ORNL/TM-2024/3600

OLCF's Advanced Computing Ecosystem (ACE) FY24 Efforts for the DOE Integrated Research Infrastructure (IRI) Program

September 2024

Approved for public release.
Distribution is unlimited.

ORNL IS MANAGED BY UT-BATTELLE LLC FOR THE US DEPARTMENT OF ENERGY

DOCUMENT AVAILABILITY

Online Access: US Department of Energy (DOE) reports produced after 1991 and a growing number of pre-1991 documents are available free via https://www.osti.gov/.

The public may also search the National Technical Information Service's National Technical Reports Library (NTRL) for reports not available in digital format.

DOE and DOE contractors should contact DOE's Office of Scientific and Technical Information (OSTI) for reports not currently available in digital format:

US Department of Energy
Office of Scientific and Technical Information
PO Box 62
Oak Ridge, TN 37831-0062
Telephone: (865) 576-8401
Fax: (865) 576-5728
Email: reports@osti.gov
Website: https://www.osti.gov/

ORNL/TM-2024/3600

NCCS

OLCF'S ADVANCED COMPUTING ECOSYSTEM (ACE) FY24 EFFORTS FOR THE DOE INTEGRATED RESEARCH INFRASTRUCTURE (IRI) PROGRAM

Sarp Oral, Rafael Ferreira da Silva, Subil Abraham, Ryan Adamson, Valentine Anantharaj, Tom Beck, Ashley Barker, Katie Bethea, Michael J. Brim, Joshua Brown, Elton Cranfill, Dan Dietz, Brian Etz, Anjus George, Jens Glaser, Gustav R. Jansen, Olga A. Kuchar, Jack Lange, Ketan Maheshwari, Zach Mayes, Bronson Messer, Ross Miller, Ryan Prout, David M. Rogers, Mallikarjun Shankar, Tyler J. Skluzacek, Veronica M. Vergara, Jordan Webb, Patrick Widener, Christopher Zimmer

Collaborators: Daniela Cassol, Ryan Coffee, Mario Cromaz, Eli Dart, Philip W. Fackler, Kjiersten Fagnan, William F. Godoy, Steven E. Hahn, Christina M. Hoffmann, Ezra Kissel, Wilko Kroeger, Takuya Kurihana, Valerio Mariani, Narasinga Rao Miniskar, Zachary Morgan, Frédéric Poitevin, Eric Pouyoul, Andrei T. Savici, Jana Thayer, Pedro Valero-Lara, Kiran Vasu, Rama K. Vasudevan, Jeffrey S. Vetter, Cong Wang, Aaron Young

September 2024

Prepared by
OAK RIDGE NATIONAL LABORATORY
Oak Ridge, TN 37831
managed by
UT-BATTELLE LLC
for the
US DEPARTMENT OF ENERGY
under contract DE-AC05-00OR22725

EXECUTIVE SUMMARY

This report highlights significant strides made by Oak Ridge National Laboratory's Oak Ridge Leadership Computing Facility (OLCF) in advancing computational research and infrastructure. Through the Advanced Computing Ecosystem (ACE) strategic initiative, OLCF has been successfully integrated with DOE's Integrated Research Infrastructure (IRI) program, establishing itself as a critical framework for enhancing scientific computing capabilities across various domains. This report outlines the activities, accomplishments, and future directions of ACE, emphasizing its role in developing cutting-edge technologies, supporting science pilots, and fostering collaborations that drive scientific innovation.

One of the key achievements of ACE in FY24 has been the establishment of a robust testbed environment that facilitates the development and validation of new computing technologies. This testbed, integrated into the larger IRI ecosystem through Energy Sciences Network (ESnet), allows for seamless collaboration among DOE's Advanced Scientific Computing Research (ASCR) facilities, including the Argonne Leadership Computing Facility (ALCF), the National Energy Research Scientific Computing Center (NERSC), and Jefferson Lab (JLab). Today the ACE testbed includes Summit through the SummitPLUS allocation program and also the Defiant system as a test and development platform, as well several other compute and storage solutions and an AI appliance. The testbed has been instrumental in enabling science pilots to experiment with innovative methodologies in a controlled environment, thereby accelerating the translation of research into practical applications.

