject-specific repositories like the National Center for Biotechnology Information (NCBI), AmeriFlux, DAAC, and Phytozome are preferred, but where those are unavailable, others, like university (Illinois Data Bank and Iowa State University DataShare) and public (Zenodo, Mendeley, and GitHub) repositories are used.
- CABBI DNA sequencers are used extensively and produce large datasets:
  - The main data-producing technology is currently the Illumina system. The system's instruments generate up to 8 terabase pairs per day, which would translate to 2 to 20 terabytes depending on the data format, retention of quality values, nature of sequence, and the compression used.
  - Current CABBI usage is around 30 terabases per year, which in raw data usually requires 50 terabytes or so in the formats usually employed.
  - Data is retrieved from the DOE facilities the JGI and EMSL. Data is also sometimes shared between CABBI partner sites.
- CABBI uses a suite of models including three ecosystem models, an economic model, and a techno-economic analysis/life cycle assessment model. The researchers using these systems collaborate frequently, sharing both data and code:
  - Some modelers make use of BioCluster, IGB's fee-based high-performance computing resource. BioCluster contains 2,824 cores and over 27.7 TB of RAM with 1.3 PB of storage on a general parallel file system (GPFS).
  - Researchers archive their simulations, including input data, code, and results, in a repository stored on IGB's file archive server.
  - The uploaded files vary in number and size depending on the scope of the simulation and the model being used. Currently, archived simulations require 2 MB to 110 MB.
  - Researchers access this server via their individual networks using a secure file transfer protocol (SFTP) client.
- CABBI has CABBI has automated sensors for collecting environmental data at field trial locations in Illinois, Iowa, and Florida. The instrumentation includes:
  - Meteorological sensors (wind, temperature, precipitation, radiation).
  - Soil sensors (soil temperature and moisture).
  - Gas sensors (concentration, flux).

- CABBI field sites sometimes have permanent infrastructure (e.g., power and networking) to support automated collection of data and operation of instruments. Other times it may be necessary to manually retrieve data using removable media:
  - Field data is collected continuously all year. Each file holds 30 minutes of data collection, and is approximately 5 MB in size. At the end of the year, field data is compressed to 70 MB in size. The meteorological and soil data each year generates approximately 30 MB of data.
  - Remote sites can make use of commercial networking, which is expensive and less reliable than the facilities provided at a university.
  - Remote access is not feasible for many of the field measurements because they are done at the plant level. It would be expensive and unfeasible to connect sensors to specific plants that will need to be removed from the field at harvest.
- The IGB, where CABBI is headquartered and housed on the Illinois campus, works within the University of Illinois's infrastructure to provide network access:
  - For internal data sharing, the IGB provides CABBI 10 TB of shared storage. Storage is accessible only using SFTP clients.
  - Globus is integrated into the biotransfer system of IGB's high-performance computing resource, BioCluster.
  - A perfSONAR node will be made available for validating end-to-end performance.
  - These networking services are expected to stay the same for the next three to four years.
- CABBI features centralized networking, data storage, and compute facilities at the University of Illinois are the main source of computer resources. These resources have connections to ESnet via the National Center for Supercomputing Applications (NCSA) on the Illinois campus.
- CABBI has access to Amazon, Google, and Microsoft cloud computing resources through the University of Illinois. To date, only AWS has been used to host a web tool developed by CABBI researchers. Additionally, cloud computation through AWS was explored in the past when a Windows operating system was required for specific software applications. It was found to be cost-prohibitive at that time, and local resources were used. For other computational needs, the available local resources are sufficient and cheaper to use.
- Cloud-based commercial file sharing and file management are used for files 150 GB or smaller. Files are shared both within CABBI with restrictive permissions, and publicly via a link from the research data website.
- Current constraints include using commercial wireless networks for data transfer from remote environmental instrumentation. Finding cheaper and more reliable options for this would benefit CABBI's science.
- As the rate of genomic data increases, increased network connectivity from CABBI to facilities like the JGI and EMSL will be necessary to continue these partnerships.
- In the future, as CABBI moves towards more remote measurements in both the laboratory and the field, network security will become more critical. Ensuring outside actors are unable to take control of equipment will be necessary not only for the protection of the equipment and data, but also for public safety. Determining the appropriate security restrictions without compromising the commitment to ensuring CABBI research is usable and reproducible by others will be challenging.
- CABBI and ESnet can share information on remote laboratory environments that leverage cellular and satellite connectivity options.

5.9.2 CABBI Facility Profile

One of the major challenges the world faces is how to provide sustainable sources of energy that meet societal needs as the population continues to grow. CABBI will develop efficient ways to grow bioenergy crops, transform biomass into valuable chemicals, and market the resulting biofuels and other bioproducts.

CABBI will integrate recent advances in agronomics, genomics, and synthetic and computational biology to increase the value of energy crops — using a "plants as factories" approach to grow fuels and chemicals in plant stems, an automated foundry to convert biomass into valuable chemicals, and ensuring that its products are ecologically and economically sustainable. This holistic approach will help reduce our nation's dependence on fossil fuels, thus increasing national security.

5.9.2.1 Science Background

CABBI is organized around three interconnected themes, which align with the DOE priority research areas of Sustainability, Feedstock Production, and Conversion. CABBI addresses the following long-term goals:

1. Provide an integrated economic and environmental framework for determining feedstock supply and its sustainability;
2. Provide a regionally adaptive yet national-scale platform for grass-based biorefining based on high-yielding feedstocks with improved environmental resilience;
3. Provide a broad set of platform microorganisms, and automated tools to engineer them, to produce value-added products from plant-produced feedstocks or substrates.

To achieve these diverse goals, many types of data are collected and used by CABBI researchers. CABBI's datasets include phenomic, genomic, transcriptomic, metabolomic, proteomic, and fluxomic data for both bioenergy feedstocks and the engineered microbes, in addition to bioeconomic, edaphic and climatic data for sustainability determination.

Figure 5.9.1: Map of CABBI partner locations.

Due to the breadth of research data types, instrumentation, and protocols at the various CABBI partner locations, the data management plan focuses on each researcher or lab group being responsible for developing data collection and storage procedures that suit their work with the requirement that data is shared within the group, are well-documented, and are backed up in multiple secure locations. During the analysis period, researchers are only required to describe their data (i.e., variables collected, temporal or spatial resolution), where it is stored, and who is responsible for it on CABBI's researcher website to encourage collaboration within the Center. Once data has been analyzed, and quality tested, researchers are then required to share datasets internally as manuscript development is undertaken. Data-sharing locations include cloud-based locations, local servers, and web applications. These are documented in Section 6. Researchers are required to publicly share supporting data once the corresponding manuscript is published. CABBI depends on public data repositories including subject-specific, DOE, and general for this, as these will ensure data is findable and accessible immediately and into the future. Researchers are asked to use non-proprietary formats for storage and to ensure descriptive and complete documentation is also submitted. When public repositories are not appropriate, supplemental files are added to the manuscript. These data and links to data in repositories are made available via CABBI's public website.

5.9.2.2 Collaborators

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
JGI (Berkeley, CA or Huntsville, AL)	Primary	JGI data portal (https or Globus) or Phytozome (https) for completed genomes	Typically files are 109-1010 bytes, datasets are 10-100 files. The 2022 JGI allocation to CABBI is 26.4 terabase pairs, conversion varies according to sequence but likely tens of terabytes and often multiple transfers are needed.	Ad hoc on completion of runs, can be daily, weekly or monthly	Not usually. Occasionally a small sequence is sent back for analysis	-Files are large and speed could be faster, but generally OK using Globus
EMSL (at PNNL, Richmond, WA)	Primary	EMSL portal is recommended but requires generation of download link that often freezes. Hard to transfer data directly to the HPC system. Have recently used physical hard drive transfer from EMSL (sneakernet)	Most data is sequence files similar in size to JGI. Fewer samples (perhaps 5-10 per year but with more data per sample). 1.6 x 1011 bytes per year predicted.	Ad hoc on completion of experiments	Not usually	Portal has some glitches. Sneakernet is a high-latency mode of transfer (though bandwidth is great).

Table 5.9.1: CABBI Collaboration Space

The data in both cases is exclusively or mostly genetic sequence data. See Case Study 1, Genomic Sequencing.

5.9.2.3 CABBI Science Use Cases

The CABBI facility is a collaboration of multiple University and National Lab groups focused on biology, biochemistry, chemistry and environmental sciences. Centralized networking, data storage and compute facilities at the University of Illinois are the main source of computer resources currently. These resources have excellent connections to ESnet via the NCSA on the Illinois campus. CABBI operations include many disciplines. Four case studies of data-intensive research are described below:

- Genomics Sequencing
- Field Data Collection
- Modeling
- iBioFAB

5.9.2.3.1 Genomics Sequencing

5.9.2.3.1.1 Instruments and Facilities

DNA sequencers are used extensively and produce large datasets. The sequencers are at the University of Illinois within CABBI, at JGI, and at multiple other sites. The main data producing technology is currently the Illumina system, particularly the NovaSeq 6000 and NovaSeq X series. These instruments generate up to 8 terabase pairs per day, which would translate to 2-20 terabytes depending on the data format, retention of quality values, nature of sequence and the compression used. Current CABBI usage is around 30 terabases per year, which in raw data usually requires 50 terabytes or so in the formats usually employed. The Center and JGI are also increasingly using the PacBio Sequel II sequencer, which produces longer DNA reads but does not have as high a data output as the Illumina system currently, though it is increasing rapidly, and the Oxford Nanopore system, which is lower than PacBio but still produces significant data, and finally the older capillary electrophoresis system, which is still widely used but has relatively low throughput. The data production rate of DNA sequencers has historically doubled every seven months. See Stephens et. al, 2015[57], and Illumina[58] and PacBio[59] websites.

5.9.2.3.1.2 Generalized Process of Science

Samples of DNA or RNA are collected from plant, microbial or ecosystem samples and the nucleic acid is extracted and purified in the laboratory. DNA is then fragmented and sequenced in a large number of short “reads” consisting of molecular fragments of the original DNA. RNA is treated differently depending on the sequencing protocol, it is reverse transcribed into DNA first, then either fragmented or sequenced as a single molecule. Once a population of molecules is prepared, a “library” is then constructed using molecular linkers made of DNA and sometimes other compounds, and the base pair sequence of the individual molecules of the sample are read either directly from the single molecule via quantum mechanical inferences (PacBio) or electrochemical interactions (Oxford Nanopore) or the molecule is amplified using the polymerase chain reaction (PCR) then sequenced by the addition of fluorescent moieties that are sequence-specific and visualization by laser-induced fluorescence (Illumina and capillary sequencing methods). For an overview of the science process (the technology descriptions are outdated and should be obtained from the links above) see Hudson, 2008:

https://doi.org/10.1111/j.1471-8286.2007.02019.x

57 https://doi.org/10.1371/journal.pbio.1002195
58 https://www.illumina.com/systems/sequencing-platforms.html
59 https://www.pacb.com/technology/hifi-sequencing/sequel-system/

5.9.2.3.1.3 Remote Science Activities

Data is retrieved from the DOE facilities the JGI and EMSL. Data is also sometimes shared between CABBI partner sites.

5.9.2.3.2 Field Data Collection

5.9.2.3.2.1 Instruments and Facilities

CABBI has automated sensors for collecting environmental data at field trial locations in Illinois, Iowa and Florida.

The instrumentation includes:

- Meteorological sensors (wind, temperature, precipitation, radiation)
- Soil sensors (soil temperature and moisture)
- Gas sensors (concentration, flux)

Illinois

These field sites are located at a University of Illinois facility, the Energy Farm, in a rural area near campus, but the facility is equipped with University networking. The instrument computers are connected to an on-site network box via Ethernet cables, and the network box is connected through underground cabling to the main facility hub and stored on a server there.

Via the University network, data is then sent to a server located in the IGB. This server has a total of 4 TB of which 1 TB is currently being used for this project. To gather for long-term storage.

Data are collected continuously year around at a 10Hz frequency. The flux data is stored in Licor greenhouse gas (GHG) files, so they are easily processed by Licor software. Each file holds 30 minutes of data collection, and is approximately 5 MB in size. These files are placed on the server in a directory for that year. At the end of the year, each day's 48 files are compressed with .tar.gz, and sorted into directories by month. Each daily compressed file is approximately 70MB. The meteorological and soil data is stored in .dat text files on the computer at the Energy Farm. Each year generates approximately 30 MB of meteorological and soil data. The data is then transferred for long-term storage to a dedicated file server located at IGB. The data uses 2.6 TB of space currently.

There are not any major upgrades planned for this system at this time.

Florida

These field sites are located at a research facility on a ranch in central Florida. Data are stored in a USB within the instrumentation and a card inside the datalogger. Both the USB and the datalogger are connected to a modem where the data is transferred via Verizon Internet every four hours. The flux data is transferred to the same IGB server as the Illinois data while the meteorological and soil data is transferred to the server at the Energy Farm facility. The file processes are identical to that at the Illinois site. The files that are generated are processed the same and generate similar file sizes.

Iowa

These field sites are located at a Iowa State University facility in a rural area near campus. Data are transferred from the towers via both flash drive and a cell network to a server located at the facility site. Data processing and file storage are similar to the process used at Illinois.

5.9.2.3.2.2 Generalized Process of Science

The raw 10Hz data is averaged to a 30-minute time step using Licor's EddyPro software with standard flux corrections. Flux data is quality-controlled and quality-assured to remove spikes. Meteorological data is gap-filled using external data sources from local automatic weather stations and gridded global data from the European Center for Medium Range Weather forecasts. Carbon, water, and energy fluxes are gap-filled using an Artificial Neural Network approach. Carbon fluxes are also partitioned into gross primary productivity and ecosystem respiration using three different models. QA/QC, gap filling and partitioning processes are performed using PyFluxPro 9. The half an hour data is also summarized to daily time steps.

Data are made available within CABBI after QA/QC has been completed in a shared Box folder, as both NetCDF and Excel file types. Processes are currently being developed to publicly share the data via AmeriFlux.

These data represent a long-term flux and weather data resource currently being used by ecosystem modelers within CABBI and to others outside the Center.

5.9.2.3.2.3 Remote Science Activities

Both Florida and Iowa State while at CABBI facilities make use of commercial networking which is expensive and less reliable than the facilities at Illinois. But remote access is not feasible for many of the field measurements because they are done at the plant level. It would be expensive and unfeasible to connect sensors to specific plants that will need to be removed from the field at harvest.

Additionally, this data is not needed in real-time as analysis is done at the end of the growing season, so transferring data manually as needed either via datalogger or USB is sufficient and more cost effective. This could potentially change in the future if there is a need to monitor plants and the ecosystem in real-time, but it's unlikely in the near term.

5.9.2.3.3 Modeling

5.9.2.3.3.1 Instruments and Facilities

CABBI uses a suite of models including three ecosystem models, an economic model, and a techno-economic analysis/life cycle assessment model. The researchers using these systems collaborate frequently sharing both data and code.

Currently, researchers work separately on their own models. Some modelers make use of BioCluster, IGB's fee-based high-performance computing resource. BioCluster contains 2824 cores and over 27.7 TB of RAM with 1.3 PB of storage on a GPFS file system. Iowa State researchers are in the process of setting up a new server with a great computation capacity. It has 1 TB of RAM and internal storage of 50TB.

At the time of publication or when benchmark results are completed, researchers archive their simulations including input data, code and results in a repository stored on IGB's file archive server. Collaborators within CABBI are able to access these stored simulations to either modify or use the results for other studies.

The uploaded files vary in number and size depending on the scope of the simulation and the model being used. Currently, archived simulations require 2 MB to 110 MB.

5.9.2.3.3.2 Generalized Process of Science

Currently, researchers access this server via their individual networks using an SFTP client. CABBI will soon begin the development of an internal web application to move away from the data access via the repository, and make it easier to locate and view the data CABBI researchers need.

This site will host common input datasets used by CABBI researchers. This will ensure they are accessing the same versions so results will be comparable. Researchers will be able to download this data directly from the website with customizations (location, grid spacing) needed for their model. The site will also host the benchmark simulation results. These will be visualized spatially as well as also offering the customizable download through the website.

5.9.2.3.3.3 Remote Science Activities

Some modeling researchers remotely access BioCluster to perform their modeling simulations.

5.9.2.3.4 iBioFAB

5.9.2.3.4.1 Instruments and Facilities

iBioFAB is an automated laboratory facility located in the IGB. The system supports rapid design, fabrication, validation/quality control and analysis of genetic constructs and organisms. Current services available are HT yeast strain engineering and synthesis of transcription activator-like effector nucleases (TALEN) libraries.

Instrumentation includes:

- Labcyte Echo acoustic liquid handler
- Tecan Evo and Fluent automated liquid handlers
- Thermocyclers
- Centrifuge
- Plate sealers and unsealers
- Media dispensers
- Shaking incubators
- High-capacity non-shaking incubator
- Plate reader
- Fragment Analyzer
- Cytation 5 throughput fluorescent microscope
- OrbiTrap LC_MS with TriVersa NaNoMate

A variety of data is generated for the various iBioFAB instrumentation. Some include data generated from instrumentation log files, Excel spreadsheets containing plate reader data, capillary electrophoresis results, and microscopy images. The number of files and the size are dependent on the analysis being performed. All data is retrieved for the computers attached to the instruments.

In the next two to five years, there are plans to integrate the liquid chromatography - MS into the automation system. This will involve significant customization.

5.9.2.3.4.2 Generalized Process of Science

No standard analysis procedure exists currently for the iBioFAB since it is used for many different projects and for the retrieval of different data. Currently, data is manually aggregated for the various instruments and emailed to the requesting researchers for analysis. There are general plans for the software team to develop standardized methods for storage and delivery in the future, but these are still in the beginning stages.

5.9.2.3.4.3 Remote Science Activities

While data collection and transfer is primarily performed at the facility and can make use of the University of Illinois network, at some point in the future, users of the facility may want to remotely operate instrumentation. This could potentially become a challenge for the current protocol.

5.9.2.4 Software Infrastructure

LabCollector: A commercial web-based laboratory information management system to facilitate the sharing of DNA parts and materials. Both local and remote CABBI researchers can access this through a web browser. The software also includes an API to allow outside applications like BioParts.org to access the selected DNA parts. (Present and beyond)

Box.com: Cloud-based commercial file sharing and file management for files 150 GB or smaller. Files are shared both within CABBI with restrictive permissions and publicly via a link from the research data website. (Present and beyond)

Institute of Genomic Biology's (IGB) Local Network Data Storage: For internal data sharing, the IGB provides CABBI 10 TB of shared storage. Storage is only accessible using SFTP clients. (Present and beyond)

Data Repositories: All data supporting CABBI publications are released either via Box or through publishing via a data repository. Subject specific repositories like NCBI, AmeriFlux, DAAC, and Phytozome are preferred, but where those are unavailable, others like university (Illinois Data Bank and Iowa State University Data Share) and public (Zenodo, Mendeley, and GitHub) repositories are used. (Present and beyond)

Sustainability Data Hub: CABBI is in the process of developing a web application/database to host a number of datasets from the Sustainability theme. Initially, this tool will only be available internally. It will include an ecophysiological traits and yields database to store quality tested field trial measurements and management data. It will also host modeling input datasets, like land use and climate change datasets, and benchmark modeling results. Access to these datasets will be standardized for the needs of each Sustainability modeling team. (Implementation in the next year)

Pathway Tools Software: CABBI is testing the use of this software to host a metabolic engineering database for non-model yeast genomics. (Implementation in the next year)

5.9.2.5 Network and Data Architecture

The IGB, where CABBI is headquartered and housed on the Illinois campus, works within the University of Illinois's infrastructure to provide network access. In the IGB building, there are dual 10 Gb connections to the campus's public network. Desktops have 1 Gb connections.

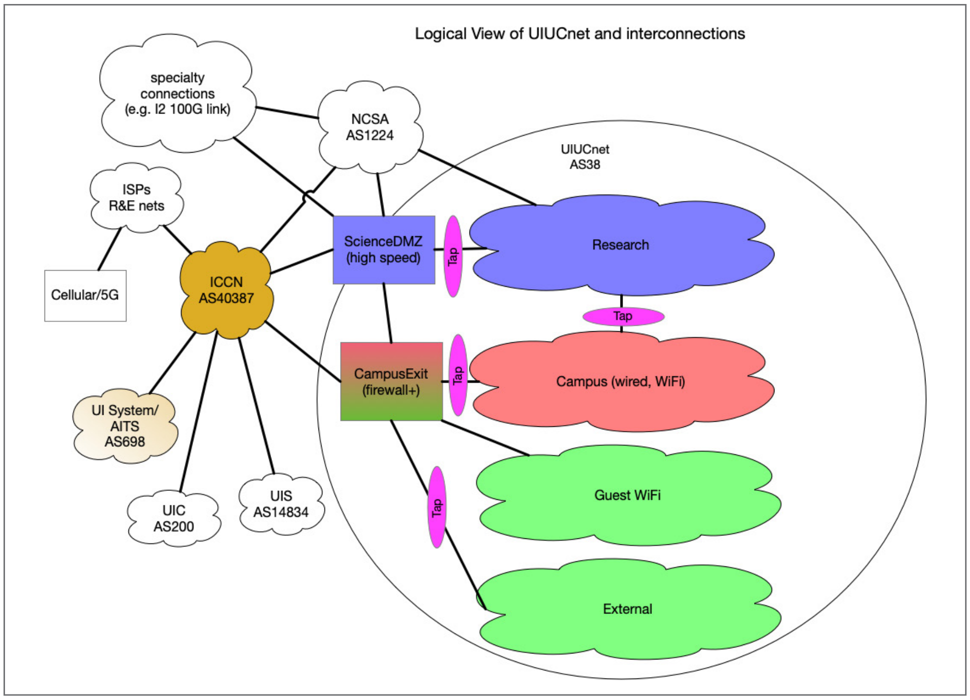

Figure 5.9.2: UIUC Campus Network

The server room within IGB has a 100 Gb connection to Campus Advanced Research Network Environment (CARNE), the campus research network (pictured in blue in Figure 5.9.2). CARNE has multiple 100 Gbps connections to various high-speed regional research networks as well as I2. All systems that connect to this research network do so at 10, 40, 50, or 100 Gbps. CARNE is the Science DMZ at the University of Illinois.

Globus is integrated into the biotransfer system of IGB's high-performance computing resource, BioCluster. Next year, a perfSONAR node will be made available.

These networking services are expected to stay the same for the next 3-4 years. Planning for longer into the future has not started at this point.

5.9.2.6 Cloud Services

CABBI has access to Amazon, Google and Microsoft cloud computing resources through the University of Illinois. To date, only Amazon Web Services have been used to host a web-tool developed by CABBI researchers. Additionally, cloud computation through AWS was explored in the past when a Windows operating system was required for specific software applications. It was found to be cost-prohibitive at that time and local resources were used. For other computational needs, the available local resources are sufficient and cheaper to use.

Cloud storage is also available through the University of Illinois' resources. Box is frequently used by CABBI researchers as described in Section 6. The University of Illinois also offers Google and Microsoft cloud storage for free of charge.

5.9.2.7 Data-Related Resource Constraints

Current constraints include the using commercial wireless networks for data transfer from remote environmental instrumentation. Finding cheaper and more reliable options for this would benefit CABBI's science.

As the rate of genomic data increases, increased network connectivity to facilities like JGI and EMSL will be necessary to continue these partnerships.

5.9.2.8 Outstanding Issues

While not a current issue, in the future as CABBI moves towards more remote measurements in both the laboratory and the field, network security will become more critical. Ensuring outside actors are unable to take control of the equipment will be necessary not only for the protection of the equipment and data, but also for public safety. Determining the appropriate security restrictions without compromising the commitment to ensuring CABBI research is usable and reproducible by others will be challenging.

5.9.2.9 Facility Profile Contributors

CABBI Representation

Matthew Hudson, *Prof. of Crop Sciences, U. of Illinois Urbana-Champaign, CABBI Science Integration Head,* mhudson@illinois.edu

Elizabeth Murphy, *Managing Director, Institute for Sustainability, Energy, and Environment (iSEE), CABBI Deputy Director of Research Strategy & Management,* eamurphy@illinois.edu

Leslie Stoecker, *CABBI Data Manager,* lensor@illinois.edu

5.10 CBI

The CBI vision is to accelerate domestication of bioenergy-relevant plants and microbes to enable high-impact, value-added fuels and coproduct development at multiple points in the bioenergy supply chain.

CBI has identified research targets for a thriving bioeconomy:

- Sustainable biomass feedstock crops using plant genomics and engineering.
- Advanced processes to simultaneously break down and convert plants into specialty biofuels.
- Valuable bioproducts, including chemical feedstocks, made from the lignin residue remaining after bioprocessing.

5.10.1 Discussion Summary

- CBI is a multi-institutional and interdisciplinary organization encompassing numerous institutions across the United States. The core DOE institutions involved are ORNL and the National Renewable Energy Laboratory (NREL). Major university partners include the University of California, University of Georgia, University of Tennessee, University of Maryland, University of Oregon, Washington State University, and Dartmouth College. CBI will accelerate domestication of bioenergy-relevant plants and microbes to enable fuels and coproduct development at multiple points in the bioenergy supply chain.
- CBI has identified research targets for a thriving bioeconomy:
  - Sustainable biomass feedstock crops using plant genomics and engineering.
  - Advanced processes to simultaneously break down and convert plants into specialty biofuels.
  - Valuable bioproducts, including chemical feedstocks, made from the lignin residue remaining after bioprocessing.
- CBI has identified four innovation targets that will enable the future bioeconomy:
  - Development of high-yielding, process-advantaged, sustainable plant feedstocks.
  - Creation of lower-cost consolidated bioprocessing to produce fermentation intermediates.
  - Valorization of lignin to atom-efficient biofuels and bioproducts.
  - Development of chemocatalytic conversion processes for biobased intermediates to produce a complete biomass-based sustainable aviation fuel (SAF).
- Hyperspectral instrumentation located both in the greenhouses and in the field is a major source of large data volumes for CBI. Each hyperspectral image is approximately 1 gigabyte of data, and the instruments can produce an image every three minutes.
- Data from various research components within CBI are shared, stored, and preserved by building upon existing integrations of resources. CBI datasets are TB to PB scale. The diverse data types in CBI can be categorized into basic data levels. The EarthData processing levels devised by NASA's Earth Observation System Data and Information System largely inspired CBI's data processing levels:
  - Level 0 data is raw, unprocessed data from original sources. Sources are either instruments or recorded observations.
  - Level 1 data is typically generated by a third-party provider on behalf of CBI.

    - Level 2 data is derived from data in Levels 0 and 1 and consists of data in standardized formats that are widely recognized within each field of study.

- CBI provides its users with a cloud-based Laboratory Information Management System (LIMS) repository based on the LabKey Server technology. The LabKey Server system houses Level 0 and Level 1 data. CBI investigators are encouraged to use the system to track and exchange their Level 0 and Level 1 data assets. The CBI LIMS can accommodate data dumps from both manual and automated sources.
- CBI uses the Advanced Plant Phenotyping Laboratory (APPL) facility at ORNL, the PY-MBMS instrumentation at NREL, and a custom-designed facility for low-cost fermentation of biomass with co-treatment located at Dartmouth College. In addition to these specialized facilities, CBI uses standard molecular biology and field experimentation techniques at its various locations. CBI currently maintains active field sites located in Davis, California; Corvallis, Oregon; Clatskanie, Oregon; Knoxville, Tennessee; and Athens, Georgia that are equipped with field-based instrumentation.
- An Azure instance is being evaluated for CBI. The purpose of the evaluation is to ascertain whether Azure can be used in a hybrid on-prem/ off-prem environment to access the data. The main use case for the Azure instance will be to surface virtual machines that are staged with specific Microsoft Windows-only applications that need to run against the main data store.
- CBI has remote devices include field stations, edge instrumentation, and space-based assets. Not all of the partner institutions have high-speed networking capabilities, and this can be problematic regarding data transfer. Often the fastest transfer modality is shipping hard drives on FedEx.
- CBI is in the process of evaluating a Kubernetes environment hosted by OLCF. It is expected that this environment will become the main hosting environment for both data and compute.
- CBI will work with ESnet to learn more about Kubernetes deployment.
- The default max storage configuration for the ORNL CADES cloud virtual machines, where a lot of development is done, comes with 80 GB of space. This is insufficient for a single day's worth of data from the APPL facility. Storage requirements can be mitigated by mounting various devices to the CADES cloud, but there is a bottleneck with both the read and write speed of the devices themselves. This has limited some data transfer activities within the ORNL environment.
- The cybersecurity landscape at some of the institutions imposes many constraints upon scientific researchers. Each institution has independent cybersecurity policies, teams, and infrastructure, with different technical and policy requirements that must be traversed:
    - At ORNL, the CADES infrastructure is logically separated from ORNL/OLCF. Hence, to move data between the two organizations, one must traverse two firewalls—each run by separate cybersecurity groups—and run through a series of DTNs and jump servers before the data can be ingested into the application database for broad access.
    - This creates layers of friction that impede seamless data transfer between indigent organizations within a single entity such as ORNL, as well as between partner institutions.
- CBI has some data-related resource constraints on the network to address. Performance to remote sites, such as NREL, can sometimes be below 1 Gbps. For large datasets (e.g., PB scale), this can take weeks to transfer. This level of friction for data that is already publicly available is a serious hindrance to the utility of ESnet.

- CBI will work with ORNL, NREL, and ESnet to address slow data-transfer problems. This will involve migration of data to ORNL resources that have access to the institutional Science DMZ and performance testing with Globus.

5.10.2 CBI Facility Profile

CBI pursues a host of new technologies to alleviate critical cost barriers to sustainable, economically viable production of biobased products and advanced biofuels. The center focuses on creating robust, high-yielding feedstock plants using genetic technology and bioengineering. CBI creates methods for high-yield production of advanced biofuels that can be blended with existing transportation fuels. Researchers also study ways to develop valuable byproducts from lignin left over after biomass processing.

The CBI team is positioned to develop breakthrough technologies and scientific advances by uniting a team of senior science leaders with a proven desire and ability to accomplish challenging goals specifically targeted to barriers facing the emerging bioeconomy.

CBI is one of four BER Centers within the DOE's SC created to expand on the foundational successes of former bioenergy research centers and to lay the scientific groundwork for a new robust, biobased economy.

5.10.2.1 Science Background

The world (e.g., climate change), biology (e.g., genomics), and the bioeconomy (still in its infancy) have experienced profound changes in the past decade. Presently, there is a demand for the emerging bioeconomy to meaningfully contribute to the displacement of petroleum-based fuels and products. To this end, the CBI has identified four innovation targets that will enable the future bioeconomy:

1. development of high-yielding, process-advantaged, sustainable plant feedstocks,
2. creation of lower cost consolidated bioprocessing to produce fermentation intermediates,
3. valorization of lignin to atom-efficient biofuels and bioproducts, and 4) development of chemocatalytic conversion processes for biobased intermediates to produce a complete biomass-based sustainable aviation fuel (SAF).

CBI achieves these innovation targets through the development and use of accelerated domestication tactics that leverage complex phenotypes in relevant perennial plants and non-model microbes, and through the application of hybrid biological and catalytic processing. CBI builds on a strong organizational structure, ongoing integrative assessments, world-class institutional capabilities, and highly experienced research and management teams to implement and accelerate these critical advances. CBI aims to create new technologies, informed by science, that can alleviate cost barriers that have thus far prevented the biofuels and bioproducts endeavor from becoming a sustainable and economically viable reality.

The CBI is a multi-institutional and interdisciplinary organization encompassing numerous institutions across the United States. The core DOE institutions involved are ORNL and NREL. Major University partners include the University of California, University of Georgia, University of Tennessee, University of Maryland, University of Oregon, Washington State University, and Dartmouth College. CBI generates data for the diverse people, science, and technologies that drive its experiments, collaborations, and objectives. Data from various research components within CBI are shared, stored, and preserved by building upon existing integrations of resources. CBI adheres to and embraces FAIR (findability, accessibility, interoperability, and reusability) data principles for the exchange and dissemination of data generated. This paradigm advances research progress, and allows for cross-validation and quality control (QC) of experimental results. Data is shared and stored while adhering to guidance, laws, and regulations intended to protect confidentiality and security. Critical data interdependencies have been identified and are managed through careful consideration of data workflows.

Data Types and Generation

The diverse data types in CBI can be categorized into basic data levels. The EarthData processing levels devised by NASA's Earth Observation System Data and Information System largely inspired CBI's data processing levels. Understanding and defining the types and sources of data generated will enable FAIR and consistent data sharing and quality across CBI, as well as externally to the broader scientific community. A standardized set of terminology is being continually developed and made accessible to the community to enable translation and common controlled language across researchers and data types. Level 0 data is raw, unprocessed data from original sources. Sources are either instruments or recorded observations. Level 1 data is typically generated by a third-party provider on behalf of CBI. Level 2 data is derived from data in levels 0 and 1 and consists of data in standardized formats that are widely recognized within each field of study. Data at levels 3 and above has undergone a comprehensive FAIRification process and is heavily curated and scrutinized for QC.

Data Storage, Sharing, and Preservation

CBI provides its users with a cloud-based LIMS repository based on the LabKey Server technology. The LabKey Server system houses Level 0 and Level 1 data. CBI investigators are encouraged to use the system to track and exchange their Level 0 and Level 1 data assets. The CBI LIMS can accommodate data dumps from both manual and automated sources. During the first five years of CBI, significant effort was put into automation and standardization of inputs. This work can be leveraged to generate interoperability between various systems used internally within research groups that are a part of CBI. Efforts are being made in terminology standardization and the auto-ingest of data using a standard process into a cloud-based LIMS to enable rapid validation of experimental results.

5.10.2.2 Collaborators

CBI is a multi-institutional entity comprising multiple institutions across the United States. The list of partners and their locations is available online.[60]

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
ORNL	Primary	Data portal	500 GB	3 times/week	N	Yes. (slow network)
NREL	Secondary	Data portal	200 GB	Biweekly	N	Yes. (poor infrastructure/ slow network)
University of California, Davis	Primary	Hard drive	250 GB	Ad hoc	N	Ad hoc transfers, no repeatable process
Other University partners	Secondary	Data portal	Varies	As needed	N	Yes (heterogenous networks)

Table 5.9.1: CABBI Collaboration Space

60 https://cbi.ornl.gov/partners/

5.10.2.3 Instruments and Facilities

CBI uses the APPL facility at ORNL, the PY-MBMS instrumentation at NREL, a custom designed facility for low cost fermentation of biomass with co-treatment located at Dartmouth College. In addition to these specialized facilities, CBI uses standard molecular biology and field experimentation techniques at its various locations. CBI currently maintains active field sites located in Davis, Ca, Corvallis, Oregon, Clatskanie, Oregon, Knoxville, TN, Athens, Georgia that are equipped with field-based instrumentation. Some of the edge devices require a substantial investment in compute and data infrastructure.

ORNL also hosts Frontier, the world's fastest supercomputer, while NREL hosts the world's greenest one. CBI uses both of these systems for its computational biology needs. There is a tremendous need to streamline the data feeds for Frontier.

Hyperspectral instrumentation located both in the greenhouses, and in the field is a major source of large data volumes. Each hyperspectral image is approximately 1 gigabyte of data, and the instruments can produce an image every 3 minutes.

5.10.2.4 Generalized Process of Science

Science begins with a proposal from a CBI PI. These ideas are consolidated into the five-year proposal that is sent to DOE. Upon approval, investigators begin their work. Each PI has operational discretion regarding how to collect and process initial data. CBI encourages, but does not require, that Level 0 data be stored in the CBI LIMS . However, both Level 0 and Level 1 data are managed at the PI's discretion. Level 2 data is centralized in the CBI LIMS and is FAIRified by the CBI data management team. This data may be published as data releases or further processed into higher level data products for consumption by the community

5.10.2.5 Remote Science Activities

Remote devices include field stations, edge instrumentation, and space-based assets. Not all of the partner institutions have high-speed networking capabilities, and this can be problematic regarding data transfer. Often the fastest transfer modality is shipping hard drives on FedEx.

5.10.2.6 Software Infrastructure

CBI uses LabKey Server as the main repository for its level 0 and level 1 data. Level 2 data is stored in a custom database built on postgresql. The system uses postgis, Python, php, and JavaScript technologies, and the codebase is hosted at code.ornl.gov

There is also an Azure instance that is being evaluated. CBI is in the process of evaluating a Kubernetes environment hosted by the Office of Leadership computing. It is expected that this environment will become the main hosting environment for both data and compute.

5.10.2.7 Network and Data Architecture

Default is 1 gigabit throughput between buildings, while some have a 10 gigabit fiber backplane.

ORNL connects to ESnet via redundant border routers, each of which currently connects to a diverse ESnet router at 100G. The expectation is that these connections will soon be upgraded to 400G connections. The ORNL border routers connect the ORNL Enterprise network and OLCF to ESnet. This connectivity is depicted in the diagram below.

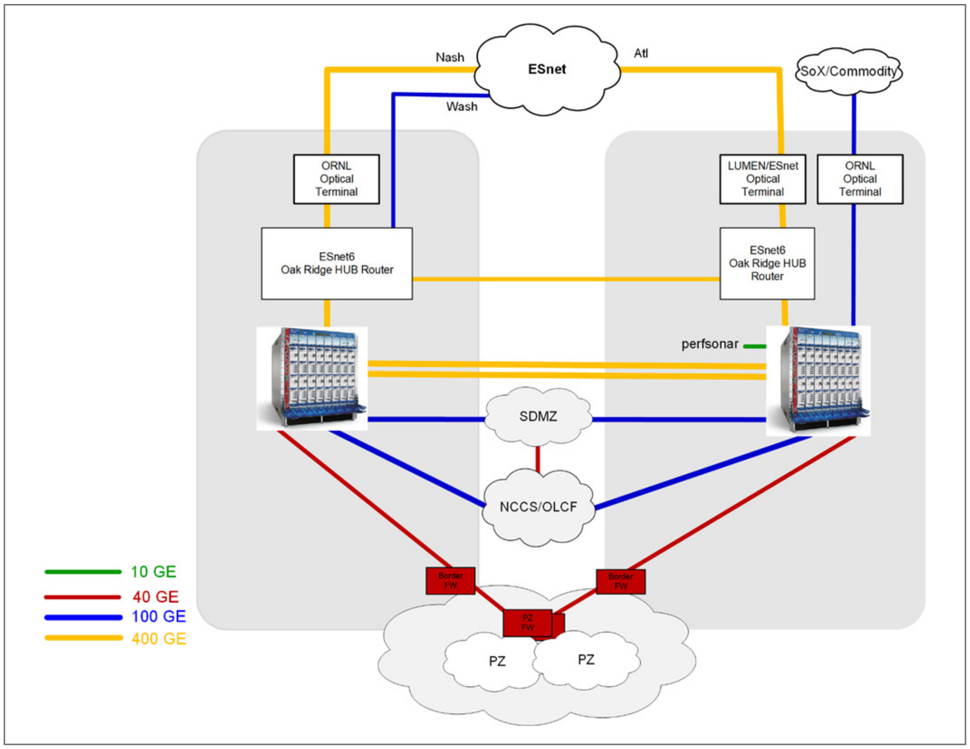

Figure 5.10.1: ORNL Network Diagram

The Enterprise firewalls are connected to the border routers at 40G currently. ORNL does use a Science DMZ architecture for high-performance data transfer. This environment connects to the border routers with 10/40/100G DTN connections available. Globus is the approved transfer method. A border perfSONAR node is connected to the border router and participates in the ESnet grid.

5.10.2.8 Cloud Services

CBI is not considering the use of commercial cloud services at this time.

5.10.2.9 Data-Related Resource Constraints

Slow or non-existent network connectivity between computers where data resides and border routers, especially on the ORNL/CADES side.

As an example, the default max storage configuration for the CADES cloud virtual machines, where a lot of development is done comes with 80 GB of space. This is insufficient for a single day's worth of data from the APPL facility. This is mitigated by mounting various devices to the CADES cloud, (example Black Pearl) but there is a bottleneck both with the read/write speed of the devices themselves, and of the fact that the data bandwidth available from a users Office PC to the CADES cloud is only 1 Gig (shared bandwidth from a single building) over copper Ethernet. While there are a few fiber channel drops in some buildings with 10 GB capacity, this is the exception rather than the rule. Exacerbating this problem is the cybersecurity landscape which society imposes upon scientific researchers. Each institution has independent cybersecurity policies/teams/ and infrastructure, with different technical and policy requirements that must be traversed. At ORNL, the CADES infrastructure is logically separated from ORNL/OLCF. Hence, in order to move data between the two organizations, one must traverse two firewalls - each run by separate cybersecurity groups - and run through a series of DTNs and jump servers before the data can be ingested into the application database for broad access. This creates layers of friction that impede seamless data transfer between indigent organizations within a single entity such as ORNL, as well as between partner institutions.

5.10.2.10 Outstanding Issues

NREL has high-performance data transfer nodes (using Globus) connected directly to ESnet, but there are potential hurdles with getting data from the DTN's to the underlying data store that is connected to Eagle/ Kubernetes. NREL has a similar issue to ORNL in that there is logical separation between the HPC compute infrastructure and commodity computing. NREL uses Amazon/S3 and glacier in addition to their on-site infrastructure for compute.

5.10.2.11 Facility Profile Contributors

CBI Representation

Gerald Tuskan, *ORNL,* tuskanga@ornl.gov

Dan Jacobson, *ORNL,* jacobsonda@ornl.gov

Stan Martin, *ORNL,* martins@ornl.gov

ESCC Representation

Susan Hicks, *ORNL,* hicksse@ornl.gov

5.11 The GLBRC

The GLBRC is a cross-disciplinary research center led by the University of Wisconsin–Madison. With Michigan State University and other collaborators, GLBRC draws on the expertise of over 400 scientists, engineers, students, and staff to develop sustainable biofuels and bioproducts.

Created in 2007 by the US DOE, GLBRC focuses on three areas of research: sustainable cropping systems, efficient biomass conversion, and field-to-product integration. Together, we are helping to replace petroleum-derived fuels and products and enable a new generation of biorefineries.

Our mission is simple: creating biofuels and bioproducts that are economically viable and environmentally sustainable. Our vision rests in developing sustainable biofuels and bioproducts from all usable portions of dedicated energy crops grown on marginal, nonagricultural lands.

5.11.1 Discussion Summary

- The GLBRC is led by the University of Wisconsin–Madison with a mission of creating biofuels and bioproducts that are economically viable and environmentally sustainable. GLBRC focuses on three areas of research:
  - Sustainable cropping systems.
  - Efficient biomass conversion.
  - Field-to-product integration.
- The mission of the GLBRC has expanded to include collaborating with the other three BRCs on a variety of shared research objectives, putting an even larger emphasis on high-speed networking to allow for transferring of shared data over the next five years of research.
- GLBRC produces heterogeneous data types, at a variety of time and spatial scales and at ever-increasing rates, with different management requirements, and has a variety of pre-existing data management practices. The key data management challenge is to leverage diverse resources and systems in support of the needs of individual research projects while optimizing information integration to promote collaboration and productive synergies across the center.
- GLBRC works with the following types of data:
  - Biomass compositional analysis.
  - DNA/RNA sequencing.
  - High-throughput (HT) chemical genomics.
  - Droplet-based microfluidics.
  - HT mass spectrometry (MS)-based protein and metabolite analysis.
  - Proteomic analysis.
  - Metabolite analysis.
  - Experimental fermentation lab (EFL).
  - Structural biology.
  - Cryo-EM and Cryo-ET (evapotranspiration) instrument access and training.
  - Imaging.
  - Separations.

- GLBRC provides researchers with access to an array of in-house analytical, computational, and business services. These include:
  - Experimental measurements and field observations.
  - Experimental designs and protocols.
  - Equipment descriptions, designs, and procedures for use.
  - Materials inventories and characterizations.
  - Algorithms and source code.
- GLBRC workflow typically involves:
  - Analysis of materials generated in field experiments.
  - Different datasets being compiled for a number of core measurements (e.g., climate, soil, emissions, and other field sensor measurements).
  - Composition and quality of materials determined through analysis at member institutions and national labs.
  - Materials being transferred for processing through a variety of platforms generating additional rounds of data from DOE national labs and institutional core facilities.
  - Additional analytic techniques possibly being applied at any stage of this process to assist with experimental goals.
  - This data being transferred to, and stored in, the GLBRC data catalog.
- Optimizing the GLBRC workflow integration process to improve collaboration and productive synergies is a key challenge, and network capabilities play a critical role.
- GLBRC routinely shares 100 GB datasets with a number of collaborators: JGI, EMSL, other BER-funded BRCs and national laboratories, and collaborators at the University of Wisconsin and Michigan State University.
- GLBRC collaboration with DOE light sources facilities requires low-latency, high-bandwidth network connections:
  - Diffraction experiments that generate crystallographic data are controlled and monitored in real time using secure virtual network computing (VNC) connections.
  - Existing network architecture is strained by the pace of data collection.
  - Datasets consisting of numerous 10 GB images are common, and many datasets are required for each solved structure. Acquisition can take minutes, and after planned upgrades dataset sizes will increase and acquisition times will decrease by orders of magnitude.
  - Even with rapid on-site data reduction, this pace of data collection will put new strain on network infrastructure and data storage for collaborations.
- The two primary institutions within the GLBRC, the University of Wisconsin-Madison and Michigan State University (MSU), have multi-hundred gigabit backbone connections from campus and high-speed intra-campus connections that range from 1 to 100 Gb between buildings. [Section 5.11]
  - The primary data center for GLBRC, located within the Wisconsin Energy Institute, has redundant 40 Gb connections to the campus backbone and 10 Gb to all equipment internally.

- High-performance data transfer is available between the primary institutions using Globus endpoints.
- New network architectures with improved transfer speeds (Science DMZ) are under investigation by the University of Wisconsin's campus networking team for potential use in the next five years of research.