ACE has been instrumental in advancing a series of science pilots using computational resources from OLCF. These pilots span a wide range of scientific domains, from basic energy sciences to biological and environmental research, demonstrating the versatility and impact of the ACE framework. By integrating cutting-edge technologies, such as AI-driven analytics and real-time data processing, with HPC, these pilots have enabled significant advancements in scientific workflows. For instance, projects like LCLStream and DELERIA have showcased the potential of ACE in optimizing experimental designs, accelerating data analysis, and driving innovations that are critical for addressing some of the most complex scientific challenges of IRI pathfinder projects.

ACE has also made significant progress in foundational technology development, particularly in areas such as data movement, interface design, and scheduling. These advancements have been key in addressing the complex requirements of modern scientific workflows, ensuring that OLCF remains at the forefront of high-performance computing (HPC). The report details several ongoing projects, including the development of the OLCF Facility API, Data Streaming technologies, and the Zambeze distributed orchestration system, which together provide a secure and flexible framework for managing computational resources across a distributed infrastructure.

In addition to technical achievements, ACE has played a central role in outreach and engagement activities. Through participation in conferences, hackathons, and training sessions, the ACE team has actively contributed to the broader scientific community. These efforts have not only facilitated the exchange of knowledge but have also strengthened collaborations with other DOE facilities and research institutions, further enhancing the impact of the ACE program.

Looking ahead, the report outlines strategic plans for FY25, which include continued support for science pilots, further integration with ASCR facilities, and the exploration of new technologies. These initiatives are designed to ensure that ACE remains a dynamic and adaptable platform, capable of meeting the evolving needs of the scientific community. The ongoing collaboration with ORNL's INTERSECT initiative and the planned technology refreshes underscore ACE's commitment to maintaining its leadership in HPC.

CONTENTS

1. INTRODUCTION

The landscape of computing systems and their underlying architectures is evolving rapidly to meet the demands of emerging applications and interconnected workflows [1]. This growing complexity necessitates a comprehensive approach to understand these advancements, evaluate their strengths and weaknesses, assess their operational impacts and capabilities, and develop solutions to address any gaps.

The **Integrated Research Infrastructure** (IRI) Program [2], spearheaded by the Department of Energy (DOE), aims to create a cohesive and interconnected research environment. IRI focuses on "integrating diverse computational resources, data infrastructures, and scientific instruments to facilitate seamless collaboration and data sharing across multiple research domains." By fostering an ecosystem where resources and data can be easily accessed and utilized by scientists, the IRI program enhances the efficiency and effectiveness of scientific research, enabling breakthroughs and accelerating innovation.

The IRI vision is to "empower researchers to meld DOE's world-class research tools, infrastructure, and user facilities seamlessly and securely in novel ways to radically accelerate discovery and innovation." This vision emphasizes new modes of integrated science, rapid data analysis, AI-enabled insights, and the integration of vast data sources to drive forward scientific progress. Core to this vision are principles such as flexibility, performance, scalability, transparency, interoperability, resiliency, extensibility, and cybersecurity. These principles ensure that the infrastructure supports innovative and secure scientific endeavors while being adaptable to future needs.

A significant aspect of IRI is its mission to democratize access to high-performance computing (HPC) and data resources. By linking distributed resources and creating a more open and collaborative environment, IRI aims to accelerate discovery and innovation, drawing new talent into the scientific community and advancing open science. The program also addresses the challenges of the exascale science era, including the data deluge from advanced source/detector technologies and observational platforms, and the application of Artificial Intelligence (AI) for science, energy, and security.

To support these ambitious goals, the Oak Ridge Leadership Computing Facility (OLCF) has consolidated various related initiatives under the **Advanced Computing Ecosystem (ACE)** umbrella (Figure 1). ACE is a strategic framework designed to support the OLCF in the development and deployment of the upcoming OLCF-6 project, as well as to contribute to DOE programs such as IRI and FASST. ACE leverages a close partnership with Oak Ridge National Laboratory's INTERSECT (Interconnected Science Ecosystem) LDRD initiative and ASCR's Energy Sciences Network (ESnet) user facility to ensure a synergistic approach to advancing scientific computing. Driven by the imperative of scientific advancement, ACE aims to provide robust science campaign support and infrastructure to enhance research capabilities.

In FY24, ACE encompassed four key areas of focus to advance the goals of IRI:

- **Science Pilots and Workflows –** Developing and implementing pilot projects and workflows that leverage integrated research infrastructure, facilitating seamless collaboration and data sharing across various scientific domains. These pilots aim to create new models of scientific research by integrating multiple user facilities and employing advanced data management and analysis techniques.
- **Testbeds –** Providing experimental platforms for testing and validating new technologies in a controlled environment, allowing for iterative development and refinement. These testbeds serve as a proving ground for innovative ideas and technologies, enabling researchers to experiment with cutting-edge solutions and evaluate their practicality and effectiveness before large-scale deployment before productization.

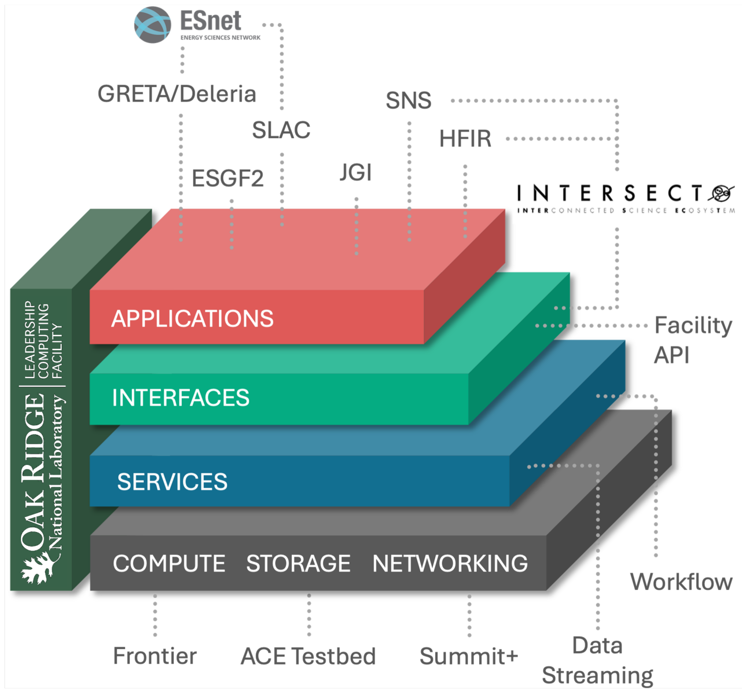

Figure 1. Advanced Computing Ecosystem (ACE) within the OLCF.

- **Research and Development of Foundational Technologies –** Advancing the core technologies that underpin next-generation computing systems, driving innovation and enhancing performance and efficiency. This includes developing new algorithms, software tools, and hardware solutions that can handle the increasing demands of modern scientific research.
- **Technology Evaluations –** Conducting thorough assessments of new computing technologies to understand their potential benefits and limitations, ensuring informed decision-making for future deployments. These evaluations are crucial for identifying the most promising technologies and integrating them into the research ecosystem effectively.

By addressing these areas, ACE aims to equip the OLCF with the tools and knowledge needed to navigate the complexities of modern computing environments, ultimately enhancing the capabilities and impact of scientific research and discovery. The integration with ORNL's INTERSECT initiative further strengthens ACE's ability to support interconnected and integrated scientific workflows, fostering an environment where collaborative science can thrive and laboratories of the future can be realized.