- GLBRC is a multi-institutional center with labs generating data from a variety of instruments. The output of the equipment varies in type and size, from spreadsheets (KB in size) to hundreds of gigabytes of sequencing data. It is important that this data is shared across teams and stored in a central repository, requiring high-speed network connectivity for the larger data types.
- The GLBRC uses a custom-built data catalog to store data and metadata along with custom-built applications that assist with collection and transfer of data from external data providers.
- Currently, the GLBRC uses cloud storage options for some smaller facility data transfers (Google Drive, OneDrive) but all compute and storage are located on campus. Over the next several years, GLBRC will be looking into using public cloud resources for applications (AWS, GCP) and to host the upcoming cross-BRC data-sharing portal, bioenergy.org.
- While transferring large amounts of data at high speed is critical to sharing and using compute resources outside of local resources, the storage of data is often an overlooked issue. GLBRC has sufficient storage space for the petabytes of research data that are expected, and has built this technology solution over a number of years of operation. A central data storage option provided by the DOE could be very beneficial.
- Data created in labs and facilities that do not have a high-speed connection to ESnet can be hard to access and share. If these datasets are large, it becomes even more difficult and requires alternative solutions, such as scheduled archival of data, or other workarounds to be able to share with collaborators outside of that lab.
- Initial testing has shown that speeds through the border of MSU's campus network can be slow, limited to 500 mb per second. Hopefully new resources, such as the Science DMZ, can help alleviate this problem.

5.11.2 GLBRC Facility Profile

Developing sustainable processes to generate the liquid hydrocarbon fuels and products currently derived from petroleum is an increasingly urgent challenge. Since 2007, researchers at the GLBRC have been addressing this challenge by generating the knowledge needed to supply these compounds from lignocellulosic biomass. GLBRC will address significant barriers in the sustainable pro- duction of specialty biofuels and bioproducts via microbial conversion of the lignocellulosic biomass of bioenergy crops grown on land that is unsuitable for food crops, with low carbon (C) stores and low conservation value (i.e., bioenergy land). In collaboration with other Bioenergy Research Centers (BRCs) and the broader scientific community, the Center will produce knowledge and technological advances and train the next generation of researchers to establish a sustainable lignocellulosic industry.

5.11.2.1 Science Background

The Great Lakes Bioenergy approach to interdisciplinary science depends on a wide range of research projects and supporting facilities having access to diverse resources and being able to integrate and leverage these resources for maximum impact. These projects are embedded in laboratories with heterogeneous data types, management requirements, and a variety of pre-existing data management practices. They are, moreover, generating data at a variety of time and spatial scales and at ever-increasing rates, even as the effort required for multidimensional data analysis may exceed the effort expended to generate the data.

Our key data management challenge is to leverage diverse resources and systems in support of the needs of individual research projects while optimizing information integration to promote collaboration and productive synergies across the Center.

5.11.2.2 Collaborators

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
JGI	Yes	Portal/Data transfer	100s GB	Daily	N	
EMSL	No	Data transfer	???	Ad hoc	N	
UW – Campus/Core Facilities	No	Data transfer	100s GB	Daily	N	
MSU - Campus/ Core Facilities	No	Data transfer	100s GB	Daily	N	
DOE National Labs	No	Data transfer	100s GB	Ad hoc	N	
Other BRCs	Yes	Data transfer	GB	Ad hoc	Y	

Table 5.11.1: GLBRC Collaboration Space

5.11.2.3 Instruments and Facilities

The specific pieces of equipment used to generate scientific data for the GLBRC are too broad to enumerate here. Generally speaking the Center works with the following types of data:

- Biomass Compositional Analysis
- DNA/RNA Sequencing
- HT Chemical Genomics
- Droplet-based Microfluidics
- HT MS-based Protein and Metabolite analysis
- Proteomic Analysis
- Metabolite Analysis
- EFL
- Structural Biology
- Cryo-EM and Cryo-ET Instrument Access and Training
- Imaging
- Separations

This data is generated at the Center's Core Facilities, programs in its member institutions, DOE national labs, and user facilities.

The two primary member institutions of the GLBRC are public universities with significant compute resources that are used to analyze the above data types.

5.11.2.4 Generalized Process of Science

To facilitate collaboration and enable rapid progress, the Center provides researchers with access to an array of in-house analytical, computational, and business services.

Our datasets include all recorded factual material and metadata that are required to validate research findings: 1) experimental measurements and field observations; 2) experimental designs and protocols; 3) equipment descriptions, designs, and procedures for use; 4) materials inventories and characterizations; 5) algorithms, source code, and more. GLBRC data cover a breathtaking scope, from small molecule identification and quantification to enzyme kinetics, cropping practices, greenhouse gas fluxes, and socio-economic models.

GLBRC data types include Next-Generation Sequencing (NGS), MS, Statistical analysis, modeling, and visualization plant phenotyping data; phospholipid fatty acid (PLFA) profiles of microbial cultures and soil microbiome samples; high-throughput environmental qPCR (HTE-qPCR); nuclear magnetic resonance (NMR); plant cell wall digestibility assays; GENPLAT biomass digestibility assays on mixtures of enzymes; Biolog phenotypic arrays; microbial growth curves; chemical genomics profiles; engine pressure, emissions, and high-speed imaging; climate, soils, topography, land use, land cover, land management, and other geospatial data; soil greenhouse gas emissions; protein structures; etc. A detailed discussion of all GLBRC data types is beyond the scope of this document.

An example workflow might include analysis of materials generated in field experiments across Wisconsin and Michigan. Land management & sustainability datasets are compiled for climate, soil, GHG emissions, ecological diversity and other field sensor measurements. Composition and quality of materials is determined through protein and metabolite analysis at member institutions and national labs. Materials are transferred between member institutions for processing through a variety of platforms generating additional rounds of proteomic, metabolomic, lipidomic, and genomic data from DOE national labs and institutional core facilities. Additional analytic techniques may be applied at any stage of this process to assist with experimental goals. This data is transferred to and stored in the Center's Data Catalog. Optimizing this data integration process to improve collaboration and productive synergies is a key challenge to the GLBRC and network capabilities play a critical role.

Another example workflow from genes to solved structures might begin with the association of a valuable enzymatic activity with a specific sequence. Hundreds to tens of thousands of proteins belonging to the same protein family might be clustered and tree structures generated to generate preliminary target lists. Additional bioinformatic analysis may be performed to help predict which similar proteins would be easier or more difficult to express by HT techniques. Attributes associated with facile protein expression and purification are few or no highly repetitive sequences, and few predicted intrinsically disordered regions. After gene synthesis of representative target enzymes, hundreds or thousands of proteins can be produced in a wheat-germ cell-free system. In favorable cases, enzyme activity assays can be performed without purifying the protein. Targets of interest are then expressed and purified in milligram quantities, submitted to HT crystallization screens. Diffraction data collection most often occurs at DOE synchrotron light sources, like APS or NSLS-II. After active site residues are confidently identified, 3D structures with substrates may be obtained.

5.11.2.5 Remote Science Activities

GLBRC is a multi-institutional center with labs generating data from a variety of instruments. The output of the equipment is varied in type and size, from small KB spreadsheets to hundreds of gigabytes of sequencing data. It is important that this data is shared across teams and stored in a central repository, requiring high-speed network connectivity for the larger data types.

The mission of the GLBRC has expanded to include collaborating with the other 3 BRCs on a variety of shared research objectives, putting an even larger emphasis on high-speed networking to allow for transferring of shared data over the next five years of research.

(Remote collection of crystallographic data at synchrotron light sources requires low-latency, high-bandwidth network connections. Diffraction experiments are controlled and monitored in real-time using secure VNC connections, or technologies that make remote and local data collection functionally equivalent. Existing network architecture is strained by the pace of data collection even now. Single-crystal macromolecular datasets with 10 GB of image data are not uncommon. Many datasets are often required for each solved structure. At the APS, acquisition of 10 GB of image data now takes about a minute. After a planned upgrade, due to be completed in spring 2024, the source will become 100-1000 times brighter, and acquisition of the same 10 GB dataset may take less than a second. Latency between datasets may be reduced to a minute. Serial crystallography applications may see sustained data rates of 1GB per second. Even with rapid on-site data reduction, this pace of data collection will put new strain on network infrastructure and data storage.)

5.11.2.6 Software Infrastructure

The GLBRC uses a custom-built data catalog to store data and metadata along with custom-built applications that assist with collection and transfer of data from external data providers.

The GLBRC uses a variety of other tools to transfer data, including Globus, rsync, SCP/SFTP, and cloud storage solutions such as Google Drive, OneDrive, and Box.com. These tools will be used over the next five years of research within the Center.

5.11.2.7 Network and Data Architecture

The two primary institutions within the GLBRC, UW-Madison and MSU, have multi-hundred gigabit backbone connections from campus and high-speed intra-campus connections that range from 1-100 Gb between buildings. The primary data center for GLBRC, located within the Wisconsin Energy Institute has redundant 40 Gb connections to the campus backbone and 10 Gb to all equipment internally.

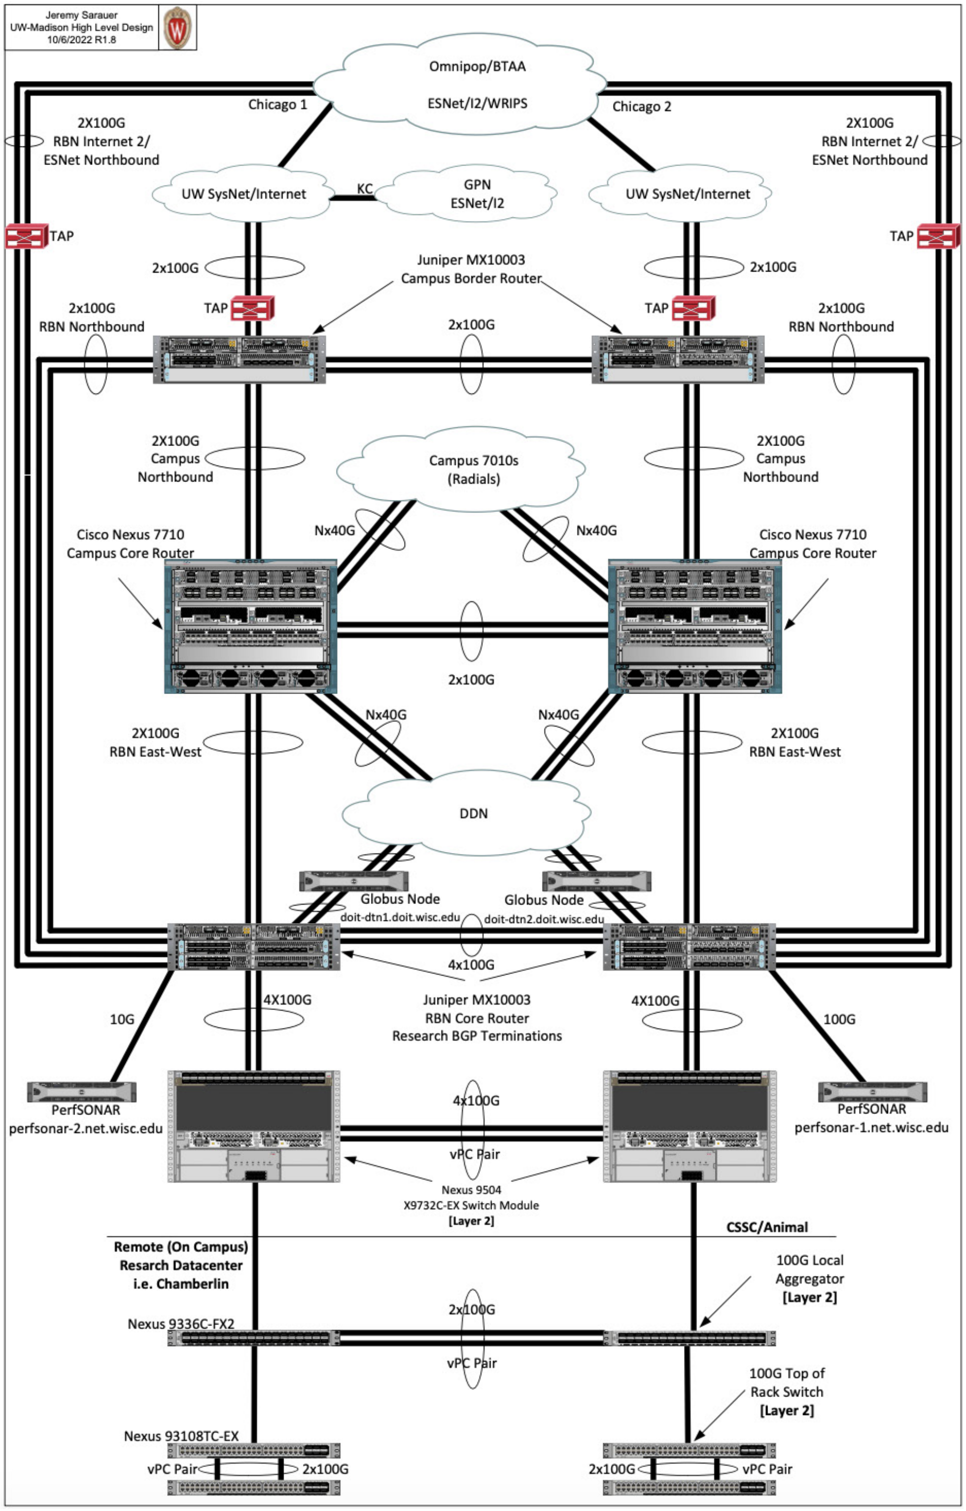

Figure 5.11.1: University of Wisconsin Network Diagram

High-performance data transfer is available between the primary institutions using Globus endpoints.

New network architectures with improved transfer speeds (Science DMZ) are under investigation by central campus for potential use in the next five years of research. See outstanding issue below regarding final mile connectivity concerns.

5.11.2.8 Cloud Services

Currently the GLBRC uses cloud storage options for some smaller facility data transfers (Google Drive, OneDrive) but all compute and storage is located on campus. Over the next several years GLBRC will be looking into utilizing public cloud resources for applications (AWS, GCP) and to host the upcoming cross-BRC data-sharing portal, bioenergy.org.

5.11.2.9 Data-Related Resource Constraints

While there is sufficient storage space for the petabytes of research data the GLBRC creates, it is only because GLBRC has built this solution. A central data storage option provided by ESnet could be very beneficial (see outstanding issues). Initial testing has shown that speeds through the border of MSU's campus network can be slow, limited to 500mb per second. Hopefully new resources, such as the Science DMZ, can help alleviate this problem.

5.11.2.10 Outstanding Issues

Storage–While transferring large amounts of data at high speed is critical to sharing and utilizing compute resources outside of local resources, the storage of data is often an issue that is overlooked. Having some sort of centralized storage resources, directly connected to ESnet would be a great resource and allow for better collaboration and sharing, along with longer-term storage outside of public repositories such as NCBI.

"Final mile" networking–Data created in labs that do not have a high-speed connection to ESnet can be hard to access and share. If these datasets are large, it becomes even more difficult and requires alternative solutions such as scheduled archival of data, or other workarounds to be able to share with collaborators outside of that lab.

5.11.2.11 Facility Profile Contributors

GLBRC Representation

Dirk Norman, *GLBRC*, dirk.norman@wisc.edu

Nick Thrower, *GLBRC*, throwern@msu.edu

Tim Donohue, *GLBRC*, tdonohue@bact.wisc.edu

Craig Bingman, *GLBRC*, cabingman@wisc.edu

Brian Fox, *GLBRC*, bgfox@biochem.wisc.edu

Yiying Xiong, *GLBRC*, yiying.xiong@wisc.edu

5.12 JBEI

The JBEI is a US DOE Bioenergy Research Center dedicated to developing advanced biofuels—liquid fuels derived from the solar energy stored in plant biomass that can replace gasoline, diesel and jet fuels.

JBEI researchers are using the latest tools in molecular biology, chemical engineering, computational and robotic technologies to transform biomass into biofuels and bioproducts.

JBEI is a research partnership combining the scientific talent, expertise, resources and support of four national laboratories, six academic institutions and one industry partner.

5.12.1 Discussion Summary

- The JBEI is a US DOE Bioenergy Research Center dedicated to developing advanced biofuels: liquid fuels derived from the solar energy stored in plant biomass that can replace gasoline, diesel, and jet fuels. JBEI partners include Brookhaven, Pacific Northwest and Sandia National Laboratories; Iowa State University; the University of California (UC) campuses at Berkeley, Davis, San Diego, and Santa Barbara; and TeselaGen Biotechnology, Inc.
- JBEI will have access to scientific infrastructure and instrumentation at LBNL and partner institutions that are well suited to address the most challenging tasks:
  - At LBNL, these include dedicated biotechnology laboratories, specialized facilities for HT structural biology and genomics analysis, high-resolution imaging systems, extensive plant biology and greenhouse research areas, and a powerful supercomputing environment.
  - UC Berkeley (UCB) facilities include a biomolecular nanotechnology facility and nuclear magnetic resonance (NMR) facility.
  - JBEI has access to combustion research, supercomputing, and microsystems facilities at Sandia, UC Davis (UCD) feedstock growth facilities, and fungal research, proteomics, and bioinformatics facilities at PNNL.
- In addition, JBEI researchers make use of national research facilities, often operated by DOE national laboratories.
- In addition to the facilities in JBEI's central location, approximately 10% of JBEI's researchers are housed in off-site locations at DOE facilities and laboratories, universities, and private industrial partners.
- JBEI routinely shares datasets that are .5 TB in size with other DOE facilities: the JGI, EMSL, ALS, CBI, CABBI, GLBRC, and the NMDC. Data transfer to some facilities could be improved, and is attributed to the tools used, as well as some of the networking abilities at remote sites.
- JBEI has extensive, fully operational research facilities and resources at Emeryville Station East (ESE). Processed data is added to the JBEI Experimental Data Depot (EDD) to aid information transfer to individual researchers and data storage and communication with the biofuels community. Capabilities include:
  - Analytical laboratory: instruments to determine reaction outcomes.
  - Biomass pretreatment laboratory: to test the effects of mechanical and chemical processing of biomass.
  - Chemistry laboratory: to develop reactions, increase scale, study the effects of high temperature and pressure, and safely handle organic solvents and other hazardous materials.

    - Fermentation laboratory: to study the physiology of microbial strains engineered for production of biofuels or enzymes during fermentation.
    - Microfluidics laboratories: microfluidic tools for HT synthetic biology, using machine learning (ML).
    - Microscopy laboratory: provides researchers with a suite of spectral microscopes.
    - Plant growth facilities: provide precisely controlled indoor greenhouse facilities for growing genetically engineered feedstocks.
    - Robotics laboratory: commercial automation equipment to enable HT research.
    - Spectroscopy laboratory: specialized instruments to detect and characterize biofuel feedstocks.
    - Synthetic biology laboratories: multidisciplinary spaces that house specialized equipment for experimentation.
    - NextGen sequencing laboratory: next-generation sequencing research efforts such as RNA-seq analysis, genome resequencing applications, and general quality assurance and quality control (QA/QC) for synthetic biology and metabolic engineering.
    - X-ray crystallography facilities: protein purification equipment for the preparation of protein samples.
    - Computing server room: houses JBEI systems for data storage, data backup, bioinformatics analysis, Docker services, and web services.
    - Access to a computing cluster maintained by LBNL IT.
- JBEI researchers, by virtue of the scope of the project, will generate data of many different types, which need to be preserved and shared with the broader scientific community:
    - JBEI data will be available to the scientific community beyond that published in peer-reviewed journals.
    - JBEI data will be preserved to the maximum degree possible through inclusion in public, long-lived databases.
    - Larger datasets will be deposited with recognized community databases.
    - The research community will be provided with other JBEI datasets through community resources such as KBase, the Experimental Data Depot (EDD), and the Inventory of Composable Elements (ICE) registry.
- JBEI research will generate multiple kinds of experimental data in a digital format that can be used by other researchers to validate research findings and potentially further analyzed:
    - Genomic data will be obtained through gene sequencing and will be created through construct design.
    - Multiple functional genomic data types (transcriptomics, proteomics, metabolomics, and fluxomics) will be generated in the characterization of engineered microbes, adapted microbes from adaptive laboratory evolution (ALE) experiments, and engineered plants.
    - Enzymatic data will be generated in the analysis of wild type and in engineering glycosyl hydrolases, lignases, and other biosynthetic enzymes
    - NMR studies of plant biomass will generate multidimensional spectroscopic data.
    - Structural studies of these same proteins will generate crystallographic and cryo-electron microscopy data, and associated atomic models.

- For DNA, RNA, and protein sequences, JBEI uses community data-exchange formats including FASTA, GenBank, and Synthetic Biology Open Language (SBOL).
- For proteomic and metabolomic data, JBEI will use community standard data formats and processing methods for compound identification and validation.
- For structural biology data, JBEI will make use of standard formats for internal workflows, including experimental data, molecular model, and geometric restraints formats.

- Local area network (LAN) and wide-area network (WAN) services for building 978 where JBEI is housed at LBNL are provided by LBNLnet. The LAN consists of two switches connected by a 10 G link with 10 G uplinks to the WAN. The primary WAN link is a 100 G link through ESnet. JBEI has a secondary 10 G link through the Corporation for Education Network Initiatives in California (CENIC).
- LBNL is a Google Workspace customer and has substantial discounts on services running in the Google Cloud Platform (GCP). JBEI is currently working on the instrumentation required to:
  - Archive data to GCS.
  - Provide access to Google's Compute Engine for projects that exceed the capabilities of local resources.
  - Deploy webapps to Google's App Engine.
- JBEI currently has 14 hardware systems dedicated to computational biology that are more than five years old. These will soon need to be replaced. Given the increasing focus on ML, JBEI should replace these with GPU-enabled systems.
- JBEI has established new inter-BRC shared research objectives that will increase the need for efficient data transfer between all four BRCs over the next five years.

5.12.2 JBEI Facility Profile

The DOE's JBEI, led by LBNL, was selected in July 2017 as one of four DOE BRCs and was awarded a total of $125 million over five years to support innovative research on biofuels and bioproducts.

JBEI partners include Brookhaven, Pacific Northwest and SNL; Iowa State University; the UC campuses at Berkeley, Davis, San Diego, and Santa Barbara; and TeselaGen Biotechnology Inc.

JBEI was among three BRCs established by DOE in 2007 to accelerate fundamental research in advanced, next-generation biofuels, and to make such technology cost-effective and widely available. The other two centers were the BioEnergy Science Center, led by Oak Ridge National Laboratory, and the GLBRC, led by the University of Wisconsin–Madison in partnership with Michigan State University.

Since its founding, JBEI has contributed many scientific achievements, including:

- Engineering bioenergy crops to increase sugar-containing polymers and decrease lignin in plant cell walls
- Developing an affordable and scalable ionic liquid pretreatment technology
- Developing microbial routes for the conversion of biomass-derived sugars into advanced, "drop-in" blendstocks for gasoline, diesel, and jet fuels

5.12.2.1 Science Background

Fossil fuels derived from oil, coal, and natural gas provide 81% of the energy consumed in the United States (US), with the established US transportation infrastructure operating nearly exclusively on petroleum-derived hydrocarbons. The global economy is quickly rebounding from the COVID-19 pandemic economic downturn, and fossil fuel consumption worldwide is once again on the rise. Burning fossil fuels contributes to climate change, and the production of many chemicals and materials from petroleum requires significant energy inputs and pollutes the environment. The development of renewable biofuels and bioproducts that reduce reliance on petroleum is critical to the energy, environmental, and economic security of the nation. Lignocellulosic biomass from non-food crops could provide a large fraction of those biofuels and bioproducts, as it has been estimated that there are approximately one billion dry tons of lignocellulosic biomass available annually in the US. To mobilize this strategic renewable carbon resource and accelerate the bioeconomy of the US, there remain significant scientific and technological roadblocks that must be addressed: 1) lack of scalable and sustainable bioenergy crops optimized for deconstruction and conversion; 2) difficulty in deconstructing and fractionating bioenergy crops into targeted bioavailable intermediates at high yields using conventional biomass pretreatments; 3) lack of a robust feedstock-agnostic pretreatment technology; and 4) lack of efficient and affordable microbial routes to produce drop-in biofuels for automotive, aviation, and diesel engines, and bioproducts that can displace, and/or be superior to, petroleum-derived products. JBEI's mission is to establish the scientific knowledge and new technologies in feedstock development, deconstruction and separation, and conversion needed to transform the maximum amount of carbon available in bioenergy crops into biofuels and bioproducts (Figure 5.12.1). When fully scaled, JBEI's technologies will enable the production of:

- gasoline, diesel and jet fuel replacements at ~$4.00 per gallon (rack price) without a bioproduct;
- gasoline, diesel and jet fuel replacements at ~$2.50 per gallon (rack price) when bioproducts are co-produced with the fuel;
- drop-in, commodity bioproducts (production of a million tons per year or more) that can compete with the same petroleum-derived molecules and that reduce biofuel prices; and
- novel bioproducts that cannot be efficiently produced from petroleum, have desirable properties, and reduce biofuel selling prices.

In doing so, JBEI will reduce the nation's dependence on fossil fuels, significantly reduce the amount of carbon added to the atmosphere by at least 70% over fossil fuels, reduce contamination of the environment, and provide the scientific tools and knowledge required to transform the bioenergy marketplace. Informed and benchmarked by rigorous techno-economic and life-cycle analyses, JBEI's research program takes a systems-level approach to maximize the conversion of biomass into biofuels and bioproducts, simultaneously optimizing the composition of bioenergy crops, the deconstruction and separation process, and the metabolism of biofuel- and bioproduct-producing microorganism. That system has the following integrated components:

- high-yielding bioenergy crops engineered for drought tolerance and disease resistance that can be readily deconstructed into sugar and aromatic intermediates and that can be nearly fully used by the engineered biofuel- and bioproduct-producing microorganism;
- an integrated, feedstock-agnostic deconstruction process using biocompatible ionic liquids that liberates at least 90% of the sugars and lignin-derived intermediates suitable for biological conversion; and
- microorganisms engineered with a metabolism that simultaneously uses the sugars and lignin-derived intermediates (e.g., aromatics) resulting from the deconstruction process (and thus match the composition of the engineered plants) and produce a variety of fuels and bioproducts at industrially relevant titers, rates, and yields (TRY).

Developing an integrated biomass-to-biofuels-and-bioproducts process requires a highly integrated research program that benefits from an embedded technology development program to build tools for all aspects of the research. These advances require continual reassessment of the techno-economics and life-cycle implications when there are discoveries or roadblocks that may require a change in the research program. The integration needed to simultaneously develop and optimize these components can only be achieved through an integrative research institute like JBEI. This integrated high-risk, high payoff approach to science aligns with and supports the DOE's strategic plans for renewable energy, bioenergy, mission innovation and sustainability.

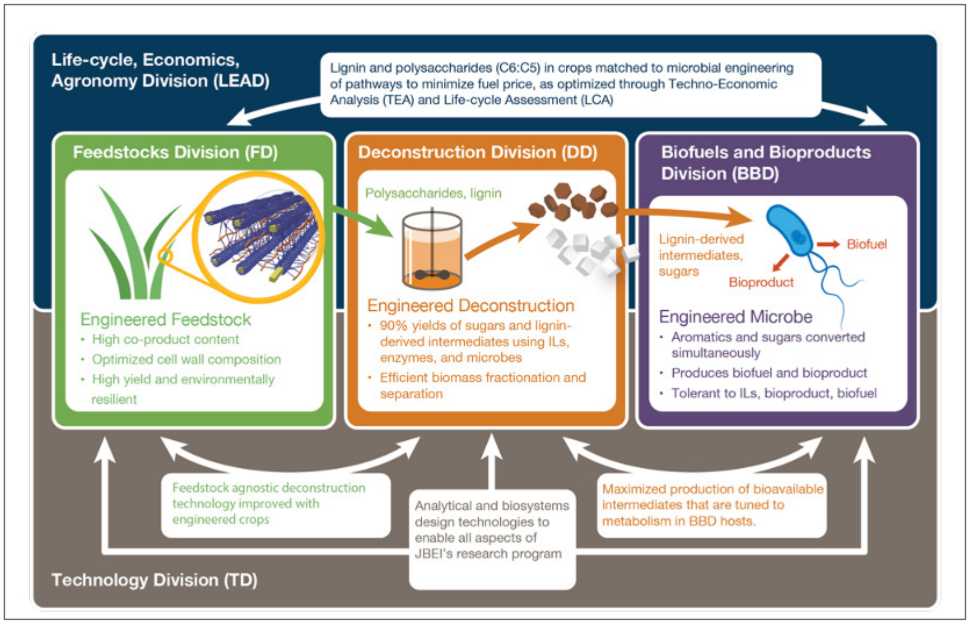

Figure 5.12.1: JBEI's integrated approach to the bioenergy research enterprise.

5.12.2.2 Collaborators

JBEI will have access to scientific infrastructure and instrumentation at LBNL and partner institutions that are well suited to address the most challenging tasks. At LBNL, these include dedicated biotechnology laboratories, specialized facilities for HT structural biology and genomics analysis, high-resolution imaging systems, extensive plant biology and greenhouse research areas, and a powerful supercomputing environment. UCB facilities include a biomolecular nanotechnology facility and NMR facility. JBEI has access to combustion research, supercomputing, and microsystems facilities at Sandia, UCD feedstock growth facilities, and fungal research, proteomics, and bioinformatics facilities at PNNL. In addition, JBEI researchers make use of national research facilities, often operated by DOE national laboratories. These facilities are described in detail below.

5.12.2.2.1 EMSL

EMSL (operated by PNNL) supports world-class research in the biological, chemical, and environmental sciences to provide innovative solutions to the nation's environmental challenges as well as those related to energy production. EMSL's distinctive focus on integrating computational and experimental capabilities as well as collaboration among disciplines yields a strong, synergistic scientific environment. Bringing together experts and state-of-the-art instruments critical to their research under one roof, EMSL has helped thousands of researchers use a multidisciplinary, collaborative approach to solve some of the most important national challenges in energy, environmental sciences, and human health. These challenges cover a wide range of research, including synthesis, characterization, theory and modeling, dynamical properties, and environmental testing. EMSL is a 200,000 sq. ft. facility that houses multiple state-of-the-art instruments and capabilities supported by trained staff. EMSL capabilities include light and electron microscopy (some of which are located in the EMSL Quiet Wing), MS for proteomics and metabolomics analysis, MS imaging instrumentation, a variety of liquid and solid-state NMR instruments as well as laboratory space for sample preparation. As a user facility, EMSL has office and lab space available for visiting scientists. Researchers in Deconstruction and Biofuels and Bioproducts Division (BBD) have ongoing projects funded through the EMSL User Program in the fields of lignin depolymerization and valorization using NMR and MS, lignocellulose deconstruction by microbial communities using omics, and metabolic flux analysis.

5.12.2.2.2 JGI

LBNL manages the JGI, DOE's genomics user facility located on the main LBNL campus. JGI was the world's first genomic user facility dedicated to the mission areas of bioenergy and the environment. It is currently the largest producer of plant and microbial genomes worldwide and focuses on user projects whose scale and complexity require the unique suite of data-generation and data- analysis capabilities found at JGI. The JGI provides users a diversity of cutting-edge experimental and computational technologies that will include massive-scale sample preparation and sequencing, single-cell genomics, biodesign and gene synthesis, and deep functional annotation of plant and microbial genomes. JBEI researchers have numerous ongoing collaborative projects resourced through the BRC pipeline in place at JGI on the topics of sequencing plants, metagenomics, de novo sequencing, gene synthesis, bioinformatics, and resequencing.

5.12.2.2.3 NMDC

The long-term vision of the NMDC is to support microbiome data exploration through a sustainable data discovery platform that promotes open science and shared ownership across a broad and diverse community of researchers, funders, publishers, societies, and other collaborators such as the teams working on microbial communities at JBEI. The NMDC mission is to work with the community to iteratively develop and pilot an integrated, open-source microbiome science gateway that leverages existing resources and enables comprehensive access to multidisciplinary microbiome data and standardized, reproducible data products. The NMDC enables the microbiome research community to decode the molecular underpinnings of fundamental biological processes, and ultimately, drive transformational discoveries. JBEI researchers are active collaborators with the NMDC team and will share data from the study of microbial communities.

5.12.2.2.4 KBase

KBase is a software and data platform designed to meet the grand challenge of systems biology: predicting and designing biological function. KBase integrates data, tools, and their associated interfaces into one unified, scalable environment, so users do not need to access them from numerous sources or learn multiple systems to perform sophisticated systems biology analyses. Users can perform large-scale analyses and combine multiple lines of evidence to model plant and microbial physiology and community dynamics. KBase enables JBEI researchers to upload their own data, add their own analysis methods, analyze their data (along with collaborator and public data), build increasingly realistic models, and share and publish their workflows and conclusions. JBEI researchers from Deconstruction, BBD, and Technology use KBase narratives to analyze RNA-seq data and will work closely with KBase researchers to integrate JBEI tools into KBase.

5.12.2.2.5 Genomic Sciences Foundational Science Focus Area Ecosystems and Networks Integrated with Genes and Molecular Assemblies (ENIGMA)

The LBNL-led ENIGMA (Ecosystems and Networks Integrated with Genes and Molecular Assemblies) SFA is focused on basic scientific research on microbial communities in groundwater and sediment. ENIGMA seeks to understand the mechanisms that drive the persistent and changeable activities of microbial communities in complex environments. Defining features of the SFA are: a quantitative approach to characterizing microbial dynamics from molecules to meters through integrated field and laboratory experimentation, an emphasis on new technology development, and the application of rigorous computational methods. The program targets DOE mission-relevant microbes and their communities, to understand the impacts of environmental stresses on microbial survival and adaptation, such as contamination in soil and water. Although ENIGMA's research mission is different from JBEI's, there are natural areas of complementarity, principally in the area of technology development. ENIGMA's study of microbial communities has led to the development of methods for analyzing metabolites, protein networks, and cellular organization. Exchange of expertise and technologies between the two programs is therefore expected and encouraged. The synergy between these two programs helps both JBEI and ENIGMA more rapidly reach their goals and also explore new research directions. The JBEI Vice President for Technology, Paul Adams, is also the Laboratory Research Manager for the ENIGMA program, and is therefore well placed to promote synergistic interactions between the two programs.

5.12.2.2.6 Microbial Community Analysis and Functional Evaluation in Soils (m-CAFEs)

m-CAFÉ is LBNL-led project is a multi-institutional SFA aimed at understanding the interactions, localization, and dynamics of grass rhizosphere communities at the molecular level (i.e., genes, proteins, metabolites). Gaining this understanding will enable prediction of community responses to perturbations and insights into the persistence and fate of engineered genes and microbes for secure biosystems design. To achieve this understanding, advanced fabricated ecosystems (EcoFABs) are used in combination with gene-editing technologies such as CRISPR-Cas and bacterial-virus approaches (i.e., phage-based methods) to interrogate gene and microbial functions in situ—addressing key challenges highlighted in recent US DOE reports. This work is integrated with the development of predictive computational models that are iteratively refined through simulation and experimentation. Simplified microbial consortia assembled in the laboratory are used in combination with studies of partially reduced, native soil-derived microbial communities to gain critical insights into engineered genes and microbes within soil microbiomes and the biology and ecology of unc ultivated microbes. This SFA is currently piloting meter-scale contained and controlled ecosystems (called EcoPODs) that enable the extension of m-CAFEs science into more complex environments. The research on microbial communities leverages the team's extensive expertise in CRISPR-Cas systems, phage biology, genome-resolved metagenomics, HT bacterial and fungal genetics, and systems biology, as well as plant mutant collections and phenotyping capabilities. Together, these efforts lay a critical foundation for developing secure biosystems design strategies, harnessing beneficial microbiomes to support sustainable bioenergy, and improving the understanding of nutrient cycling in the rhizosphere. The knowledge gained and approaches pioneered by m-CAFEs can be extended to other ecosystems and will advance microbiome research toward a more predictive and causative science.

5.12.2.2.7 Agile BioFoundry (ABF) (LBNL, SNL, PNNL, NREL, ANL, LANL, ORNL, INL, Ames)

The Agile BioFoundry (ABF), supported by the DOE's Office of Energy Efficiency and Renewable Energy (EERE), Bioenergy Technologies Office (BETO), aims to reduce biomanufacturing process development and commercialization timelines and costs, as well as increase biomanufacturing carbon and energy efficiencies. This is accomplished through a collaborative effort distributed across nine national laboratories, leveraging their aggregate infrastructure and domain expertise. Specifically, the ABF brings together an Integrated Design/Build/Test/Learn cycle, Integrated Analytics (Techno-Economic Analysis, Life Cycle Analysis), Host Onboarding (further development of industrially relevant organisms for use at the ABF), and Process Integration and Scaling. In the course of its efforts, the ABF develops and generates software, methods, and information that can broadly enable biomanufacturing efforts in industry, academia, and the nonprofit and government sectors[61].

61 http://agilebiofoundry.org

5.12.2.2.8 ALS — LBNL

ALS, supported by the DOE's Office of Basic Energy Sciences, is the world's brightest soft X-ray synchrotron radiation source, one of the most prolific synchrotrons and a powerful source of hard X-rays for structural studies. It is home to seven macromolecular crystallography beamlines, one small angle scattering beamline. These eight beamlines are run by four groups: the Berkeley Center for Structural Biology (BCSB, beamlines 5.0.1, 5.0.2, 5.0.3, 8.2.1, and 8.2.2), the Structurally Integrated for Life Sciences (SIBYLS) group (beamline 12.3.1), the Molecular Biology Consortium (MBC, beamline 4.2.2), and the UCB/UCSF consortium (beamline 8.3.1). These ALS structural biology beamlines are state-of-the-art mature resources, serving collectively ~800 users annually. The National Center for X-ray Tomography operates beamline 2.1 for soft X- ray microscopy of biological materials and soft matter. JBEI researchers will have access to the ALS beamlines through the merit-based peer-reviewed general user program for protein-structure determination and other measurements of plant cells and cellulosic materials[62].

5.12.2.2.9 Advanced Biofuels Process Demonstration Unit (ABPDU) — LBNL

The ABPDU, located at LBNL, is a state-of-the-art facility for testing and developing emergent biofuels and bioproducts technologies. This 15,000-square-foot facility is available to BRCs, DOE-supported researchers, academic institutions, nonprofit research organizations, and companies involved in biofuels and bioproducts research and development. Capabilities at this newly built center include: chemical catalysis, thermochemical pretreatment, enzymatic biomass saccharification, biomass extraction, biofuel and bioproduct fermentation, enzyme fermentation, downstream processing, and analytical technologies. The ABPDU was built and is operated with funds from the Bioenergy Technologies Office (BETO) within the DOE Office of EERE, including start-up funds from the American Recovery and Reinvestment Act (ARRA). When elements of the JBEI R&D program reach an appropriate level of technical maturity, JBEI researchers and the ABPDU team will draft and agree on statements of work to use the facilities at the ABPDU to evaluate, optimize and scale-up the relevant technology at the process demonstration stage of production (50 – 200 L scale of biomass deconstruction and 50-300 L scale fermentation[63]).

5.12.2.2.10 Berkeley Center for Structural Biology (BCSB) — LBNL

The BCSB provides experimental facilities for the collection of X-ray diffraction data at the ALS. The Center operates 5 experimental end stations (beamlines): 5.0.1, 5.0.2, 5.0.3, 8.2.1, and 8.2.2. Beamlines 5.0.1, 5.0.2 and 5.0.3 are located on sector 5.0, which uses a 2 Tesla wiggler to generate a high intensity X-ray beam. Beamline 5.0.2 is wavelength tunable and optimized for multiwavelength anomalous dispersion experiments. Beamlines 5.0.1 and 5.0.2 are monochromatic (12.7 keV) and are ideally suited for the collection of native data for structure solution by molecular replacement, and selenium anomalous scattering experiments. All three beamlines are equipped with fast pixel array or CCD-based X-ray detectors, and Oxford cryostream low temperature systems for crystal cooling. Additionally, all three beamlines have automated sample changing robots to facilitate rapid screening of crystals. Beamlines 8.2.1 and 8.2.2 are located on sector 8.2, which uses a 5 Tesla superbend magnet to generate a high intensity X-ray beam. Both beamlines are wavelength tunable and optimized for multiwavelength anomalous dispersion experiments. Both beamlines are equipped with fast readout CCD-based X-ray detectors, and Oxford cryostream low temperature systems for crystal cooling. Additionally, both beamlines have automated sample changing robots to facilitate rapid screening of crystals. A new microfocus beamline (2.0) will come on line in the near future and provide the highest brightness macromolecular crystallography X-ray source at the ALS. These beamlines are accessed by JBEI researchers through the ALS peer-reviewed general user program, or through BCSB discretionary time. Crystallographic experiments can be used for optimizing protein design, and understanding and engineering conformational changes in enzymatic processes[64].

62 https://www-als.lbl.gov/
63 http://abpdu.lbl.gov/
64 https://bcsb.als.lbl.gov/

5.12.2.2.11 Biomolecular Nanotechnology Center (BNC)

The Biomolecular Nanotechnology Center (BNC) is an 11,500 sq. ft. class 1,000/10,000 cleanroom facility located in Stanley Hall on the UC Berkeley campus. The BNC is a unique fabrication and experimentation facility specializing in BioMEMS (biomedical micro-electromechanical systems) and microfluidic devices. The facility has a complete lithography microfabrication as well as soft lithography and glass or polymer bonding capabilities. The BNC features a full range of deposition, etching, metrology, and microscopy equipment as well as facilities for performing biological experiments. The center focuses on microfluidic processing of glass and polymer materials. Experimentation on proteins, nucleic acids, cells, and tissues are encouraged. JBEI turns to the LBNL Molecular Foundry and the BNC for assistance in the development and testing of new bioproducts that are precursors to new materials[65].

5.12.2.2.12 Center for Integrated Nanotechnologies (CINT) — Los Alamos and Sandia

The microfluidics work at JBEI may need access to equipment housed in DOE nanotechnology centers. The Center for Integrated Nanotechnologies (CINT) is a DOE / SC Nanoscale Science Research Center operating as a national user facility devoted to establishing the scientific principles that govern the design, performance, and integration of nanoscale materials. Through its Core Facility and Gateways to both Los Alamos and SNL, CINT provides access to tools and expertise to explore the continuum from scientific discovery to the integration of nanostructures into the micro and macro worlds. Unique to CINT are Discovery Platforms™, modular, micro-laboratories designed and batch fabricated expressly for the purpose of integrating nano and micro length scales and for studying the physical and chemical properties of nanoscale materials and devices that can be used for the study of enzymes and microbes relevant to JBEI's bioenergy mission[66].

5.12.2.2.13 Central California 900 MHz NMR Facility — UC Berkeley

The 900 MHz NMR facility at UCB is administered by the California Institute for Quantitative Biosciences (QB3). Instruments include state-of-the-art 900 MHz and 600 MHz spectrometers, each equipped with a cryoprobe, as well as 800 MHz and 500 MHz spectrometers. JBEI researchers access the facility through a collaborative agreement with the director, Professor David Wemmer. This facility provides invaluable support to organic synthesis capabilities as well as biomass and lignin analytics[67].

5.12.2.2.14 College of Natural Resources Biological Imaging Facility — UC Berkeley

The College of Natural Resources Biological Imaging Facility. The laboratory is equipped with state-of-the-art confocal and standard fluorescence microscopy, micro-technique equipment, and provides training in digital image processing and analysis with two full-time instructors. Two other imaging centers exist at UCB. The Molecular Imaging Center has multiphoton, laser scanning confocal and deconvolution microscopy; and the Electron Microscopy Laboratory has transmission and scanning electron microscopy, including an Hitachi environmental scanning electron microscope that will be used to study the structural motifs present in plant cell walls and how they are disrupted during IL pretreatment[68].

65 http://qb3.berkeley.edu/bnc/
66 http://cint.lanl.gov/
67 http://qb3.berkeley.edu/nmr/
68 http://microscopy.berkeley.edu/

5.12.2.2.15 Combustion Research Facility (CRF) — Sandia

The Combustion Research Facility (CRF) is a 75,000-square-foot facility with 36 labs. It is home to about 100 scientists, engineers, and technologists who conduct basic and applied research aimed at improving the control and use of combustion processes. CRF scientists pioneered the use of laser-based diagnostics, now essential for the study of combustion. There is an existing CRF-JBEI memorandum of understanding (MOU) to evaluate the advanced biofuels produced at JBEI in the engine labs at the CRF. This collaboration has generated several publications, including one on the testing of isopentanol in an Homogeneous Charge Compression Ignition engine30 that serves as the foundation for JBEI's work on isopentanol as a fuel. JBEI will continue to work with the CRF to identify new candidate fuel molecules that can be produced biologically[69].

5.12.2.2.16 Electron and X-ray Microscopy Facilities — UC Berkeley

JBEI researchers access in-house electron microscopes including a JEOL 1400 for sample screening, and a JEOL 3200 for high-resolution cryo-EM/tomography work. JBEI is also able to conduct basic electron microscopy (EM) and room-temperature electron tomography work at the UC Berkeley Electron Microscopy Laboratory. As a member of the Bay Area Cryo-EM consortium JBEI has access to a FEI Krios electron microscope equipped with a K3 direct electron detector, to perform high-resolution electron tomography and/or single particle analysis cryo-EM work. The LBNL Molecular Foundry provides access to a basic Focused Ion Beam Scanning Electron Microscopy (FIB-SEM) microscope and regularly solicits proposals for microscope access.

5.12.2.2.17 IMRL and MANTL — Sandia

The 140,000-square-foot Integrated Materials Research Laboratory (IMRL) and 80,000- square-foot Micro- and Nano-Materials Research Laboratory (MANTL) house most of the advanced materials research and development functions at SNL. With about 350 researchers, these facilities integrate research from the atomic scale, through the development of electronic devices, to full-scale mechanical components. The experimental work is augmented by advanced computer modeling and simulation techniques. JBEI has ongoing collaborations with SNL staff in the analysis of samples using ITC-DSC, TGA, TMA, and AFM. JBEI will have access to these facilities to evaluate new bioproducts that are precursors to new materials.