This report will discuss these four key areas in detail as well as the lessons learned and challenges encountered, workforce development initiatives, and outreach activities undertaken as part of the ACE program. Through these efforts, we aim to highlight the progress made and the future directions for advancing the capabilities of scientific computing in support of groundbreaking research.

Engagement and Collaboration with the IRI Management Council. Since the IRI program interconnects not only the ASCR facilities but the broader DOE SC complex, we are keeping our efforts deeply aligned with the IRI program and its governance and management objectives. We are actively collaborating with the IRI Management Council to ensure our efforts are closely coordinated with the overarching goals of the IRI program where aligned and its pathfinder science projects requirements. Our active involvement includes occupying key leadership positions, such as co-chairs of the recently launched technical subcommittees (Table 1). This strategic collaboration enables us to contribute effectively to the IRI's mission, fostering integrated research and facilitating the successful implementation of its innovative scientific applications and workflows.

IRI Management Council	**ORNL Representatives**
Executive Committee	Mallikarjun Shankar
Leadership Group	Sarp Oral, Rafael Ferreira da Silva
Technical Subcommittee: Outreach and Engagement	Rafael Ferreira da Silva (co-chair), David M. Rogers
Technical Subcommittee: Interfaces	Ryan Prout (co-chair), Addi Malviya Thakur
Technical Subcommittee: TRUSTID	Ryan Adamson (co-chair), Carl Bai

Table 1. ORNL Representatives in IRI Management Council and Technical Subcommittees.

2. SCIENCE PILOTS AND WORKFLOWS

The Science Pilots and Workflows initiative aims to develop and implement pilot projects that leverage ACE's capabilities to facilitate seamless collaboration and data sharing across various scientific domains. This initiative focuses on creating and testing new models of scientific research by integrating multiple user facilities and employing advanced data management and analysis techniques. The goal is to demonstrate how interconnected resources and innovative workflows can accelerate scientific discovery and enhance research efficiency. In this section, we describe our efforts with applications spanning a range of science programs from the DOE's Office of Science, highlighting the practical impact and potential of these pilots in advancing the frontiers of knowledge.

2.1 Basic Energy Sciences (BES)

2.1.1 LCLStream

External Collaborators: Frédéric Poitevin, Jana Thayer, Ryan Coffee, Cong Wang, Valerio Mariani, Wilko Kroeger, LCLS, SLAC National Accelerator Laboratory

Pilot Description. The LCLStream pilot project carries out training of a generalist AI model able to interpret experimental X-ray data from a stream of detector events (Figure 2). These streams are produced from either archived or live experiments run at SLAC's Linac Coherent Light Source (LCLS) and LCLS-II beamlines. Because of its size and complexity, the model requires hundreds to thousands of Summit nodes (thousands of GPUs) for hours at a time as O(terabyte) batches of training data becomes available. The resulting AI model functions as a shared vision backbone to produce feature maps that enable multiple downstream prediction tasks such as hit classification, Bragg peak segmentation and image reconstruction. It will be useful to inform a variety of data analysis queries from experimentalists.

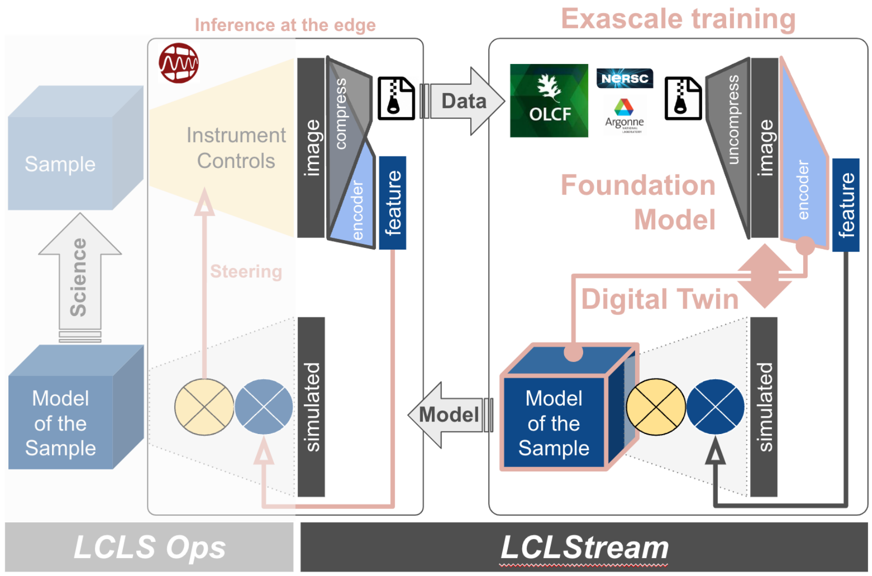

Figure 2. Overview of the LCLStream pilot project, showcasing the end-to-end workflow for training a generalist AI model to interpret experimental X-ray data from SLAC's LCLS and LCLS-II beamlines. The workflow includes data streaming, preprocessing, scaling, and on-site modeling, with integration across facilities and compute resources at SLAC and OLCF. ***Courtesy of SLAC.***

Challenges. In addition to its high-performance computing needs on ACE, this pilot depends on developing new technologies that will be critical for expanding LCLS-II analysis capabilities in several directions: (1) Optimizing the event building, data reduction and pre-streamed data processing; (2) Scaling up the AI model training to petabytes of detector data; (3) Augmenting the data stream with compute-intensive, on-site modeling and simulation; and (4) Experimental design for steering and optimizing the various instrumentation.

Reliably implementing these technologies requires close collaboration between facilities, application scientists and instrument scientists. Specifically, pre-processing of the data before being streamed requires speed-up of the data reduction and calibration pipelines, up to real-time speeds. Scaling up image data-streams requires not only increasing physically available site-to-site bandwidth, but also the mechanism for allocating bandwidth, and cross-facility identity management. Running compute-intensive analysis requires on-demand compute capability, which necessitates a method for scheduling high-priority reservations for compute resources at leadership-level HPC scales. Instrument steering requires considering the problem of secure, authenticated, and trustworthy two-way API interactions.

Leveraging ACE. An end-to-end streaming demo has been designed and prototyped using a series of software and components for data source management[1], data analysis (event-building and data reduction), and flexible streaming[2] across SLAC and OLCF sites. On the computational front, a highly efficient AI dataset storage mechanism[3] has been implemented to optimize data storage for AI inference. Job control and orchestration are managed through a dedicated API, enabling the execution of stream-processing jobs on high-performance computing resources such as Summit and Defiant. The entire workflow is cohesively integrated[4], supporting documentation of the installation, initialization, and execution phases, ensuring a streamlined and reproducible process.

A new software called Masked Autoencoder for X-ray Image Encoding (MAXIE) has been released, offering support for models ranging from hundreds of millions to billions of parameters. MAXIE's training is computationally intensive, requiring a dataset comprising roughly 1 trillion visual tokens (16-by-16 image patches). To facilitate this process, data segmentation, shared memory utilization, and parallelization strategies are employed to enable high-compute-intensity jobs and enhance data loading efficiency. Significant project milestones include the transfer of 286 TB of image data to Summit's Alpine2 filesystem, the successful demonstration of the image loading pipeline in streaming mode from SLAC to ORNL at rates of 10 Mb/sec and local benchmarking at 70 GByte/second with 60 MPI ranks, and the development of a comprehensive tutorial for the MAXIE image autoencoder. Additionally, a cross-laboratory documentation effort on streaming from LCLS is being developed[5].