5.12.2.2.18 Molecular Foundry — LBNL

The Molecular Foundry is a DOE-funded nanoscience research facility that provides users access to cutting-edge expertise and instrumentation in a collaborative, multidisciplinary environment. The Molecular Foundry at LBNL supports research and training using state-of-the-art instruments, techniques, and expertise to pursue the integration of biological components into nanotechnology. Two of its focus areas are especially germane to bioenergy: biological nanostructures and organic and biomolecular synthesis. JBEI will turn to the Molecular Foundry for assistance in the development and testing of new bioproducts that are precursors to new materials[70].

69 http://crf.sandia.gov/
70 http://foundry.lbl.gov/

5.12.2.2.19 National Energy Research Scientific Computing Center (NERSC) — LBNL

NERSC provides one the most powerful open basic-science supercomputing environments in the world. NERSC delivers billions of hours of computer time on advanced HPC systems to researchers who analyze extensive genomics and bioinformatics databases, simulate complex fluidics, chemical reactions, and biological systems. The latest computing platform at NERSC is Perlmutter, an HPE Cray EX supercomputer, is a heterogeneous system with both GPU-accelerated and CPU-only nodes. Its projected performance is three to four times that of NERSC's current flagship system, Cori. In addition to this computing resource, NERSC has a 120 Pbyte of high-performance storage capacity. NERSC's HPC experts excel in migrating cluster-sized computational workflows to HPC. Researchers directly affiliated with JBEI have access to NERSC's facilities to develop advanced, multiscale models of cellulosic deconstruction, IL pretreatment, and for large-scale bioinformatic analysis of JBEI experimental data[71].

5.12.2.2.20 Oxford Tract Greenhouse Facility, College of Natural Resources — UC Berkeley

The Oxford Tract resources include mist propagation beds and rooms, workspace, supplemental lighting, and temperature and lighting controls. The Genetics and Plant Biology (GPB) (North) greenhouse opened in 1993 and offers 10,000 sq. ft. of climate-controlled greenhouse space, and shared head house and laboratory workspace. The Oxford Research Services (RES) (South) greenhouse has 23,100 sq. ft. of greenhouses, four lath houses, four shared head houses, assigned laboratory, office and storage space, and six growth chambers. The Insectary greenhouse opened in 2002 contains 4400 sq. ft. of greenhouse and 2200 sq. ft. of lath house space in support of insect biology research[72]. JBEI will use the greenhouses for growing sorghum and poplar to phenotype plants and provide biomass samples to teams at JBEI, the other BRCs, and external collaborators.

5.12.2.2.21 QB3 Computational Genomics Resource Laboratory (CGRL) — UC Berkeley

The mission of the QB3 Computational Genomics Resource Laboratory (CGRL) is to serve as a focal point for scientists interested in comprehensive analysis of next-generation sequencing data. The CGRL provides: computational infrastructure for data analysis; training in analytical tools for next-generation sequence data; and project- specific consultation on experiment design and analysis. To equip researchers with state-of-the-art tools and knowledge on evolving trends in computational biology and bioinformatics, the CGRL organizes seminars by leading experts in the field. In addition, there are workshops on specific methods and topics given by experts in the subject. The CGRL also provides one-on-one consultation on specific projects. CGRL coordinates with the existing next-generation sequencing UCB core facilities - Vincent J. Coates Genomics Sequencing Laboratory (GSL) and the Functional Genomics Laboratory (FGL) - to keep ahead of new developments. JBEI users of these sequencing facilities have the option of seamlessly transferring data and carrying out secondary and tertiary analysis of their data utilizing the CGRL computing infrastructure[73].

5.12.2.2.22 Spallation Neutron Source (SNS) — Oak Ridge

JBEI researchers access beamlines at the Spallation Neutron Source (SNS) at ORNL through standard user proposal mechanisms. Neutron reflectometry is being used to study the physical and chemical interactions of ligninases and mediators with lignin. JBEI has discovered that ligninases are effective for degrading lignin only for a certain range of concentrations, and JBEI is using neutron reflectometry to unravel the mechanisms behind that behavior. This work supports development of optimized ligninase cocktails for lignin breakdown[74].

---

71 http://www.nersc.gov/
72 https://nature.berkeley.edu/oxford-facility/about
73 http://qb3.berkeley.edu/cgrl/
74 https://neutrons.ornl.gov/sns

5.12.2.2.23 University of California Agriculture and Natural Resources

Work conducted at the University of California Agriculture and Natural Resources Kearney and West Side Research and Extension Centers (KARE and WCREC) will be supported as part of the subcontract with UC Davis.

The 330-acre Kearney Agricultural Research and Extension Center is the University of California's largest off-campus agricultural research facility. Kearney's research facilities encompass a state-of-the-art greenhouse, a postharvest laboratory, a mosquito control laboratory, multiple insectaries, and academic and administrative offices. Its controlled field studies occupy 260 acres of orchards, vineyards and fields. The Kearney Agricultural Research and Extension Center has 170,000 square feet of facilities consisting of 45 offices, 40 laboratories and 30 greenhouses. Other available resources include a postharvest lab, drying equipment, cold storage, a weighing lysimeter and a weather station. Conferences, workshops, seminars and other educational activities are accommodated by a 240-seat conference facility and smaller meeting rooms. Office and laboratory space is available on an as-needed basis for temporary and visiting scientists.

The West Side Research and Extension Center (WSREC) has one laboratory equipped for pathology or limited wet chemistry work, including a fume hood. Two other labs can be used for entomology work or other microscope work, or are suitable for plant or soil sample preparation or measurement work. Soil and plant grinding areas are available in the headhouse area. Multiple storage areas are available for equipment and supplies storage, plus specialized storage for gas-powered portable equipment or equipment such as neutron probes that require locked storage spaces. The center has a 400-square-foot lath house for transplants and outdoor plant growing, two greenhouse spaces (one approximately 1,800-square-foot and the other approximately 1,200 square feet) and both are currently used primarily for drying of soil and plant samples.

5.12.2.2.24 UC Davis Genome Center — UC Davis

The UC Davis Genome Center, of which Ronald is a faculty member, is the technological antenna in the genomics arena for the campus, providing state-of-the-art facilities and expertise in genomics, proteomics, metabolomics and bioinformatics. The bioinformatics group provides expertise and infrastructure to carry out acquisition, curation and distribution of complex datasets, as well as to develop and perform computations, analyses and simulations addressing a wide variety of biological questions from genomics and systems biology. The Bioinformatics Core's Director is available to assist with project objectives. The Proteomics Core provides state-of-the-art analytical proteomic services. In addition to highly trained staff, the facility has three tandem mass spectrometers capable of providing sensitive, accurate and fast proteomic analysis of complex biological samples, including a Thermo Q-exactive mass spectrometer with a Proxeon nano-spray source and Easy-LC (liquid chromatography) II high-performance liquid chromatography (HPLC). They also perform AA analysis and Edman protein sequencing[75].

5.12.2.2.25 UC Davis Core Greenhouse Facilities — UC Davis

The College of Biological Sciences supports 17 greenhouses, three headhouses, a screenhouse, outdoor benches, and two lath houses at the Orchard Park facility suitable for switchgrass and sorghum growth. The College of Biological Sciences supports two greenhouses and one headhouse at the Core Greenhouse Complex. The greenhouse staff provides the services needed for successful plant growth, including pots and media, daily watering (seven days/week including holidays), inspections, light intensity control and pest control.

75 http://proteomics.ucdavis.edu.

5.12.2.2.26 UC Davis Field and Laboratory Research — UC Davis

UC Davis has excellent large field facilities for the measurement of crop phenotypes under true-to-life crop production environments. Soils at Davis are deep alluvial clay- loams, with excellent soil fertility and water-holding capacity. The environmental conditions (temperature, daylength, rainfall) are typical of a Mediterranean climate and representative of the large agricultural regions of California (San Joaquin Valley and the Sacramento Valley) and the world (e.g., South Africa, Australia, Argentina, Chile, Southern Europe, N. Africa) which have similar growing conditions. The lack of rainfall during most of the season (May-October) enables careful control of water applications, and studies of water deficits for crop production that are impractical in temperate regions with summer rains that interfere with deficit water experimental designs. These are high-yielding environments, so the genetic potential of crops is often maximized here, particularly heat-loving C4 crops. All of the infrastructure for crop management, tillage equipment, experimental planters, harvesters, sprayers, and the expertise to manage this equipment exists at UC Davis, and is primarily housed in the Plant Sciences Field Research Unit on Campus or the UCD College of Agriculture and Environmental Sciences. While some specific field equipment may be needed specific to this grant for the switchgrass, PIs on campus have developed the capabilities for planting and harvesting small- and large-plots with hundreds of treatments per experiment over many years of field research. Yields will be measured with experimental crop harvesters that measure small-plots (e.g., 1 m to 15 m in length, 1 m width). Different machines are usually needed for sorghum and for switchgrass since many sorghum varieties are very tall in stature. Varied approaches to irrigation are feasible: drip, sprinkler, check-flood or furrow, depending upon the experimental design, plot size, and the needs of the researcher. The infrastructure and water supply for drip irrigation or sprinkler or flood irrigation techniques is available on campus, and the expertise to manage those facilities exists in the persons of the PI, field station staff, and project scientists. Field Evapotranspiration (ET) is available on campus, and detailed measurements of ET are feasible. Greenhouse space is available for generation of plants for transplanting, and laboratory space is available for sample processing, drying, grinding, and other analyses. Machine-driven transplanters (generally used for tomatoes) are available for use in the biofuel project, if needed. The forage lab on campus in the Plant and Environmental Sciences building has capabilities for seed and sample handling, processing, in vitro gas analysis, Near-Infrared Spectroscopy, wet chemistry analysis of Neutral Detergent and Acid Detergent Fiber, respectively). Other analyses for soil and plant tissue for protein, minerals, and other plant characteristics are available from the UC Davis Analytical Lab, although it is expected that most detailed biofuel analyses will be conducted at JBEI.

5.12.2.2.27 UC Davis Growth Chamber Facilities — UC Davis

The Controlled Environment Facility maintains a total of 151 plant growth chambers in two separate locations on campus. Chamber controls include light intensity, temperature, relative humidity, photoperiod and irrigation systems. There are also many specialized units, including 15 for precise control of carbon dioxide levels.

5.12.2.2.28 Vincent J. Coates Genomics Sequencing Laboratory — UC Berkeley

The Vincent J. Coates GSL at the University of California, Berkeley, is a core research facility managed by the California Institute for Quantitative Biosciences (QB3), a cooperative effort between the state of California, private industry, venture capital, and the UC Campuses at Berkeley, San Francisco, and Santa Cruz. Partnered with the Functional Genomics Laboratory (FGL) for Library Preparation, the GSL in Stanley Hall provides Next-Generation Sequencing on Illumina's HiSeq and MiSeq platforms, as well as Illumina Array services on the HiScan and Biospecimen processing. The GSL's mission is to make massively parallel DNA sequencing, Illumina Array Services, and Biospecimen processing affordable to UC Berkeley and the greater UC community. The GSL is equipped with an Illumina HiSeq2500 in Rapid Mode, an Illumina HiSeq4000, two Illumina MiSeq systems, and an Illumina HiScan for arrays. The GSL can also prep PacBio RSII libraries for sequencing at a sister UC Core until an instrument can be acquisitioned[76].

---

76 https://qb3.berkeley.edu/facility/genomics/

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
JGI	Primary	Data portal	200-500 Gb	Biweekly	N	Slow network
EMSL	Primary	Data portal and data transfer	100-200 Gb (more if includes EM images)	Monthly	N	Slow network
ALS	Primary	Data portal and data transfer	100-200 Gb	Ad hoc	N	Slow network
CBI	Both	Data transfer	200-500 Gb (envisioned)	Ad hoc/Monthly	Y	Slow network
CABBI	Both	Data transfer	200 -500 Gb (envisioned)	Ad hoc/monthly	Y	Slow network
GLBRC	Both	Data transfer	200-500 Gb (envisioned)	Ad hoc/monthly	Y	Slow network
NMDC	Secondary	Data transfer	50 Gb	Ad hoc	N	Slow network

Table 5.12.1: JBEI Collaboration Space

5.12.2.3 Instruments and Facilities

JBEI has extensive, fully operational research facilities and resources at Emeryville Station East (ESE), a facility located in the hub of the San Francisco Bay Area biotech corridor at 5885 Hollis Street in Emeryville, CA. The facilities are managed day-to-day by professional staff and are available to all JBEI researchers. Each facility and resource is described in detail below.

The JBEI team brings together research from a variety of scientific disciplines under one roof for a truly integrated approach to maximize productivity and increase the likelihood of success in every facet of the JBEI research program. JBEI will be headquartered in ESE. In 2008, the fourth floor of ESE was leased by LBNL and built into a state-of-the-art 65,000-square-foot research facility. The space was built and sized with capabilities to provide room for shifting research needs and new technologies. It houses over $14 million of equipment that will be dedicated to facilitating JBEI's proposed research program. LBNL has a long-term lease with the landlord and pays all lease costs. There is no additional expense over and above approved burden rates for lease or space costs.

ESE was renovated and fully equipped in 2008 with a total investment of $14.879 million. The 65,000 sq. ft. facility with 239 desks and 173 lab benches requires no renovation or alteration and has ample space to house JBEI personnel, equipment and support space through the duration of the performance period. ESE will house the vast majority (~90%) of JBEI personnel and activities. JBEI activity not housed in the ESE (~10%) will occur at partner institutions for work such as field trials. These facilities are described in detail above for each partner. There is space designated at JBEI to house partner scientists when they are in-residence at ESE. LBNL's Guest House is available to provide short-term accommodations to visiting JBEI partners and collaborators. The Guest House is conveniently located on the LBNL campus and features 57 tastefully appointed guest rooms.

Currently, JBEI's access to ESE is managed under LBNL's Covid-19 access and safety protocols. These protocols are revised based on continual monitoring of the local COVID-19 case rate, CDC guidance and local public health direction. Protocols are approved by DOE and JBEI has developed a comprehensive internal COVID-19 portal. JBEI is currently able to provide all bench researchers 24/7 access.

JBEI's combined laboratories contain 15 chemical fume hoods (including one walk-in fume hood); ten LABCONCO Purifier Logic+ Class II Type A2 Biosafety Cabinets; two 37 °C and two 30 °C warm rooms totaling 500 square feet; three 4 °C cold rooms totaling 445 square feet; two industrial deionizing water systems, three Thermo Scientific Barnstead Nanopure water polishers, two Elga Purelab water polishers and an Elga Chorus Reverse Osmosis water polisher system; 21 Thermo Revco -80 °C freezers; house vacuum, purified nitrogen, and clean, dry air.

JBEI's dedicated Laboratory Operations group provides dishwashing, autoclaving, and media preparation services; strain archive and retrieval; central shipping and receiving; and storerooms for glassware, chemicals, and common laboratory supplies. The warm and cold rooms, freezers, refrigerators, plant growth room, and growth chambers are constantly monitored and alarmed for temperature, power consumption and/or humidity by Klatu Systems Traxx Motes. All major equipment is covered for on-site preventative maintenance and repairs by manufacturer's service agreements.

Robotics Laboratory is equipped with Beckman Coulter Laboratory Automation Workstations (Biomek FXp [equipped with two liquid-handling pods, Spectramax Multimode Microplate Reader, Sigma 6K centrifuge, Cytomat hotels, random- access ambient to 70 °C incubator, random-access stacker-carousel, vacuum filtration unit, shaker-Peltier, ELISA microplate washer, tip washer, and bubble- paddle mixer] and Biomek NXP); Coastal Genomics Nimbus Select Workstation; Liquid Handler Robots (Beckman Coulter Biomek FX and NX); Labcyte Echo 550 Acoustic Liquid Dispensers; Molecular Devices QPix 460; Formulatrix Mantis liquid handler; Agilent ZAG DNA Analyzer, Agilent BenchCel Microplate Handler, VCode Microplate Barcode Labeler. Two Opentrons OT-2 liquid handlers and one Dispendix I.DOT One Liquid Handler. A Biorad CFX96, a Bio-Rad CFX384, and a Bio-rad CFX-Opus384 RT-PCR. An Eppendorf 5392 Heat Sealer and an Illumina MiSeq sequencing platform and an Inscripta Onyx digital genome engineering platform.

Biomass Pretreatment Laboratory is equipped with three Globe reactor systems; one Agilent 1290 HPLC and one Agilent 1200 HPLC; three Parr Reactors; YSI 2950 Analyzer; Mettler Toledo Infrared (IR) Moisture Analyzer; Retsch PM100 Ball Mill; two Thomas Wiley Model 4 Knife Mills a Labman Automated Grinding and Dispensing Robot; a SpexSamplePrep 6875 Freezer Mill; a Mettler Toledo TGA/DSC 3+ Thermal Analyzer; an Across International GFC1200 Muffle Furnace; a Tosoh HLC-8320 GPC and a Thermo Dionex ASE 350 Solvent Extractor.

Microscopy Laboratories (2) is equipped with a Zeiss LSM 710 Laser Scanning Microscope; Bruker Optics RFS27 MultiRAM FT-Raman Spectrometer / RamanScopeIII Raman Microscope; Bruker Optics Tensor 27 FTIR / Hyperion 3000 FPA Imaging Microscope; Leica DM4000 Upright Microscope; Leica MZ16 CMO Microscope; and a Thermo Fisher EVOS M7000 Imaging System.

Spectroscopy Laboratory is equipped with an Agilent 1100 LC with a G7102A ELSD; Beckman Coulter DTX880 Multimode Detector, Bruker Optics MPA FT-NIR (near-infrared) Spectrometer, Bruker Optics Vertex 70 FTIR / HTS-XT 96-well Plate Reader.

Analytical Facility is equipped with an Agilent 1290 LC / Thermo Orbitrap Exploris 480 LC-MS; Agilent 1290 LC / 6550 Q-TOF MS; two Agilent 1290 LC / 6460 QqQ MS; an Agilent 1260 LC / SCIEX 4000Q-Trap MS; an Agilent 1200 LC / 6210 TOF MS; an Agilent 1290 LC / 6545 QTOF LC/MS; an Agilent 1260 / LC 6520 Q-TOF LC/MS; two Agilent LC/MSD iQ systems; an Agilent 1200 Preparative LC; an Agilent 7890A GC / 5975C MSD; one Agilent 8890 GC-FID; two Agilent Intuvo 9000 GC/ 5977B MSD; one Agilent Intuvo 9000 GC / 7000D MS/ MS; one Agilent 6980 GC/ 5973 MSD; a Thermo Trace GC Ultra / Polaris Q MS; a Thermo Trace GC Ultra / Polaris Q MS / CDS Analytical 5250 Pyroprobe; a Shimadzu GC-2010 Plus /GCMS-QP-2010 Plus / CDS 6200 Pyroprobe; and a Shimadzu GC-2014 GC and a Bruker ultrafleXtreme MALDI-TOF MS.

Synthetic Biology Laboratories is equipped with a BD Biosciences FACSAria II Cell Sorter; two Cytiva ÄKTA Explorer fast protein liquid chromatography (FPLC) protein purification systems; a Cytiva ÄKTA Pure 25 M protein purification system and an ÄKTA Avant protein purification system; two Thermo Scientific Dionex

ICS-5000 HPAEC-PAD systems and one Thermo Scientific ICS-6000 system; Berthold Technologies TriStar2 LB 942 Multimode Microplate Reader; Plate Readers (models Tecan Infinite F200 PRO, three Biotek Synergy H1, three Molecular Devices Spectramax M2; one Molecular Devices Spectramax Plus and one Molecular Devices FilterMax M5); Beckman Coulter Spectrophotometers (models DU-800 and DU-640); Beckman Coulter Centrifuges (models Allegra 25R (two), J-15R (four), AVANTI JE, Avanti J-25, Avanti J-20XP, GS-15R, and J2-MC) and Ultracentrifuges (models Optima L-90K,L8-80M and TLX); and 30 Life Technologies Veriti™ Dx 96- Well Thermal Cyclers. A Nanotemper Monolith N.115 Thermophoresis instrument.

Cultivation Facilities, a dedicated fermentation laboratory is equipped with two Sartorius BIOSTAT B-plus reactors and one Sartorius BIOSTAT B Twin reactor; Kuhner RAMOS (Respiration Activity Monitoring System); m2p-labs BioLector 1/ Robolector and one m2p-labs Biolector Pro microbioreactor; Thermo ONIX Vg Gas Prima db mass spectrometer, and two Agilent 1100 LCs. Facilities for anaerobic microbiology include two Coy Laboratory Products Type B (two- person) anaerobic chambers and an anaerobic gassing station.

Plant Growth Facilities includes a 400 square feet light- and temperature-controlled plant growth room; five Percival AR-100L3 growth chambers; one Percival AR41L2 growth chambers; three Percival PGC-105-LED growth chamber; five Percival CU41L4 plant tissue culture growth chambers, and two Conviron PGR15 plant growth chambers; a Taiwan Hipoint FH-2300 plant growth Chamber; a Taiwan Hipoint SS-2300 growth chamber; custom controlled atmosphere growth chamber for 13C labeling of plants; two Thermo Fisher Scientific MaxQ 5000 lighted; refrigerated shaking incubators for plant cell culture; and vertical and horizontal laminar-flow hoods for plant tissue culture.

Computing Server Room houses JBEI systems for data storage, data backup, bioinformatics analysis, Docker services, and web services. JBEI IT staff maintain 15 hardware hosts and about 65 virtual hosts. The primary storage system provides about 100 TB of high-speed disk storage. Primary storage gets replicated to a secondary system co-located at the UC Davis campus data center. The secondary system provides about 200 TB of archive storage. JBEI IT manages several mission critical web applications, for example EDD and DIVA.

Computing Cluster maintained by LBNL IT, and is located in a large server room on the main LBNL site. The cluster consists of 120 nodes, where each node has dual Intel X5550 2.66Ghz 4-core processors (8 cores/node) for a total of 960 cores. Each node has 24 GB of RAM, and they are connected by a QDR 40 Gb per second Infiniband interconnect. The cluster has 40 GB of disk storage attached. In addition, the cluster is connected to a shared global Lustre parallel file system that provides 1 PB of scratch space. JBEI researchers also have access to the LBNL Lawrencium midrange computing cluster, which consists of approximately 900 nodes providing an aggregate of 12,000 cores and 40 TB of RAM.

5.12.2.4 Generalized Process of Science

Unlike most academic research laboratories, JBEI investigators are not assigned space. Rather, space is managed centrally and assigned dynamically as researchers leave and join. Researchers are assigned lab and desk space by the Laboratory Operations Director based on research needs. This creates a research environment where scientists are integrated daily by the very nature of their workspace assignments. Cross pollination of ideas and integration occur daily at the work site rather than primarily in social areas outside of workspaces. Social areas are equipped with white boards and interaction space to support serendipitous research discussions across the Center.

Analytical Laboratory. The JBEI Analytical Laboratory provides researchers with a powerful array of instruments to rapidly and accurately determine reaction outcomes with high sensitivity. JBEI uses the Analytical Lab to characterize feedstocks, engineered microbes, enzymes, and advanced biofuels. The Analytical Lab provides walk-up and expert use of GC-MS, proteomics, HPAEC, microscale thermophoresis, metabolomics LC-MS, capillary electrophoresis MS, and MS imaging instruments.

JBEI performs HT metabolite profiling and enzyme assays using a novel and unique integration of acoustic printing and Nanostructure-Initiator Mass Spectrometry (NIMS), performed on an Bruker ultrafleXtreme MALDI-TOF MS. Enzyme assays use a JBEI library of mass-tagged substrates targeting specific glycan-modifying reactions, and metabolite assays leverage JBEI's extensive metabolite standards library.

Processed data from JBEI mass spectrometers are added to the JBEI EDD to aid information transfer to individual researchers, data storage and communication with the biofuels community. The instrumentation is integrated with a computational system hosting Mascot, Mascot Distiller (de novo peptide sequencing), MassHunter Qualitative Analysis (Agilent), MassHunter Quantitative Analysis, Skyline Proteomics, iPython notebooks and Scaffold for data analysis. JBEI developed and implemented custom software, which aids experimental approaches using stable isotopes and tandem MS, to facilitate unknown metabolite identification. JBEI will use strong cation exchange (SCX) and high-pH reverse phase chromatography systems for separation of complex protein mixtures to aid global proteomic analyses as needed.

Biomass Pretreatment Laboratory. In the JBEI Biomass Pretreatment Laboratory, researchers test the effects of mechanical and chemical processing of biomass on sugar yield, as well as generate hydrolysates and lignolysates for biofuel and bioproduct conversion. Grass, wood, and other biomass samples are broken down by using automated mills, including a Labman robotic grinder for HT sample processing, and sieves are used to separate the particles by size. Uniformly-sized particle batches are subjected to treatment with ionic liquid (IL) to determine optimal solids loading, temperature, time, and other parameters. JBEI uses the pretreated biomass, in an integrated feedstocks-to-biofuels pipeline, to determine which set of process conditions most efficiently produces the sugars and lignin-derived intermediates amenable to transformation into advanced biofuels and products. This lab also houses three 1-L Globe reactors, three 1-L Parr reactors, moisture analyzer, Karl-Fischer titration system to determine water content in ILs, Retsch PM100 Ball Mill, two Thomas Wiley Model 4 Knife Mills, a SpexSamplePrep 6875 Freezer Mill; a Mettler Toledo TGA/DSC 3+ Thermal Analyzer , an Across International GFC1200 Muffle Furnace, a Tosoh HLC-8320 GPC and a Thermo Dionex ASE 350 Solvent Extractor.

Chemistry Laboratory. The JBEI Chemistry Laboratory provides advanced instrumentation for researchers to develop reactions, increase scale, study the effects of high temperature and pressure, and safely handle organic solvents and other hazardous materials. Chemistry Laboratory facilities include three chemical fume hoods (including one large walk-in hood). A 300 mL Parr reactor will be used for biomass conversion reactions at up to 350°C and up to 3000 psi; an oil-heated 200 mL Buchiglasuster Reactor Parr unit for reactions at up to 300°C and up to 2000 psi; and, an integrated six-reactor Parr system for investigating multiple reactions simultaneously at up to 300°C and up to 3000 psi. The Chemistry Laboratory also contains a large glass Soxhlet extraction system, a large-capacity still, vacuum and convection ovens, variable-temperature cooling baths, digital temperature-control and data-acquisition systems, specialty glassware for synthesis, vacuum/inert gas manifolds for performing air sensitive reactions, a rotary evaporator, and a high-capacity vacuum pump and filtration system. The Chemistry Laboratory is used to prepare bulk samples of pretreated biomass, to synthesize precursors and substrates for biochemical reactions, and to separate advanced biofuels and chemicals.

Fermentation Laboratory. The physiology of microbial strains engineered for production of biofuels or enzymes (such as cellulases) during fermentation is studied in the JBEI Fermentation Laboratory. The lab is equipped with: (a) a Kuhner RAMOS (Respiration Activity Monitoring System) shaker that is capable of simultaneously measuring the respiratory rates of up to eight microbial cultures; (b) one Sartorius Biostat B Twin reactor and two Sartorius BIOSTAT® B+ twin bioreactor systems, each capable of simultaneously running up to four fermentations (at 2-L or 10-L scale), and each outfitted with high- precision peristaltic pumps for controlling feed rates, and controllers for regulating gas mixes; (c) m2p-labs BioLector 1/ Robolector and one m2p-labs Biolector Pro microbioreactor system to perform HT fermentations; and (d) analytical instruments such as two Agilent 1100 HPLC units for measuring fermentation byproducts, and a process Mass Spectrometer for online gas composition analysis. These tools provide critical insights to growth conditions and process development for optimal microbial production of biofuels and bioproducts.

Microfluidics Laboratories. The JBEI Microfluidics Laboratories provides state-of-the-art microfluidic tools for HT synthetic biology, using ML to fully automate the design-build-test-learn (DBTL) cycle. Employing expertise from industry and academia to design, fabricate and prototype devices in-house, JBEI harnesses the capabilities developed by the Feedstocks, Deconstruction, BBD to accelerate project completion at an unprecedented rate. These micro electro-mechanical systems (MEMS) devices interface with emerging automated liquid-handling technologies that are capable of dispensing nanoliter volumes of reagents (e.g., the Labcyte Echo, and I-DOT from Cellink), drastically increasing the scale at which methods such as CRISPR, gene assembly, and protein pathway engineering can be performed. In addition, Sandia supports systems-integration activities for automated and integrated microfluidic platforms for HT screening applications.

Microscopy Laboratory. The JBEI Microscopy Laboratory provides researchers a suite of spectral microscopes to examine engineered feedstocks and study their deconstruction. A Zeiss LSM-710 confocal laser scanning microscope with Ar, He, diode, and tunable (488 – 640 nm range) lasers allows in situ study of biomass degradation. A Bruker Hyperion 3000 FT-IR microscope equipped with a focal plane array produces chemical images in the 800 – 4000 cm-1 range in minutes, while a Bruker FT-Raman spectrometer with attached FT-Raman microscope permits high spectral resolution in precisely-selected regions of a sample. JBEI uses these instruments and the addition a Thermo Fisher EVOS M7000 Imaging System, to study the microscopic structure and composition of engineered feedstocks, examine the effects of pretreatment or enzymatic hydrolysis (including in situ studies) on plant cells, and develop new analytical methods.

Plant Growth Facilities. JBEI's on-site plant-growth room and plant-growth chambers provide precisely controlled indoor greenhouse facilities for growing genetically- engineered feedstocks. Nineteen (19) independent chambers comprising a combined 1494 cubic feet allow plant growth under computer-controlled temperature, humidity, and daylight cycle. JBEI has developed a specialized chamber for growing plants in a 13CO2 environment, necessary for NMR analysis of biomass. A plant-growth room provides an additional 400 square feet of light- and temperature-controlled space. JBEI uses the plant-growth facilities to develop engineered sorghum and poplar plants with improved biofuel and bioproduct phenotypes.

Robotics Laboratory. JBEI's Robotics Laboratory provides state-of-the-art commercial automation equipment to enable HT research. For routine use, equipment is available to JBEI and ABF researchers on a walk-up or scheduled basis. For more complex applications, a staff laboratory automation engineer provides assistance with method development and troubleshooting. The Robotics equipment includes an in-house DNA construction service powered by multiple liquid handler units, programmable through the JBEI- developed custom PR-PR scripting language; these liquid handlers are: (1) Labcyte Echo 550 acoustic liquid dispensers for precise transfer of nanoliter droplets, (2) four Biomeks for quantifying feedstock recalcitrance and enzyme activity, automated proteomic/metabolomic sample preparation workflows, directed evolution of ionic liquid-tolerant cellulases, and measurement of advanced biofuels, and (3) Hamilton Nimbus for general purpose liquid handling. There is also other specialty equipment: (4) Molecular Devices QPix 460 for colony picking, (5) Formulatrix Mantis for bulk reagent dispensing, (6) Agilent ZAG DNA fragment analyzer, (7) Agilent microplate labeling platform, and (8) Labman solids handler for grinding and dispensing biomass and soil samples. Auxiliary equipment includes Veriti thermal cyclers, plate sealers, a laminar-flow hood, incubator shakers, microscope, centrifuges and semi-automated electronic pipettes for protein expression, cell culture and other applications.

Spectroscopy Laboratory. The JBEI Spectroscopy Laboratory provides multiple specialized instruments to detect and characterize biofuel feedstocks. Advanced instrumentation used in the Spectroscopy Lab includes an Agilent 1100 LC with a G7102A with a G7102A Evaporative Light Scattering Detector; a circular dichroism (CD) spectrometer (Jasco J-815) for biophysical characterization of proteins and enzymes; a Beckman Coulter DTX880 Multimode Detector, a Wyatt Technology DyaPro Plate reader and a Thermo Scientific Spectronic 200 to determine the identity and quantity of plant components; NIR and mid-infrared (MIR) spectrometer (Bruker MPA NIR); a Jobin Yvon FluoroLog 3 fluorometer; a HT dynamic light-scattering 96-well plate reader to test protein and lignin size; and a gel permeation chromatography system equipped with refractive index, ultraviolet (UV) and viscometric detector (EcoSEC HLC-8320GPC) for lignin molecular weight analysis. JBEI researchers use the Spectroscopy Laboratory to determine the composition of plant feedstocks and to understand the detailed effects of pretreatments on plant-cell components.

Synthetic Biology Laboratories. JBEI's Synthetic Biology Laboratories are large-scale, multidisciplinary spaces combining the five scientific divisions. Centralized around three rooms and comprising 21,000 square feet, the "SynBio" Labs house much of the specialized equipment and facilities required to advance JBEI's mission with specialized rooms for routine microscopy, plant cell culture, microfluidics, microarray, anaerobic growth, strain archive, seed storage, media preparation, glass wash and autoclave, and chemical storage, as well as three 4 °C cold rooms, and two each of 30 °C and 37 °C warm rooms. The labs contain all the equipment necessary for molecular and synthetic biology, such as laminar-flow hoods (7), -80 °C freezers (21), thermal cyclers, and incubator shakers. With the latest techniques and instruments, JBEI uses its Synthetic Biology Labs to engineer model plants for biofuels, deconstruct biomass, and apply synthetic biology to produce advanced biofuels.

NextGen Sequencing Laboratory. JBEI will use an Illumina MiSeq sequencer for next-generation sequencing research efforts such as RNA-seq analysis, genome re- sequencing applications, and general QA/QC for synthetic biology and metabolic engineering workflows. The system is connected to the JBEI data storage system and JBEI EDD for experiment metadata tracking, data sharing, processing, and storage. To enable data analysis, JBEI has various visualization tools such as Jupyter Notebooks, Arrowland, and components in the DOE KBase to aid data analysis.

X-ray Crystallography Facilities. Standard protein purification equipment (Akta) is available at JBEI for the preparation of protein samples. Two JBEI robotic crystal screening systems (Phoenix, from Art Robbins) are located close to the ALS, permitting screens for 800 different crystallization conditions to be set up in approximately four hours. Crystals are screened and diffraction data collected using the BCSB beamlines at the ALS, using either BCSB discretionary time or general user time obtained through the standard ALS peer-reviewed application process.

JBEI's warm and cold rooms, freezers, refrigerators, plant growth room, and growth chambers are constantly monitored and alarmed for temperature, power consumption and/or humidity by Klatu Systems Traxx Motes. Major equipment is covered for on-site preventative maintenance and repairs by manufacturer's service agreements.

JBEI's combined laboratories contain 15 chemical fume hoods (including one walk-in fume hood); ten LABCONCO Purifier Logic+ Class II Type A2 Biosafety Cabinets; two 37 °C and two 30 °C warm rooms totaling 500 square feet; and three 4°C cold rooms totaling 445 square feet; two industrial deionizing water systems; three Thermo Scientific Barnstead Nanopure water polishers; two Elga Purelab water polishers; and an Elga Chorus water polisher; twenty-one Thermo TSX Series -80 °C freezers; and house vacuum, purified nitrogen, and clean, dry air.

All major equipment is connected to the JBEI data storage system and integrated with the JBEI EDD and a computational system hosting Mascot, Mascot Distiller (de novo peptide sequencing), Skyline proteomics, Agilent MassHunter Qualitative, Agilent MassHunter Quantitative, iPython notebooks and Scaffold for data analysis.

5.12.2.5 Remote Science Activities

In addition to the facilities in JBEI's central location, approximately 10% of JBEI's researchers are housed in off-site locations. Below JBEI's off-site facilities are detailed.

5.12.2.5.1 University of Adelaide

Prof Jenny Mortimer's laboratory is co-located with seven other plant molecular biology research groups in the Plant Research Centre, part of the School of Agriculture, Food and Wine on the Waite Campus at the University of Adelaide. This provides access to a large amount of shared equipment, which has been particularly important in the new lab. The Mortimer Lab has a large PC2 laboratory space, which includes quarantine facilities, four laminar-flow cabinets for tissue culture, controlled condition growth cabinets equipped with near-sunlight LEDs, and a fume hood. In addition to the standard molecular biology equipment (PCR, Q-PCR, protein analysis, etc.), the Lab also have extensive glycan analytical capabilities and microscopy facilities. At the Waite campus, the Lab is fortunate to be co-located with national facilities including the NCRIS-funded Australian Plant Phenomics Facility, the Bioplatforms Australia Metabolomics platform at the Australian Wine Research Institute (AWRI), the South Australian Genomics Centre (SAGC) Waite node, and the Adelaide Glycomics facility.

5.12.2.5.2 Brookhaven National Laboratory (BNL)

Dr. Chang-Jun Liu has a laboratory of ~2000 sq. ft. located in the Biology Department, Brookhaven National Laboratory Campus. The Laboratory is fully equipped with instruments for molecular genetics, biochemistry, and structural biology studies. The space includes a main wet experiment room, two bio-analytic rooms (hosting a ThermoFisher Q-Exactive plus ultra-high performance liquid chromatography UPLC-MS system, an Agilent LC-MS system, a GC-FID, a GC-MS, and a FPLC system), a laminar-flow tissue culture room, a reagent preparation room, a dark room, an isotope radioactivity room, a microscopy-bioimaging room, and a temperature-controlled cold room. Dr. Liu's lab has four offices with a total of ~800 sq. ft. housing the PI and current research staff. Growth facilities include two temperature- and humidity-controlled walk-in growth chambers (50 sq. ft. each), two plant tissue culture chambers, and a newly built computer-controlled greenhouse (3000 sq. ft.). Field plot resources include an ~2-acre field located within the campus of BNL and used for agricultural experimentation.

5.12.2.5.3 Georgia Institute of Technology

Prof. Sankar Nair's group operates a total of approximately 5,000 sq. ft. of laboratory facilities in four buildings on the Georgia Tech campus. These laboratories are dedicated for synthesis and characterization of nanoporous materials and membranes and experimental evaluation of their applications in multicomponent liquid and gas-phase separations. Materials and membrane synthesis (including graphene oxide, inorganic zeolites, and metal-organic frameworks) is carried out in two 800 ft2 laboratory spaces which have been established for Nair's research in the School of Chemical & Biomolecular Engineering. These laboratories are located in the Environmental Science & Technology (ES&T) Building and the Carbon Neutral Energy Solutions (CNES) Building respectively, and are equipped with a total of 5 fume hoods, 10 workbenches, and utilities. Three 400 ft2 laboratories in the Renewable Bioproducts Institute (RBI) building contain comprehensive facilities for fabrication of membranes on flat, tubular or hollow fiber surfaces. The Nair group laboratory space has facilities that house comprehensive equipment for measurement of adsorption and membrane permeation properties and for demonstration of adsorption-based and membrane-based separation processes. These are described in the Equipment section. The RBI building houses a comprehensive facility for chemical analysis of mixtures and materials relevant to bioproducts research. The facility houses numerous instrumental analysis techniques including gas and liquid chromatography, MS, infrared spectroscopy, emission spectroscopy and capillary electrophoresis. Classical wet chemistry techniques including titrimetry, gravimetry and calorimetry are also used routinely. Georgia Tech has state-of-the-art central cost center facilities for materials characterization (electron and surface microscopy/microanalysis, NMR/electronic/vibrational spectroscopies, X-ray diffraction, and mass spectroscopy). The usage fees of cost center facilities is built into the material and substantial (M&S) budget. Machine and electronics shops are available in the School of Chemical & Biomolecular Engineering and in the Georgia Tech Research Institute (GTRI). Graduate students and postdocs in Nair's group are located in two office areas in the ES&T building, one office area in the CNES building, and one office area in the RBI building. There is ample space for the personnel associated with the proposed work.

5.12.2.5.4 Iowa State University

The biorenewables area continues to be critically important to Iowa State University and the State of Iowa, and has been identified as a top priority of the university. A robust research infrastructure also exists on the Iowa State campus, including a state-of-the- art, 33,000-square-foot, $32 million Biorenewables Research Laboratory (BRL). In this building, the Center for Biorenewable Chemicals (CBiRC) has administrative space, enabling equipment, and experimental facilities of over 9,000 square feet, including space for shared reactor, fermentation and analytical instrumentation. Dr. Shanks has a dedicated laboratory and student office space in the BRL. Dr. Shanks' laboratory space consists of 3200 square feet of space located in the Biorenewables Research Laboratory building. The Shanks group has graduate student and postdoctoral scholar desk space adjacent to the lab. The lab space itself is equipped with twelve flow hoods, various equipment and glassware required for catalyst synthesis (e.g., rotovap, balances, vacuum oven, synthesis ovens, nitrogen glove box, Schlenk line) and air, gas, water and electrical supply. The lab includes eight HT 75 ml Parr reactors (2 of the reactors being Hastelloy) and two 100-ml stainless steel Autoclave reactors with operating pressures up to 10 atm. There are also three stainless steel and two glass flow reactor systems. Sample analysis can be performed with a Waters Acquity UPLC with - photo diode array and evaporative light scattering detectors, Agilent GC/MS with FID, 2 GCs (one with an autosampler), microGC, and Metrohm 798 MPT Titrino automatic titrator. Material characterization equipment includes a Micromeritics ASAP 2020 automatic micropore & chemisorption analyzer, Micromeritics 2090 dynamic chemisorptions apparatus, Bruker Tensor 27 FTIR, Malvern Zetasizer, and Perkin Elmer TGA 7 thermogravimetric analyzer. User facilities are available at Iowa State University for X-ray diffraction, SEM/TEM, solution and solid-state NMR, and thermogravimetric analysis (TGA)/MS. These capabilities are critical for JBEI's chemistry-biology integration objectives in achieving a greater range of compounds directly from pretreated carbon streams as well from microbial and in planta derived compounds.

5.12.2.5.5 LLNL

All LLNL JBEI-paid researchers will be located at the JBEI central facility. Dr. Singh and his group have access to the following facilities at LLNL.

Advanced Manufacturing Laboratory: The Advanced Manufacturing Laboratory (AML) is a new 14,000-square-foot facility where scientists and engineers work side-by-side with partners to develop new materials and technologies. The AML houses some of the most sophisticated and capable equipment in the field of advanced/additive manufacturing, some of which are not yet commercially available. AML's facilities include equipment for direct ink writing, powder bed fusion, electrophoretic deposition, projection microstereolithography, and laser-based processes such as two-photon lithography and selective laser melting. Additional resources include material evaluation and characterization equipment and access to high-performance computing (HPC) modeling and simulation capabilities.

Center for Micro and Nanotechnology (CMNT): The CMNT facility houses a 6,400 ft2 class 100 cleanroom (class 10 capable) with an additional 1,000 ft2 Class 1000 space, and 6000 ft2 of other laboratories. Most equipment can accommodate 6" wafers, although some are limited to 4". Equipment available for lithography includes Picotrack spin/develop system for multiple wafers, several individual spinners for photoresist, polyimide, SU-8 etc.; electrostatic spray coater for coating of 3D patterned substrates; contact aligners (Karl Suss, EVG, OAI) capable of sub 2 μm resolution with backside alignment; maskless aligner (Heidelberg) capable of sub 2 μm resolution; and a Raith External Border Gateway Protocol (EBPG) electron beam writer. Equipment for physical vapor deposition includes electron beam evaporators for metal and dielectric deposition; electron beam evaporator dedicated for fabrication of Josephson Junctions; large sputter chamber for coating of multiple 6" wafers; confocal sputterer with true co-sputtering and up to 500oC heating capability; and thermal evaporator. For chemical vapor deposition, the laboratory includes a thermal oxidation furnace; low pressure chemical vapor deposition for silicon nitride; plasma enhanced chemical vapor deposition for SiO2, SiNx, poly-Si, and amorphous SiC; and atomic layer deposition system. For wet etching, the laboratory has Numerous wet chemical benches to support hydrogen fluoride, metal etch, alkali (eg KOH), sulfuric, piranha, and other chemistries as required. For dry etching, the laboratory has a Bosch Deep Reactive Ion Etcher for deep Si etching, capable of loading up to 3 x 4" wafers simultaneously; an Oxford Inductively Coupled Plasma (ICP) etcher for up to 4" III-V semiconductors including GaN and GaAs, with a heated stage capable of up to 400oC; a ULVAC ICP etcher for semiconductor, dielectric, and metal etching; a Nordson oxygen plasma asher; and XeF2 vapor etcher. For dicing, the laboratory has dicing saws, one with polygon cutting capability (up to 60 sides, near circular cutting). For thermal processing, the lab has a Rapid thermal annealer capable of up to 1350oC with Ar, N2, O2, H2, and NH3, gas mixing capability, vacuum capability; numerous ovens and furnaces for baking, curing, and dopant diffusion; and high temperature 1760oC furnace for dopant activation (eg: for SiC). For metrology, the lab has numerous contact and contactless tools including Dektak profilometer, Veeco white light interferometer, atomic force microscopy, Toho Nanospec, spectroscopic ellipsometer; numerous optical microscopes; Photoluminescence with mapping, cryogenic, and pulsed capability; and FEI Apreo scanning electron microscope with energy dispersive X-ray spectroscopy (EDS) capability. For bonding/packaging, the lab has wafer bonding capability; flip chip bonders; and wire bonders. For device testing, the lab has a full suite of electrical characterization equipment including probe station, semiconductor parameter analyzer, CV meter, Deep Level Transient Spectroscopy (77 to 700K), Hall Effect (50 to 400K), high-voltage power supplies.

5.12.2.5.6 PNNL

While ~90% of PNNL's JBEI researchers will be located at the JBEI central facility, JBEI will use the facilities located at the Fungal Biotechnology Center (FBC), operated at PNNL and funded by the DOE, that addresses barriers to fungal fermentation in the primary areas of morphology control, genomics, proteomics, fungal hyper-productivity, and feedstocks-to-products via fungal-based consolidated bioprocesses. The facility contains two biological safety cabinets centrifuges, incubators, cell disruptors, UV/Vis spectrophotometer, thermal cyclers, gel electrophoresis, image analyzers, microscopes (standard, dissecting and confocal) and other equipment necessary for shake flask cultures, biological sample processing, and genetic engineering of bacteria, yeast or filamentous fungi. A Thermo LTQ mass spectrometer and Agilent nanoflow HPLC system dedicated to global proteomic analysis are also available. JBEI's Fungal Biotechnology team will leverage the expertise and tools developed at the FBC to expedite the realization of hosts capable of producing high levels of recombinant cellulases and ligninases.

5.12.2.5.7 Northwestern University

The Broadbelt group has office space totaling 700 square feet, with 10 desks and 10 monitor stations. A 280-square-foot office is available for Prof. Broadbelt. The following programming languages and environments are installed on the computers: Matlab, C++, Perl, Python, in-house code for pathway generation and dynamic modeling (network generation, ODE creation and solution) and in-house code for stochastic modeling (structure generation and kinetic Monte Carlo). These resources will be used for establishing the Pickaxe tool for the JBEI retrobiosynthesis platform.

5.12.2.5.8 Sandia National Laboratories

While ~95% of Sandia's JBEI researchers are slated to work in the centralized JBEI facility in Emeryville, JBEI will use a number of specialized facilities and equipment available at Sandia (Livermore and Albuquerque) as described below.

Applied Biosciences Laboratories (ABL). ABL, located in Livermore, CA, comprises 15 individual laboratories for a total of 12,000 sq. ft. All research rooms (except for the Chip Fabrication Technologies Laboratories) are BSL–2 certified. The laboratories available for JBEI support include the following:

Microfluidic Devices Laboratories (~2100 sq ft). These laboratories are primarily used for development of microfluidics-based platforms. They are equipped with home-built devices for these purposes, as well as off-the-shelf equipment for construction and validation of the devices. Capabilities include: fabrication of stationary phases (bonded silica and porous polymer monoliths) in capillaries and on chip; optical detection of analytes in capillaries and on chip using UV absorbance or laser-induced fluorescence; fabrication of high-pressure electroosmotic pumps; and design and fabrication of compact high-voltage power supplies. JBEI has over 15 chip-stations for immunoassays, capillary electrophoresis, and electrochromatography using UV-detectors, microscopes and laser-induced fluorescence detectors.

Chip Fabrication Technologies Laboratories. A range of additional equipment is located in these labs including fabrication facilities (1000 sq. ft.) for fused silica, glass and plastic microchips. Polymer replication equipment includes a 60-ton vertical injection molding machine, three Carver presses for embossing and bonding of microfluidic devices, a three-bay nickel sulfamate electroplating machine, nanoimprint lithography tool, and a diverse range of polymer films and pelleted resin thermoplastics.

Rapid Prototyping and Engineering Capabilities. Sandia has a strong rapid prototyping capability with multiple 3D printers capable of processing polymers and metals, design software (CAD) for solid modeling including Pro Desktop, Pro Engineer and Solid Works, and structural and thermal analysis software. In addition, JBEI has multiple Computer Numerical Control Machines (CNC), three-, four-, and five-axis CNC machining centers, a lathe, high-speed machining and micro machining electrical discharge machining (EDM), jeweler's lathes, CNC lathes, milling machines and jig grinders. There are electronic prototype fabrication capabilities including chassis wiring/assembly, PCB loading & rework, ESD protection, PCB conformal coating, as well as multilayer (up to 12) printed circuit board fabrication in both solder mask and silk screen, quick turnaround print and etch and in-house cross section analysis.

Microsystems Science and Applications (MESA) Complex. Two unique facilities are co- located within the 65,000-square-foot MESAfab complex including: (1) a 34,500-square-foot Silicon Fab for 6-inch silicon wafer processing with Class 1 bays and (2) a 30,400-square-foot MicroFab with Class 10 and Class 100 bays for compound semiconductor material processing and silicon wafer post-processing. The Silicon Fab has processing expertise in both complementary metal–oxide–semiconductor (CMOS) and MEMS technologies. Packaging and assembly include interconnecting ICs and/or other components to a plastic or ceramic substrate, and subsequently integrating that package into a system.

5.12.2.5.9 TeselaGen Biotechnology, Inc.

TeselaGen biotechnology occupies 1200 square feet of office space in San Francisco, CA. The space is fully wired for Internet access and provides up to 15 software developers with a spacious, comfortable and healthy environment for advanced software development. TeselaGen uses Amazon Web Services for server hosting and ORACLE 12c as a relational database management system (RDBMS). TeselaGen exploits Node.js as a backend webserver, as well as a fully JavaScript enabled front end development stack that includes Bookshelf/Knex, Ext.js as well as custom designed tools and services.

5.12.2.5.10 University of California, Berkeley (UCB)

The Doudna Laboratory is located on the second and fifth floors of the Innovative Genomics Institute (IGI) at UC Berkeley. The lab is fully equipped for biochemistry, structural and molecular biology research, with instrumental setups for FPLC and HPLC chromatography, gel electrophoresis, and centrifugation. There are 12 full benches, 24 desks, and two chemical fume hoods on the fifth floor. There is also a tissue culture facility equipped with laminar-flow hoods on the second floor. The lab contains four Akta chromatography systems, in addition to an X-ray crystallization room and room dedicated to experiments with radioactive materials. The laboratory is equipped with tabletop centrifuges; ultracentrifuges for protein purification; a Typhoon phosphorimager; Tecan and Biotek fluorescence plate readers; gel electrophoresis apparatus; gel dryers; thermocyclers; a Zeiss AxioExplorer fluorescent microscope; Speed-Vac concentrators; bacterial shakers; crystallization incubators; and a Milli-Q plus water purification system. Lab personnel have access to graphics and computing workstations for diffraction data analysis and structure refinement. The Doudna research group has ten PC and Macintosh computers and workstations, five of which are dedicated to running laboratory equipment. Molecular graphics and computing workstations are located in the core computing facility adjacent to the main lab space. Networking and computing support including maintenance and troubleshooting of hardware, software, and stored data is provided by Campus Shared Services. Computational resources are also available on Savio, a high-performance computing cluster available to the Berkeley community. Dr. Doudna has an office on the fifth floor of the IGI in close proximity to her laboratory. The office suite includes offices housing a senior scientific collaborator and a project manager; a small reception area; and a large administrative workspace for two assistants, as well as a copier and a fax machine. The offices and workstations are equipped with standard furnishings and with five networked computers with high-speed Internet access.

The UC Berkeley DNA sequencing facility[77] performs Sanger sequencing, genotyping, and cell line authentication. The Berkeley MacroLab[78] is composed of four full-time staff members and equipment for automated gene cloning, protein expression, purification, and crystallography with a particular focus on multi-subunit protein complexes. Cell culture studies are aided by the Tissue Culture Facility, which performs all standard tissue culture techniques and provides cells as requested. Access is also provided to proteomics and MS facilities[79]. Three separate facilities are located on campus with expertise in protein identification and the detection and analysis of post-translational modifications. Services include analysis of complex samples by phosphoenrichment protocols, the Multidimensional Protein Identification Technology (MudPIT) technique, and quantification of relative protein levels in samples.

Fee-for-service facilities available to UC Berkeley PIs The Berkeley PIs have access to numerous fee-for-service facilities. The Tissue Culture and Media Core Facility in Barker Hall, the building adjacent to the lab, is equipped with CO2 incubators, cryogenic storage tanks, and a Coulter Z2 cell sorter, and provides cell cultures, media, and supplements as well as assistance and instruction. The Genomic Sequencing Laboratory has three Illumina 2G Genome Analyzers; a Proteomics Mass Spectrometry facility has a Thermo Finnigan LCQ and a Thermo Fisher LTQ ion trap mass spectrometer, as well as HP 6890 & HP 5973 GC/MS, Agilent 1100 LC/MS, and Applied Biosystems DEpro MALDI-TOF instruments. The MacroLab has a Biomek NXP, 2000 and 3000, Caliper LC90 microfluidic gel system, two ÄKTA Explorer 100 FPLC systems, HT FPLC ÄKTAxpress 4-pack, Mosquito crystal drop setting robot, and a Fluidigm Topaz microfluidics workstation for crystallizing proteins. The Biological Imaging Facility, located in Koshland Hall, offers all aspects of modern biological light microscopy, including confocal and deconvolution microscopy; computer image processing and analysis; and training.

5.12.2.5.11 University of California, Davis (UCD) Facilities

Ronald Laboratory personnel will be located at UCD. UCD will provide access to its experimental fields and state-of-the-art greenhouses. UCD also maintains the UC Agricultural Experiment Station (AES) as part of a state-funded UC research program dedicated to the testing and development of economically viable and environmentally sustainable agricultural production systems. DOE invested $3 million in ARRA funds to build new greenhouses and install additional growth chambers specifically for JBEI feedstocks development. A special dew chamber is available for studying pathogen response in sorghum.

77 https://ucberkeleydnasequencing.com/home
78 https://macrolab.qb3.berkeley.edu/
79 http://qb3.berkeley.edu/pmsl

Ronald's Laboratory at the UCD consists of a 2500 sq. ft. of newly renovated laboratory with sufficient bench space to accommodate 25 researchers. The lab has all equipment required to perform state-of-the-art experiments in cell and molecular biology. Equipment dedicated to the laboratory includes the full range of incubators, freezers, centrifuges (low-speed, and refrigerated superspeed- and tabletop ultra-centrifuge) and microcentrifuges, two flow hoods, assorted water baths, autoclave, scintillation counter, spectrophotometer, electroporator, water purification system, two PCR temperature cyclers, DNA/RNA and protein electrophoresis and blotting equipment, speed vacuum concentrator. A pathogen isolation laboratory is available for select agent work. Ronald is a member of the UCD Genome Center which has polysaccharide MS, HT genotyping, proteomic, and bioinformatic capabilities). Offices located adjacent to the laboratory are provided to the PI, postdocs and students.

5.12.2.5.12 University of California, San Diego (UCSD)

Laboratory. The Palsson and Feist Labs / Systems Biology Research group are fully equipped for both wet-lab and dry-lab research and have one large laboratory facility with four major areas. There are four rooms dedicated to wet-lab work – two large rooms (rooms 409, 412) with molecular biology bench space, major equipment (bioreactors, centrifuges, HPLC) and desk space for wet-lab researchers, and two small rooms for general cell culture (room 411) and BSL2 (room 410) work. Five rooms are dedicated to computational research (rooms 413, 414, 415, 417, 418) serving as offices for personnel and housing the computer workstations. Furthermore, JBEI is equipped with an Illumina MiSeq which will be used for DNA sequencing experiments proposed.

The wet-lab rooms house the automated Adaptive Laboratory Evolution (ALE) technology that JBEI has developed and that JBEI proposes to expand through a purchase using JBEI funds. In the current year, funds will be used to purchase an automated laboratory evolution machine that will be BSL2 certified and housed next to existing ALE equipment in the lab in the Bioengineering Department at UCSD. This machine was developed over the past five years at UCSD and the majority of cost (approximately 90%) is for the purchase of a Tecan Fluent automated liquid-handling platform with an Infinite M200 reader from Tecan, Inc. The other costs are for heating and agitating culturing blocks for multi-mL cultures, which is an essential part of the machine's functionality.

Computer. The Palsson computational group is equipped with 35 workstations, 14 shared computational servers dedicated to processor intensive calculations (both Windows and Linux-based), and three laptops. JBEI also has several workstations in the wet lab; some running lab equipment and some for routine use by personnel. JBEI has a dedicated server running a SimPheny database for metabolic reconstructions. Specialized software available for research use includes Matlab, LINDO, Simpheny, Mathematica, GAMS, and Tomlab.

Office. The Palsson group occupies a total of nine offices (rooms 413-415, 417-419, 426, 428-429). In addition, the PI occupies a conference room (416) for up to 14 people. JBEI also has nine more offices and floor space in the integrative Cal-IT2 (California Institute for Telecommunications and Information Technology) building for computational work. All offices are set up with wired and wireless networking and telephone. Administrative support and office space is provided by the department. All permanent staff on this project will have designated space.

5.12.2.5.13 University of California, Santa Barbara (UCSB)

O'Malley Laboratory personnel working on JBEI will be located at UC Santa Barbara. Dr. Michelle O'Malley's laboratory in UCSB's Department of Chemical Engineering is located in Elings Hall on the UCSB main campus and occupies approximately 2,300 square feet of dedicated wet-lab bench space and adjacent seating. The O'Malley lab has all the relevant equipment for microbial culture (i.e. flasks, shakers, and appropriate incubators), anaerobic microbial culture, spectrophotometers, TECAN fluorescent plate reader, autoclave, refrigerators, freezers, nucleic acid and protein electrophoresis equipment, centrifuges, thermocyclers, imaging equipment, microscopes (inverted light, and with fluorescence capabilities), and designated equipment/instruments for handling radioactive reagents. Each bench space is equipped with gas, air, and vacuum hookups to facilitate sterile microbial culture technique. The lab also contains a laminar-flow hood to control containment of hazardous and/or noxious materials. Importantly, the lab is outfitted with CO2 piping throughout to facilitate experiments and cell manipulation in the absence of oxygen at each bench, including the laminar-flow hood; an anaerobic chamber has also recently been installed in the lab. The lab has as an apparatus to allow batch-scale growth of anaerobic microbes in Hungate tubes and serum bottles using the Pressure Transducer Technique (PTT) and an inverted light microscope with attached camera to visualize anaerobic microbes in direct culture. A YSI analyzer and temperature-controlled HPLC (with autosampler and multi-angle light scattering detector) are available in the lab to monitor substrate (sugar) consumption and the production of liquid metabolites, and a gas chromatograph is available to evaluate fermentation gas composition. A QiaCube instrument in the lab facilitates automated nucleic acid and protein preparations as needed. In addition to yeast/bacterial culture, the lab is skilled in basic microbiology for anaerobic microbe extraction from fecal materials, including hand-casting of anaerobic roll tubes and anaerobic isolation of colonies, which does not require an anaerobic chamber.

Equipment and instrumentation in the neighboring Valentine lab will be available throughout this project, and include four Shimadzu gas chromatographs (models 14B and 8A, equipped with a flame ionization detectors or thermal conductivity detectors [TCDs)], Agilent 3000 micro TCD, Peak Labs gas chromatograph (PP1) equipped with a reduction gas analyzer (for H2 and CO), Hewlett Packard 5890 series II gas chromatograph equipped with a flame ionization detector, Monitor Instruments linear cycloid mass spectrometer, Seabird CTD/Rosette with 6 × 6L bottles, Sentinel Workhorse 300 kHz Acoustic Doppler Current Profiler (ADCP), oxygen respirometer for microbial cultures, oxygen optode for laboratory or field application, New Brunswick BioFlo fermentor with a 14 L vessel, a custom hydrogen stripping fermentor, hungate-style gassing stations for anaerobic microbiology, Coy labs anaerobic chamber, Autoclave Engineers impelled and jacketed pressure vessel for microbial cultivation at 100+ atm.

The PI will also have access to two Finnigan Delta+XP isotope ratio mass spectrometers (equipped with a GC combustion interface, an elemental analyzer interface, and a gas bench interface) through the Marine Science Institute Analytical Laboratory (awarded from NSF equipment grant DBI-0200461, J. Schimel - PI). This equipment is housed and maintained at the Marine Science Institute's Analytical Laboratory at UCSB, and is available to UCSB researchers on a recharge basis.

Several other shared facilities are located in the California NanoSystems Institute (CNSI) within Elings Hall and the Materials Research Laboratory (MRL), which are located in adjacent buildings to Engineering II. Additionally, the Neuroscience Research Institute (NRI) houses several confocal microscopes (including spinning disk and 2-photon), and TEM/AFM/SEM microscopes along with necessary support instrumentation. The CNSI and MRL have resources for confocal light and electron microscopy, nucleic acid analysis (Qbit, Tapestation, Bioanalyzer), next-generation sequencing (Ion Torrent, Illumina NextSeq, Illumina MiSeq), circular dichroism, and fluorimetry, as well as facilities for protein sequencing, mass spectroscopic identification of proteins/protein modifications, and microarrays; all available on a recharge basis. High-performance computing (HPC) is readily accessible through affiliation with the California NanoSystems Institute on the UCSB campus. The laboratory regularly uses these resources for de novo transcriptome assembly via TRINITY, among several other NGS analysis pipelines. Although UCSB houses several internal genomic resources (e.g., Ion Torrent, Illumina NextSeq) in core labs, it is sometimes necessary to use off-campus resources to obtain next-generation sequencing data in a timely fashion. For this purpose, the UC Davis Genome Center is available for use at internal UC recharge rates. Resources that may be relevant to this project include RNAseq (Illumina HiSeq 2500, Illumina MiSeq, PacBio RSII), with necessary library preparation, and optional bioinformatics support. The sequencing pipeline also includes thorough quality control checks for RNA Integrity Number (RIN) score and nucleic acid content.

5.12.2.6 Software Infrastructure

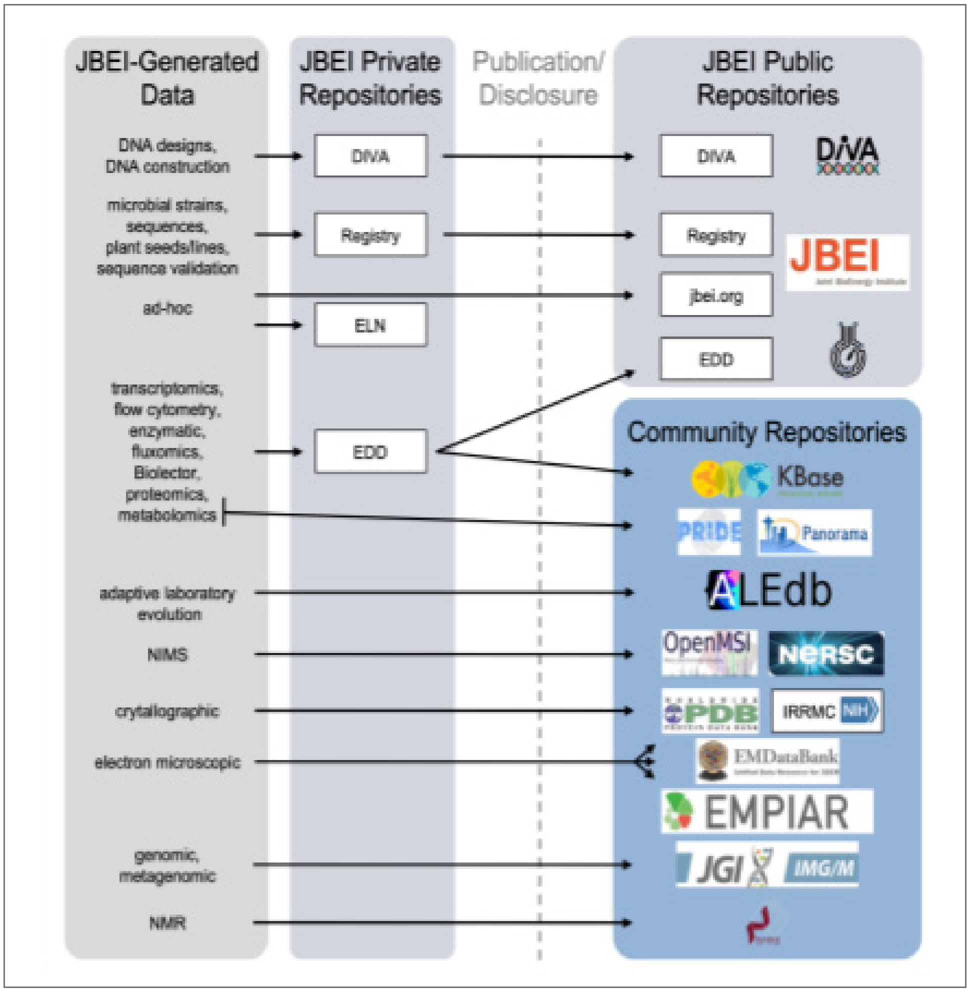

Figure 5.12.2: JBEI Data Types

JBEI researchers, by virtue of the scope of the project, will generate data of many different types, which need to be preserved and shared with the broader scientific community. JBEI will adhere to the guidance provided in the DOE Statement on Digital Data Management to maximize the impact of JBEI generated data. The goal is to make JBEI data available to the scientific community beyond that published in peer-reviewed journals. JBEI data will be preserved to the maximum degree possible through inclusion in public, long-lived databases. The appropriate metadata will be archived with datasets to make it possible for other researchers to validate the data and make use of it in their own work.

Where possible, data will be included in digital format as part of peer-reviewed publications. Larger datasets will be deposited with recognized community databases. Where practical the research community will be provided with other JBEI datasets through community resources such as the DOE Knowledge Base (KBase), or through public-facing sections of JBEI-maintained databases such as the Experimental Data Depot (EDD) and the ICE Registry. Other data will be made available publicly through the JBEI website as downloadable files. By going beyond solely providing data through publications JBEI aims to provide valuable datasets to the community, including negative results where feasible, and thus enabling alternative analyses of JBEI data, and knowledge discovery. A summary of the data types, their management and dissemination are shown in Figure 2.

JBEI data will have a broad impact in the broader research community and the public domain. Examples will include collections of genomic, transcriptomic, proteomic and metabolomic data for engineered plants and microbes, data from field trials, as well as biochemical and structures of lignocellulolytic enzymes and glycosyl transferases. Analysis of this data will inform the development of improved methods for biosystems design. Metabolic flux data, in particular that informed by 13C labeling experiments, will be used to test and develop improved metabolic models, and algorithms for metabolic modeling. JBEI strain and gene construct data in combination with characterization data will provide a community resource that describes parts for biosystems design and other bioengineering applications. JBEI structural biology data, provided through community databases, will enable the modeling of related structures for other biological applications.

5.12.2.6.1 Data Types and Sources

JBEI research will generate multiple kinds of experimental data in a digital format which can be used by other researchers to validate research findings and potentially further analyzed. Genomic data will be obtained through gene sequencing, and will be created through construct design. Multiple functional genomic data types (transcriptomics, proteomics, metabolomics, and fluxomics) will be generated in the characterization of engineered microbes, adapted microbes from ALE (Adapted Lab Evolution) experiments, and engineered plants. Enzymatic data will be generated in the analysis of wild type and engineering glycosyl hydrolases, lignases, and other biosynthetic enzymes. NMR studies of plant biomass will generate multidimensional spectroscopic data. Structural studies of these same proteins will generate crystallographic and cryo-electron microscopy data, and associated atomic models.

5.12.2.6.2 Content and Format

For DNA, RNA, and protein sequences, JBEI uses (and its software tools support) community data-exchange formats including FASTA, GenBank, and SBOL[80].

JBEI will continue to participate in community standards events and activities, and will move towards the adoption of new synthetic biology experiment data standards as they emerge.

For proteomic and metabolomic data, JBEI will use community standard data formats and processing methods for compound identification and validation. Raw data from targeted proteomic and metabolomic experiments will be processed via Skyline and distributed to the community through the Panorama web portal. Panorama captures both the Skyline data and methods necessary for subsequent validation experiments. NIMS data collected via MS imaging methods will be analyzed and shared through the OpenMSI imaging platform at NERSC. Flux data will be provided in Systems Biology Markup Language (SBML) format.

For structural biology data JBEI will make use of standard formats for internal workflows, including experimental data, molecular model, and geometric restraints formats. JBEI will use the widely accepted MTZ format for crystallographic diffraction data, the CCP4/MRC map format for 3-dimensional cryo-EM reconstructions, and the Common Intermediate Format (CIF) for geometric restraints. For deposition of experimental data to community resources JBEI will make use of the mmCIF standard. Paul Adams is the chair of the wwPDB mmCIF working group and thus well positioned to petition for extension data items to be added to the mmCIF dictionary to accommodate any new JBEI specific data.

5.12.2.7 Network and Data Architecture

LAN and WAN services for building 978 where JBEI is housed at LBNL are provided by LBNLnet. The LAN consists of two switches connected by a 10G link with 10G uplinks to the WAN. It has a network mask of 255.255.252.0, giving us approximately 1000 usable Internet protocol (IP) addresses. The primary WAN link is a 100G link through ESnet. JBEI has a secondary 10G link through CENIC.

80 http://sbolstandard.org/

5.12.2.8 Cloud Services

LBNL is a Google Workspace customer and has substantial discounts on services running in GCP. JBEI is currently working on the instrumentation required to:

1. Archive data to GCS.
2. Provide access to Google's Compute Engine for projects that exceed the capabilities of local resources.
3. Deploy webapps to Google's App Engine.

5.12.2.9 Data-Related Resource Constraints

We currently have 14 hardware systems dedicated to computational biology which are over five years old. These will soon need to be replaced. Given the increasing focus on ML, JBEI should replace these with GPU-enabled systems.

5.12.2.10 Outstanding Issues

We have established new inter-BRC Shared Research Objectives that will increase the need for efficient data transfer between all 4 BRCs over the next five years:

Inter-BRC	Bioenergy Biodesign (Theme 1)	Prototype development of strain libraries, HT assays and workflows for data generation and integration with AI–ML using bioenergy-relevant proteins	Use AI–ML to guide strain and enzyme development to show improved substrate conversion to product
Inter-BRC	Feedstocks-to-Fuels (Theme 7)	Complete at least one iteration of a multi-BRC feedstocks-to-fuels pipeline and establish baseline performance	Compare multi-BRC feedstocks-to-fuels pipelines for 3-4 feedstocks to assess performance
Inter-BRC	Plant Microbiome (Theme 3)	Compile BRC plant and soil microbiome resources in shared and discoverable formats	Identify core bacterial and fungal microbiomes across BRC bioenergy crops
Inter-BRC	Plant Syn Bio (Theme 2)	Initiate catalog of gene regulatory elements for inter-BRC sharing and develop testing and implementing technologies for plant synthetic biology	Pooling of BRC transformation tools across complementary bioenergy crop species and tools for multigene delivery
Inter-BRC	Bioenergy crops (Theme 4)	Initiate cross-use of phenotyping platforms for AI–ML and feedstock improvement (greenhouse- and field-based)	Advance cross-use of phenotyping platforms (both greenhouse- and field-based).
Inter-BRC	Data Sharing (Theme 6)	Develop FAIR inter-BRC data products portal approach and work with BRC researchers to curate data to populate the portal on bioenergy.org	Create inter-BRC imaging data products repository and connect it to the data portal on bioenergy.org
Inter-BRC	Lignin (Theme 8)	Maximize yields from lignin deconstruction and separation by optimizing RCF/OCF	Determine the effect of targeted bioenergy crops on potential for product recovery from lignin

Inter-BRC	Enabling Tech (Theme 10)	Enhance and standardize analytical capabilities for collecting data from fabricated ecosystems (i.e., EcoFABs) for the study of bioenergy crops	Generate ML-AI-enabled bioenergy research analytics for feedstocks, deconstruction, and microbial conversion
Inter-BRC	TEA-LCA-LU (Theme 9)	Use Feedstock-to-Fuel pipeline data as basis for biorefinery simulations/TEA	Harmonize life-cycle assessment modeling inputs and common environmental performance metrics for biofuel and bioproduct production routes across the BRCs
Inter-BRC	Training & Outreach	Accelerate training of a diverse bioenergy research workforce and increase public awareness of bioenergy	Accelerate training of a diverse bioenergy research workforce and increase public awareness of bioenergy

Table 5.12.2: inter-BRC Shared Research Objectives

We recommend that JBEI has a separate discussion about the potential use of ESnet to support these inter-BRC efforts.

5.12.2.11 Facility Profile Contributors

JBEI Representation

Jay Keasling, *LBNL,* JDKeasling@lbl.gov

Blake Simmons, *LBNL,* basimmons@lbl.gov

Hector Garcia-Martin, *LBNL,* hgmartin@lbl.gov

Henrik Scheller, *LBNL,* hscheller@lbl.gov

Corinne Scown, *LBNL,* cdscown@lbl.gov

Mark Kulawik, *LBNL,* mkulawik@lbl.gov

Arthur Panganiban, *LBNL,* ahpanganiban@lbl.gov

ESCC Representation

Rune Stromsness, *LBNL,* rstrom@lbl.gov

Richard Simon, *LBNL,* rsimon@lbl.gov

5.13 E3SM

The E3SM project is an ongoing, state-of-the-science Earth system modeling, simulation, and prediction project that optimizes the use of DOE laboratory resources to meet the science needs of the nation and the mission needs of DOE. In this context, "laboratory resources" include current and future personnel, programs, and facilities. Collectively, they represent a unique combination of proficiencies in science, engineering, leadership computing, and information technologies—the expertise required to construct, maintain, and advance an Earth system modeling capability needed by the country and DOE.

5.13.1 Discussion Summary

- E3SM is a multi-laboratory Earth system modeling, simulation, and prediction project that optimizes how existing DOE laboratory resources are used to develop a leading-edge climate and Earth system model. The science questions the project aims to address are centered around water cycle changes and impacts, the evolution of the human-Earth systems and polar processes, and sea-level rise and coastal impacts.
- Earth system modeling is data intensive. Data transfer can be needed at every phase of the model development cycle for E3SM, including:
  - Running the model.
  - Archiving and publishing the simulation.
  - Post-processing and analyzing the simulation.
- Running E3SM generates a large amount of data:
  - Some data generated during the model development phase is one-off and not saved.
  - For production-level runs, the project requires a standard workflow to evaluate, archive, and publish the data. Storage allocation is available at NERSC and LLNL.
  - Moving data is important for post-processing and analyzing the simulations. For simulations from a group of ensemble runs, or for very long simulations, generated data are sometimes distributed across multiple machines, while the results need to be concatenated or evaluated against each other.
  - All data transfer has been done using Globus and via ESnet.
- The E3SM project is a collaborative effort across eight national laboratories (including LLNL, ANL, LANL, PNNL, SNL, ORNL, LBNL, and BNL) and more than 10 academic institutions. Data transfer is essential to share data with collaborators and the broader community.
- Members have access to E3SM data from at least one E3SM-supported facility. E3SM projects also make production runs available publicly through the DOE's ESGF for external projects. ESGF publication currently goes through the LLNL-ESGF node, with plans to also use the ANL ESGF node and ORNL ESGF node.
- The following facilities play a more important role in model development and in conducting production simulations and evaluation:
  - Chrysilas and Anvil at ANL owned by E3SM.
  - Compy at PNNL (Shared by E3SM/RGMA/ESMD).
  - NERSC (Cori and Perlmutter).
  - ACME1: Archival system at LLNL with ESGF connection.

- ALFC: Theta, Cooley, Polaris, Aurora (first use expected in 2024).
  - OLCF: CADES, Summit, Frontier, Andes.
  - ANL: the Joint Laboratory for System Evaluation (JLSE).
  - LANL: Darwin, Badger, Chicoma/Rome, Chicoma/GPU.
  - LBNL: Lawrencium clusters, GCP via LBNL.
  - LLNL: Livermore Computing HPC systems.
  - PNNL: Research Computing.
  - Sandia National Laboratories (SNL): Center for Computing Research and HPC systems.
- The next project phase will focus on supporting model development and conducting simulation campaigns on GPU-based machines as well as maintaining its excellent performance on traditional central processing unit (CPU) architectures. In addition to the existing pre-exascale GPU systems installed at the leadership computing facilities (LCFs) and at NERSC, E3SM is targeting two exascale systems: OLCF's Frontier and ALCF's Aurora.
- Major simulations are conducted on supported production-level machines (Chrysalis and Anvil at ANL, and Perlmutter at NERSC, Summit at OLCF, Theta at ALCF, etc.).
- ACME1 has 10 PB designated to archive E3SM:
  - As of summer 2022, on ACME1, the E3SM volume reached more than 4 PB, comprising 2.9 PB of formal archive, and total ESGF publication volume of another 1.2 PB.
  - Estimated combined archive and publication growth is another 2.5–3.0 PB per year, having gained 1 PB in just the last five months.
- Every E3SM production simulation is required to be archived uniformly using Zstash, moved to NERSC HPSS for centralized deep storage, and then sent to the ACME1 server maintained at LLNL for data publication:
  - The massive data movements across platforms are done through Globus DTNs via ESnet.
  - The E3SM science campaign simulation results are being processed and published through the BER-supported ESGF, where all the datasets are cataloged and publicly discoverable through the ESGF search portal.
- Via the LBNL IT department, the E3SM project has set up a GCP cluster that is being used for running the E3SM nightly testing suite. Only test results (in the form of a single XML file) need to be exported from the GCP cluster.
- The massive data movements across platforms are done through Globus data transfer via ESnet. This software and network connection is the central piece that the project relies on.
- ESnet has adequately addressed the present need of data transfer in E3SM. A period over the last year when the Globus transfer performance from ANL to NERSC HPSS dropped significantly was noticed. The simulation team had to get a workaround to first transfer to NERSC's Community File System (CFS) then to HPSS. The problem was resolved after several weeks/months, which suggested diagnosing and fixing this type of issue currently is time consuming.
- Archiving data on the NERSC HPSS system plays an important role for data-intensive projects, such as E3SM. E3SM does realize there is a hardware bottleneck to put and retrieve data from tape systems. It would be more user friendly to make the connection stable and efficient.

5.13.2 E3SM Facility Profile

The E3SM project is a multi-laboratory project with the goal to develop a leading-edge climate and Earth system model designed to address US DOE mission needs and specifically targets DOE Leadership Computing Facility resources. The E3SM model development and simulation campaigns are motivated by addressing the most critical challenges facing our nation and DOE. The project is designed to accelerate the development and application of a fully coupled, state-of-the-science earth system model for scientific and energy mission applications. Scientific development of the system will be dictated by three science drivers that broadly cover the foundational science for advancing Earth system prediction. The science questions the project aims to address are centered around water cycle changes and impacts, the evolution of the human-Earth systems and polar processes, sea-level rise and coastal impacts. Using the model as a tool, simulations and projection experiments are conducted to address targeted overcharging science questions.

5.13.2.1 Science Background

Earth system modeling is data-intensive. Data transfer can be needed at every phase of the model development cycle, including running the model, archiving and publishing the simulation, post-processing and analyzing the simulation. These process are detailed as follows:

Particularly, running E3SM generates a large amount of data. Some data generated during the model development phase are one-off, they are being evaluated and archived at facilities most convenient to access to the project members. For production-level runs, the project requires a standard workflow to evaluate, archive and publish the data. More than 4 PB of official simulation has been archived with one copy at NERSC HPSS and another on an designated archival system maintained at LLNL. Moving data around is also important for post-processing and analyzing the simulations. For simulations from a group of ensemble runs, or for very long simulations, they are sometimes distributed across multiple machines, while the results need to be concatenated or evaluated against each other. In this case, either native simulation or the immediate files need to be moved to one location for future post-processing and analysis. All of the massive data transfer has been done using Globus and via ESnet.

When running the model, a set of input files including initial condition files, boundary condition datasets, mapping files, etc. needs to be staged on a machine. Current collection of this data for all possible configurations is ~ 5 TB. The project supports a centralized input data directory served via HTTP at ANL's Laboratory Computing Resource Center to provide both project and public access to these datasets. When configuring a run, the model case control infrastructure, CIME checks if all needed input files are present and downloads exactly what is needed for the simulations to be run. Typically all the users on a machine will use one shared input data directory, and all files will be downloaded to this directory. This has the advantage that each file is only downloaded once per machine even if needed by multiple users. If the file already exists on a machine, it will not be downloaded again. Larger size data transfer can be involved to perform some specialized simulation runs, for instance nudging simulations that require Interpolated analysis fields (including multiple three-dimensional variables in variable temporal resolution) from reanalysis datasets (e.g., ECMWF, NCEP, MERRA, and etc.) or fields from an existing model run to targeted model resolution. This input data needs to be generated and (or) transferred to a machine that performs the simulations.

5.13.2.2 Collaborators

E3SM project is a collaborative effort across eight national laboratories (including LLNL, ANL, LANL, PNNL, SNL, ORNL, LBNL and BNL) and more than ten academic institutions. The project involves around 50 FTEs spread over more than 80 scientists.

Data transfer is essential to share data with collaborators and the broader community. Members on the E3SM project and collaborators who have collaboration agreements have access to E3SM data from at least one of

E3SM supported facilities, as detailed in the table in this section and additional local resources for lab employees are available (see a list in session 3). E3SM projects also make production runs available publicly through DOE's ESGF for external projects, as well as the broader science community. For now all the ESGF publication goes through the LLNL-ESGF node. There is a plan to port the data publication workflow and execute it through the ANL-ESGF node and the ORNL-ESGF node for simulations generated at those locations.

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
Chrysilas and Anvil at ANL owned by E3SM	No, sometimes users backup simulation to its 4PB archival system	3 PB of disk spaced shared between Chrysalis and Anvil Data transfer	e.g., 80 GB per simulation year. 12 TB per 165 historical low-resolution (1 degree) run.[81]	Ad hoc, once a production run is completed	No, data send to NERSC HPSS and LLNL for two copies of archive	Data transfer to HPSS is in general slow.
Compy at PNNL Shared by E3SM/RGMA/ ESMD	No	1 PB of E3SM designated disk And 0.75 PB shared disk space data transfer	Same as above	Ad hoc, once a production run is completed	Same as above	Same as above
NERSC Cori (to-be-retired) Perlmutter	Store a primary copy of E3SM production run on HPSS	HPSS Data portal Data transfer Disk E3SM quota: 1.8 PB of disk, 9 PB HPSS storage	Same as above	Daily	No	Data transfer from HPSS to different end points are also slow.
ACME1: Archival system at LLNL with ESGF connection	Store a second copy of E3SM production run on spinning disk	9 PB of disk	Same as above	Daily	No	Data transfer from NERSC HPSS is slow.
ALFC: Theta, Cooley, Polaris, Aurora [expected in 2023]	Store a second copy of E3SM production run on HPSS	HPSS Data portal Data transfer	Same as above	Ad hoc	No	Data transfer from HPSS to HPSS end points are slow.
OLCF: CADES, Summit, Frontier, Andes	Store a second copy of E3SM	HPSS Data portal Data transfer	Same as above	Ad hoc	No	N/A

Table 5.12.1: JBEI Collaboration Space

81 Based on the configuration of the simulations, each present production run could range from 2 TB to 80 TB

5.13.2.3 Use of Instruments and Facilities

The facilities play a more important role in model development and in conducting production simulations and evaluation are listed in the table in session 2.

Other than the facilities provided in the table in session 2, additional institutional facilities are available to E3SM staff and collaborators, listed as follows:

- ANL: The Joint Laboratory for System Evaluation (JLSE)
- LANL: Darwin, Badger, Chicoma/Rome, Chicoma/GPU
- LBNL: Lawrencium clusters, GCP via LBNL
- LLNL: Livermore Computing HPC systems
- PNNL: Research Computing
- SNL: Center for Computing Research and HPC systems

The next project phase will focus on supporting model development and conducting simulation campaigns on GPU-based machines as well as maintaining its excellent performance on traditional CPU architectures. In addition to the existing pre-exascale GPU systems already installed at the LCFs and at NERSC, E3SM is targeting two exascale systems: OLCF's Frontier and ALCF's Aurora.

5.13.2.4 Process of Science

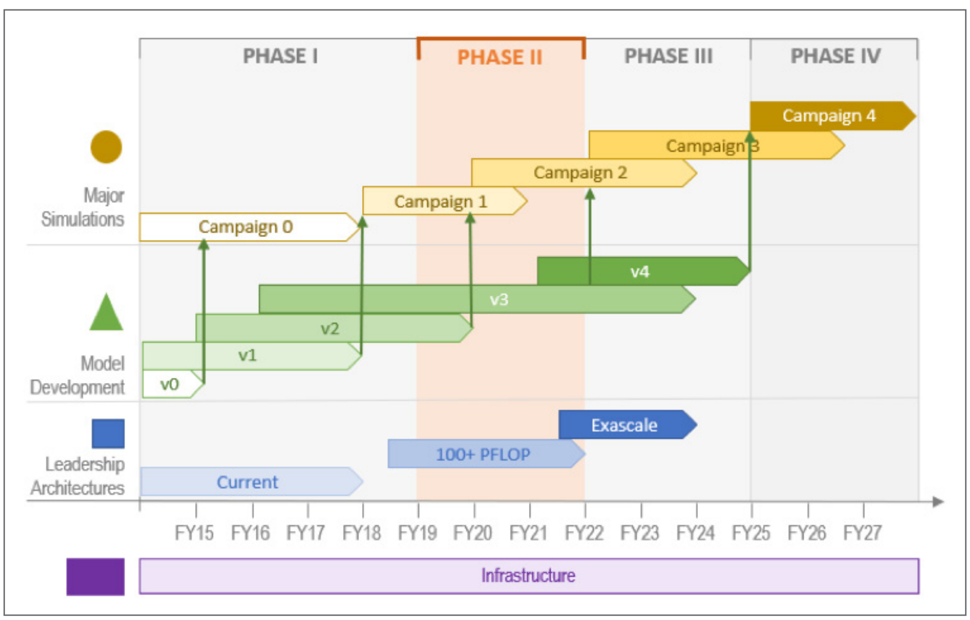

Figure 5.13.1: E3SM Project Life Cycle

Data Production and Management in E3SM:

Major simulations are conducted on supported production-level machines (Chrysalis and Anvil at ANL, and Perlmutter at NERSC, Summit at OLCF, Theta at ALCF, etc.). ACME1 is a prototype machine for data analysis and publication. Every E3SM production simulations are being transferred from where the data is generated to NERSC HPSS for deep storage and then to ACME1 for another archive copy on spinning disk and for downstream data publication and distribution.

ACME1 has 10 PB designated to archive E3SM on spinning disk and a data post-processing and publication workflow is set up on this machine to directly extract data from the archive, processing and publishing the data to the ESGF node at LLNL. As of August 2022, on ACME1, the E3SM volume reached more than 4 PB, comprising 2.9 PB of formal archive, and total ESGF publication volume of another 1.2 PB (1.2 PB in E3SM project space and 18 TB to CMIP). The new E3SM v2 simulations, the initial ~ 0.5 PB from Water Cycle simulation has been published (~ 0.3 PB) in its native format to ESGF. CMIP6 submission will follow. Estimated combined archive and publication growth is another 2.5-3.0 PB/year, having gained 1 PB in just the last five months. Other than the ESGF archive, this simulation campaign data is archived and publicly available at NERSC HPSS.

E3SM deploys a unified Python analysis environment on all the supported platforms. However, simulations were distributed over different platforms depending on the architecture of the machines required by the simulation configuration, as well as available allocation and demand of the machines. There were cases that one simulation or portions of a group of large ensemble runs was distributed among different machines. Therefore, data movement is critical for analysis. In addition, it is often needed to cross-compare candidate models with existing CMIP models, in this case the analysis can be done with data being transferred to ACME1, where holds the complete archive of ESGF's CMIP data archive mounted to allow comparing E3SM simulations to peer climate models. There is a CMIP data archive on NERSC, which is managed by RGMA FSA.

Data availability for science community data users: Every E3SM production simulation is required to be archived uniformly using zstash, moved to NERSC HPSS for centralized deep storage and to ACME1 server maintained at LLNL for data publication. The massive data movements across platforms are done through Globus DTNs via ESnet. The E3SM science campaign simulation results are being processed and published through the BER-supported ESGF, where all the datasets are cataloged and publicly discoverable through the ESGF search portal. Nearly all of the native model output from the production experiments are directly available from the E3SM project space under ESGF. The experiments that were conducted following CMIP protocols are processed and submitted to the CMIP project under ESGF.

5.13.2.5 Remote Science Activities

No additional information to report for this section.

5.13.2.6 Software Infrastructure

The massive data movements across platforms are done through Globus data transfer via ESnet This software and network connection is the central piece that the project relies on. E3SM also develops and maintains tools for data post-processing and management. In most instances, raw model output needs to be remapped into standard grids for generating final data products for analysis and for distributing to the science community (e.g., participating in CMIP activities). For instance, NCO (netCDF [network Common Data Form] Operators) are command-line operators to manipulate and analyze output from E3SM component model output (EAM, ELM, MPAS-SI, MPAS-O). The most common uses of NCO are generating climatologies, performing regridding, and time-series extraction. E3SM_to_CMIP[82] is the tool to rewrite E3SM output into the format that is compliant to CMIP standard. A workflow tool DataStateMachine[83] is developed to operationally stage raw model output, rewrite, and publish the data to ESGF.

zstash[84] is an HPSS long-term archiving solution for E3SM. Standard tar files (~500 GB each) are generated locally before transfer model output to HPSS. Checksums (md5) are computed on-the-fly during archiving and verified on extraction. Metadata is stored in a sqlite3 index database, enabling faster retrieval for target files.

82 https://github.com/E3SM-Project/e3sm_to_cmip
83 https://github.com/E3SM-Project/datasm
84 https://github.com/E3SM-Project/zstash

Zstash has parallel extraction and verification for increased performance. Most E3SM production simulations have been archived on NERSC's HPSS using zstash. The ability for initiating Globus data transfer was recently enabled in zstash to streamline the process of archiving long production simulations run on machines without HPSS. Every E3SM production simulation is required to be archived uniformly using zstash, moved to NERSC HPSS for centralized deep storage and to ACME1 server maintained at LLNL for data publication.

5.13.2.7 Additional Network and Data Architecture Requirements

There are no additional data architecture requirements to relay at this time.

5.13.2.8 Use of Cloud Services

Via LBNL IT department the E3SM project has set up a GCP cluster that is being used for running E3SM nightly testing suite. Other cloud-based functions the project uses day-to-day includes instances such as Atlassian tools: Confluence and Jira hosted at Atlassian.net for documentation and task tracking. Nightly test results are displayed on cdash.org. The code repository is managed through github.com. E3SM uses a web-based video/audio teleconferencing service from GoTo Meeting.

5.13.2.9 Data-Related Resource Constraints

Archiving data on HPSS systems plays an important role in a data-intensive project as E3SM. E3SM does realize there is a hardware bottleneck to put and retrieve data from tape systems. It would be more user-friendly to make the connection stable and efficient.

5.13.2.10 Outstanding Issues

ESnet has adequately addressed the present need of data transfer in E3SM. It was noticed there was a period of time over last year when the Globus transfer performance from ANL to NERSC HPSS dropped significantly. The simulation team had to get a workaround to first transfer to NERSC cfs file system then to HPSS. The problem was resolved after several weeks/months, which suggested diagnosing and fixing this type of issue currently is time consuming. Maybe more efficient diagnostic tools/approaches can be helpful to timely addressing problems of this kind.

5.13.2.11 Facility Profile Contributors

E3SM Representation

Robert Jacob, ANL, jacob@anl.gov

Jill Zhang, LLNL, zhang40@llnl.gov

Dave Bader, LLNL, bader2@llnl.gov

ESCC Representation

Linda Winkler, ANL, winkler@mcs.anl.gov

Corey Hall, ANL chall@anl.gov

Gianni Carlo Contardo, LLNL, gianni@llnl.gov

5.14 AmeriFlux Network

AmeriFlux is a network of PI-managed sites measuring ecosystem CO2, water, and energy fluxes in North, Central, and South America. It was established to connect research on field sites representing major climate and ecological biomes, including tundra, grasslands, savanna, crops, conifer, deciduous, and tropical forests. As a grassroots, investigator-driven network, the AmeriFlux community has tailored instrumentation to suit each unique ecosystem. This "coalition of the willing" is diverse in its interests, use of technologies, and collaborative approaches. As a result, the AmeriFlux network continually pioneers new ground.

5.14.1 Discussion Summary

- AmeriFlux is a network of more than 500 PI-managed sites measuring ecosystem CO2, water, and energy fluxes in North, Central, and South America:
  - The project was established to connect research on field sites representing major climate and ecological biomes, including tundra, grasslands, savanna, crops, and conifer, deciduous, and tropical forests.
  - Data produced by the network provides information and new knowledge central to the BER mission to understand complex biological and environmental systems, across spatial and temporal scales, by joining theory, observations, experiments, and models.
- AmeriFlux Core Sites, and the National Ecological Observation Network (NEON), typically upload datasets less than 100 GB in size for yearly data. Noncore sites have similar data volumes, but the cadence of upload (as well as mechanism) varies depending on access to data.
- AmeriFlux data products are publicly and freely accessible:
  - Most data users are based at universities, research institutions, or government agencies worldwide.
  - Over the last seven years, about 5,000 users globally have downloaded AmeriFlux data, totaling around 28,000 unique downloads.
  - Users can typically download data less than 100 GB in size, and in some cases may be interested only in a very small subset of data.
- Data transfers from measurement facilities to data repositories involve a preprocessing step to convert the data to a required format, and then upload it to a data processing facility via a web portal for QA/QC checks. Some measurement facilities transfer raw high-frequency data to NERSC through secure copy protocol (SCP).
- AmeriFlux data's life cycle involves:
  - Data collection at each site, using a full suite of instruments.
  - Processing of high-frequency observations into half-hourly time series.
  - Preparation of the data in a standardized format.
  - Submission to the AmeriFlux website: a web data portal and a semi-automated data pipeline to accommodate data submissions, quality-checking, tracking, and releasing.
- AmeriFlux Core Sites are flux towers whose managing PIs and staff have agreed to deliver timely, high-quality, continuous data to the AmeriFlux database. These 49 sites represent nearly every major bioregion and vegetation type in the contiguous USA, span 15 states and multiple climate regions, and comprise contrasts of demography and disturbance.

- AmeriFlux is considering allowing a subset of sites to transfer (near real-time) raw high-frequency data to AMP. There is demand among data users for this transfer, as well as taking advantage of 5G capabilities for near real-time data transfer. AMP would process the high-frequency data to half-hourly data products using standardized protocols and codes. This will drastically increase the network demands, e.g., 1.5 to 5 GB per month per site, and also the computation demand, e.g., ~0.5 to 3 hours processing time for one site-month of data.
- AmeriFlux sites are at a mix of locations, some more urban with ample infrastructure and others in more remote areas where power is derived from solar panels and batteries. Similarly, networking ranges from those with wired connections to those where available options are cellular/satellite connections. Ideally, all data can be sent from the field sites to teams' institutions, where various data checks and processing occur. Practically, measurements are often recorded to a data logger or computer on-site. Then the site teams retrieve the data via physical visits or remote connection (e.g., cellular modem, radio transfer, Ethernet). Many site teams manage multiple field sites and, in rare cases, up to 40 to 50 field sites, e.g., the National Ecology Observatory Network (NEON). Usually, data transfer from field sites to designated institutions poses the most significant challenge in the data's life cycle.
- The AmeriFlux data is stored in Berkeley Lab servers and in NERSC's CFS. An additional copy is stored in the high-performance storage system (HPSS) tape data storage, using metadata files that incorporate effective versioning, file naming, and backup strategies.
- In the past five years, the project has generated ~25 TB of data (primarily high frequency). In the next five years, AmeriFlux expects the project to generate on the order of 50–100 TB of data given the current workflow and project needs.
- Due to cost constraints, the use of cloud resources is somewhat limited. AmeriFlux has test instances on AWS but sees limited usage at this point as the costs are prohibitive for an instance that runs 24/7 instead of owning machines directly. This does impact redundancy to a certain extent.
- As data is processed to a half-hour resolution at each site team's institution before submitting to AmeriFlux, the data transfer is less constrained by the network, except for some sites that choose to transfer raw high-frequency data. If future scientific needs change and require more sites to transfer high-frequency data to AMP, network infrastructure can be an issue.
- Currently, there are no major outstanding data-related resource constraints. However, AmeriFlux anticipates that as real-time data needs come to the forefront, both networking and storage space demands will scale up significantly. There is a discussion within the community that AmeriFlux may slowly shift towards more real-time data processing and usage, where APIs will see much heavier usage in the next five years. In addition, tools like real-time data exploration with visualization will drive up both networking and data needs.
- It is recommended that AmeriFlux and ESnet discuss some of the lessons learned from remote sensor networks that rely on sensor and satellite connectivity.

5.14.2 AmeriFlux Network Facility Profile

The AmeriFlux network is a community of sites and scientists measuring terrestrial carbon, water, and energy fluxes between ecosystems and the atmosphere and a suite of ancillary meteorological and soil data. AmeriFlux has more than 500 sites across the Americas, each operated by individual PIs and teams who produce and share high-quality eddy covariance data. These sites represent various climate and ecological conditions, natural disturbances, and land management. Data produced by the network provides information and new knowledge central to the DOE BER mission to understand complex biological and environmental systems, across spatial and temporal scales, by joining theory, observations, experiments, and models.

The network was launched in 1996, after an international workshop on flux measurements in La Thuile, Italy, in 1995, where some of the first year-long flux measurements were presented. Early support for the network came from many sources, including the US DOE's Terrestrial Carbon Program, the DOE's National Institute of Global Environmental Change (NIGEC), NASA, NOAA, and the US Forest Service. The network grew from about 15 sites in 1997 to more than 110 active sites registered today. Sixty-one other sites, now inactive, have flux data stored in the network's database. In 2012, the US DOE established the AmeriFlux Management Project (AMP) at LBNL to support the broad AmeriFlux community and the AmeriFlux sites.

5.14.2.1 Science Background

In 2012 DOE established the AmeriFlux Management Project (AMP) at LBNL to support the broad AmeriFlux community and sites. AMP team members are mostly based in the Climate and Ecosystem Sciences Division (CESD) and the Computational Research Division (CRD) at LBNL but also include collaborators from multiple institutions. AMP works with hundreds of AmeriFlux scientists to ensure the quality and availability of the continuous, long-term ecosystem measurements necessary to understand these ecosystems and to build effective models and multisite syntheses. AMP's primary goals are to (1) support a set of long-term Core sites to provide high-quality, long-term data (details below); (2) enhance the quality of AmeriFlux data and its usability and discoverability by a broad community (details below); (3) expand the network's impact for basic research and earth system model improvement; (4) foster innovation and discovery.

AmeriFlux Core Sites are flux towers whose managing PIs and staff have agreed to deliver timely, high-quality, continuous data to the AmeriFlux database. They receive direct operation and instrumentation support from the AMP, unlike other AmeriFlux sites. These Core sites represent nearly every major bioregion and vegetation type in the contiguous USA, span 15 states and multiple climate regions, and comprise contrasts of demography and disturbance. Currently, 49 AMP-supported Core sites exist, of which 33 and 15 sites have been operating for more than 10 and 20 years, respectively.

AmeriFlux data's life cycle begins with data collection at each site, using a full suite of instruments at a high frequency, operated by individual site teams. The site teams collect and process the high-frequency observations into half-hourly time series, prepare the data in a standardized format, and submit it to the AmeriFlux website[85], maintained by AMP. Over the years, the AMP team built a web data portal and a semi-automated data pipeline to accommodate data submissions, quality-checking, tracking, and releasing. Upon submission, the processing pipeline is initiated and performs a series of quality assessment and quality control (QA/QC) checks. If the submitted data pass the QA/QC checks, the resulting BASE data product is published, i.e., made publicly available on the AmeriFlux website. The BASE data product is a format-standardized and quality-controlled data product.

The AMP team further uses the ONEFlux (Open Network-Enabled Flux) codes to process the BASE into the FLUXNET data product, i.e., fully processed and gap-filled data products with additional value-added variables for scientific research. The ONEFlux processing codes are jointly developed by a global collaboration that includes AMP. Both BASE and FLUXNET data products are published in a standard format, bundled with sites' metadata, versioned, and assigned with a DOI. The data users can register an account, search, and download data through the AmeriFlux website. Over the last seven years, about 5,000 users globally have downloaded the AmeriFlux data, totaling around 28,000 unique downloads. The data were widely used in a wide range of uses, e.g., multisite synthesis, benchmarking with RS and Land surface models, and education.

85 https://ameriflux.lbl.gov/

5.14.2.2 Collaborators

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
D. Baldocchi, University of California Berkeley	No	N/A				
D. Papale, University of Tuscia	No	N/A				
M. Humphrey, University of Virginia	No	N/A				
C. Hanson, Oregon State University	No	N/A				
Core site institutions, multiple locations across the US (~49 sites)	Yes (subset)	data portal	< 20 MB per site per year for half-hourly data 20-60 GB per site per year for archiving raw data for a subset of sites	Quarterly	No, the processed data is made publicly available on the website	
National Ecological Observation Network (NEON) across the US (48 sites)	Yes (subset)	data portal	similar as above	Quarterly	similar as above	
Other AmeriFlux sites, various locations across Americas (varied, ~400 sites)	Yes (subset)	data portal	similar as above	Varied, from yearly to one-time submission	similar as above	
Data users worldwide (varied, ~5000 users)	Yes, secondary for individual use	data portal	Varied, 25-90 GB, or smaller if a subset	One-time download, or multiple times if needed	No	

Table 5.12.1: JBEI Collaboration Space

Data user: AmeriFlux data products are publicly and freely accessible. The data users can register an account and download data themselves through the AmeriFlux website. Most data users are based at universities, research institutions, or government agencies worldwide. Although there is no direct and official collaboration between AMP and most data users, one of the AMP's major goals is to provide continuous and quality-assured data to all sorts of users, research, and applications. And, AMP developed many web-based user-interface features for facilitating data use.

5.14.2.3 Use of Instruments and Facilities

AmeriFlux data's life cycle begins with data collection at each field site, using a full suite of automatic instruments. AmeriFlux has more than 500 sites (i.e., both active and historical) across the Americas, each operated by individual PIs and teams. These sites are located in various natural or managed ecosystems, e.g., forests, savannas, grasslands, croplands, shrublands, wetlands, and lakes. The types of instruments may vary from site to site but typically include the following:

- Tower structure for instrument mounting: Types might vary from scaffolding walk-up tower, triangular tower, to tripod.
- Data acquisition system: data logger, computer
- Flux measurement: sonic anemometer for three-dimensional wind velocities and sonic temperature, infrared gas analyzer for water vapor and trace gas concentrations (e.g., CO2, CH4).
- Meteorological and soil measurement: Typical sensors used at each site include, but are not limited to wind speed, wind direction, solar radiation of various types, air temperature, water temperature, soil temperature, relative humidity, precipitation, snow depth, barometric pressure, soil moisture, soil electrical conductivity, and soil heat flux.

The data acquisition systems record the data streams continuously (at 10-20 Hertz for flux measurements and 1-0.1 Hertz for others). Then, the site teams retrieve the data, process them into a half-hour time series, prepare them in a standardized format, and submit them to the AmeriFlux website. Data files are ASCII text files using a CSV (comma-separated value) format, typically < 20 MB for a year of data at a site. The data upload frequency varies among sites, from quarterly updates (i.e., add three months of data every calendar quarter) to one-time submission (i.e., submit all data (e.g., 3-4 years) of a project at once). A subset of sites also submits their raw high-frequency data for backup purposes, typically 20-60 GB for a year of data at a site.

The AMP team has deployed the following cyber-infrastructure components to manage data and information.

- **Ingest:** Data are submitted to AmeriFlux via the AmeriFlux website. Larger files (such as high-frequency data) are submitted via secure File Transfer Protocol (FTP) directly to storage at NERSC. All data uploads are logged along with provenance information.
- **QA/QC and processing:** The AMP team has developed a code suite that checks the quality of data sent to AMP. Half-hourly/hourly data files are checked for issues such as time shifts, thresholds, consistency checks, etc., and flagged issues are communicated with the site teams using Jira with plots and summaries. This is an iterative process and may require multiple back-and-forths between the AMP team and the site teams. All communications are tracked, and the data provenance is maintained. Once the requirements are met, data is released for public consumption, enhanced processing is further done, and then released as other value-added data products. Besides the more regular half-hourly/hourly measurements, metadata and also other kinds of time-series data with irregular frequencies also go through the same QA/QC and processing process before being released to the public.
- **Data backup:** The AmeriFlux data is stored in Berkeley Lab servers and in NERSC's CFS. An additional copy is stored in the HPSS tape data storage, using metadata files that incorporate effective versioning, file naming, and backup strategies.

In the past five years, the project has generated ~25 TB of data (primarily high-frequency). In the next five years, AmeriFlux expects the project to generate on the order of 50–100 TB of data given the current workflow and project needs.

In terms of near-term plans (current technology horizon), there are discussions about allowing a subset of sites to transfer (near real-time) raw high-frequency data to AMP. There is demand among data users for this transfer, as well as taking advantage of 5G capabilities for near-real-time data transfer. AMP would process the high-frequency data to half-hourly data products using standardized protocols and codes. This will drastically increase the network demands, e.g., 1.5-5 GB per month per site, and also the computation demand, e.g., ~0.5-3 hours processing time for one site-month of data, using a commercial PC.

5.14.2.4 Process of Science

AMP team built a web portal and a semi-automated data pipeline to accommodate data submissions, quality-checking, tracking, and releasing. Upon submission, semi-automated BASE processing pipeline is initiated and performs quality assessment and quality control (QA/QC) checks. If the submitted data pass the QA/QC checks, the resulting BASE data product is published, i.e., made publicly available on the AmeriFlux website. The BASE data product is a standard, quality-controlled, half-hourly/hourly data product.

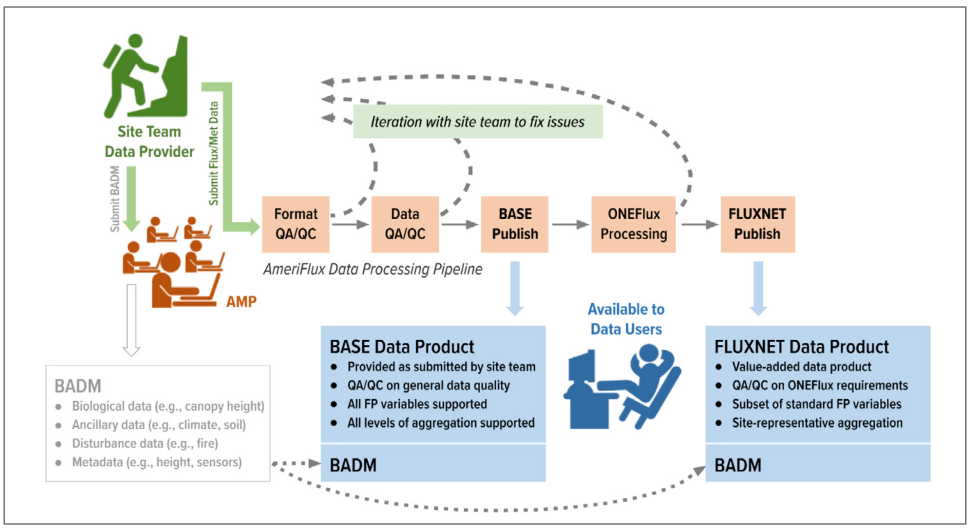

Figure 5.14.1: An overview of the data processing pipeline. Data submitted to AMP are processed using the pipeline to generate BASE and FLUXNET Data Products.

The AMP team further uses the ONEFlux (Open Network-Enabled Flux) codes to process the BASE into the FLUXNET data product, i.e., fully processed and gap-filled data products with additional value-added variables for scientific research. The ONEFlux processing codes are jointly developed by a global collaboration that includes AMP. The computational costs for ONEFlux processing are highly dependent on the number of variables and the number of data years. More complex processing runs on HPC resources for sites with more data years, and variables have ranged between one and a half days (even up to six days in some instances), while the shorter sites with a year's worth of data may complete in under an hour. The timings listed here are raw computational times and do not account for the time while waiting in the batch queue. Both BASE and FLUXNET data products are published in a standard format, bundled with sites' metadata, versioned, and assigned with a DOI. The data users can register an account, search, and download data through the AmeriFlux website. Over the last seven years, about 5,000 users globally have downloaded the AmeriFlux data, totaling around 28,000 unique downloads. The data were widely used in many uses, e.g., multisite synthesis, benchmarking with RS and Land surface models, and education.

In terms of near-term plans (current technology horizon), there are discussions about allowing a subset of sites to transfer (near real-time) raw high-frequency data to AMP. There is demand among data users for this transfer, as well as taking advantage of 5G capabilities for near-real-time data transfer. AMP would process the high-frequency data to half-hourly data products using standardized protocols and codes. This will drastically increase the network demands, e.g., 1.5-5 GB per month per site, and also the computation demand, e.g., ~0.5-3 hours processing time for one site-month of data, using a commercial PC.

5.14.2.5 Remote Science Activities

AmeriFlux data's life cycle begins with data collection at more than 500 field sites across the Americas, each operated by PIs and teams. These sites are located in various natural or managed ecosystems, e.g., forests, savannas, grasslands, croplands, shrublands, wetlands, and lakes. The types of instruments may vary from site to site but typically include the following:

- Tower structure for instrument mounting: Types might vary from scaffolding walk-up tower, triangular tower, to tripod.
- Data acquisition system: data logger, computer
- Flux measurement: sonic anemometer for three-dimensional wind velocities and sonic temperature, infrared gas analyzer for water vapor and trace gas concentrations (e.g., CO2, CH4).
- Meteorological and soil measurement: Typical sensors used at each site include, but are not limited to wind speed, wind direction, solar radiation of various types, air temperature, water temperature, soil temperature, relative humidity, precipitation, snow depth, barometric pressure, soil moisture, soil electrical conductivity, and soil heat flux.

The data acquisition systems record the data streams continuously (at 10-20 Hertz for flux measurements and 1-0.1 Hertz for others). The site teams retrieve the data via physical visits (e.g., portable drive) or remote connection (e.g., Wi-Fi, cellular modem, radio transfer, or Ethernet). Then, the site teams archive and process the data at each institution before submission to AmeriFlux.

In the next phase, AmeriFlux anticipates the number of sites to continue growing and start submitting data to AmeriFlux. But, there will also be a portion of currently active sites that will become inactive and stop collecting data. The total number of active sites that submit data regularly (e.g., quarterly) should be stable or grow only gradually.

There are discussions about potential plans that a subset of sites will transfer (near real-time) raw high-frequency data to AMP and AMP processes the data to half-hourly data products using standardized protocols and codes. It will drastically increase the network demands, e.g., 1.5-5 GB per month per site. Also, not every site currently has access or sufficient bandwidths to networks to allow the real-time transfer of raw data from field sites to research institutions.

5.14.2.6 Software Infrastructure

The software involved is separated into multiple stages, from where measurements occur to when a data product is released.

- Measurement facilities (AmeriFlux sites): This is managed by multiple site teams, and the preparation and processing of data are usually done using a mix of commercial software (e.g., ReddyProc), and self-developed software tools. To date, more than 500 measurement sites are registered with the network, and each team uses their commercial/self-developed software based on their own needs.

- Data transfers from measurement facilities to data repositories: Once data is prepared in the required format, these are uploaded to the data processing facilities via a web portal for QA/QC checks. Some measurement facilities transfer raw high-frequency data to NERSC through SCP.
- Data processing and QA/QC: Software here is mostly written in-house and managed by the AMP team using Python. AmeriFlux is shifting from a Windows-based stack to a *nix-based one, and the software stack is now mostly open-source. Currently, AmeriFlux relies heavily on PostgreSQL and Flask. The only exception is Jira (used for process tracking), which is commercial. On occasion, an Azure Cloud instance is used for testing work of the legacy Windows stack.
- Serving data to users: This involves Web services via a mix of Flask and older .NET stack. The front-end Web UI is on a WordPress stack. The data provision is accomplished via FTP, managed through the department.

In the project's next stage, AmeriFlux foresees the amount of data flow will increase, and there will be a pivot to real-time data transfers. High-frequency data will likely become a more prominent use case. In addition, there is a gravitation towards real-time data consumption with tools that support exploration and visualization.

5.14.2.7 Additional Network and Data Architecture Requirements

AmeriFlux sites are at a mix of locations, some more urban with ample infrastructure and others in more remote areas where power is derived from solar panels and batteries. Similarly, networking ranges from those with wired connections to those where available options are cellular/satellite connections. Ideally, all data can be sent from the field sites to teams' institutions, where various data checks and processing occur. Practically, measurements are often recorded to a data logger or computer on-site. Then the site teams retrieve the data via physical visits or remote connection (e.g., cellular modem, radio transfer, Ethernet). Many site teams manage multiple field sites and, in rare cases, up to 40-50 field sites, e.g., NEON. Usually, data transfer from field sites to designated institutions poses the most significant challenge in the data's life cycle.

As data is processed to a half-hour resolution at each site team's institution before submitting to AmeriFlux, the data transfer is less constrained by the network, except for some sites that choose to transfer raw high-frequency data. If future scientific needs change and require more sites to transfer high-frequency data to AMP, network infrastructure can be an issue.

5.14.2.8 Use of Cloud Services

Due to cost constraints, the use of Cloud resources is somewhat limited at this point. AmeriFlux has test instances on AWS but see limited usage at this point as the costs are prohibitive for an instance that runs 24/7 instead of owning machines directly. This does impact redundancy to a certain extent. However, AmeriFlux is looking into ways to engineer backup solutions. Code repositories are less relevant to the science but critical to the development processes, where AmeriFlux leverages GitHub and CircleCI for continuous integration.

5.14.2.9 Data-Related Resource Constraints

Currently, there are no major outstanding data-related resource constraints. However, AmeriFlux anticipates that as real-time data needs come to the forefront, both networking and storage space demands will scale up significantly. There is a discussion within the community where AmeriFlux may slowly shift towards more real-time data processing and usage, where APIs will see much heavier usage in the next five years. In addition, tools like real-time data exploration with visualization will drive up both networking and data needs.

5.14.2.10 Outstanding Issues

As mentioned above, the varied infrastructure of individual research sites presents a significant hurdle in enabling near real-time data streams from the collection (i.e., sensor locations) to end-use. Remote sites in hard-to-reach places lack access to common data transmission services, such as Ethernet, cellular, or telecommunication, and remain outstanding challenges.

5.14.2.11 Facility Profile Contributors

AmeriFlux Network Representation

Project PI: Margaret Torn, *LBNL, University of California, Berkeley,* mstorn@lbl.gov

Housen Chu, *LBNL,* hchu@lbl.gov

You-Wei Cheah, *LBNL,* ycheah@lbl.gov

ESCC Representation

Rune Stromsness, *LBNL,* rstrom@lbl.gov

Richard Simon, *LBNL,* rsimon@lbl.gov

5.15 NGEE

The ESS program advances an integrated, robust, and scale-aware predictive understanding of terrestrial systems and their interdependent biological, chemical, ecological, hydrological, and physical processes. ESS is part of the Earth and Environmental Systems Sciences Division (EESSD) within the U.S. DOE's BER Program. The ESS program seeks to develop an integrated framework using a systems approach to elucidate the complex processes and controls on the structure, function, feedback, and dynamics of terrestrial ecosystems that span from the bedrock through the rhizosphere and vegetation to the atmospheric interface.

The Next-Generation Ecosystem Experiments, focusing on the understudied, yet globally important, Arctic and tropical biomes, are designed to advance the predictive power of Earth system models through understanding of the structure and function of Arctic terrestrial ecosystems, as well as the predictive understanding of tropical forest responses to global changes across scales.

5.15.1 Discussion Summary

- The NGEEs, focusing on the Arctic and the tropics, are designed to advance the predictive power of Earth system models through understanding of the structure and function of Arctic terrestrial ecosystems, as well as the predictive understanding of tropical forest responses to global changes across scales.
- The NGEE Arctic project seeks to understand how surface and subsurface processes and properties are interconnected across permafrost-dominated tundra ecosystems through a series of collaborative investigations across a gradient of remote, permafrost landscapes in coordination with landowners from multiple Native corporations.
- NGEE–Tropics aims to fill the gaps in knowledge of tropical forest-climate system interactions, and develop a predictive understanding of how tropical forest carbon balance and climate system feedbacks will respond to changing environmental drivers. The overarching goal of NGEE–Tropics is to develop a greatly improved predictive understanding of tropical forests and Earth system feedback to changing environmental drivers.
- Several NGEE initial field studies are underway at key sites, with measurements that address:
  - Impacts of thawing permafrost on greenhouse gas emissions.
  - Impacts of wildfire and thermokarst disturbance on tundra carbon balance.
  - Whether the Arctic is becoming wetter or drier.
  - Whether the Arctic is greening or browning.
  - Forest carbon cycle—hydrology interactions.
  - Nutrient limitations in tropical secondary forests.
  - Plant functional diversity response to climate change.
  - Regional variation in the causes of tree mortality.
- NGEE Arctic and NGEE–Tropics will each deliver an ecosystem model extending from bedrock to the top of the vegetative canopy/atmospheric interface that can be modeled at the scale of a high-resolution earth system model grid cell (~10 x 10 km2 grid size).
- The NGEE projects, and their local teams, collaborate on data collection, sharing instrumentation and data protocols, doing data processing and data quality control, and generating data products, which are ultimately stored in data-archive facilities. It is anticipated that for long-term preservation, this data is also synchronized with the ESS-DIVE data repository.

- NGEE Arctic is led by ORNL with partnerships at LBNL, LANL, Brookhaven National Laboratory (BNL), and the University of Alaska Fairbanks, as well as multiple Alaskan Native communities.
- NGEE Arctic also collaborates with:
  - National Aeronautics and Space Administration (NASA) Arctic and Boreal Vulnerability Experiment (ABoVE) (a memorandum of understanding is in place between NASA ABoVE and NGEE Arctic).
  - Permafrost Carbon Network.
  - Tundra Trait Team.
  - Alfred Wegener Institute (AWI).
  - Study of Environmental Arctic Change (SEARCH).
  - International Land Model Benchmarking (ILAMB) project.
  - E3SM.
  - ARM.
  - Barrow Arctic Science Consortium.
  - Ukpeagvik Iñupiat Corporation (UIC) Science, and Sitnasuak, Mary's Igloo, Council, and Bering Strait Native corporations.
  - UAF Northwest Campus.
- NGEE–Tropics researchers collaborate with both field (e.g., research groups from Manaus, Brazil) and modeling groups (e.g., E3SM) on tropical forests research questions.
- NGEE–Tropics field sites typically produce observations that can vary in size from a megabyte a few times a month, to a gigabyte over the course of a year. The frequency of sharing this depends heavily on data availability, network connectivity, and local staff availability to retrieve sensor data. Data is shared through a variety of tools, with cloud services being popular along with other more traditional methods like Globus and SCP.
- NGEE–Tropics uses data from collaborator repositories (e.g., NASA, NOAA, and NCAR), which can grow to TB scales and typically must use Globus for transfer.
- NGEE–Tropics uses Google Earth Engine for some aspects of the workflow, which results in MB to TB scales of transfer using web tools for transfer (e.g., Google Drive, etc.).
- The activities and field campaigns for NGEE Arctic at remote field sites in Arctic Alaska target uncertainties in model initialization, parameterization, and evaluation, and include the physical and remote monitoring of multiple surface and subsurface processes, ranging from abiotic variables such as temperature, moisture, and rain and snowfall, to biological variables from microbial activity to plant community composition and physiology.
- Much of the NGEE project model development and use occurs at large computing centers such as OLCF and NERSC [Section 5.15]:
  - Model analysis is either on workstations or via Jupyter Notebooks running on large computing resources.
  - Datasets are stored in public repositories, and backed-up into tape regularly and automatically.
  - All publicly available datasets will be synchronized with ESS-DIVE.

- The NGEE Arctic data ecosystem leverages several partnerships:
  - NGEE Arctic leverages the high-performance computing cluster and development area using the Compute and Data Environment for Science (CADES) infrastructure at the Oak Ridge Leadership Computing Facility (OLCF). This infrastructure hosts an E3SM land model container and tools to allow simulations in a common computing environment with datasets needed to run site- to regional-scale simulations.
  - NGEE Arctic uses ESS-DIVE for long-term data archival.
  - The NGEE Arctic project data archive at ORNL—an interim storage location before permanent data archival at ESS-DIVE—currently has 233 datasets with 161 available to the public, 62 available to the team only, and 10 in the planned (forthcoming) stage. The project archive currently has 3.5 TB of data.
  - The EMSL and JGI user facilities have been instrumental in improving the understanding of subsurface microbial communities across the Arctic tundra.
- NGEE Arctic data workflow is described as follows:
  - Researchers are responsible for the collection and management of data until it is ready for submission to the project's archive and the data management team (DMT).
  - Researchers using the project archive create metadata using the online metadata editor (OME) or email to initiate submission.
  - Data (raw, processed, etc.) is submitted via OME upload, email, Globus, Dropbox, Google Drive, or other file-transfer system.
  - Researchers using other data-type specific archive facilities like AmeriFlux, JGL, or GitHub will follow their workflow submission protocols.
  - Data is shared with the public as CCby4, with some datasets reserved for team members until eventually released to the public with a manuscript publication.
  - Digital Object Identifiers (DOIs) are assigned using the DOE Office of Scientific and Technical Information (OSTI) application programming interface.
  - Data is currently stored on an NGEE Arctic server at ORNL located in Oak Ridge, Tennessee.
  - Data will initially be mirrored and eventually migrated to the DOE ESS-DIVE.
  - The NGEE Arctic project archive at ORNL currently has 233 datasets with 161 available to the public, 62 available to the team only, and 10 in the planned (forthcoming) stage.
- NGEE–Tropics workflows involve:
  - Receiving and processing field data from multiple sites.
  - Harmonizing this data in preparation for experimental analyses.
  - Synthesis analyses (multiple types of data combined).
  - Feeding analyzed data into simulations/model runs on HPC resources.
  - Synchronizing all publicly available datasets (currently 109) with ESS-DIVE, with automatic synchronization of new datasets
- The data life cycle of field NGEE instruments is as follows:
  - Light computation happens at time of collection (in the field) to normalize results.

- At the point of retrieval, data goes through additional processing to check for data integrity and completeness.
- Data analysis typically happens at host institutions.
- Many abacus variants are possible: dozens of KB-level files, millions of KB-level files, thousands of MB-level files, dozens of gigabyte-level files, all in many possible storage organization structures/directories.

- For remote-sensing (RS) data use cases in NGEE (satellites or flown in aircraft/UAVs), heavy processing needs to happen close to where the data was collected:
  - Later analyses are conducted using GIS-type tools or numeric array data analysis tools.
  - RS data is large in nature, now easily going into PB and EB ranges, but almost always only parts of the data are needed for studies.
- NGEE simulation data is generated by running models like the Functionally Assembled Terrestrial Ecosystem Simulator (FATES) and E3SM for experiments that are more general (e.g., long-term climate predictions) or specialized (e.g., parameter sensibility analyses):
  - Data generated by many of these runs becomes valuable for further analysis.
  - Some of these datasets can be deposited with data archives such as the project archive, but many are too large to be moved in their entirety, and must be accessed parts at a time.
- Moving data from field sites to analysis and archival locations, as well as accessing large, remote datasets, will continue to be a crucial part of projects like NGEE. New tooling to facilitate data transfers from edge locations with unreliable networking without creating additional problems like file duplication or file corruption caused at transfer time is already a pressing need. Easier access to large remote datasets, e.g., remotely accessing portions of interest from a large dataset as if it was a local file in an HPC system, or from within a Jupyter Notebook, will likely become common requirements for conducting data processing and analysis.
- Data movement uses FTP, SCP, Dropbox/GDrive, Globus.org, and interfaces between these (e.g., creating a GDrive endpoint in Globus to connect a system that cannot run GDrive to a system that cannot run Globus). Simple tools like SCP will continue to be critical for conditions of limited compute power, like field locations equipped with only dataloggers.
- Local storage capacity is a recurring constraint when it comes to analyzing remote sensing or simulation output data for NGEE. Data corruption caused by unreliable connections is also a common issue for transferring data from inhospitable locations (not necessarily in a sense of "checksum corruption," but more of interrupted transfers that cannot be resumed later when connectivity is restored).
- Having ways to communicate with instruments in remote locations, e.g., using satellite uplinks such as Starlink or long-range radio to offer local coverage (maybe mesh-like) for connectivity, would allow many new forms of instrumentation of field sites.
- NGEE Arctic notes that a data team must become an integral, well-funded, component of large projects, with the appropriate toolboxes to facilitate interaction among data and models. This requires skills in data handling, preservation, and access by empiricists and modelers, all of which must be trained and integrated.
- It is recommended that ESnet and LANL collaborate on ways to effectively share data with other NGEE collaborators.

5.15.2 NGEE Facility Profile

The NGEE Arctic project seeks to understand how surface and subsurface processes and properties are interconnected across permafrost-dominated tundra ecosystems. In Phases 1-3 (2012 - 2024), NGEE Arctic conducted a series of collaborative investigations across a gradient of remote, permafrost landscapes in coordination with landowners from multiple Native Corporations. These landscapes include: polygonal tundra underlain by continuous, cold permafrost on the coastal plain at the Barrow Environmental Observatory (BEO) outside of Utqiagvik, Alaska and tundra hillslopes underlain by discontinuous, warmer permafrost on the Seward Peninsula, Alaska, that span a range of coastal and interior climates and glaciation histories. A pan-Arctic Phase 4 for NGEE Arctic is currently being proposed.

Related research experiments and modeling efforts on Alaska's North Slope and Seward Peninsula are designed to answer the following six science questions (Figure 5.15.1):

- How does the structure and organization of the landscape control the storage and flux of carbon and nutrients in a changing climate?
- What will control rates of carbon dioxide (CO2) and methane (CH4) fluxes across a range of environmental conditions?
- How do plant functional traits change across environmental gradients, and what are the consequences for carbon, water, and nutrient fluxes?
- What determines shrub distribution across the tundra landscape, and how will shrub distribution and atmospheric feedback change in the future?
- Where, when, and why will the Arctic become wetter or drier, and what are the pan-Arctic implications?
- What controls the vulnerability of Arctic ecosystems to disturbance, and how do disturbances alter the structure and function of these ecosystems?

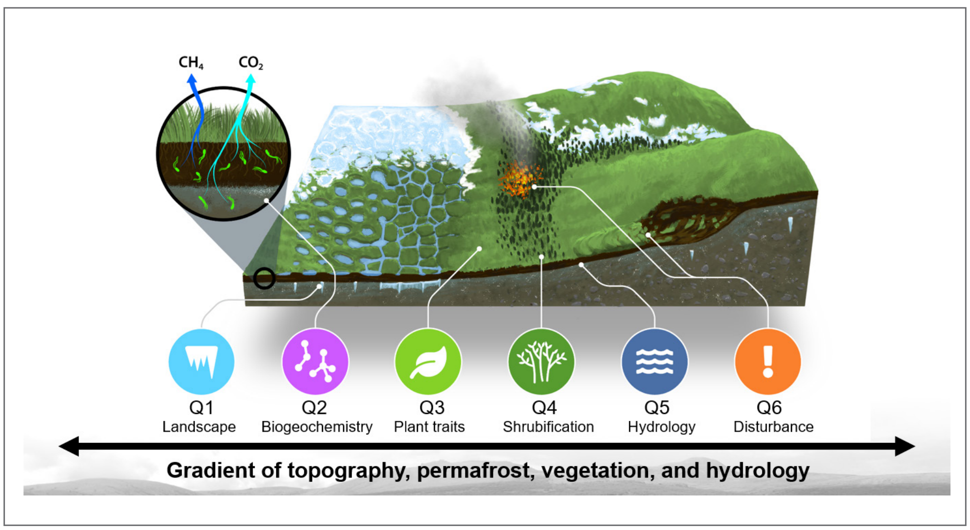

Figure 5.15.1: NGEE Arctic Science Questions. In Phase 3, science questions encompassed landscape organization (Q1), greenhouse gas production (Q2), plant functional traits (Q3), vegetation dynamics (Q4), dynamic hydrology (Q5), and disturbance (Q6).

The overarching goal of NGEE–Tropics is to develop a greatly improved predictive understanding of tropical forests and Earth system feedback to changing environmental drivers. To attain this goal, NGEE–Tropics will deliver a state-of-the-art, process-rich tropical forest ecosystem model that accurately represents forest structure and function and provides robust projections of tropical forest responses to global change.

5.15.2.1 Science Background

5.15.2.1.1 NGEE Arctic

The goal of NGEE Arctic is to support the DOE BER mission to advance a robust predictive understanding of Earth's climate and environmental systems by delivering a process-rich ecosystem model, extending from bedrock to the interface between the vegetative canopy and the atmosphere, in which the evolution of Arctic ecosystems in a changing climate can be modeled at the scale of a high-resolution earth system model grid cell. To achieve this ambitious goal, NGEE Arctic has brought together a team of more than 100 scientists spanning four national laboratories and the University of Alaska Fairbanks. In Phases 1-3, a series of collaborative investigations was conducted across a gradient of permafrost landscapes in coordination with landowners from multiple Native Corporations. These landscapes include polygonal tundra underlain by continuous, cold permafrost on the coastal plain at the BEO outside of Utqiagvik, Alaska, as well as tundra hillslopes underlain by discontinuous, warmer permafrost on the Seward Peninsula, Alaska, that span a range of coastal and interior climates and glaciation histories.

In Phase 1 (2012–2014), the modeling approach was driven by the recognition that the current generation of Earth system models fails to capture many processes controlling the cycling of carbon, water, nutrients, and energy in the Arctic tundra. Initial model improvements emphasized the unique aspects of cold, continuous permafrost environments and improved mechanistic model representation of coupled surface- and subsurface thermal hydrology, biogeochemistry, and snow. In Phase 2 (2015–2018), the multiscale modeling approach informed hypotheses about which reductions in model complexity were possible while retaining predictive power, and in turn, which processes required the use of fine-scale model parameterizations to inform intermediate-scale models. NGEE Arctic also confronted improved models with unique aspects of hillslope environments across warmer, discontinuous permafrost environments. In Phase 3 (2019–2024), NGEE Arctic observations were leveraged, synthesis products, and model scaling over space and time to inform the conceptualization and parameterization of six integrated model (IM) modules in an NGEE Arctic fork of the E3SM repository (E3SM Version 2.1.0). These modules encompass new model conceptualization and parameterizations, including new tundra-specific plant functional types, snow distribution and dynamics along tundra hillslopes, and redox biogeochemistry impacts on greenhouse gas production. Furthermore, continued observations and model experiments indicate that intensifying impacts of climate change on extreme events, disturbance, erosion, and abrupt thaw can affect carbon-climate feedbacks on tundra processes ranging from potential ice-rich permafrost subsidence to wildfire impacts.

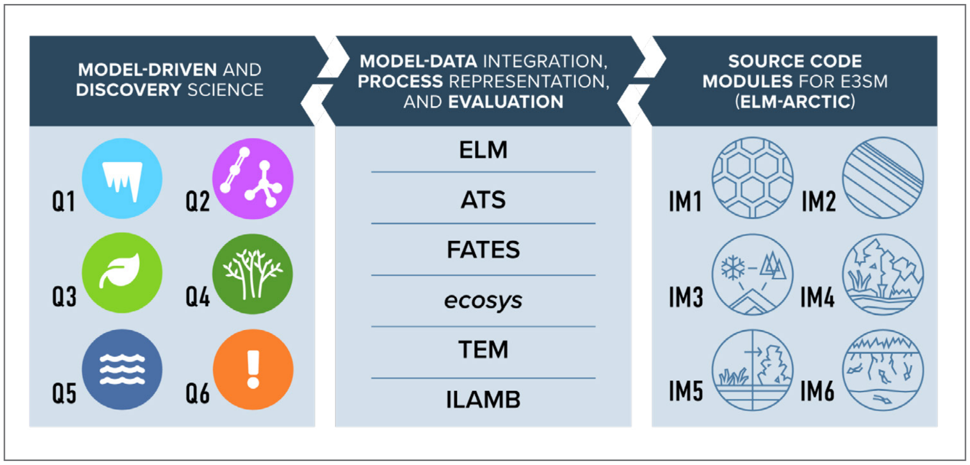

Figure 5.15.2: NGEE Arctic Scientific Questions and Integrated Modeling Modules. NGEE Arctic Science Questions in Phase 3 encompassed landscape organization (Q1), greenhouse gas production (Q2), plant functional traits (Q3), vegetation dynamics (Q4), dynamic hydrology (Q5), and disturbance (Q6). These questions informed the development of six integrated model (IM) modules in ELM ('ELM-Arctic') spanning improved representation of fractional inundated area (IM1); improved representation of surface and subsurface distribution of water (IM2); improved representation of interactions among snowpack, terrain, and vegetation (IM3); expanded and improved representation of Arctic tundra plant functional types (IM4); new implementation of dynamic tundra biogeography and fire disturbance (IM5); and improved representation of tundra biogeochemistry (IM6).

5.15.2.1.2 NGEE–Tropics

The Next-Generation Ecosystem Experiments–Tropics, or NGEE–Tropics, is a ten-year, multi-institutional project funded by the US DOE), SC, BER. NGEE–Tropics aims to fill the critical gaps in knowledge of tropical forest-climate system interactions. The overarching goal of NGEE–Tropics is to develop a predictive understanding of how tropical forest carbon balance and climate system feedback will respond to changing environmental drivers over the 21st Century.

In Phase 2 of NGEE–Tropics (ending FY2025), assessment of what is known about tropical forest ecosystems, and how well these processes are represented in earth system models, is being performed. Uncertainty assessments are informing field and laboratory studies that will provide new, high-priority data for forest ecosystem models. Several initial field studies are underway at key sites, including pilot studies in Brazil, Panama, and Puerto Rico, and a select group of ForestGEO sites across the tropics. Measurements at these sites address:

- Forest carbon cycle—hydrology interactions
- Nutrient limitations in tropical secondary forests
- Plant functional diversity response to climate change
- Regional variation in the causes of tree mortality

The overarching goal of NGEE–Tropics is to develop a predictive understanding of how tropical forest carbon balance and climate system feedback will respond to changing environmental drivers over the 21st Century. NGEE–Tropics' grand deliverable is a representative, process-rich tropical forest ecosystem model, extending from bedrock to the top of the vegetative canopy-atmosphere interface, in which the evolution and feedbacks of tropical ecosystems in a changing climate can be modeled at the scale and resolution of a next-generation Earth System Model grid cell (~10 x 10 km2 grid size).

Our fundamental approach to advancing understanding and model representation of tropical forests in FATES is a strong coupling of model development and evaluation with experiments and observations. This deliberate and focused integration of models and experiments (termed ModEx) ensures that model development is informed by the latest empirical knowledge, and that field measurements are explicitly designed to target gaps in process understanding or parameterization, thereby addressing substantial uncertainty in E3SM-FATES. This approach requires that NGEE–Tropics develop model testbeds—which bring site-specific meteorological and plant trait data to drive the model together with other observations to test the model—to evaluate model performance at individual, community and regional scales.

Three Research Focus Areas (RFAs, described in Figure 5.15.3) are groups collecting experimental data at different spatial scales, from leaf, to canopy and ecosystem, to regional and global scales. This data is collected at multiple field sites, with two of the main sites for the project being located in Panama, at the Barro Colorado Island Research Station (Smithsonian Tropical Research Institute), and in Brazil, at the Manaus ZF2 Research Station managed by the National Institute for Amazon Research (Instituto Nacional de Pesquisas da Amazônia [INPA]), and many additional field sites spread throughout the Tropics, including in Puerto Rico. NGEE–Tropics and the local teams collaborate on the data collection, sharing instrumentation and data protocols, doing data processing and data quality control, and also generating data products, which are ultimately stored in the NGEE–Tropics Data Archive hosted at NERSC. For long-term preservation, this data is also synchronized with the ESS-DIVE data repository for BER. Data collected by RS platforms, including airborne and satellite data, are also critical observational data for the project, and derived data products are also stored in the archive. This data is then used to drive and validate model runs, with many of the results being stored in disk and tape systems on NERSC and other HPC centers, as well as the archive for published results of these runs.

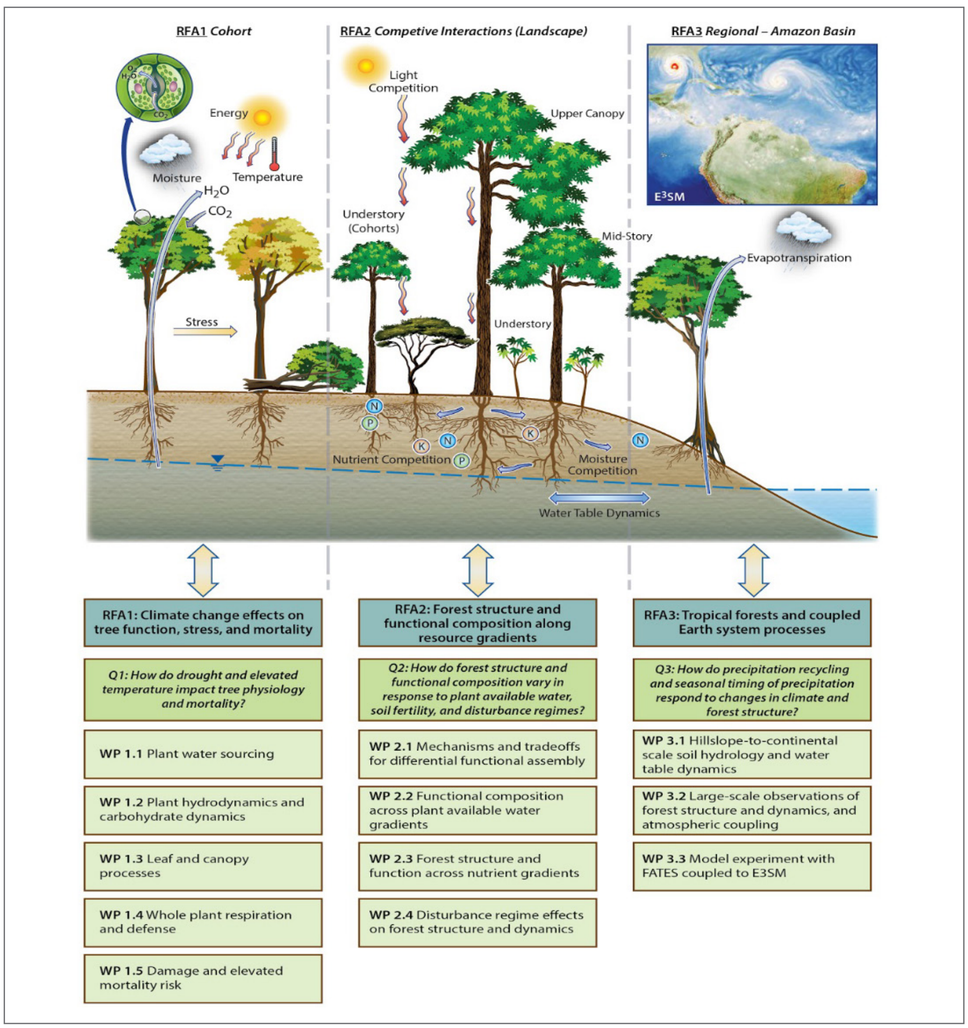

Figure 5.15.3: NGEE–Tropics Research Areas

5.15.2.2 Collaborators

5.15.2.2.1 NGEE Arctic

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
Globus at ORNL	Secondary copy for transfer	Direct transfer from user computer through Globus to NGEE server	Over 2GB	Monthly	No	Glitchy since new users are not familiar. LANL cannot use it at all. Maintenance and upgrades are not controllable
Dropbox	Secondary copy for transfer	Direct from user		Almost never since using Globus	No	Not usually any issues. Still cannot use it for large files since researcher typically has free version.

Table 5.15.1: NGEE Arctic Collaboration Space

NGEE Arctic is led by ORNL, with partner institutions at LBNL, LANL, BNL, and the University of Alaska Fairbanks, as well as several Alaskan Native Communities.

NGEE Arctic also collaborates with:

- NASA Arctic and Boreal Vulnerability Experiment (ABoVE) (a memorandum of understanding is in place between NASA ABoVE and NGEE Arctic)
- Permafrost Carbon Network
- Tundra Trait Team
- AWI
- SEARCH
- ILAMB project
- E3SM
- ARM
- Barrow Arctic Science Consortium
- UIC Science, Sitnasuak, Mary's Igloo, Council, and Bering Strait Native Corporations
- UAF Northwest Campus

Data workflow within the project data archive:

- Researchers are responsible for the collection and management of their data until it is ready for submission to the project's archive and the DMT.
- Researchers using the project archive create metadata using the OME or email to initiate submission. Data (raw, processed, etc.) submitted via OME upload, email, Globus, Dropbox, Google Drive, or other file-transfer system.
- Researchers using other data-type specific archive facilities like AmeriFlux, JGL, or GitHub will follow their workflow submission protocols.
- Data is shared with the public as CCby4 with some datasets reserved for team members until eventually released to the public with a manuscript publication.

- DOIs are assigned using the DOE OSTI API service.
- Data is currently stored on NGEE Arctic server at ORNL located in Oak Ridge, Tennessee
- Data will initially be mirrored and eventually migrated to the DOE ESS-DIVE, which is a data repository for Earth and environmental science data funded by DOE's SC, BER, ESS program.

5.15.2.2.1 NGEE–Tropics

NGEE–Tropics researchers collaborate with both field and modeling groups on tropical forests research questions. One key set of collaborators are team members and other researchers at field sites, many of which are in remote locations. For example, NGEE–Tropics has a large field presence at a research forest near Manaus, Brazil, with several sets of observational streams being collected.

We also collaborate with Earth system modeling groups such as the DOE E3SM project, and both develop land models and Earth system models, as well as use the output from such models for analysis. Much of the model development and use occurs at large computing centers, while much of the model analysis is either on workstations or via Jupyter Notebooks either running on large computing resources or local workstations.

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
Field Site: Forest Management Laboratory, Manaus, Brazil. Team of about 5 staff and 10-40 students at different degree stages.	Primary/ Secondary (NGEE–Tropics functions as primary for some of the raw data and most of the derived data products)	Hard drive, Dropbox/GDrive, SCP, Globus.org	10 MB (a few times/ month) 100 GB (1-2 times/year)	Ad hoc	Yes, in transformed data product form	Collaborators unfamiliar with tools like SCP; issues with Dropbox/GDrive (slow uploads, duplicate files, inconsistent copies), and unreliable network
Field Site: Large Biosphere-Atmosphere Experiment, Manaus, Brazil (also multiple field sites throughout the Amazon forest). Team of about 6 staff and occasional students and external collaborators.	Same as above	Same as above	10 GB	2-4 times/ month	Same as above	Same as above
Field Site: Smithsonian Tropical Research Institution, multiple sites in Panama, coordinated by a team based at Barro Colorado Island. Team of about 4 staff and 10-20 students at different degree stages.	Same as above	Same as above	10MB-100GB	Regular for some instruments (a few times/ month), and ad hoc for campaigns	Same as above	Same as above

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
Multiple other field sites throughout the Tropics – 10s of collaborators spread geographically	Primary	Same as above	10MB-1TB	Ad hoc (campaigns and research collaboration), usually one-time only interaction	Same as above	Same as above
NASA and NOAA Data Repositories (e.g, DAACs and other services, multiple locations in the US)	Primary	Data portals	10s MB to 100s GB	Ad hoc (for specific experiments)	No	Local storage not always adequate
Model Data Repositories (e.g., NCAR in Boulder, CO)	Primary	Data portals, SCP/rsync, and Globus.org	10s MB to 10s TB	Ad hoc (for specific experiments)	Rarely	Local storage not always adequate
Google Earth Engine	Secondary	GDrive and browser upload	10s MB to 100s MB (upload), 10s GB (download)	Ad hoc (for specific experiments)	Frequently	Data transfer methods limited and not reliable

Table 5.15.2: NGEE–Tropics Collaboration Space

Workflows usually involve receiving and processing (including QA/QC) field data from multiple sites, harmonizing this data in preparation for experimental analyses, synthesis analyses (multiple types of data combined), and for being fed into simulations/model runs.

Field data is usually collected as:

- Continuous sampling (longer-term instrumentation, ranging from a few months to decades being deployed in the field collecting data); or,
- Field campaigns (more limited in scope and time, but usually more data intensive)

Other important activities for the project include:

- Field logistics support (e.g., remote maintenance of instrumentation, sometimes with remote access to dataloggers).
- Access to large datasets stored in other data centers for:
  - Use within simulation runs in HPC systems
  - Analyses conducted in smaller clusters or laptops (e.g., using Jupyter Notebooks)
  - Data exchange with commercial compute-capacity providers such as Google Earth Engine to access, combine, process, and retrieve RS data products.

Research data products for the project are all stored within the NGEE–Tropics Data Archive, developed at LBNL and hosted at NERSC (using Spin infrastructure). Datasets are stored in CFS and backed-up into tape regularly and automatically. All publicly available datasets within the NGEE–Tropics Data Archive have been synchronized with ESS-DIVE and new datasets are synchronized shortly after they are published. This is only possible because the NGEE–Tropics Archive now maintains full compatibility with ESS-DIVE data and metadata guidelines. Availability within ESS-DIVE guarantees long-term availability while project users and collaborators interacting with the NGEE–Tropics Archive still have access to tailored support for their data-archival needs and for following the project's data policy and DOE data availability requirements. Data deposited with the NGEE–Tropics archive automatically receives a DOI (issued by OSTI).

5.15.2.3 Use of Instruments and Facilities

5.15.2.3.1 NGEE Arctic

NGEE Arctic is supported by the Biological and Environmental Research Program in the DOE's SC. Colleagues at DOE's ARM North Slope of Alaska atmospheric observatory in Utqiagvik, Alaska, and at the University of Alaska Fairbanks Northwest Campus in Nome, Alaska, facilitate work at remote field sites. Native corporations Ukpeagvik Iñupiat Corporation (UIC) Science, the Bering Straits Native Corporation, the Council Native Corporation, Mary's Igloo Native Corporation, and the Sitnasuak Native Corporation permit us to conduct research on the traditional homelands of the Iñupiat. NGEE Arctic leverages the high-performance computing cluster and development area using the Compute and Data Environment for Science (CADES) infrastructure at the OLCF that hosts an ELM model container and tools to allow simulations in a common computing environment with datasets needed to run site- to regional-scale simulations. NGEE Arctic uses the Environmental System Science Data Infrastructure for a Virtual Ecosystem (ESS-DIVE) for long-term data archival. The NGEE Arctic project data archive[86] at ORNL — an interim storage location before permanent data archival at ESS-DIVE — currently has 233 datasets with 161 available to the public, 62 available to the team only, and 10 in the planned (forthcoming) stage. The project archive currently has 3.5 TB of data. The EMSL and JGI user facilities have been instrumental in improving the understanding of subsurface microbial communities across the Arctic tundra.

5.15.2.3.2 NGEE–Tropics

Table 5.15.2 lists facilities and locations relevant to the project. As to instruments, three main types can be identified/categorized as data sources for NGEE–Tropics: field instruments, RS instruments, and simulation results.

For field instruments, only basic computation happens at time of collection (still in the field, using limited datalogger processing capabilities mostly for computing data aggregates and issuing quality flags). At the point of retrieval, data might go through additional light processing to check for data integrity and completeness. Only when generating derived data products and conducting data analysis heavier processing will happen, usually back at host institutions. Many variants are possible in this case: dozens of KB-level files, millions of KB-level files, thousands of MB-level files, dozens of GB-level files, all in many possible storage organization structures/directories.

For RS data, imagery collected with instruments either onboard satellites or flown in aircraft/UAVs, heavy processing needs to happen close to where the data was collected, since most of the information to generate even basic versions of the data are strictly dependent on the instruments and their configuration parameters. Later analyses are conducted using GIS-type tools or numeric array data analysis tools. RS data is large in nature, now easily going into PB and EB ranges, but almost always only parts of the data are needed for studies.

Simulation data is generated by running models like FATES and E3SM for experiments that are more general (e.g., long-term climate predictions) or specialized (e.g., parameter sensibility analysis). Data generated by many of these runs become valuable for further analysis. Some of these datasets can be deposited with data archives such as the project archive, but many are too large to be moved in their entirety, and must be accessed parts at a time.

---

86 https://data.ngee-arctic.ornl.gov/data/#fq=datasource%3A%22NGEE%20Arctic%22&q=*%3A*

Phase 1 of the project involved limited deployment of instrumentation in field sites. New field deployments in Phase 2 were severely limited by COVID, but have been ramping back up significantly in 2022 and plans continue to see this increase for the remainder of Phase 2. As an example, the project now likely maintains the largest site instrumented with tree sapflow sensors in the world (about 90 trees monitored with sampling every second). Additional measurements, especially as identified by model development needs, are planned to continue to be made at these field sites in the next five years. As these new measurements become more well established and their value proven, long-term deployments can be planned. As examples for experimental data, micrometeorological records can date as far back as the 1950s with fairly reliable records; forest inventories have been conducted regularly for likely over 100 years; and, more recently, the eddy covariance technique for measuring carbon, water, and energy fluxes, has been used consistently to monitor ecosystems continually for as long 30 years now. Other types of measurements will likely prove to be as useful as these and become standard measurements for these types of research field sites.

5.15.2.4 Process of Science

5.15.2.4.1 NGEE Arctic

Variation in landscape structure and organization—including ridges, valleys, and watershed basins—provide the organizing framework for integrating process studies. Multiscale research activities organized around these components are designed to help determine whether the Arctic is, or in the future will become, a negative or positive feedback to anthropogenically-forced climate change. These activities include landscape heterogeneity, soil biogeochemistry, plant traits, shrub dynamics, watershed hydrology, and disturbance.

- Landscape heterogeneity considers how vegetation, biogeochemistry, and hydrology are influenced by regional landforms such as ridges, valleys, and drainage basins. Variations in landscape form and function are simplified for input into models while still retaining fine-scale features with the strongest influence on climate feedback.
- Soil biogeochemistry focuses on the influence of temperature, moisture, and nutrients on microbial dynamics and soil organic matter decomposition. Understanding controls over CO2 and CH4 fluxes across a range of environmental conditions has been critical for improving predictions of net carbon exchange in Arctic systems.
- Plant traits provide an important link among plant communities, tundra biogeochemistry, and biophysical feedback to the atmosphere. Improved understanding of tundra plant traits, above- and belowground, have led to new approaches for scaling plant function from leaf to landscape using remotely sensed data and models.
- Shrub dynamics focuses on advancing the capability to determine shrub distributions and predict their impacts on climate, now and in the future. The development of dynamic tundra vegetation models, including new shrub and moss functional types, has facilitated process-rich simulations over the twenty-first century and the identification of structural and parameter priorities for dynamic predictions.
- Watershed hydrology determines the spatial distribution and temporal dynamics of soil saturation and inundation and, in turn, surface and subsurface hydrological and biogeochemical cycling and local- to regional-scale energy balance. Observations have informed key model improvements in prediction of past, present, and future distributions of snow, soil moisture and inundation, and surface and subsurface water flow across the Arctic landscape.
- Disturbances such as wildfire, thermokarst, and thermal erosion are expected to increase in frequency and magnitude over the coming decades due to a warmer climate and the intensification of human activities in the Arctic, and have the potential to cause profound changes by disrupting processes that regulate climate feedback. Observations and

model simulations have targeted knowledge gaps and interactions among multiple forms of disturbance.

This comprehensive suite of NGEE Arctic process studies and observations is being strongly linked to model development and application requirements for improving process representation, initializing multiscale model domains, calibrating models, and evaluating model predictions. A fundamental challenge for the NGEE Arctic modeling activity is to relate new process knowledge gained at fine- and intermediate spatial scales to states and fluxes relevant for integration in earth system models. Consequently, a nested hierarchy of models is being engaged at fine-, intermediate-, and global scales, connecting process studies to models and models to each other in a quantitative upscaling and downscaling.

5.15.2.4.2 NGEE–Tropics

Moving data from field sites to analysis and archival locations, as well as accessing large, remote datasets will continue being a crucial part of projects like NGEE–Tropics. New tooling to facilitate data transfers from edge locations with unreliable networking without creating additional problems like file duplication or file corruption caused at transfer time are already a pressing need. Easier access to large remote datasets, e.g., remotely accessing portions of interest from a large dataset as if it was a local file in an HPC system or from within a Jupyter Notebook will likely become common requirements for conducting data processing and analysis.

5.15.2.5 Remote Science Activities

5.15.2.5.1 NGEE Arctic

NGEE Arctic is a model-driven, multiscale, pan-Arctic research program that builds on a foundation of model-data integration from a decade of observations across Arctic Alaska. NGEE Arctic emphasizes iterative collaboration among interdisciplinary teams of empiricists and modelers to incorporate experiments into models (i.e., a 'ModEx' approach), underscored by a foundation of open science and data sharing (DOE SC, EESSD). NGEE Arctic field research activities span a north-south gradient of sites - a northern location near Utqiagvik on the North Slope of Alaska is underlain by cold, discontinuous permafrost, and more southerly locations outside of Nome on the Seward Peninsula of Alaska are underlain by warmer, discontinuous permafrost. Permits to work at the Barrow Environmental Observatory (BEO), as well as across five hillslopes on the Seward Peninsula, are obtained annually from the Native Communities who have lived on and stewarded the land since time immemorial. NGEE Arctic Phase 1 began on the BEO, and the project added additional sites on the Seward Peninsula in Phases 2 and 3. As the project moves into Phase 4, it will extend to the pan-Arctic, but with a strong modeling, rather than observational, approach. As Phase 3 in Arctic Alaska begins to wrap up by October 2024, equipment and infrastructure temporarily installed at remote field sites is being decommissioned.

The activities and field campaigns for NGEE Arctic at remote field sites in Arctic Alaska target uncertainties in model initialization, parameterization, and evaluation, and include the physical and remote monitoring of multiple surface and subsurface processes, ranging from abiotic variables such as temperature, moisture, and rain and snowfall, to biological variables from microbial activity to plant community composition and physiology. For example, micrometeorological instrumentation (weather stations, Figure 5.15.4a), including various sensors for both surface and subsurface measurements, are deployed to understand diurnal, seasonal, and annual variation in abiotic factors. Wells, piezometers, and surface water collection enable characterization of water chemistry and flow rates. Soils and plants are sampled during snow-free seasons from permafrost to the top of the vegetative canopies to monitor edaphic conditions, permafrost ice content, microbial community composition and activity, isotopes, and dynamic carbon and nutrient cycling (Figure 5.15.4b). Fluxes of carbon dioxide and methane are sampled from the soil at a small, point scale using hand-held, battery-powered instrumentation and chambers, and from the ecosystem at a larger scale using eddy covariance towers (Figure 5.15.4c). In turn, snow surveys, where researchers access remote field sites in using either helicopters or snow machines, quantify snow characteristics and thickness with snow pits and magnaprobes, as well as subsurface properties with ground-penetrating radar and electrical resistivity tomography (Figure 5.15.4e). Aerial imagery is collected using kites, unoccupied aerial systems (drones, Figure 5.15.4f), airplanes, and satellites. Much of the aerial RS work is done in collaboration with partners at NASA ABoVE.

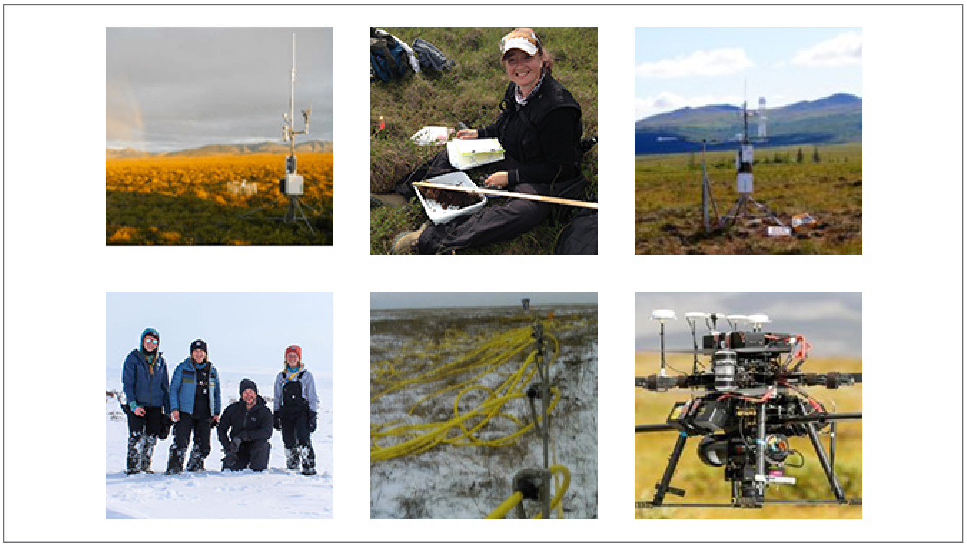

Figure 5.15.4: Activities and field campaigns for NGEE Arctic at remote field sites in Alaska. (a) Weather station at the Council Site on the Seward Peninsula of Alaska, which uses trickle charge batteries powered by solar panels, and a nearby precipitation gauge. (b) Sampling of plants and soil at the Kougarok Hillslope on the Seward Peninsula of Alaska. (c) Eddy covariance tower at the Council Site on the Seward Peninsula of Alaska, battery-operated and powered by solar panels, small wind turbines, or a methanol fuel cell. (d) Snow campaign at the Teller Watershed on the Seward Peninsula of Alaska. (e) Electrical resistivity tomography (ERT) at the Barrow Environmental Observatory (BEO) outside of Utqiagvik, Alaska. (f) Quadcopter unoccupied aerial system (UAS) collecting vegetation imagery.

5.15.2.5.2 NGEE–Tropics

NGEE–Tropics is inherently distributed in nature: the team and collaborators are spread throughout the globe; field sites are in remote tropical locations; large remote sensing and simulation datasets are hosted at many different institutions. This is likely to continue and intensify in the future.

5.15.2.6 Software Infrastructure

5.15.2.6.1 NGEE Arctic

Data movement to the archive uses an FTP upload through the Online Metadata Editor or for large files there is an option of Globus, lab file-transfer system, or Dropbox. Sometimes files are simply emailed.

The NGEE Arctic team hosted an ESS project-wide 'field to model'[87] workshop that facilitated understanding among field ecologists on how to access and understand E3SM land model code via Jupyter Notebook and a Docker platform.

5.15.2.6.2 NGEE–Tropics

Data movement uses FTP, SCP, Dropbox/GDrive, Globus.org and interfaces between these (e.g., creating an GDrive endpoint in Globus to connect a system that cannot run GDrive to a system that cannot run Globus).

87 https://drive.google.com/file/d/1pzA2SE-RlPUvpvHpp9lFN1uI5qOsslMW/view?usp=share_link

We also believe that simple tools like SCP will continue to be critical for conditions of limited compute power like field locations equipped with only dataloggers.

5.15.2.7 Additional Network and Data Architecture Requirements

5.15.2.7.1 NGEE Arctic

ORNL connects to ESnet via redundant border routers, each of which currently connects to a diverse ESnet router at 100G. The expectation is that these connections will soon be upgraded to 400G connections. The ORNL border routers connect the ORNL Enterprise network and OLCF to ESnet. This connectivity is depicted in the diagram below.

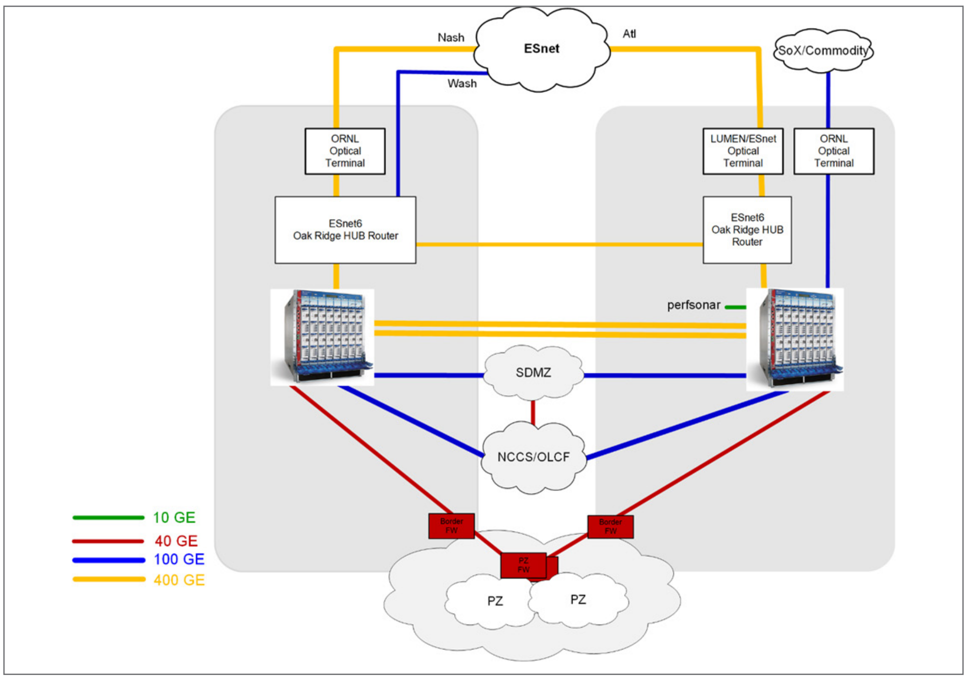

Figure 5.15.5: The Enterprise firewalls are connected to the border routers at 40G currently. ORNL does use a Science DMZ architecture for high-performance data transfer. This environment connects to the border routers with 10/40/100G DTN connections available. Globus is the approved transfer method. A border perfSONAR node is connected to the border router and participates in the ESnet grid.

5.15.2.7.2 NGEE–Tropics

Examples of bottlenecks and limitations for the project were listed in the examples above, as well as examples of services that might be beneficial to address these.

5.15.2.8 Use of Cloud Services

5.15.2.8.1 NGEE Arctic

Globus as a free service is used for file transfer (although problematic for colleagues at LANL, as noted above) and Google Drive is used for file sharing and collaboration although not typically to house data or data transfer since the project uses a free version.

5.15.2.8.2 NGEE–Tropics

Two main uses of cloud services are ongoing for NGEE–Tropics:

- Our data archive is hosted in a cloud-like environment on NERSC (called Spin);
- Several team members use GCP, via their Google Earth Engine tool for analysis of large spatial and remotely sensed data.

We also make extensive use of cloud storage services, mainly a combination of Dropbox and Google Drive, depending on what collaborators prefer to use.

5.15.2.9 Data-Related Resource Constraints

5.15.2.9.1 NGEE Arctic

We have similar constraints to those listed below by NGEE–Tropics. In addition, the ability to transfer and download data in real-time from remote field sites that are hard to access is a constraining factor. NGEE Arctic tackled this problem in one remote location on the Seward Peninsula of Alaska by installing a repeater on a mountaintop that allowed signal transmission from an eddy covariance tower in a remote watershed to a separate mountaintop hosting a satellite uplink. But this is time and resource-intensive. Another constraint is people - there is excitement that data handling, preservation, and access by empiricists and modelers is increasingly a focus of large, national-lab worthy programs. However, a data team must become an integral, well-funded, component of large projects, with the appropriate toolboxes to facilitate interaction among data and models.

5.15.2.9.2 NGEE–Tropics

Local storage capacity is a recurring constraint when it comes to analyzing remote sensing or simulation output data. Data corruption caused by unreliable connections are also a common issue for transferring data from inhospitable locations (not necessarily in a sense of "checksum corruption", but more of interrupted transfers that cannot be resumed later when connectivity is restored).

5.15.2.10 Outstanding Issues

5.15.2.10.1 NGEE Arctic

Given the remote nature of the Arctic, outstanding issues will always be related to installation of equipment and infrastructure, and then accessing ongoing data collection in 'real time' to understand the biology but also understand and assess problems with equipment or data streams. This is particularly relevant to biomes such as the Arctic where harsh winter conditions can limit physical site access for more than half of the year, or where trips for equipment maintenance that require helicopter support are quite expensive. Aerial and satellite remote sensing are also important components, but remote access, as well as the path of satellites across the globe, can hinder year-round understanding in the high latitudes. However, the dynamics of biological processes on the 'shoulder seasons' of spring and fall are important to monitor and appropriately capture in Earth system models. The NGEE Arctic project has innovated in the space of methanol fuel cells to power data loggers over winter when solar panels are not an option (e.g., Krassovski et al. 2020), as well as satellite relay stations to relay data collected from eddy covariance systems and meteorological stations to cloud storage. However, more work is needed, and a potential path forward is described below.

5.15.2.10.2 NGEE–Tropics

Considering connectivity to edge networks (e.g., field sites), having ways to communicate with instruments in remote locations – e.g., using satellite uplinks such as Starlink or long-range radio to offer local coverage (maybe mesh-like) for connectivity would allow many new forms of instrumentation of field sites.

5.15.2.11 Facility Profile Contributors

NGEE Representation

Arctic

Stan Wullschleger, *ORNL,* wullschlegsd@ornl.gov

Terri Velliquette, *ORNL,* velliquettet@ornl.gov

Susan Heinz, *ORNL,* heinzsl@ornl.gov

Colleen Iversen, *ORNL,* iversencm@ornl.gov

Tropics

Charles Koven, *LBNL,* CDKoven@lbl.gov

Gilberto Pastorello, *LBNL,* gzpastorello@lbl.gov

Sandy Chin, *LBNL,* schin@lbl.gov

ESCC Representation

Rune Stromsness, *LBNL,* rstrom@lbl.gov

Richard Simon, *LBNL,* rsimon@lbl.gov

Susan Hicks, *ORNL,* hicksse@ornl.gov

5.16 Coastal Efforts

The COMPASS program aims to improve fundamental scientific understanding, model representation, and predictive capacity of coastal systems. COMPASS is a multi-institutional effort funded by the Earth and Environmental System Science Division of the US DOE's SC. COMPASS will help dramatically enhance predictive understanding of coastal systems, including their response to short- and long-term changes.

5.16.1 Discussion Summary

- The Coastal Observations, Mechanisms, and Predictions Across Systems and Scales (COMPASS) program aims to improve fundamental scientific understanding, model representation, and predictive capacity of coastal systems. The project comprises two parts: a field study and a coastal modeling study:
  - COMPASS-FME (Field, Measurements, and Experiments) focuses on field studies and associated process and ecosystem modeling of two coastal interfaces.
  - COMPASS-GLM (Great Lakes Modeling) focuses on modeling and analysis of coastal systems in the Great Lakes Region.
- COMPASS-FME seeks to advance a scalable, predictive understanding of the fundamental biogeochemical processes, ecological structure, and ecosystem dynamics that distinguish coastal terrestrial–aquatic interfaces (TAIs) from the purely terrestrial or aquatic systems to which they are coupled. FME will focus on overarching long-term science questions:
  - What fundamental mechanisms control the structure, function, and evolution of coastal TAIs?
  - How do these fundamental mechanisms interact across spatial scales, and what interactions are most important to improving predictive models?
- COMPASS-FME is a multi-institutional project led by PNNL, and partners with ORNL, ANL, LANL, and LBNL, as well as the Smithsonian Environmental Research Center, the University of Toledo, and Heidelberg University. FME generates data across seven different sites in the Chesapeake Bay and Western Lake Erie regions. This data provides broad spatial and temporal coverage of ecosystem processes. Data is remotely logged, collected through telemetry, and compiled for use by experimentalists and modelers to discern patterns of change and guide model improvements.
- COMPASS-GLM aims to enhance understanding of freshwater coastal systems, especially how they respond to climate warming, land use land cover change (LULCC), and other perturbations at watershed-to-regional scales. GLM will focus on overarching long-term science questions:
  - What multiscale mechanisms govern the structure, function, and dynamics of coastal systems at different spatial and temporal scales?
  - How do coastal systems respond to natural and anthropogenic influences?
  - Is it possible to generalize new process knowledge and predictive skill gained at a small number of sites or regions across the observed diversity of coastal systems?
- COMPASS-GLM modeling activities will center on two tasks. The first is focused on coupled, integrated atmosphere–land–lake regional modeling in the Great Lakes Region. The second will focus on high-resolution modeling of the Portage River watershed. GLM will share model output across tasks to drive the high-resolution watershed model with output from the atmospheric model to understand how the watershed responds to individual events and future climate conditions.

- COMPASS-FME data from sensors amounts to datasets that are kilobytes to megabytes in size and are collected constantly. These are typically shared via cloud sharing services, and stored either at the facilities from the project partners or shared cloud resources.
- COMPASS-GLM model datasets will be shared between two multi-institutional task-level teams for atmosphere–land–lake regional modeling (Great Lakes Region) and high-resolution watershed modeling (Portage River Basin). The models are GB scale, and are stored on HPC resources at ANL, LBNL, and ORNL. Globus is used for data sharing, along with other tools such as web portals and SCP.
- COMPASS-FME operates seven real-time environmental sensor clusters (e.g., synoptic nodes) at sites along Chesapeake Bay and Lake Erie. Each node contains data loggers and associated communication infrastructure, including cell phone and radio telemetry, as well as hardened solar power infrastructure. Each node produces about 5 Gb raw data per year, with 40 Gb of processed data annually. Current instrumentation includes:
  - Water quality measurements.
  - Meteorological parameters.
  - Soil measurements
  - Tree health.
- COMPASS-FME aims to make data available through a collaborative network. The workflow will develop data distribution infrastructure and data products, and is designing a system that will facilitate interactivity with the data, as well as facilitating access. Starting with cloud services for data sharing and storage has revealed several areas of friction:
  - Performance when uploading large datasets is not predictable.
  - Permissions to release and share data with collaborators is not consistent, and often does not integrate with institutional identities.
  - Data changes are challenging to track over time.
- A common document-sharing platform supported/certified by all DOE laboratories will be helpful for collaborative work in COMPASS and potentially improve productivity.
- COMPASS-GLM has access to HPC resources at facilities at the National Energy Research Scientific Center (NERSC) for modeling. All the data products are stored on and can be accessible from NERSC's HPSS. Future simulations will be performed on COMPASS-HPC, a midrange high-performance computing facility at PNNL. Data transfer and sharing between researchers across scales use Globus, which allows efficient, secure, and reliable data transfer across all supercomputers at DOE HPC facilities as well as personal computers.
- COMPASS generates two broad categories of data, and makes heavy use of Jupyter Notebooks to document computational workflows including data sources, codes used, and all associated metadata. To that end, both efforts will generate, and work with:
  - Observational and experimental data generated by field and laboratory-based experiments.
  - Simulation data generated from multiscale modeling activities.
- COMPASS-GLM raw data output from WRF-FVCOM is approximately 200 MB over the course of a day. This is a combination of over 20 variables from the sensors. These are combined into multiday datasets that are approximately 10 GB in size. Raw data is shared via Globus, and used in modeling activities on HPC resources. The output files from the transient daily integrated hydrology simulations in ATS for the Portage basin (~1,000 km2) for 15 years of simulation are about 150 GB.

- Most major data transfer operations for COMPASS, both incoming and outgoing, make use of the Globus data mobility software. Smaller data transfers may use other protocols such as SFTP or https.
- COMPASS-FME uses data across existing repositories: ARM, ESGF, ESS-DIVE, MSD-LIVE, NASA DAAC, USGS, NOAA, Great Lakes Observing System, NERR, NEON, AmeriFlux, LTERs, MarineGEO, the Smithsonian's Coastal Carbon Research Coordination Network (CCRCN) for coastal wetland carbon stocks, NOAA's National Estuarine Research Reserve System-Wide Monitoring Program Database, and others.
- COMPASS desires a secure method of interacting with commercial cloud storage providers (Dropbox, Google Drive, etc.), or a similar in-house service, to help ease transition into using more sustainable utilities when working with collaborators. Key features of this storage would be the ability for all collaborators—at national labs, research institutes, or universities—to access project-related files such as documents, presentations, and small (say, 1 GB or less) datasets.
- Federated authentication mechanisms among partner institutions would greatly ease data transfer and other collaboration efforts. This has been a particular problem in bringing the COMPASS-HPC computer online, as the registration and approval process was quite slow.
- The COMPASS compute resource will have an on-premises Kubernetes deployment to facilitate containerized software deployment. It also hosts an Open OnDemand web UI to provide Jupyter hosting and other software-as-a-service products. Other DOE HPC user facilities may be engaged depending on the needs of researchers and the scale of their work.
- There is a risk that some COMPASS datasets may exceed the current per-dataset capacity at commonly used affiliated data archives, such as ESGF and ESS-DIVE:
  - In these cases, the archives are used to index the data, for which the researchers must still locate and secure long-term storage.
  - Similar constraints may appear for HPSS as large amounts of data from the regional climate modeling for long-term historic simulations and future projections are generated.
- There is also an extremely high demand for COMPASS computational resources for regional models for long-term simulations and projections. One of the reasons is that the input/output (I/O) reading and writing take a good portion (20 to 30%) of the entire calculation time. This cost can be significantly reduced by splitting the model output on each processor. However, it also takes wall clock time (instead of computing allocations) to stitch these split files back to a file that can be processed and analyzed for science problems
- The majority of data-related COMPASS-FME issues surround data sharing and tracking:
  - Cloud services, like Google Drive, do not have enough flexibility with permission options and leave the project vulnerable to data loss. For example, users must be given edit access to upload data, but this also gives users permission to delete or change data as well.
  - It is desirable to track dataset usage and version history to ensure that people are aware of how up-to-date a dataset is.
  - Sharing dataset availability and contact information associated with a dataset to the team is challenging. Often, people do not know whom to contact about locating a specific dataset, due to difficulty in searching.

5.16.2 Coastal Efforts Facility Profile

The COMPASS project comprises two relatively independent but collaborative parts: a field study and a coastal modeling study. COMPASS-FME focuses on field studies and associated process and ecosystem modeling of two coastal interfaces. COMPASS-GLM focuses on modeling and analysis of coastal systems in the Great Lakes Region.

COMPASS-FME seeks to advance a scalable, predictive understanding of the fundamental biogeochemical processes, ecological structure, and ecosystem dynamics that distinguish coastal terrestrial-aquatic interfaces (TAIs) from the purely terrestrial or aquatic systems to which they are coupled. To achieve the COMPASS vision, FME will focus on overarching long-term science questions:

- What fundamental mechanisms control the structure, function, and evolution of coastal TAIs?
- How do these fundamental mechanisms interact across spatial scales, and what interactions are most important to improving predictive models?

The two-year COMPASS-GLM pilot study aims to enhance predictive understanding of freshwater coastal systems, especially how they respond to climate warming, LULCC, and other perturbations at watershed-to-regional scales. This includes local climate and weather feedbacks that affect wintertime lake effect snows and summertime convective storms. To achieve the COMPASS vision, GLM will focus on overarching long-term science questions:

- What multiscale mechanisms govern the structure, function, and dynamics of coastal systems at different spatial and temporal scales?
- How do coastal systems respond to natural and anthropogenic influences?
- Is it possible to generalize new process knowledge and predictive skill gained at a small number of sites or regions across the observed diversity of coastal systems?

5.16.2.1 Science Background

5.16.2.1.1 COMPASS-FME

Coastal Observations, Mechanisms, and Predictions Across Systems and Scales - Field, Measurements, and Experiments (COMPASS-FME) is an Environmental Systems Science pilot project. This multi-institutional project is led by PNNL, and partners with ORNL, ANL, LANL, and LBNL, as well as the Smithsonian Environmental Research Center, the University of Toledo, and Heidelberg University.

Our objective for the pilot study is a new predictive understanding of both saline and freshwater coastal terrestrial–aquatic interfaces (TAIs) that can be applied and evaluated at new sites, providing a foundation for broader representation of coastal TAIs within earth system models. A key outcome of the FME study is to develop a measurements- and mechanisms-informed coastal module for ecosystem biogeochemistry. Such a modeling framework — encompassing the interactions of water, soil, and plants — does not currently exist to bridge the land and water components of earth system models at coastal interfaces. This limits understanding of how coastal regions — which are home to critical ecosystems and an outsized fraction of the human population and play an outsize role in modulating biogeochemical fluxes — will respond to climate trends, shifts in land use, and other environmental changes.

We generate spatial and temporal data across seven different sites in the Chesapeake Bay and Western Lake Erie regions. These data provide broad spatial and temporal coverage of ecosystem processes. Much of the data is remotely logged, and collected through telemetry. This data is compiled for use by experimentalists and modelers to discern patterns of change along the study transects, and to guide model improvements and inform new hypotheses.

5.16.2.1.2 COMPASS-GLM

The two-year COMPASS-GLM pilot study aims to enhance predictive understanding of freshwater coastal systems, especially how they respond to climate warming, LULCC, and other perturbations at watershed-to-regional scales. This includes local climate and weather feedbacks that affect wintertime lake effect snows and summertime convective storms. During the pilot study, modeling activities will focus on two tasks: the first task is focused on coupled, integrated atmosphere–land–lake regional modeling in the Great Lakes Region, the second task will focus on high-resolution modeling of the Portage River watershed.

An important component of this integrated study is to share model output across tasks, so that, for example, it is possible to drive the high-resolution watershed model with output from the atmospheric model in order to understand how the watershed responds to individual events, and how that response might change in future climate conditions.

5.16.2.2 Collaborators

5.16.2.2.1 COMPASS-FME

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
Pacific Northwest National Laboratory - Marine and Coastal Research Laboratory (MCRL)	Primary	Google Drive	5 kb	Ad hoc	No	Need to store raw instrument data somewhere
Pacific Northwest National Laboratory - Richland	Primary	Google Drive	5 kb	Ad hoc	No	Need to store raw instrument data somewhere
Pacific Northwest National Laboratory - JGCRI	No	Google Drive	N/A	Ad hoc	N/A	Organization difficult on Google Drive, multiple copies of same data
University of Toledo	No	Google Drive	N/A	Ad hoc	N/A	NA
Smithsonian Environmental Research Center	Primary	Google Drive	20 mb / month	5-15 minute transfer intervals	No, mirrored from source	Google Drive permission issues
Argonne National Laboratory	No	Google Drive	N/A	Ad hoc	N/A	No
Los Alamos National Laboratory	No	Google Drive	N/A	Ad hoc	N/A	N/A
Heidelberg University	No	Google Drive	N/A	Ad hoc	N/A	N/A

Table 5.16.1: COMPASS-FME Collaboration Space

5.16.2.2.2 COMPASS-GLM

Model datasets will be shared between two multi-institutional task-level teams for atmosphere–land–lake regional modeling (Great Lakes Region) and high-resolution watershed modeling (Portage River Basin). With the objective of integrating research from regional to watershed scale, there is a need for quick and efficient sharing of model-data between two teams. In the table below, a list of institutions across which team members from two teams will be sharing model datasets.

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
Argonne National Laboratory (ANL)	Yes, all the datasets are backup on The HPSS on NERSC	Data transfer	10-100GB for the case study, 10days of data that is needed by ATS is 1.2GB. 10 days of raw WRF output is 15 GB.	Ad hoc	Yes, through shared directory on HPC. Can also use Globus	N/A
Oak Ridge National Laboratory (ORNL)	Yes	Data transfer	10-100GB	Ad hoc	Yes, through shared directory on HPC	N/A
Pacific Northwest National Laboratory (PNNL)	Yes	Data transfer	10-100GB	Ad hoc	Yes, through shared directory on HPC. Can also use Globus	Access to the PNNL COMPASS computer (e.g., via ssh) has limited access to computer resources.

Table 5.16.2: COMPASS-GLM Collaboration Space

As a case study, COMPASS-GLM considers sharing of the output datasets from an atmosphere–land–lake regional model at ANL for the Great Lakes Region to drive a high-resolution spatially distributed surface-subsurface integrated hydrology model at ORNL for the Portage river basin.

5.16.2.3 Use of Instruments and Facilities

5.16.2.3.1 COMPASS-FME

COMPASS-FME operates 7 real-time hierarchically designed environmental sensor clusters called synoptic nodes at sites along Chesapeake Bay and Lake Erie. Each node contains data logger infrastructure that can host a suite of environmental sensors for measuring ModEx derived meteorological physical parameters for biogeochemistry and carbon cycle modeling. Current instrumentation includes:

- Water Quality measurements via AquaTroll600 or EXO2 multiparameter sonde (up to 4 per site) - pH, conductivity, temperature, dissolved oxygen, chlorophyll content, turbidity to be added later (4 per node)
- Meteorological parameters via ClimaVUE met station- PPT, PAR, temperature, windspeed, pressure, (1 per node)

- Soil – Volumetric Water Content, Electrical Conductivity, Temperature at various depths (16 per node)
- Tree Health – Sap flow measurements for assessing tree response using Granier method (24 per node)
- Each node contains 3 Cr1000X data loggers and associated communication infrastructure including cell phone and radio telemetry as well as hardened solar power infrastructure. Each node produces about 5 Gb raw data per year; 40 Gb of processed data annually. All nodes share a common program and deployment architecture limiting operational burdens and allowing for site expansion.

Our network objective is to make data available to researchers and modelers through a collaborative network with the Smithsonian Environmental Research Center. The data workflow will push data from Smithsonian Institute to DOE? Operated networking utilities. COMPASS-FLM is in the process of developing data distribution infrastructure and data products. From initial work, several types of network issues are anticipated:

1. As an initial measure, COMPASS-FLM has been pushing data to Google Drive as more robust products are being developed. This experience has elucidated potential issues with this platform with respect to access control and data integrity.
   - There is no easy mechanism for non-enterprise account Google users to get official access to the system. Non-enterprise users (non-DOE institution) collaborators use personal Google accounts for access) making access control difficult and unverified.
   - Specific use access often has to be set at full edit access for researchers making data integrity hard to maintain without outside copies For sensor data, raw data is mirrored/controlled from outside source to control this issue.
   - Data changes are hard to track.
2. Our long-term data plan involves allowing R-based scripts to be run by research scientists on XY utilities. COMPASS-FLM needs to anticipate access or query from a variety of researcher-used tools such as Git, R-Studio, Python, and other high-level analysis computational tools. Because this is cross-institutional and collaborative research, outside researchers will need meaningful access to query the data lake is anticipated.

5.16.2.3.2 COMPASS-GLM

High-performance computing (HPC) facilities at NERSC are used for the Great Lakes modeling at regional and watershed scales. HPC clusters include Cori (set to retire in April 2023) and the newest supercomputer, Perlmutter, an HPE Cray EX supercomputer, features both GPU-accelerated and CPU-only nodes. Its projected performance is three to four times that of NERSC's current flagship system, Cori. The system was installed in two phases. Phase 1, which includes the system's GPU-accelerated nodes and scratch file system, has been available for early science campaigns since the summer of 2021. 2022's Phase 2 installation added CPU-only nodes. three state-of-the art modeling tools: Weather Research and Forecasting (WRF), Finite Volume Community Ocean Model (FVCOM), and Advanced Terrestrial Simulator (ATS) (Coon et al., 2019) are installed on the NERSC HPC clusters. Simulations are conducted using these models with multiple nodes on Cori and Perlmutter. All the data products from COMPASS regional-scale modeling are stored on and can be accessible from NERSC HPSS. While the ATS simulations so far have been performed on NERSC, a majority of the future simulations will be performed on COMPASS-HPC[88] a midrange high-performance computing facility at PNNL. Data transfer and sharing between researchers across scales (regional to watershed) use Globus, which allows efficient, secure, and reliable data transfer across all supercomputers at DOE HPC facilities as well as personal computers.

88 https://compass.pnnl.gov/Compass/COMPASSHPC

5.16.2.4 Process of Science

5.16.2.4.1 COMPASS-FME

Our project generates two broad categories of data: (1) observational and experimental data generated by short- and long-term field experiments, laboratory-based experiments, and routine monitoring efforts; and (2) simulation data generated from multiscale modeling activities. Modeling activities heavily rely on observational and experimental data to parameterize, calibrate, and validate numerical models. The experimental data created in this project include but are not limited to ‘omics measurements including genomic, transcriptomic, MS proteomic, and metabolomics (MS and nuclear magnetic resonance) data; a suite of ionic, stable isotopic, and -ray absorption spectroscopy (XAS) measurements to trace water sources and bulk geochemical properties of soils (e.g., particle size, water retention curves, redox state) and vegetation (e.g., exposure to ions); and quantitative characterization of dissolved organic matter via GC-MS and LC-MS-MS at MCRL. Continuous field observations include but are not limited to distributed, high-frequency (1–30 minutes) monitoring data of water level, meteorology, soil moisture, specific conductance, temperature, dissolved oxygen, redox potential, and tree stem sap flow across the Chesapeake Bay and Great Lakes regions.

Most of the model simulations are performed with the coastal midrange computing facility and DOE and PNNL supercomputers. While the raw model outputs can be massive, they are typically post-processed to produce results needed to address science questions and generate plots to be used in journal publications. There are multiple value-adding steps that can reduce the dataset size. Simulation data is usually well structured and can be reproduced given all the input files and specific simulator change set or version. Because all input files and settings will be archived, preserving the raw model outputs may not be necessary. The team makes heavy use of Jupyter Notebooks (https://jupyter.org/) to document computational workflows including data sources, codes used, and all associated metadata.

5.16.2.4.2 COMPASS-GLM

The COMPASS-GLM includes two domain scales in their modeling activities. At the Great Lakes regional scale, there is model output from regional climate model WRF standalone, hydrodynamic model FVCOM standalone, and fully coupled WRF-FVCOM. At the smaller watershed scale, there is model output from the Advanced Terrestrial Simulator (AT), a spatially resolved integrated hydrology and reactive transport code. The output of WRF-FVCOM over the land is at 4 km spatial resolution and hourly temporal resolution. Selected variables from WRF-FVCOM output are being used as input for ATS, including precipitation, air temperature, wind speed, solar radiation, and relative humidity (calculated based on surface pressure, air temperature and water vapor). This transfer of data from WRF to drive ATS is the topic of the case study, as it is representative of other model coupling activities.

The output of the regional climate model — fully coupled WRF-FVCOM includes many 3D variables in addition to the near surface variables needed for ATS forcing. The watershed modeling team has been downloading select outputs of WRF model from NERSC, transforming data in ATS-consistent coordinate reference system as needed to create ATS forcings and running ATS model. While the format of raw output of WRF-FVCOM is NetCDF, and that of ATS is HDF5, they can be easily converted to any other format depending on the purpose of the use of the data. COMPASS-GLM has post-processed the model output and retrieved the most frequently used variables and made them ready to share within the project team. On the other hand, COMPASS-GLM also backs up the raw output in NERSC HPSS so it can be used for other scientific studies as well that may require more variables that are less frequently used. Thus the general workflow is to subset, remap, then convert the selected data from WRF used to drive ATS.

5.16.2.5 Remote Science Activities

5.16.2.5.1 COMPASS-FME

We use data across existing repositories (ARM, ESGF, ESS-DIVE, MSD-LIVE, NASA DAAC's, USGS, NOAA, Great Lakes Observing System, NERR, NEON, AmeriFlux, LTERs, MarineGEO, the Smithsonian's CCRCN for coastal wetland carbon stocks, NOAA's National Estuarine Research Reserve System-Wide Monitoring Program Database of nationwide coastal wetland sites and water quality data, etc.). The research will leverage RS data products that provide spatial distribution of plant functional types, aboveground biomass, algal blooms, and water chemistry available from NASA, NOAA, and USGS. COMPASS-FME will start to assemble a metadatabase of relevant datasets focusing on biogeochemistry, ecology, and hydrology likely to be necessary and/or valuable for the overall coastal work proposed here. The metadatabase will benefit the broader coastal community.

5.16.2.5.2 COMPASS-GLM

The raw output from WRF-FVCOM for a day (with 24 hours of output) is about 1.5Gb. Such output includes more than 20 variables at near surface and vertical profiles. Not all of them are needed by ATS. Therefore, COMPASS-GLM retrieved the needed variables (air temperature, wind speed, relative humidity, solar radiation, and precipitation) for the event and the days before the event. Ten days of these variables at hourly temporal resolution take ~1.2GB. The raw data for 10 days takes 15 GB. These model simulations are generated on Cori NERSC, and are shared with team members with access to Cori. COMPASS-FLM can also use Globus (global.app.org) to share data with users. Globus would take a few minutes to transfer the data with such size. Processing of WRF-FVCOM output datasets to generate ATS inputs can be done either on a local machine or Cori/Perlmutter at NERSC itself. In this case study, datasets were downloaded onto a local machine using rsync. The generated ATS input files were uploaded back on Cori/Perlmutter NERSC for ATS-ISSHM simulations. The output files from the transient daily integrated hydrology simulations in ATS for the Portage basin (~1000 km2) for 15 years of simulation are about 150GB.

5.16.2.6 Software Infrastructure

5.16.2.6.1 COMPASS-FME

- PostgreSQL Database, NoSQL/Mongo database for reporting purposes,
- R Shiny for data model simulations/runs,
- multithreaded file-transfer system for large files, multiprocessor to prevent system-locks when one very large job is taking place,
- "Round Robin" file management node which can handle multiple connections from outside (e.g., redis)

5.16.2.6.2 COMPASS-GLM

- Globus to move data from different HPC systems like NERSC and COMPASS Computer.
- HPSS for data archive and long-term storage.
- Python libraries for generation of ATS input files from WRF-FVCOM using Jupyter Notebooks and for geospatial analysis
- Python-based Watershed Workflow modules for setting up ATS model – generation of mesh, synthesizing soil, geology, land cover, and other data products for high-resolution ATS simulation

- Software programs and third-party libraries supporting ATS installation[89]
- Libraries that can handle NetCDF data processing, such as NCO command, and NCAR Command language (NCL), which has functions to read, process, and visualize NetCDF (e.g., WRF output) data.

5.16.2.7 Additional Network and Data Architecture Requirements

Most major data transfer operations, both incoming and outgoing, make use of the Globus data traffic software. Smaller data transfers may use other protocols such as SFTP or https.

A secure method of interacting with commercial cloud storage providers (Dropbox, Google Drive, etc.) would ease transition into using more sustainable utilities.

Federated authentication mechanisms among partner institutions would greatly ease data transfer and other collaboration efforts.

5.16.2.8 Use of Cloud Services

5.16.2.8.1 COMPASS-FME

There will be some AWS deployment for visualization software.

Google Drive is used for small-scale dataset storage and propagation. Google Docs is used for collaborative editing.

The COMPASS compute resource will have an on-premises Kubernetes deployment to facilitate containerized software deployment. It will also host an Open On-Demand web UI to provide Jupyter hosting and other software-as-a-service products.

Other DoE HPC user facilities may be engaged depending on the needs of researchers and the scale of their work.

5.16.2.8.2 COMPASS-GLM

NERSC HPSS is where COMPASS-GLM will store the model output as well as other data products, such as processed observation data including in situ, satellite, and buoy over lakes.

Argonne certificated BOX is used for GLM regional modeling team's collaborative writing (e.g., scientific manuscripts, reports), material collection, sharing and co-editing (e.g., group meeting slides). Similarly, ORNL certified Microsoft SharePoint is used for GLM watershed hydrology modeling team's co-editing manuscripts and presentations. The GLM team also uses Slack, Google Drive, and other ad-hoc resources, sometimes creating some confusion about where certain datasets or documents are.

5.16.2.9 Data-Related Resource Constraints

5.16.2.9.1 COMPASS-FME

There is a risk that some datasets may exceed the current per-dataset capacity at commonly-used affiliated data archives, such as ESGF and ESS-DIVE. In these cases the archives are used to index the data, for which the researchers must still locate and secure long-term storage.

89 https://github.com/amanzi/amanzi/blob/master/INSTALL_ATS.md

5.16.2.9.2 COMPASS-GLM

In addition to the constraints listed by FME, similar constraints may appear for HPSS as COMPASS-FLM generates a large amount of data from the regional climate modeling for long-term historic simulations and future projections. Meanwhile, there is also an extremely high demand for computational resources for regional models for long-term simulations and projections. One of the reasons is that the input/output (I/O) reading and writing take a good portion (20-30%) of the entire calculation time. This cost can be significantly reduced by splitting the model output on each processor. However it also takes wall clock time (instead of computing allocations) to stitch these splitted files back to a file that can be processed and analyzed for science problems.

5.16.2.10 Outstanding Issues

5.16.2.10.1 COMPASS-FME

The majority of data-related COMPASS-FME issues surround data sharing and tracking. First, Google Drive does not have enough flexibility with permission options and leaves us vulnerable to data loss. For example, everyone must be given edit access to upload data, but this also gives users permission to delete or change data as well. Additionally, because Google Drive is currently used for project documents and data, data often gets lost or is put into sub-folders. COMPASS-FME anticipates that the COMPASS-HPC will alleviate some of these issues. Second, COMPASS-FME would like to be able to track dataset usage and version history to ensure that people are aware of how up-to-date a dataset is. Lastly, there are issues sharing dataset availability and contact information associated with a dataset to the team. Often, people do not know whom to contact about a dataset since they are difficult to find on the drive.

5.16.2.10.2 COMPASS-GLM

For larger modeling related datasets and outputs, spaces on HPC systems (NERSC, COMPASS-HPC), as mentioned in the previous sections, have been effective. However, for documents like meeting notes, presentation slides from science talks and conferences, which are of interest to the whole GLM team, there is not a common tool supported by all institutions. For example, Argonne and ORNL use BOX and SharePoint, respectively, for internal document sharing. These files cannot be made available to collaborators from other institutions. Some project-wide documents are on Google Drive; however, not all labs support Google platform for official purposes. A common document sharing platform supported/certified by all DOE laboratories will be helpful for collaborative work and potentially improve productivity.

5.16.2.11 Facility Profile Contributors

Coastal Efforts Representation

FME

Vanessa Bailey, *PNNL,* vanessa.bailey@pnnl.gov

Xingyuan Chen, *PNNL,* Xingyuan.Chen@pnnl.gov

Stephanie Pennington, *PNNL,* stephanie.pennington@pnnl.gov

Kurt Maier, *PNNL,* khm@pnnl.gov

Roy Rich, *Smithsonian Environmental Research Center,* richr@si.edu

Michael Degan, *PNNL,* michael.degen@pnnl.gov

GLM

Robert Hetland, *PNNL,* robert.hetland@pnnl.gov

Saubhagya Rathore, *ORNL,* rathoress@ornl.gov

Jiali Wang, *ANL,* jialiwang@anl.gov

ESCC Representation

Surya Singh, *PNNL,* surya.singh@pnnl.gov

Lewis Allen, *PNNL,* lewis.allen@pnnl.gov

Dylan Merrigan, *LANL,* dylan@lanl.gov

5.17 CASCADE

Risks of extreme weather events pose some of the greatest hazards to society and the environment as the climate system changes due to anthropogenic (i.e., human-caused) warming. If the severity of extreme climate events continues to increase, this will constitute one of the most stressing forms of change for society and the environment.

The CASCADE project advances the nation's ability to identify and project climate extremes and how they are affected by environmental drivers. CASCADE will focus on four main research areas:

- Description, detection, and designation of extremes in data and models
- Understanding and simulating the physical mechanisms that drive variability and change in the spatiotemporal characteristics of extreme events.
- High-performance computing to detect and predict changes in weather extremes.
- Experiments to explore how uncertainty in observations and models affect the understanding of extremes

5.17.1 Discussion Summary

- The CASCADE project looks to identify and project climate extremes and how they are affected by environmental drivers. CASCADE will focus on four main research areas:
  - Description, detection, and designation of extremes in data and models.
  - Understanding and simulating the physical mechanisms that drive variability and change in the spatiotemporal characteristics of extreme events.
  - High-performance computing to detect and predict changes in weather extremes.
  - Experiments to explore how uncertainty in observations and models affects the understanding of extremes.
- The CASCADE project is collaborative work at LBNL, the University of California, Berkeley, and University of California, Davis campuses, Indiana University, and Iowa State University. CASCADE also coordinates with:
  - Earth system modeling efforts.
  - Land, ocean, and atmosphere diagnostics projects.
  - Stakeholder-driven science projects and scientific communities, including CMIP, SAMSI, and ARTIP.
- CASCADE is designed to address three major and interrelated scientific questions:
  - How have changes in the physical behavior of the coupled system altered the chances of encountering and the nature of extreme climate events?
  - To what degree are anthropogenic drivers responsible for altering the statistics, properties, and projections of extreme climate events?
  - What are the dominant uncertainties in detecting, attributing, and modeling extremes, how do these uncertainties affect the first two questions, and how can uncertainties be reduced?
- Climate model output is collected from several sources, each of which may result in TB datasets being retrieved using Globus primarily:
  - The Program in Climate Model Diagnostics and Intercomparison (PCMDI).

- The Rutherford Appleton Laboratory in the UK.
- The ESGF.

- CACADE leverages 14 PB of HPSS storage at NERSC, and routinely does analysis using containers that operate Jupyter:
  - 5 PB is devoted to retrospective simulations of the 20th century with and without effects of anthropogenic climate change.
  - 1.5 PB is devoted to Model Intercomparison Projects (MIPs) from the most recent Coupled Model Intercomparison Project (CMIP6).
- None of the core CASCADE partner facilities provide mechanisms to analyze the data "server side," so it is necessary to use networks to transfer the data to local systems (e.g., Cori, etc.) where data can be operated on.
- It would be desirable to have the ability for moderate-scale HPC centers to remotely mount NERSC filesystems. This, combined with some sort of ability to preemptively cache data locally at the HPC sites before a job makes use of the data, would enable a distributed super-facility model to support the CASCADE workflow.
- The CASCADE SFA is unlikely to make use of cloud services any time soon, largely because of costs required to host the magnitudes of data that can be operated on.
- CASCADE creates and collects climate model simulations of the past, present, and future climate. Analysis of these simulations is typically focused on understanding how the statistics and properties of extremes could change in a warmer climate. Supercomputers at NERSC, OLCF, ALCF, IU (Big Red 200), and PNNL (Compy) are used to produce climate model simulations and analyze climate model outputs. Storage allocations at these facilities vary, but are typically terabytes to petabytes in size. CMIP7 will see a four to five times increase in data volumes, when related.
- DOE LCFs prioritize large-scale computation over moderate-scale analysis. Thus, the CASCADE team is increasingly using medium-scale computational resources external to the LCFs when doing moderate-scale data analysis tasks:
  - This results in repeated data transfers (e.g., to/from NERSC, or from external data repositories to external HPC systems like IU UITS described previously).
  - These transfers need to be done manually, and they often take O(1) days even when using transfer tools, like Globus, that nearly achieve 1 GBPs transfer speeds.
  - Given that large-scale computations—like simulations—will increasingly be done on LCFs, with the data stored/archived at those LCFs, this will likely continue to be a bottleneck.
- Server-side analysis is desirable for CASCADE computation, but moderate-scale analyses (e.g., using O(10)-O(100) nodes) have historically had very long queue waiting times at NERSC, when compared with small jobs (O(1) node) or large jobs (O(1,000) nodes) with similar wall clock time requirements. These wait times can be upwards of several days, which is less than the network transfer time of the datasets.

5.17.2 CASCADE Facility Profile

CASCADE SFA is a coordinated research program, funded under the DOE's SC, that aims to improve the understanding of climate extremes while accounting for a wide variety of sources of uncertainty in the drivers, mechanisms, and statistics of these critical phenomena.

The CASCADE project is a multidivisional, collaborative work at LBNL, drawing upon expertise of scientists in the lab's CRD and Climate and Ecosystem Sciences Division as well as the University of California, Berkeley, and UCD campuses. CASCADE scientists collaborate with related projects at LBNL and across BER's climate modeling efforts. These projects include Earth system modeling efforts; land, ocean, and atmosphere diagnostics projects; and stakeholder-driven science projects. The resulting connections and related projects ensure tight integration of observations, experiments, and modeling of extreme climate events. CASCADE is also active in national and international scientific communities, including CMIP, SAMSI, ARTIP, etc.

CASCADE is designed to address three major and interrelated scientific questions:

1. How have changes in the physical behavior of the coupled system altered the chances of encountering and the nature of extreme climate events?
2. To what degree are anthropogenic drivers responsible for altering the statistics, properties, and projections of extreme climate events?
3. What are the dominant uncertainties in detecting, attributing, and modeling extremes, how do these uncertainties affect the first two questions, and how can these uncertainties be reduced?

5.17.2.1 Science Background

The project creates and collects climate model simulations of the past, present, and future climate. Analysis of these simulations is typically focused on understanding how the statistics and properties of extremes could change in a warmer climate.

CASCADE is led out of LBNL with key partners and co-investigators at the University of California campuses at Berkeley and Davis, Indiana University (IU), Iowa State University, and the Program for Climate Model Diagnosis and Intercomparison (PCMDI) at LLNL. Five members of the CASCADE team are employed by, or affiliated with, IU, which hosts an HPE Cray EX system (Big Red 200[90]) along with several other high-performance computing systems. Use of these systems is free for these five members of CASCADE, and so these team members frequently use these systems for CASCADE-related research.

We are currently using 14 PB of HPSS storage for this purpose. Greater than 5 PB is devoted to retrospective simulations of the 20th century with and without effects of anthropogenic climate change. Another 1.5 PB is devoted to MIPs from the most recent Coupled Model Intercomparison Project (CMIP6) as described below. All the data is in the public domain with no rights retained. CASCADE operates under a data management plan following DOE standards for the SC (available on request).

Climate Model Output is collected from several sources, including (a. the Program in Climate Model Diagnostics and Intercomparison (PCMDI); (b. the Rutherford Appleton Laboratory in the UK); and (c. the ESGF). Much of the data used by CASCADE originates in multimodel ensembles constructed for successive rounds of the CMIP [Touzé-Peiffer et al, 2020[91]]. CMIP provides climate projections that support climate science worldwide, decision and policy-makers communities, in its objective to understand past, present and future climate changes. CMIP and its associated data infrastructure have become essential to the IPCC and other international and national climate assessments. The simulations from CMIP6 (Eyring et al., 2016[92]) are an integral input to the research conducted by CASCADE, and planning for CMIP7 by the World Climate Research Programme is now actively underway. The interval between CMIPs[93] is approximately 6 to 7 years, and CMIP7 will be underway by the latter part of this decade.

90 https://kb.iu.edu/d/brcc

91 Touzé-Peiffer, L, Barberousse, A, Le Treut, H. The Coupled Model Intercomparison Project: History, uses, and structural effects on climate research. WIREs Clim Change. 2020; 11:e648. https://doi.org/10.1002/wcc.648.

92 Eyring, V., Bony, S., Meehl, G. A., Senior, C. A., Stevens, B., Stouffer, R. J., and Taylor, K. E.: Overview of the Coupled Model Intercomparison Project Phase 6 (CMIP6) experimental design and organization, Geosci. Model Dev., 9, 1937–1958, https://doi.org/10.5194/gmd-9-1937-2016, 2016.

93 https://www.wcrp-climate.org/wgcm-cmip

Modern climate model data conforms to a set of detailed, prescriptive, and homogeneous standards[94]. As a result, CASCADE does not need to engage in transformations of the data to use it effectively – climate models from around the world conform to this standard.

Most analytical tools are constructed on the Python and R platforms. The CASCADE packages and toolkits have been developed via Jupyter Notebooks and extensions to R are described on the CASCADE website[95]

5.17.2.2 Collaborators

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
PCMDI (LLNL) Livermore, CA	Secondary	Data transfer	??	Ad hoc (activity driven by release of new CMIP data)	N	No
ESGF nodes[95] and International contributors[96].	Primary	Data transfer using Globus and wget	50 TB	Ad hoc (same as above)	N	Yes – several of the international servers connected to ESGF have exceptionally slow connections (e.g., the Chinese Academy of Sciences systems)
Rutherford Appleton Laboratory (RAL) (UK) Chilton, Oxfordshire, England We use this server to access PRIMAVERA[97] simulations from the EU	Primary	Data transfer using Globus	50 TB	Ad hoc (same as above)	N	No (RAL provided a very HT connection upon request – several 100 megabit per second.)
Indiana University/ University Information Technology Services (UITS) Supercomputers for academic research at IU[98]	Secondary	Data transfer using Globus	10TB	Ad hoc	Y, sometimes - by Globus	Transferred data reside on the UITS slate-scratch filesystem, which has strict data purging policies (data is deleted after 30 days). Data transfers are frequently repeated b/c of this.

Table 5.16.2: COMPASS-GLM Collaboration Space

94 These standards are described at https://pcmdi.llnl.gov/CMIP6/Guide/dataUsers.html.
95 https://cascade.lbl.gov/software_data/
96 https://esgf.llnl.gov/nodes.html
97 https://www.google.com/maps/d/u/0/viewer?mid=1cPCfAqvlJTcYPmuOfWju-WMnRZOC14wx&ll=33.92683654370459%2C-20.713510749999955&z=2
98 https://www.primavera-h2020.eu/modelling/our-simulations/
99 https://kb.iu.edu/d/alde

5.17.2.3 Use of Instruments and Facilities

Supercomputers at NERSC, OLCF, ALCF, IU, and the "Compy" system at PNNL to produce climate model simulations and analyze climate model outputs. CASCADE uses the storage and networking capabilities collocated with the supercomputers at these facilities to store the output from these simulations and to transfer it to other analytical facilities.

The output consists of files of typically a few GB in size – fortunately very, very few violate the upper bound of 68 GB on files that can be archived with htar. Most of the files confirm to the CMIP output standards (§5.17.2.1), and hence contain a 2D or 3D field as a time series for a single variable prognosed or diagnosed by the climate model.

In order to illustrate the scale of a typical dataset, CASCADE will discuss the holdings from the High-Resolution Model Intercomparison Project (HighResMIP; Haarsma et al., 2016[100]), on the 23 approved MIPs[101] in CMIP6 (Eyring et al., 2016). The total dataset size is 937 TB and is distributed across 8147 directories at the lowest level in the directory hierarchy containing 1044653 files.

As discussed in (§5.17.2.1), CASCADE is in a period between the conclusion of CMIP6 and the commencement of CMIP7. During this period, modeling centers are focused on the development of the next generations of their earth system models, and appreciably fewer computing resources will be devoted to creating large multimodel ensembles (MMEs) over the next two years (2023-2025). Therefore, the rate of growth of climate model data has slowed appreciably and will remain relatively low until the MMEs are started for CMIP7.

Over the next five years, it is reasonable to assume that the earth system models now under development will feature doubled model horizontal resolution relative to versions of these same models employed in CMIP6. If no other parameters of the CMIP MME change, this implies a growth in the volume of CMIP7 by 2 ^ 2 or roughly a factor of 4 to 5.

Considering the timeframe beyond five years for the purposes of strategic planning, it appears highly likely that the data volumes for CMIP7 and subsequent CMIPs will grow to hundreds of petabytes and may exceed an exabyte in under a decade. This will put climate model output at parity with comparable amounts of satellite and radar data in the public domain collected by American and European national scientific and meteorological agencies (Overpeck et al., 2011[102]).

5.17.2.4 Process of Science

None of the major data storage repositories in the CASCADE research area have server-side analytic capabilities, so the scientific workflow relies on fast networks to transfer the data to local systems (e.g., Perlmutter, etc.).

For a recent example of a sophisticated use of external data, see "A framework for detection and attribution of regional precipitation change: Application to the United States historical record"[103], in particular Figure 1 showing how the CASCADE team used models to test multiple hypotheses and the supplemental file[104] showing the model runs that were analyzed. The process of testing multiple hypotheses involved transferring over 50,000 years of climate model data from 6 different experiments and 12 different climate models–which exist across a federated system of data servers called the Earth System Grid Federation[105]– to NERSC filesystems.

The data were then reduced using custom-written R code designed by CASCADE team members to statistically model the spatiotemporal extreme value statistics of the simulated data. In other cases, which often

100 Haarsma, R. J., et al.: High Resolution Model Intercomparison Project (HighResMIP v1.0) for CMIP6, Geosci. Model Dev., 9, 4185–4208, https://doi.org/10.5194/gmd-9-4185-2016, 2016.

101 https://www.wcrp-climate.org/modelling-wgcm-mip-catalogue/modelling-wgcm-cmip6-endorsed-mips

102 Overpeck, JT, et al. 2011. Climate data challenges in the 21st century. Science, 331(6018): 700–702. DOI: https://doi.org/10.1126/science.1197869

103 https://link.springer.com/article/10.1007/s00382-022-06321-1

104 https://static-content.springer.com/esm/art%3A10.1007%2Fs00382-022-06321-1/MediaObjects/382_2022_6321_MOESM1_ESM.pdf

105 https://esgf.llnl.gov/

also rely on fast networks to transfer data to local systems, data is reduced either using custom-written code (often written in Python) or using existing tools developed by the CASCADE team[106], such as the Toolkit for Extreme Climate analysis.

Over the next two years, the CASCADE team will continue to port data analysis codes to work on GPU-enabled systems. Over the next five years, the tools, techniques, and workflows will evolve to accommodate analyses on diverse, high-resolution grids used in Global Cloud Resolving Models (GCRMs). Beyond that timeframe, a fundamental paradigm shift in the tools, techniques, and workflows will occur when the resolution becomes high enough that a single two-dimensional slice of model output (e.g., data on a latitude-longitude grid) becomes too large to fit in memory or data needed for a data-reduction process no longer fit in temporary storage (e.g., scratch) space, necessitating network transfers to be an integral part of the data analysis process.

5.17.2.5 Remote Science Activities

We currently access and share large climate datasets with the ESGF, the Program for Climate Model Diagnosis and Intercomparison (PCMDI), and the National Center for Atmospheric Research (NCAR). Over the near term (two to five years) CASCADE does not anticipate any major changes in these collaborations. Over the longer term (> five years) the nodes within ESGF may and probably will change. However, due to the persistent need to provide backward-compatible access to data from prior CMIPs, CASCADE anticipates that ESGF is a stable federation on this timescale.

We currently generate and stage data for analysis using high-performance computing capabilities provided by a number of institutions such as NERSC, the OLCF, Argonne Leadership Computing Facility (ALCF), Indiana University (IU), and PNNL. CASCADE anticipates that these collaborations will continue and potentially expand over near (two to five years) and longer term(> five years) time scales.

5.17.2.6 Software Infrastructure

The software required by the team could be classified into three groups: data generation, data management, and data analytics. The data management software is used to move data in between facilities and remote centers as well as between various file systems on site. Tools such as Globus, SCP, wget, curl, and cp are typical. In a typical data-generation scenario the team will make use of numerical weather prediction systems such as WRF to recreate an extreme event for the purposes of a detailed study. In the typical data analysis workflow the team will make use of the usual suite of data analytics software such as Python, Numpy, Jupyter, NetCDF, and HDF5. The team develops and makes use of specialized parallel software called Toolkit for Extreme Climate Analysis (TECA) based on CUDA, C++ threads, and Message Passing Interface (MPI) in order to process data on supercomputing facilities. Increasingly CASCADE relies on GPU computing technologies such as CUDA, HIP, and OpenMP. Both the data generation and data analytics use cases rely on data being staged locally at a center's internal HPSSs. The data management use cases rely heavily on internal and external network performance. While the specific software may change over time CASCADE does not anticipate major changes in the "stage and process" pattern of use.

5.17.2.7 Additional Network and Data Architecture Requirements

At the present time, CASCADE scientific workflows involve single-shot, large transfers between storage repositories and compute centers (e.g., NERSC); requirements are not particularly specialized. The current paradigm centers on the fact that the data analyzed by the CASCADE team (e.g., climate model outputs) are small enough in magnitude that logical units of data analysis, such as data from one or a few variables of data from a single climate model simulation, can at least fit in temporary storage if not permanent server-local storage.

106 https://cascade.lbl.gov/software_data/

If dataset volumes become so large–e.g., because of increases in spatial resolution–that data required for a reduction process cannot fit in temporary storage, then a fundamental change in the algorithms might be needed, necessitating workflows that rely integrally on network transfers.

5.17.2.8 Use of Cloud Services

The CASCADE SFA is unlikely to make use of cloud services any time in the near future, largely because of costs required to host the magnitudes of data that CASCADE operates on.

5.17.2.9 Data-Related Resource Constraints

Because DOE LCFs prioritize large-scale computation over moderate-scale analysis, the team is increasingly utilizing medium-scale computational resources external to the LCFs when doing moderate-scale data analysis tasks; this results in repeated data transfers (e.g., to/from NERSC, or from external data repositories to external HPC systems like IU UITS described above). These transfers need to be done manually, and they often take O(1) days even when using transfer tools, like Globus, that nearly achieve 1 GBPs transfer speeds. Given that large-scale computations–like simulations–will increasingly be done on LCFs, with the data stored/archived at those LCFs, this will likely continue to be a bottleneck.

Server-side analysis would be a nice solution, but moderate-scale analyses (e.g., using O(10)-O(100) nodes) have historically had very long queue waiting times at NERSC, when compared with small jobs (O(1) node) or large jobs (O(1000) nodes) with similar wall clock time requirements. These wait times can be upwards of several days, which is less than the network transfer time described above.

If there were some sort of ability for moderate-scale HPC centers to remotely mount NERSC filesystems—combined with some sort of ability to preemptively cache data locally at the HPC sites before a job makes use of the data—that would enable the sort of 'distributed super-facility' that has been discussed for a number of years.

5.17.2.10 Outstanding Issues

None to report at this time.

5.17.2.11 Facility Profile Contributors

CASCADE Representation

William D. Collins, *Lawrence Berkeley National Lab and UC Berkeley,* wdcollins@lbl.gov

Travis A. O'Brien, *Indiana University and Lawrence Berkeley National Lab, obrienta@iu.edu*

ESCC Representation

Rune Stromsness, *LBNL,* rstrom@lbl.gov

Richard Simon, *LBNL,* rsimon@lbl.gov

5.18 CATALYST

CATALYST represents a sustained commitment by the DOE and the University Corporation for Atmospheric Research (UCAR) to perform foundational research toward advancing a robust understanding of modes of variability and change using models, observations and process studies.

5.18.1 Discussion Summary

- CATALYST is joint work between the DOE and the University Corporation for Atmospheric Research (UCAR) to perform research toward understanding modes of variability and change using models, observations, and process studies by using:
  - The DOE E3SM, the NCAR Community Earth System Model (CESM), and a hierarchy of simpler models.
  - Coupled Model Intercomparison Project (CMIP) multimodel datasets.
  - ML methods.
  - Numerous observational datasets.
- CATALYST features four interrelated research objectives (ROs):
  - Address the limits of predictability of modes of variability using Earth system models (E3SM and the Community Earth System Model) and ML methods.
  - Use of a hierarchy of models to address relevant processes and feedbacks related to modes of variability and how they interact with each other.
  - Benchmark model representations, examine the role of changes, and address the likelihood and predictability of tipping points and irreversible changes.
  - Investigate the relationships between the modes of variability and high-impact events (e.g., flash droughts and precipitation extremes, atmospheric rivers, tropical cyclones).
- Simulations are carried out on a number of supercomputers, including the NCAR/University of Wyoming, Cheyenne, and NERSC's Cori and Perlmutter, and others:
  - Outputs from these simulations are often transferred between the computing sites for analysis, depending on specific aspects of the simulations involved.
  - The preferred tool for data transfer is Globus.
- The typical process for the use by CATALYST of models for knowledge discovery involves the following steps:
  - Experimental design is created.
  - Resources (computing, storage, and so on) are determined, applied for, and allocated at a computing center.
  - Model is executed.
  - The output is analyzed and archived.
  - Output is made available via the CDG or ESGF as appropriate using data mobility tools and portal systems.
  - Results are published.
- The most significant change to current CATALYST practices will be reengineering of the CESM workflow, to enable the creation of single-field time series format data as the simulation

is ongoing. This will enable the global user community to have easier and more efficient access to CESM results.

- CATALYST and CESM will continue to rely on Globus, the ESGF and CDG, and their follow-on projects to publish and deliver model output to the user community.
- HPC facilities involved in the model creation have petabytes of storage available for the input to the model generation process, as well as to store output for later sharing:
  - Simulation output generated on Cheyenne can range in size from tens of megabytes to many terabytes, depending on the specific configuration of the model.
  - Large-scale simulations can generate data collections consisting of 100,000s of files and 1,000–2,000 PB of output.
  - Output is shared with well-established community portals (e.g., Climate Data Guide [CDG] and ESGF).
- The global climate model output analysis community (currently thought to be in the hundreds) will retrieve model output from well-established community portals (e.g., CDG and ESGF). It is expected that these downloads could range from megabytes to terabytes in size, depending on the breadth of model that is retrieved.
- CATALYST does acknowledge a perpetual "shortage" of disk space. This can be attributed to model output increases, but also needing a better strategy to manage available disk resources across project partners.
- CATALYST is not currently using cloud services for model output storage and distribution. However, the CESM project has engaged in early and preliminary studies of the utility of AWS:
  - Some older datasets (CESM1 Large Ensemble, for example) have been "Zarr-ified" (e.g., compressed using the Zarr tool) and placed on AWS as tests.
  - Analysis of the usability of AWS-hosted data for the user community is ongoing.
  - The cloud may become an important resource in the future, but CATALYST will continue to rely on CDG and ESGF for the distribution of model output.

5.18.2 CATALYST Facility Profile

CATALYST performs foundational coordinated research in a team-oriented collaborative effort aimed at advancing a robust understanding of modes of Earth system variability and change using models, observations and process studies.

CATALYST performs foundational coordinated research in a team-oriented, collaborative environment. It aims to advance a robust understanding of the modes of Earth system variability and change by utilizing the DOE E3SM, the NCAR CESM, CMIP multimodel datasets, a hierarchy of simpler models, ML methods, and numerous observational datasets. CATALYST research is centered around modes of variability (MOV) on subseasonal to decadal timescales.

5.18.2.1 Science Background

Four ROs address key science questions that focus on modes of variability. RO1 sets the groundwork for the rest of the ROs and addresses the limits of predictability of modes of variability on subseasonal to decadal timescales using Earth system models (E3SM and the Community Earth System Model) and ML methods. To inform those predictions, RO2 uses a hierarchy of models to address relevant processes and feedback related to modes of variability and how they interact with each other. RO3 then benchmarks model representations of modes of variability in Earth system models, examines the role of external forcing in changes of modes of variability and their interactions, and addresses the likelihood and predictability of tipping points and irreversible changes. RO4 then investigates the relationships between the modes of variability, which are addressed in the other ROs, and high-impact events (e.g., flash droughts and precipitation extremes, atmospheric rivers, tropical cyclones) and how these might change in future climate.

CATALYST research uses the E3SM, along with the Community Earth System Model, Coupled Model Intercomparison Project multimodel datasets, a hierarchy of simpler models, ML methods, and numerous observational datasets.

Under the auspices of CATALYST, simulations are carried out on a number of supercomputers, including the NCAR/University of Wyoming Cheyenne, and NERSC's Cori and (soon) Perlmutter, and others.

Output from these simulations are often transferred between the computing sites for analysis, depending on specific aspects of the simulations involved. The preferred tool for data transfer is Globus.

5.18.2.2 Collaborators

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
Global climate model output analysis community (100s of people)	Secondary	Data portals (CDG and ESGF)	From megabytes to terabytes	Ad hoc; continuously	No	Not really – much improved compared to a few years ago

Table 5.18.1: CATALYST Collaboration Space

5.18.2.3 Use of Instruments and Facilities

Site	HPC name and type	Peak performance	Processors	Memory	Disk storage	Archival capacity
NCAR-Wyoming Super-computing Center (NWSC)	Cheyenne, SGI ICE XA Cluster (through 2023)	5.34 Petaflops	145,152	313 TB	38 PB	90 PB
	Derecho, HPE Cray EX cluster (online end 2023)	19.87 Petaflops	323,712	692 TB	~80 PB	~200 PB

Table 5.18.2: Simulation output generated on Cheyenne can range in size from tens of megabytes to many terabytes, depending on the specific configuration of the model. Large-scale simulations can generate data collections consisting of 100,000s of files and 1,000–2,000 PB of output.

5.18.2.4 Process of Science

The typical process for the use by CATALYST of the CESM for knowledge discovery involves an experimental design created by either an individual scientist, small NCAR group of scientists, or one of the CESM Working Groups (a collection of scientists and others with a common interest). Once the design is finalized and the necessary resources (computing, storage, and so on) are determined, the project applies for those resources at the computing center. Once those resources are allocated, then the model is executed at the center, the output is analyzed and archived, made available via the CDG or ESGF as appropriate, and papers are written and submitted to various science journals detailing what was learned from the experiments.

The same basic process for model simulations conducted at NCAR is accomplished at other computing centers. Experiments are designed, resources allocated, the simulations run, post-processed, analyzed and made available via the ESGF and CDG. Some of the original model output and post-processed data from these simulations may be transferred to NCAR, but because all of the CATALYST computing resources are associated with nodes in ESGF, it is not necessary to transfer all of the data just for the purpose of making them publicly available.

5.18.2.5 Remote Science Activities

Site	HPC name and type	Peak performance	Processors	Memory	Disk storage	Archival capacity
National Energy Research Super-computing Center (NERSC)	Cori, Cray XC40 (to end 2022)	30 Petaflops	2,388 Intel Xeon "Haswell" nodes, 9,688 Intel Xeon Phi KNL nodes	1 PB	40 PB	100+ PB
	Perlmutter, HPE Cray EX cluster (fully online late 2022)	3x-4x Cori	1536 nodes GPUs, 3072 nodes CPUs	2+ PB	~80 PB	100+ PB

Table 5.18.3: NERSC Capabilities

5.18.2.6 Software Infrastructure

The most significant change to current CATALYST practices will be the nearly complete reengineering of the CESM workflow, to enable the creation of single-field time-series format data as the simulation is ongoing. This will enable the global user community to have easier and more efficient access to CESM results. The CATALYST and CESM will continue to rely on Globus, the ESGF and CDG and their follow-on projects to publish and deliver model output to the user community. Other projects may be incorporated into the ESGF to enable data format changes (to other binary formats, from netCDF to GIS-compatible formats, for example) and the ability to extract, subset, and additionally process the model results. Whatever tools Earth Systems Grid and/or ESGF make available will be exploited by the CATALYST.

5.18.2.7 Additional Network and Data Architecture Requirements

CATALYST has not had any major issues with data architecture, other than a perpetual "shortage" of disk space, which only means that CATALYST's resources should be better managed.

5.18.2.8 Use of Cloud Services

CATALYST is not currently using cloud services for model output storage and distribution, however, the CESM project has engaged in early and preliminary studies of the utility of AWS. Some older datasets (CESM1 Large Ensemble, for example) have been "Zarr-ified" and placed on AWS as tests. Analysis of the usability of AWS-hosted data for the user community is ongoing.

The cloud may become an important resource in the future, but CATALYST will continue to rely on CDG and ESGF for the distribution of model output.

5.18.2.9 Data-Related Resource Constraints

As mentioned above, other than the usual space "shortage", CATALYST has not been faced with any network-related constraints. Globus has been working extremely well for bulk transfers between NCAR and NERSC, and CATALYST has been provided with sufficient data resources (as designed in CATALYST's computing proposals) so at the current time, we are not facing major problems.

5.18.2.10 Outstanding Issues

None to report at this time.

5.18.2.11 Facility Profile Contributors

CATALYST Representation

Gerald Meehl, *National Center for Atmospheric Research (PI),* meehl@ucar.edu

Jadwiga (Yaga) Richter, *National Center for Atmospheric Research (Co-PI),* jrichter@ucar.edu

Nan Rosenbloom, *National Center for Atmospheric Research (Program Manager),* nanr@ucar.edu

Gary Strand, *National Center for Atmospheric Research,* strandwg@ucar.edu

5.19 Multifacility Workflows in BER and Relationship to the IRI Initiative

Dedicated research communities often gather around DOE User Facilities and, over years, refine and evolve them as scientific understanding deepens and expands. While each experimental or observational User Facility is tailored to the pursuit of scientific discovery within the domain of a particular SC Program, their distinct workflows share common patterns that remain constant across the years. These common patterns offer opportunities for general solutions that enhance efficiency, avoid redundancy, and strengthen resilience of existing workflows across all DOE SC User Facilities. They also reveal nascent potential to create entirely new workflows by bringing together different User Facilities in harmonious operation toward a shared scientific objective. This case study discusses some of the ways BER facilities collaborate with each other, with ASCR HPC facilities, and across other DOE SC programs, while relying on ESnet as a common provider of data mobility solutions. These observations also compliment the ongoing efforts of the DOE SC IRI initiative.

5.19.1 Discussion Summary

- Workflows that require coupling capabilities across facilities are becoming more common within BER :
    - The FICUS effort is one notable example, but it is expected that the frequency and depth of this type of multi-analytic couplings will increase over time.
    - Several workflows that depend on data replication across sites to ensure that data is available are at the HPC facilities. This involves coordination of data movement and job management. The JGI and NMDC are using tools like JAWS to handle data movement across sites to facilitate these needs for data replication workflows.
- Data transmission amounts between coupled BER facilities vary from megabytes and gigabytes when dealing with the sharing of raw sensor or observational data, to multiple terabytes to petabyte scales when moving large models between HPC resources. Tools like Globus are critical to managing the data mobility aspects, along with other methods such as SCP, HTTP portals, and the use of cloud sharing and storage.
- JAWS is a workflow service designed to move the JGI away from bespoke and/or hand coded workflows, and to protect the production analysis pipelines from outages and/or downtimes at partner HPC facilities.
    - This has enabled the JGI to move some of its analysis runs off NERSC, and into clusters such as LBNL IT's Lawrencium Cluster and EMSL's (PNNL) Tahoma cluster, as well as private cloud clusters in AWS.
    - It is expected that most of the JAWS workflows will operate within LBNL, split between NERSC and a JGI-dedicated cluster hosted within LBNL IT, with only a small fraction going off site to PNNL and AWS.
    - As a consequence, ESnet may not see much traffic between LBNL and other sites due to JAWS. The general trajectory of data growth will be tied to the growth of JAMO and NMDC data.
- Some examples of coupled facility models within BER are as follows:
    - The JGI has a number of strategic user facility partners that operate in a multifacility workflow paradigm.
        - The JGI uses JAWS, Nextflow, Snakemake, and some additional systems for workflow management.
        - JAWS is the only resource able to submit work across distributed computing facilities.

        - Globus is leveraged for data movement.
        - NERSC is the primary repository for more than 14 PB of active and archival data.
    - KBase enables systems biology analyses, and works with the JGI to facilitate seamless access to data to facilitate modeling and analysis. KBase and the JGI partner through data-oriented APIs, with JGI services at LBNL and KBase production services at ANL.
    - The NMDC and JGI, through the JAWS effort, have partnered to use the services for NMDC analysis workflows. Initially, JAWS will handle NMDC workflows that operate at NERSC, but as the collaboration expands, other clusters with NMDC allocations will be integrated into JAWS. The NMDC links data and metadata from the same samples generated at the JGI and EMSL facilities.
    - The EMSL and JGI are partners in the FICUS program, which provides researchers with access to the world-class resources of multiple user facilities through a single proposal.
- The challenges with enabling BER multifacility workflows are significant. Most are not bracketed by technology but are sociological, policy, or related to deeply scientific constraints.
- It is anticipated that multifacility efforts will be expanded to include DOE BES Light-source facilities and other non-BER user facilities:
    - In addition, there are already many examples of BER facilities coupling with computational capabilities provided by ASCR. This connectivity should only increase over the coming years driven by demands from the scientific community.
    - Efforts are underway to make information flow across the facilities more seamless, such as those to standardize sample metadata and methods of exchange. It is anticipated that BER will place even more emphasis on this in the coming years.
    - It is envisioned that an ecosystem will exist, beyond the five-year time frame, where common interfaces are defined across facilities that will enable scientists to more easily take advantage of the instruments and computational capabilities across the DOE-SC complex in a seamless manner.
- Multifacility BER workflows can take several forms:
    - In one model, a researcher has a common physical sample that must be sub-sampled, and sent for measurements at different facilities using different capabilities (e.g., sequencing, MS, imaging). These results are then synthesized by the researcher to gain a better understanding of structure, function, dynamics, etc.
    - Similarly, a researcher might use the capabilities at multiple facilities to understand mechanistic behavior across space and time, at a larger study level.
    - In another model, staff at the facilities interact directly to improve the process of science for the community, for example by the transfer of data between facilities (e.g., provide genomics data to use in the analysis of proteomics data), or by providing surge capacity with computing resources. In each of these cases, computational and data capabilities are critical.
- Many BER user facilities are already designed around a remote model:
    - Users send physical samples to facilities for analyses and access results over data portals.
    - Multifacility models will expand on these by requiring coordination and information flow between facilities.

- Facilities already have a mature set of tools for data management and providing remote data access. The main challenge in enabling multifacility workflows in the coming years will be improving how these systems connect to each other and exchange information. Data movement is in many ways the most easily addressed given the current capabilities of tools like Globus. The major challenges will be in developing standards around metadata, data exchange, authentication and authorization, policies, and tooling to support all of these areas
- Most of the challenges for multifacility BER workflows are not at the networking layer, but the network could enable some models:
  - For example, the ability to create overlay networks across ESnet between facilities could allow distributed resources to be coupled together more securely.
  - HTTP-based communications play an important role, especially in the context of pulling datasets from existing data repositories, reference databases, etc., and bringing them on to compute systems where workflows are run.
  - Allowing for high-performance data-movement capabilities natively over standard protocols like HTTP, instead of through proprietary tools and protocols. Tools like Globus are very useful but ultimately require a fair amount of custom configuration for each site involved, and advanced capabilities require site licenses.
  - Hybrid commercial cloud and DOE HPC facility workflows, and the ability to enable a virtual network to facilitate orchestration of resources.
- It is expected that demands will increase for sensor networks with artificial intelligence (AI capabilities) as well as self-driving labs capabilities for BER use cases. To meet the needs of rapidly expanding field data from DOE Earth Systems sensors and instruments, it will be critical to enable data acquisition, transformation, and analytics workflows that can integrate real-time data streams with predictive models. Live data streams from sensors can then be used to train models using a hybrid of AI/ML and deterministic techniques, which can, in turn, generate real-time predictions and feedback for sensors out in the field.
- It is very likely that clouds over various forms will play a role in future BER multifacility workflows:
  - This could take the form of surging to commercial clouds to meet targeted turnaround times for analysis, hosting public datasets, or other web services.
  - This may also include cloud-like models provided by a future integrated research infrastructure.
  - Maintaining a replicated instance of data services on cloud resources for failover, redundancy, or backup should be considered to enable broader access, especially to resources/users outside the DOE complex.
  - Enabling analysis capabilities in the cloud for very large datasets already stored in the cloud (bringing compute to the data).

5.19.2 Multifacility Workflows in BER Profile

The BER research infrastructure includes observational and experimental User Facilities, each designed to advance mission-science. The three ASCR DOE SC high-performance computing User Facilities, and the ESnet high-performance research network User Facility, serve all SC Programs and the global research community. ESnet interconnects DOE SC research infrastructure through seamless exchange of scientific data.

Dedicated research communities gather around all DOE User Facilities and, over years, refine and evolve them as scientific understanding deepens and expands. While each experimental or observational User Facility is tailored to the pursuit of scientific discovery within the domain of a particular SC Program, their distinct workflows share common patterns that remain constant across the years. These common patterns offer opportunities for general solutions that enhance efficiency, avoid redundancy, and strengthen resilience of existing workflows across all DOE SC User Facilities. They also reveal nascent potential to create entirely new workflows by bringing together different User Facilities in harmonious operation toward a shared scientific objective.

The DOE SC IRI initiative, started in 2020, embraces the strategic imperative of seamless multi-Facility workflows. The IRI is helping to establish a new paradigm within which DOE SC User Facilities and their research communities work together to improve existing capabilities and create new possibilities by building bridges across traditional silos. The IRI 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. Participants discovered three common patterns:

- The **Time-Sensitive** pattern is characteristic of low-latency workflows requiring real-time, or near-real-time, response across more than one Facility or resource. An example is HPC analysis of data received from an ongoing experiment and used to redirect that experiment in real-time.
- The **Data Integration-Intensive** pattern is characteristic of workflows combining diverse datasets to deepen and expand context. An example is combined analysis of data received from more than one Facility or resource, such as experiment and simulation, while tracking metadata and provenance to ensure reproducible scientific results.
- The **Long-Term Campaign** pattern is characteristic of workflows requiring sustained access to more than one Facility or resource, at scale, for years or decades. An example is very large Facilities or resources designed, built, and operated over years and decades by expansive international scientific collaborations.

BER case studies that exhibit many of these patterns are inherently tied to a workflows that spans multiple Facilities. They often require high throughput connection of experimental Facilities with HPC platforms. These facilities are typically designed to capture, compute, store, categorize, and share data from a wide range of experiments, sophisticated data management systems, and use multi-modal data from a variety of different data sources, such as simulations, experiments, and observations. Examples include:

- Observational data acquired globally using diverse sensors and aggregated within a number of collaborating facilities
- Research that uses real time simulation and analysis of experimental data to steer running experiments toward the most fruitful avenues of exploration, potentially using AI/ML
- Multi-decadal studies that create intense and valuable models that expand over time

This case study is a prime example of how the future of BER can be adapted to the work being performed by IRI. In particular, when reviewing aspects of the BER use case studies through the IRI criteria, several core observations can be made:

- There are limited time-sensitive use cases within BER, and most of the time these are limited to facilities that are trying to perform high-throughput experiments (e.g., numerous "runs" for samples, and requiring a way to share the raw data with processing and storage facilities). Most experiments and facilities are not capable of doing real-time analysis currently, due to sensor locations being remote or there not being a pressing need to see instantaneous results to influence a following run.

- BER operates on multi-modal data constantly, and as a result most of the facilities and experiments are intrinsically data-intensive. This can be seen for some facilities that are designed to capture, store, categorize, and share results from a wide range of experiments, as well as long-running observational efforts that capture results from 100s of instruments and sensors.
- BER data, from observations, sequences, or historic models, remains useful over time. A number of facilities are being developed to help better categorize and serve this information, and many long-running efforts report use of historic models and datasets over time.

5.19.2.1 Science Background

The key feature in this cross-cutting case study is how these scientists will couple capabilities across facilities to enable new discoveries. This is already happening with programs like FICUS but it is expected that the frequency and depth of these couplings will increase over time.

Additionally, there are several workflows that depend on data replication across sites to ensure that data is available at the HPC facilities - this involves coordination of data movement and job management. JGI and NMDC are using (or planning to use) tools like JAWS to handle data movement across sites to facilitate this.

5.19.2.2 Collaborators

User/ Collaboration and location	Do they store a primary or secondary copy of the data?	Data access method, such as data portal, data transfer, portable hard drive, or other? (please describe "other")	Avg. size of dataset? (report in bytes, e.g., 125GB)	Frequency of data transfer or download? (e.g., ad hoc, daily, weekly, monthly)	Is data sent back to the source? (y/n) If so, how?	Any known issues with data sharing (e.g., difficult tools, slow network)?
JGI	Primary	Data Portals, Globus	15PB	Daily	N	N
EMSL	Primary	Portals, Globus, FTP	2 MB – 2 GB	Ad hoc	N	N/A
KBase	Primary & Secondary	Portals, Globus	Total ~1 PB	Ad hoc 10 GB to 1 TB per week	N	KBase provides tooling for sharing
NMDC	Primary	Portals, Globus (future)	N/A (pilot)	TBD (pilot)	TBD	TBD
ESS-Dive	Primary & Secondary	Portals, Globus	Total - 20 TB (expected to grow to 1 PB)	Daily	N	Replication to DataONE and NCEAS - peering with NSF. Potential large file transfers limited to 50GB
BRCs	Varies	Varies	Varies	Ad hoc	Varies	
Individual SFAs	Varies	Varies	Varies	Ad hoc	Varies	
General User Community	Varies	Varies	Varies	Ad hoc	Varies	

Table 5.19.1: Multifacility Workflows in BER Collaboration Space

The JGI: JGI has a number of strategic user facility partners that operate in a multifacility workflow paradigm. JGI uses the JAWS, Nextflow, Snakemake, and some additional systems for workflow management. JAWS is the only resource that is able to submit work across distributed computing facilities. Globus is leveraged for data movement. NERSC is the primary repository for more than 14 PB of active and archival data.

DOE System Biology Knowledgebase (KBase): KBase is a critical partner in enabling systems biology analyses, such as metabolic modeling, that go well beyond genome annotation and analysis. The JGI and KBase seek to enable users to seamlessly access and analyze data on scales ranging from single genes and individual genomes to metagenomes to systems-level modeling and understanding. JGI and KBase are designated DOE SC Public Reusable Research (PuRe) Data Resources are authoritative providers of data or capabilities in their respective subject area. KBase and JGI partner through data oriented APIs, with JGI services at LBNL and KBase production services at ANL. Data volume for these APIs are currently several orders of magnitude smaller than the backend services for JAMO and JAWS, and are not reported as a major integration with JGI in this document.

NMDC: In 2017, a stakeholder workshop involving JGI leadership and hosted by the American Society for Microbiology brought together representatives from academia, industry, government, and philanthropic funding agencies to conceptualize the NMDC. These efforts serve as the foundation of a community-driven national effort aimed to develop standards, processes, and infrastructure for an integrated microbiome data ecosystem. NMDC and the JAWS team have partnered to use JAWS services for NMDC analysis workflows, initially JAWS will handle NMDC workflows that operate at NERSC. In the future, as the collaboration expands, other clusters with NMDC allocations will be integrated into JAWS. The NMDC links data and metadata from the same samples generated at the JGI and EMSL facilities.

EMSL: EMSL collaborates with their user community to understand the role of molecular processes in controlling the function of biological and ecological systems across spatial and temporal scales. Through collaborative programs like FICUS, EMSL provides capabilities in multimodal molecular measurements, data analytics, and production computing, combined with the capabilities of partner facilities. EMSL also shares resources for data storage with JGI and NMDC, providing resiliency for JGI data access and serving as a distributed data source for NMDC. EMSL is the primary repository for data generated within the user facility, which currently totals over 33 PB of active and archived data.

FICUS: The FICUS program was created in 2014 to accelerate ambitious user research projects requesting resources from JGI and EMSL. FICUS represents a unique opportunity for researchers to bring to bear the capabilities of more than one DOE user facility on a research project, and generate datasets beyond what each of these facilities could generate by itself. The FICUS program has expanded and continues to expand to include the NEON, the Bio-SANS beamline, the APS, the ALS, and the ARM user facilities, as well as engaging data resources such as KBase, ESS-Dive, and NMDC.

The Environmental System Science Data Infrastructure for a Virtual Ecosystem (ESS-DIVE): is a data repository for Earth and environmental sciences. ESS-DIVE collects, stores, manages, and shares Earth and environmental systems data created through research sponsored by the US DOE. The volume, complexity, and diversity of Earth and environmental science data make information difficult to capture, store, verify, analyze, and share. ESS-DIVE is a consumer of data generated from workflows running at various DOE facilities (including the FICUS facilities). Additionally ESS-DIVE will also be working to enable data from sensor-network driven workflows (see Self-Driving Field Laboratories section in ESS-DIVE case study).

5.19.2.3 Use of Instruments and Facilities

The FICUS program, which involves sequencing at JGI and Omics at EMSL, has been brought to bear on particular projects in a multifacility fashion. It is anticipated that these types of multifacility efforts will expand to include DOE Light-source facilities and other non-BER User facilities. In addition, there are already many examples of BER facilities coupling with computational capabilities provided by ASCR. This connectivity should only increase over the coming years driven by demands from the scientific community. There are efforts underway to make information flow across the facilities more seamless such as efforts to standardize sample metadata and methods of exchange. It is anticipated that BER will place even more emphasis on this in the coming years. Beyond five years it is envisioned that an ecosystem where there are common interfaces defined across facilities that will enable scientists to more easily take advantage of the instruments and computational capabilities across the DOE-SC complex in a seamless manner.

Beyond five years it is expected that sensor Networks with AI capabilities as well as self-driving labs will be more common. To meet the needs of rapidly expanding field data from DOE Earth Systems sensors and instruments, it will be critical to enable data acquisition, transformation, and analytics workflows that can integrate real-time data streams with predictive models. Live data streams from sensors can then be used to train models using a hybrid of AI/ML and deterministic techniques, which can, in turn, generate real-time predictions and feedback for sensors out in the field.

5.19.2.4 Process of Science

Multifacility workflows can take several forms. In one model, a researcher has a common physical sample that they subsample and send for measurements in different facilities using different capabilities (e.g., sequencing, multiple omics, imagings). These results are then synthesized by the researcher to gain a better understanding of structure, function, dynamics, etc. Similarly, a researcher might use the capabilities at multiple facilities to understand mechanistic behavior across space and time, at a larger study level. In another model, staff at the facilities interact directly to improve the process of science for the community, for example by transfer of data between facilities (e.g., provide genomics data to use in the analysis of proteomics data), or by providing surge capacity with computing resources. In each of these cases, computational and data capabilities are critical, both to drive initial data analysis, integrating across datasets, and making data available through the community. Forms of this already occur today but it is expected these models will increase dramatically over the next 5 to 10 years.

5.19.2.4.1 JGI Analysis Workflow Service

JAWS is a workflow service designed to move JGI away from bespoke and/or hand coded workflows, and to protect the production analysis pipelines from outages and/or downtimes at partner HPC facilities. This has enabled JGI to move some of its analysis runs off NERSC, and into clusters such as LBNL IT's Lawrencium Cluster, EMSL's (PNNL) Tahoma cluster, as well as private cloud clusters in Amazon Web Services.

Depending on the services supported, the central dispatcher (JAWS Central) and the service that resides with the local batch scheduler (JAWS Site) use either Globus (GridFTP) or AWS S3 for transporting source and output data, as well as reference data.

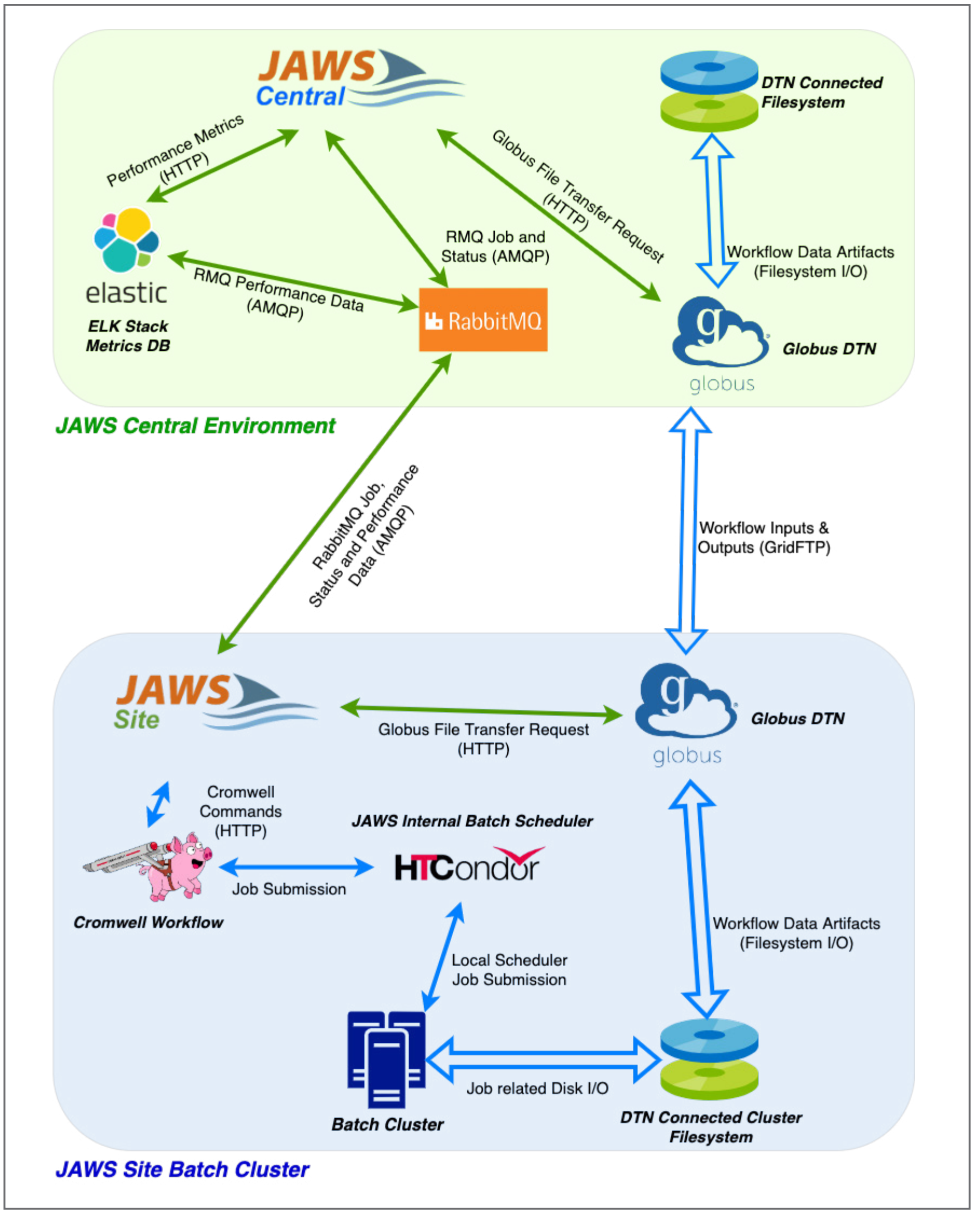

Figure 5.19.2: Example of JAWS Central and JAWS Site services and protocols

It is expected that most of the JAWS workflows will operate within LBNL, split between NERSC and a JGI-dedicated cluster hosted within LBNL IT, with only a small fraction going off site to PNNL and AWS. As a consequence, ESnet may not see much traffic between LBNL and other sites due to JAWS. The general trajectory of data growth will be tied to the growth of JAMO data + NMDC Data.

Discussions with EMSL have indicated that they are interested in creating a local instance of JAWS at PNNL for EMSL. This is especially useful to EMSL because NMDC has consolidated their multisite, distributed workflow into JAWS and NMDC researchers could then use a JAWS instance for running local and remote jobs, while EMSL researchers not in NMDC could reuse the services for local jobs. There is also interest in developing an EMSL JAMO site, so that data management could be integrated between EMSL and JGI. Those efforts are still at an early stage.

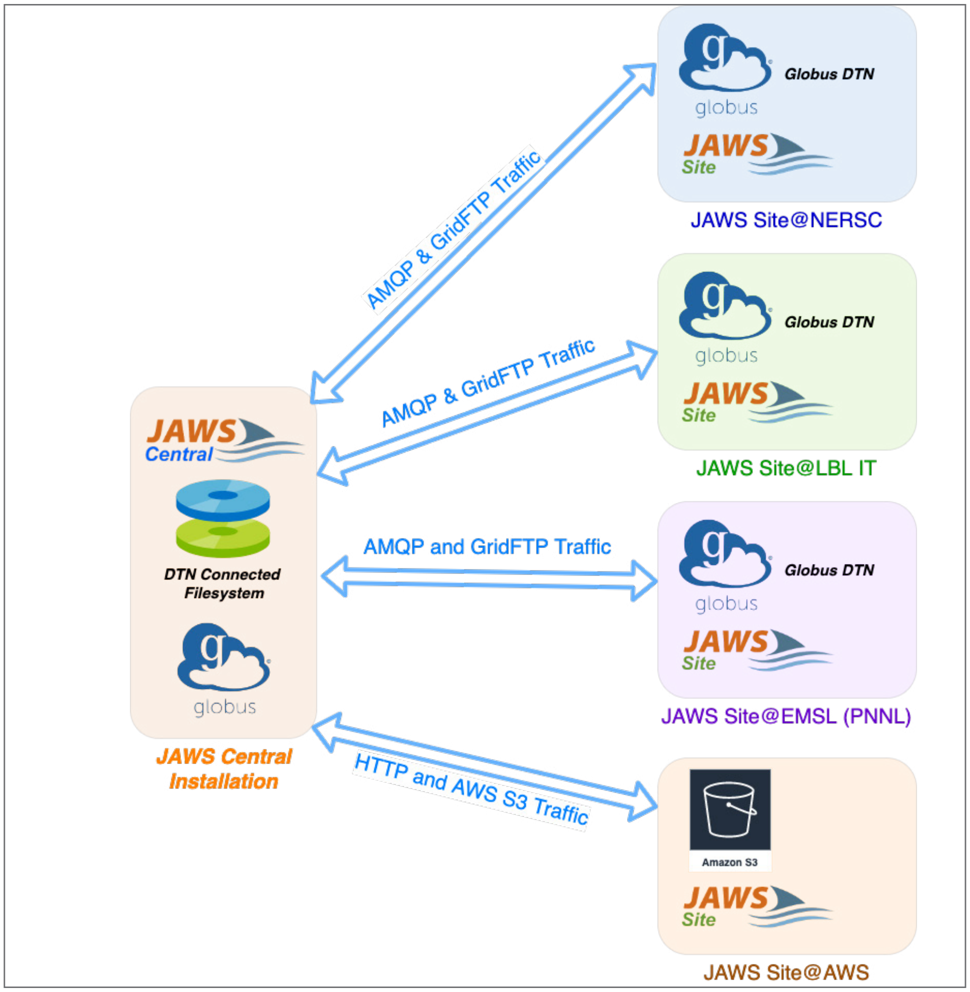

Figure 5.19.3: Current JAWS Site and Protocol Interactions

5.19.2.5 Remote Science Activities

Many of the BER user facilities are already designed around a remote model. Users send physical samples to facilities to be measured and access results over data portals. Multifacility models only expand on these by requiring coordination and information flow between facilities.

An example of an internal LBNL multifacility workflow would be JAWS. Another example of a multifacility workflow is NMDC, which spans multiple DOE laboratories, overlapping with LBNL. Through the overlap in LBNL staff with NMDC, it became clear that JAWS would be a good solution for NMDC's cross site workflows and that consolidating efforts would be the best use of finite resources. As a consequence, the JAWS team engaged with NMDC staff to configure JAWS to support NMDC workflows at LBNL, as well as planning for the expansion of NMDC workflows to EMSL. This is a prime example of DOE user facilities consolidating their multisite data and compute infrastructures to reduce redundant work and incompatibility.

Another remote science activity that is in early stages of discussion is the possibility of having an EMSL extension of the JGI JAMO service. This would provide EMSL with the metadata management services of JGI, consolidating data artifacts and metadata standards across 2 DOE sites. This service would also support ingress/egress of data between EMSL and LBNL using Globus, as well as automatically distributing archived data between the EMSL and NERSC HPSS instances.

5.19.2.6 Software Infrastructure

Individual facilities each have a mature set of tools for data management and providing remote data access. The main challenge in enabling multifacility workflows in the coming years will be improving how these systems connect to each other and exchange information. Data movement is in many ways the most easily addressed given the current capabilities of tools like Globus. The major challenges will be around developing standards around metadata, data exchange, authentication and authorization, policies, and tooling to support all of these areas.

5.19.2.7 Additional Network and Data Architecture Requirements

Most of the challenges are not at the networking layer but the network could enable some models. For example, the ability to create overlay networks across ESnet between facilities could allow distributed resources to be coupled together more securely.

HTTP-based communications play an important role, especially in the context of pulling datasets from existing data repositories, reference databases etc. and bringing them on to compute systems where workflows are run.

Beyond Five Years

High-performance data-movement capabilities natively over standard protocols like HTTP. Tools like Globus are very useful but ultimately require a fair amount of custom configuration for each site involved, and advanced capabilities require site licenses.

Hybrid Cloud / Facility workflows and the ability to enable a virtual network across these workflows to facilitate orchestration.

5.19.2.8 Use of Cloud Services

Currently the JGI JAWS service supports compute clusters in AWS that scale out on-demand for specific workflows that require the reliability of AWS services. Additionally there are small teams within JGI that start/stop their stacks in AWS to run specific analysis - some of these groups are consolidating their work into JAWS to minimize redundant effort. Some of the managed services that JGI depends on (such as Google Mail/Calendar/Apps, Cloudfront and the LBNL Laboratory Information System) are hosted in the cloud by the service provider.

In the future, JGI may move some of the core JAWS infrastructure into a cloud provider to maximize availability of services such as JAWS Central. A broad category of cloud usage is to spin up "on demand" cloud analysis platforms that do not merit ongoing, permanent installations within the JGI infrastructure. In addition, GCP is a major partner of LBNL, and it may make economic sense to migrate cloud-based compute clusters to GCP (or other cloud vendor). A cloud service that has high-speed peering with ESnet, and lower egress fees would be a compelling option for JAWS (and JGI in general).

5.19.2.9 Data-Related Resource Constraints

A broad category of Data-Related Resource Constraints is the constant need to replicate data among different sites. Analysis tools often require reference data in order to function, as well as a general need for project teams to have their data readily accessible for local access, ironically this problem had been solved in the 1980's by distributed filesystems with a homogenous global namespace (Andrew File System) albeit with performance far below HPC standards. Using a hodgepodge of tools such as rsync and Globus, multisite collaborations cobble together cross site replication using open source and free tools, however this is a deeply inefficient piecemeal solution to a pervasive problem - especially when the same data is replicated to multiple sites, with each site subject to local disk space constraints. Collaborations such as the teams around the Large Hadron Collider have also built their own distributed data management platforms, but these services are generally operated solely for the members of that collaboration and a significant administrative undertaking for smaller collaborations to take on.

Another constraint that arises constantly is negotiating firewall exceptions for services. There are a large collection of services that are required for most workflows, each of which typically requires a firewall exception to be negotiated individually with site cybersecurity teams. In some cases, these exceptions may take weeks to work through. A method that the JAWS team has been using to avoid multiple firewall exceptions for different services has been to use message queueing protocols as the transport for API requests - a single message queueing system is reused for multiple API services. This works well when a collaboration has the ability to deeply modify their software stack, however it cannot be used in all (or even most) situations. Having some form of flexible routing between trusted services at different sites would significantly simplify the effort of operated multisite services.

5.19.2.10 Outstanding Issues

The challenges with enabling multifacility workflows are significant and most are not bracketed by technology but are sociological, policy, deeply scientific, etc.

5.19.2.11 Facility Profile Contributors

Multifacility Workflows in BER Representation

Kjiersten Fagnan, *LBNL,* kmfagnan@lbl.gov

Nigel Mouncey, *LBNL,* nmouncey@lbl.gov

Georg Rath, *LBNL,* gbrath@lbl.gov

Stephen Chan, *LBNL,* sychan@lbl.gov

Boris Sadkhin, *ANL,* bsadkhin@anl.gov

Adam Arkin, *LBNL,* aparkin@lbl.gov

Shane Canon, *LBNL,* SCanon@lbl.gov

Roy Kamimura, *LBNL,* royk@lbl.gov

Emiley Eloe-Fadrosh, *LBNL,* eaeloefadrosh@lbl.gov

Shreyas Cholia, *LBNL,* scholia@lbl.gov

Lee Ann McCue, *PNNL,* leeann.mccue@pnnl.gov

Douglas Mans, *PNNL,* douglas.mans@pnnl.gov

Deb Agarwal, *LBNL,* daagarwal@lbl.gov

Charuleka Varadharajan, *LBNL,* cvaradharajan@lbl.gov

ESCC Representation

Linda Winkler, *ANL,* winkler@mcs.anl.gov

Corey Hall, *ANL,* chall@anl.gov

Rune Stromsness, *LBNL,* rstrom@lbl.gov

Richard Simon, *LBNL,* rsimon@lbl.gov

Surya Singh, *PNNL,* surya.singh@pnnl.gov

Lewis Allen, *PNNL,* lewis.allen@pnnl.gov

6 Focus Groups

A core component of the ESnet requirements review process displaced by the COVID-19 pandemic was the opportunity to hold impromptu conversations with colleagues. These could occur during the oral case study review period (and involve topics being presented or stumbled upon), but were also equally likely to occur before, during, or after the physical meeting. The importance of these interactions cannot be overstated, as they often resulted in cross pollination of ideas, collaboration, or other forms of interaction fostered by the organization of the attendees and subject matter. Facilitating these types of interactions was a high priority, despite the challenges of conducting a fully distributed review process.

6.1 Purpose and Structure

In October and November 2022, the BER requirements review team convened two virtual focus groups. The general plan for these meetings was to:

- Gather small groups of case study authors during predefined time periods, using virtual tools.
- Prepare the groups by having them review outlines of their case studies and research focus areas (if they were unfamiliar).
- Structure a conversation to review areas of research, and then seed discussion with a set of topics found to be common across all case studies in the 2022 BER requirements review.

During these focus sessions, the BER requirements review team acted as a moderator for the conversation, but let discussion flow organically toward topics of mutual interest. The goals were to:

- Allow emerging projects and facilities to ask questions of the established BER community, to better prepare for the future.
- Facilitate discussion on known problems and solutions that will guide the process of science, and support from ethnology, in the coming years.
- Establish BER practices that span the different parts of the BER program area.

6.2 Organization

The BER requirements review featured 19 case study groups. The optimal way to organize focus groups was to offer two events, and invite all parties that were available to attend; this organizational assumption acknowledged the fact that not all participants could attend both. Similar discussion topics were available at each event, but the chosen topics could differ drastically depending on participation. The events were as follows:

- Focus group 1 was held on October 27, 2022.
- Focus group 2 was held on November 1, 2022.

The agenda for each event was designed to be simple and dedicated to keeping a majority of the event available for attendee discussion. Each began with a brief introduction from the BER requirements review team and an overview of each meeting's purpose. The remainder of the time was allocated to discussion topics. These were defined prior to the meeting (and shared with attendees) by the requirements review team. All topic areas were pulled directly from observations made by case study authors. The topics were as follows:

1. **Challenges in Multi/Coupled Facility Workflows:** Are there complications in pairing BER experimentation with ASCR computing and storage via ESnet?

2. **Facility Upgrades and Changes to Data Volume and Rate:** How are users modifying workflows to prepare, and how can networks be adjusted to fit the new requirements?
3. **Data Storage Locality, Quantity, and Mobility:** Storage needs are outpacing capabilities to generate and analyze BER data. What is being done, and are there suggestions to improve the national data architecture?
4. **Computation:** Real-Time, Near Real-Time, or Offline: discussing any use case that requires computation during the process of experimentation and simulation. Focus can be on limitations to the current practice, and thoughts on how to improve.
5. **Software:** analysis, sharing, simulation, etc.: current, near-term, and future needs and development opportunities that may leverage network resources.
6. **Local Networking Support:** discussing ways that networking (either at a DOE facility or your local institution) has been challenging for your process of science. What can ESnet do to help?

A piece of polling software was used to gauge the relative interest in each topic area during the meeting. This was done to gain an understanding of what mattered to those who were represented in the room. The interest could be based on things they wanted to hear more about (potentially from other attendees), things they were concerned with implementing, or things they felt they could share experience with. Each focus group came to different conclusions about what topics mattered most, and as a result, each focus group's conversation flowed more naturally toward the strengths and weaknesses of those who attended.

6.3 Outcomes

The following sections highlight the areas of discussion and relevant findings and recommendations that emerged during the talks. Some are directly related to the structured conversation, but others came out of natural discussions that may have strayed from the topic areas. During this period of discussion, several notable items were brought up. These are detailed in the following subsections.

6.3.1 Data Storage Locality, Quantity, and Mobility

A common theme from the ESnet Requirements Review Program is the observation that most research groups have learned to effectively use networks, particularly as capacities have increased on backbones, regionals, and campuses, but that storage remains a critical bottleneck. This takes the form of two observations:

- Storage is a constant struggle to allocate, and it may not be located close to other parts of the workflow; this requires data movement several times.
- Data-movement tools are much better, but still require a level of tuning on some institutional and laboratory resources that interact with large filesystems.
- Not everyone has access to DTN resources by default (even those these are prevalent at most, if not all DOE-SC facilities), and may not know to ask how to get access.

Several participants noted that when moving information from institution to institution, observed speeds are high and ease of use has improved incredibly over the past several years. There still exists a "last foot" problem, though, where some tools are better run on laptops, and moving larger datasets can still be a challenge. Globus has simplified this from an application layer perspective, but there is nothing that can be done when running over wireless connections. Adapting workflows to use DTNs to move data "closer" is really the best mechanism, followed by streaming or bulk downloading to a laptop resource as a final step.

Globus use remains high, although not all facilities are users either due to security policy, licensing issues, or observations that performance is not measurably better. With respect to the latter, mounting storage remains a challenge since it requires lots of tuning to see high performance between facilities.

New workflows to better stream, store, process, and share data will be needed in the future as we create longer and longer paths between observation and result. This, along with some of the observations in the multifacility topic, point toward needing more intelligent APIs to join resources that span facility boundaries.

6.3.2 Facility Upgrades and Changes to Data Volume and Rate

Some larger facilities and experiments note they will see data rates increase in the coming years. ESGF (and all users of CMIP resources) expects to see an explosion as new models are released. These will require fast networks between the major HPC facilities, as well the locations of some of the ESGF nodes. It is expected that some simulations will be PB scale, approaching EB, and there will be hundreds of these in the coming years.

Improvements to GPU hardware have made it such that data compression is much better, which has simplified processing (and network transfer) for some model and simulation activities. These new hardware packages will be available on new DOE-SC systems, and will revolutionize some of the ways that people interact with data.

6.3.3 Challenges in Multi/Coupled Facility Workflows

A core problem for addressing multifacility use cases remains a lack of a cohesive identify, authorization, and authentication strategy. It is still the case that each facility (and lab) has different ways of managing this, which complicates attempting to build multifacility workflows. For example, use of a sequencer at the JGI, and storage and computation at ALCF, result in the use of multiple user identities. Managing multiple accounts can lead to poor user behavior: sharing credentials, or not practicing safe cybersecurity practices.

Participants acknowledge that despite the friction, the use of multifacility workflows will remain high in the future. Coupling computing to an instrument facility, or combining multimodal data, is still of high research interest. Tools like Jupyter have simplified the language of workflows, and remain popular.

Some participants note that when a coupled facility is unavailable (or has limited resources), the entire workflow could suffer. As an example, if a system that provides computation is down for maintenance, or has a long wait time, it may be impossible to migrate the work to another resource unless the steps to make the workflow portable were taken. If it were easier to move to other HPC facilities, or on temporal resources that may be located on ESnet, that could offer a layer of protection for some time-dependent use cases.

As real-time sensor data becomes more widely available from remote locations using tools like LEO satellites, it will be necessary to link analysis and storage to accept and process this data in a more real-time fashion. It is expected that this use case could challenge some aspects of the multifacility model not in terms of data volume, but more quantity of resources that will have to be accessed.

6.3.4 Software: Analysis, Sharing, Simulation, Etc.

Some facilities want to upgrade their software and workflow mechanisms, but lack development resources to do so. There is a desire to have more portability and also address long-standing issues with storage performance and interactive feel for real-time applications. All of this is compounded by the fact that analysis routines are regularly designed to consume more data (e.g., volume and quantity), and both the hardware and software must scale to accommodate this. Some workflows must be converted to rely on data streams (versus bulk data movement) to better handle the larger data volumes in these analysis workflows.

6.3.5 Computation: Real-Time, Near Real-Time, or Offline

Computational needs are sufficient for most participants. Interactive jobs (e.g., those that are operated via Jupyter) are easily run on cloud-like computing platforms, and meet the needs of most users. Those that require access to midrange clusters, or HPC, can use institutional resources or allocations at the major HPC centers to accommodate their computing needs.

Most participants do some commercial cloud computing, but cost remains the biggest obstacle to use. ESnet's peering with cloud providers has simplified some of the data workflow, but the use of DOE-SC resources for computing remains popular and necessary. Participants acknowledge that having more information on the ways that ESnet peers with cloud providers would be useful information when negotiating contracts.

6.3.6 Local Networking Support

A number of participants have taken advantage of ESnet's Science Engagement Team to assist with data mobility issues. These range from helping to improve international connectivity to ESGF nodes, to helping to tune DTNs, to advising on the best way to deal with remote resources that have minimal connectivity options. Participants have noted that there can be a simplified "how to ask for help" resource that is posted to the ESnet website and Fasterdata, along with continued outreach to user groups.

There is interest in understanding some of ESnet's lessons learned using LEO satellites and 5G connectivity. Many remote sensor projects could benefit from this approach, as they build a standard model for pulling data from regions with poor connectivity.

For some collaborations with a strong international presence (e.g., ESGF), high performance remains elusive. At some ESGF nodes located in the Asia-Pacific region, the network performance can be measured in Kbps values, which makes downloading large data sets challenging. There is a rich international ecosystem of connectivity (e.g., multiple 100 Gbps) across the Pacific, but local connectivity within a facility (or security policy) may be limiting performance. ESnet will continue to work with communities to characterize and understand these observations.

List of Abbreviations

ABF	Agile BioFoundry
ABL	Applied Biosciences Laboratories
ABPDU	Advanced Biofuels Process Demonstration Unit
ACCESS	Advanced Cyberinfrastructure Coordination Ecosystem: Services & Support
ADC	ARM Data Center
ADCP	Acoustic Doppler Current Profiler
AES	Agricultural Experiment Station
AI	artificial intelligence
ALCF	Argonne Leadership Computing Facility
ALE	Adapted Lab Evolution
ALS	Advanced Light Source
AML	Advanced Manufacturing Laboratory
AMP	AmeriFlux Management Project
ANL	Argonne National Laboratory
ANU	Australian National University
APPL	Advanced Plant Phenotyping Laboratory
APS	Advanced Photon Source
ARCO	analysis-ready cloud-optimized
ARM	Atmospheric Radiation Measurement
ARRA	American Recovery and Reinvestment Act
ASCR	Advanced Scientific Computing Research
ASR	Atmospheric System Research
ATS	Advanced Terrestrial Simulator
AUV	Autonomous Underwater Vehicle
AWI	Alfred Wegener Institute
AWRI	Australian Wine Research Institute
AWS	Amazon Web Services
BBD	Biofuels and Bioproducts Division
BCSB	Berkeley Center for Structural Biology
BEO	Barrow Environmental Observatory
BER	Biological and Environmental Research
BETO	Bioenergy Technologies Office
BNC	Biomolecular Nanotechnology Center
BNL	Brookhaven National Laboratory
BRC	Bioenergy Research Centers
BRL	Biorenewables Research Laboratory
CABBI	Center for Advanced Bioenergy and Bioproducts Innovation
CAD	computer-aided design
CADES	Compute and Data Environment for Science

CARNE	Campus Advanced Research Network Environment
CASCADE	Calibrated and Systematic Characterization, Attribution, and Detection of Extremes
CBI	Center for Bioenergy Innovation
CCRCN	Coastal Carbon Research Coordination Network
CD	Circular dichroism
CDG	Climate Data Guide
CEDA	Centre for Environmental Data Analysis
CENIC	Corporation for Education Network Initiatives in California
CESD	Climate and Ecosystem Sciences Division
CESM	Community Earth System Model
CGRL	Computational Genomics Resource Laboratory
CIF	Common Intermediate Format
CINT	Center for Integrated Nanotechnologies
CMIP	Coupled Model Intercomparison Project
CMNT	Center for Micro and Nanotechnology
CMOS	complementary metal–oxide–semiconductor (correct)
CNC	Computer Numerical Control
CNES	Carbon Neutral Energy Solutions
CNSI	California NanoSystems Institute
COMPASS	Coastal Observations, Mechanisms, and Predictions Across Systems and Scales
CPU	central processing unit
CRD	Computational Research Division
CRF	Combustion Research Facility
CS	Computing Sciences
CSA	Computing Sciences Area
CSV	comma-separated value
DBTL	design-build-test-learn
DKRZ	German Climate Computing Centre (Deutsches Klimarechenzentrum)
DMT	data management team
DNS	domain name system
DOE	Department of Energy
DOI	Digital Object Identifiers
DQR	Data Quality Report
DSL	digital subscriber line (correct)
DTN	Data Transfer Node
DYAMOND	Dynamics of the Atmospheric general circulation Modeled On Non-hydrostatic Domains
EB	exabyte
EcoFAB	fabricated model microbial ecosystems
EcoPOD	meter-scale contained and controlled ecosystems
EDD	Experimental Data Depot
EDM	electrical discharge machining

EDS	energy dispersive X-ray spectroscopy
EERE	Energy Efficiency and Renewable Energy
EESA	Earth and Environmental Sciences Area
EESSD	Earth and Environmental Systems Sciences Division
EFL	Experimental Fermentation Lab
ELM	E3SM Land Model (ELM)
EM	electron microscopy
EMSL	Environmental Molecular Sciences Laboratory
ENA	Eastern North Atlantic
ENIGMA	Ecosystems and Networks Integrated with Genes and Molecular Assemblies
EPICS	Experimental Physics and Industrial Control System
ERT	electrical resistivity tomography
ES&T	Environmental Science & Technology
ESCC	ESnet Site Coordinators Committee
ESE	Emeryville Station East
ESGF	Earth System Grid Federation
ESIP	Earth System Information Partners
ESS	Environmental Systems Science
ET	evapotranspiration
FACE	free-air CO2 enrichment
FAIR	Findable, Accessible, Interoperable and Reusable
FATES	Functionally Assembled Terrestrial Ecosystem Simulator
FBC	Fungal Biotechnology Center
FGDC	Federal Geographic Data Committee
FGL	Functional Genomics Laboratory
FIB-SEM	Focused Ion Beam Scanning Electron Microscopy
FICUS	Facilities Integrating Collaborations for User Science
FME	Field, Measurements, and Experiments
FPLC	fast protein liquid chromatography
FTP	File Transfer Protocol
FVCOM	Finite Volume Community Ocean Model
GaA	gallium arsenide
GaN	gallium nitride
GC	gas chromatography
GCP	Google Cloud Platform
GCRM	Global Cloud Resolving Models
GCS	Google Cloud Storage
GHG	greenhouse gas
GiB	gigabyte
GLBRC	Great Lakes Bioenergy Research Center
GLM	Great Lakes Modeling

GPB	Genetics & Plant Biology
GPFS	general parallel file system
GPU	graphical processing unit
GSL	Genomics Sequencing Laboratory
GSRM	global storm resolving models
GTRI	Georgia Tech Research Institute
GUI	graphical user interface
HPC	high-performance computing
HPLC	high-performance liquid chromatography
HPSS	high-performance storage system
HT	high throughput
ICE	Inventory of Composable Elements
ICP	Inductively Coupled Plasma
IGB	Institute for Genomic Biology
IGI	Innovative Genomics Institute
IL	ionic liquid
ILAMB	International Land Model Benchmarking
IM	integrated model
IMRL	Integrated Materials Research Laboratory
IOC	input/output controllers
IP	internet protocol
IPCC	Intergovernmental Panel on Climate Change
IPSL	Institut Pierre-Simon Laplace
IR	infrared
IRI	Integrated Research Infrastructure
IRI-ABA	IRI Architecture Blueprint Activity
ISP	Internet service provider
IU	Indiana University
JAMO	JGI Archive and Metadata Organizer
JAWS	JGI's Analysis Workflow Service
JBEI	Joint BioEnergy Institute
JGI	Joint Genome Institute
JLSE	Joint Laboratory for System Evaluation
JSON	JavaScript Object Notation
KARE	Kearney and West Side Research and Extension
KNL	Knight's Landing
KVM	Kernel-based Virtual Machine
LAN	local area network
LANL	Los Alamos National Laboratory
LBNL	Lawrence Berkeley National Laboratory
LC	liquid chromatography

LCF	Leadership Computing Facility
LEO	Low Earth Orbit
LES	large eddy simulation
LIMS	Laboratory Information Management System
LLNL	Lawrence Livermore National Laboratory
LULCC	Land Use Land Cover Change
M&S	materially and substantially
MB	megabyte
MBC	Molecular Biology Consortium
MCA	Multi Channel Analyzer
MCRL	Marine and Coastal Research Laboratory
MEMS	micro electro-mechanical systems
MIP	Model Intercomparison Projects
MIxS	Minimum Information about any (x) Sequence
ML	machine learning
MOU	memorandum of understanding
MOV	modes of variability
MPI	Message Passing Interface
MRDP	Modern Research Data Portal
MRL	Materials Research Laboratory
MS	mass spectrometry
MSU	Michigan State University
NASA	National Aeronautics and Space Administration
NCAR	National Center for Atmospheric Research
NCBI	National Center for Biotechnology Information
NCEAS	National Center for Ecological Analysis and Synthesis
NCI	National Computational Infrastructure
NCL	NCAR Command language
NCSA	National Center for Supercomputing Applications
NEON	National Ecological Observation Network
NERSC	National Energy Research Scientific Computing Center
NetCDF	network common data form
NGEE	Next Generation Ecosystem Experiments
NGS	next-generation sequencing
NIMS	nanostructure-initiator mass spectrometry
NIR	near-infrared
NMDC	National Microbiome Data Collaborative
NMR	nuclear magnetic resonance
NOAA	National Oceanic and Atmospheric Administration
NREL	National Renewable Energy Laboratory
NRI	Neuroscience Research Institute

NSF	National Science Foundation
OLCF	ORNL Leadership Computing Facility
OME	online metadata editor
ORNL	Oak Ridge National Laboratory
OSTI	Office of Scientific and Technical Information
PB	petabyte
PCMDI	Program in Climate Model Diagnostics and Intercomparison
PCR	polymerase chain reaction
PI	principal investigator
PNNL	Pacific Northwest National Laboratory
POM	particulate organic matter
POSIX	Portable Operating System Interface
PTT	pressure transducer technique
QA/QC	quality assurance and quality control
QC	quality control
RAL	Rutherford Appleton Laboratory
RAMOS	Respiration Activity Monitoring System
RBI	Renewable Bioproducts Institute
RDBMS	relational database management system
RES	Research Services
RF	radio frequency
RFA	Research Focus Areas
RO	research objectives
RS	remote sensing
SAF	sustainable aviation fuel
SAGC	South Australian Genomics Centre
SAIL	Surface Atmosphere Integrated Laboratory
SBC	Structural Biology Center
SBML	Systems Biology Markup Language
SBOL	Synthetic Biology Open Language
SBR	Subsurface Biogeochemical Research
SCP	secure copy protocol
SDFL	Self-Driving Field Laboratories
SDM	Sequence Data Management
SDSC	San Diego Supercomputing Center
SEARCH	Study of Environmental Arctic Change
SFA	Science Focus Area
SFTP	secure file transfer protocol
SGFL	Self-Guiding Field Laboratory
SGP	Southern Great Plains
SiC	silicon carbide

SLA	service level agreements
SNL	Sandia National Laboratories
SNS	Spallation Neutron Source
SRA	Sequence Read Archive
SVM	science virtual machine
TAI	terrestrial-aquatic interfaces
TB	terabyte (correct)
TBS	tethered balloon systems
TCD	thermal conductivity detector
TES	Terrestrial Ecosystem Science
TRY	titers, rates, and yields
UAS	uncrewed aerial system
UC	University of California
UCAR	University Corporation for Atmospheric Research
UCB	University of California, Berkeley
UCD	University of California, Davis
UCSB	University of California, Santa Barbara
UCSD	University of California, San Diego
UI	user interface
UIC	Ukpeagvik Iñupiat Corporation
UITS	University Information Technology Services
UPLC	Ultra-High Performance Liquid Chromatography
US	United States
UV	ultraviolet
VM	virtual machine
VNC	virtual network computing
VPN	virtual private network
WAN	wide-area network
WDL	Workflow Description Language
WHONDRS	Worldwide Hydrobiogeochemistry Observation Network for Dynamic River Systems
WRF	Weather Research and Forecasting
WSREC	West Side Research and Extension Center
XSEDE	Extreme Science and Engineering Discovery Environment

23-SN-35040