On the FY24 collaboration SLAC's Jana Thayer has stated that "SummitPlus and IRI have enabled new directions in our research. None of the large-scale training we are exploring would have been possible without the use of our SummitPlus allocations. We've developed a new architecture in which we pre-process data on the S3DF to convert to a compact and user-friendly format and stream to Oak Ridge for analysis or training. This is also new and is a pattern that will work well across sites. We've also learned a lot about possibilities for parallel streaming and live analysis that will be useful. All of our future efforts to develop digital twins and do real-time analytics will benefit from the work we've done with ORNL this last year. We've gained a lot of benefit for LCLS and ILLUMINE through these efforts and hopefully it has been beneficial for Oak Ridge as well."

[1] https://github.com/frobnitzem/lclstream
[2] https://gitlab.com/frobnitzem/nng_stream
[3] https://code.ornl.gov/99R/local_sampler
[4] https://code.ornl.gov/99R/streamrun
[5] https://github/lcls-users

New Capability. This project has successfully established an integrated workflow for end-to-end data streaming across SLAC and OLCF facilities. This capability includes the development of a new architecture for pre-streaming data processing and a new image autoencoder software that can support models containing billions of parameters.

Next Steps. Recent discussions have underscored the critical need for ongoing, close collaboration between LCLS, S3DF, and leadership computing facilities to enhance these components. While the microservices developed so far work effectively together to achieve streaming, further integration with facility-deployed infrastructure will be essential for improving their reusability across different sites and ensuring reliability in production environments. Key improvements needed include eliminating manual steps during installation, startup, and execution (such as starting servers and creating SSH tunnels), reducing site-specific adjustments required for interoperability across various experiment and computing facility APIs, standardizing security technologies like mutual TLS authentication, addressing the "co-scheduling" challenge for coordinating beam time, network bandwidth, and compute resources simultaneously, and providing seamless access to fast, open-source data movement services (e.g., XRootD or FTS3) for managing site-to-site file transfers.

These microservice prototypes hold significant value as they offer a foundation for future application projects and integration with facility-managed infrastructure. The development pathway for the IRI ecosystem is clear: facilities must provide interfaces—such as service deployment, authentication, networking, and resource reservations—that simplify component complexity and enhance usability for facility users. Science applications, in turn, need to adapt their workflows and infrastructures to leverage these advancements while effectively communicating challenges with facilities and developers. The successful execution of the LCLStream pilot project's SLAC-to-OLCF data pipeline for generalist AI on X-ray detector events marks an important first step in this ongoing exploration and development process.

2.1.2 IMAGINE-X

Pilot Description. The OLCF is collaborating with the Spallation Neutron Source (SNS) on DOE BRaVE-funded research to deliver a dynamic nuclear polarization enabled neutron crystallography instrument, IMAGINE-X, at the High Flux Isotope Reactor (HFIR). One of the strengths of neutron crystallography lies in its sensitivity to hydrogen atoms, which are important in biology but mostly invisible to X-rays. However, neutron fluxes are orders of magnitude lower than photon fluxes at synchrotron sources, which makes it necessary to grow very large protein crystals, on the order of millimeters. Dynamic nuclear polarization (DNP) addresses this limitation by using a magnetic field and embedded paramagnetic centers to polarize the hydrogen atoms in a cryogenic sample and boost the hydrogen scattering cross section. IMAGINE-X will accelerate structure solution and drug discovery for biopreparedness by improving of the signal-to-noise (S/N) ratio of neutron data by an order of magnitude, enabling the study of crystals that are radically smaller than has previously been possible, while making use of new AI driven analysis and simulation methods. This breakthrough will produce over 100-fold gains in performance for neutron diffraction analysis of biological systems, accelerating the development of new therapeutics for disease, and understanding the control of enzymes with designed catalytic and ligand binding behaviors.

Challenges. One of the primary challenges involves processing the experimental data directly from the 50 new Anger cameras that will be installed at the experiment. Every neutron will result in a timestamped event with spatial coordinates and detector indices. To integrate IMAGINE-X with the data acquisition (DAQ) capabilities used at other experiments such as DEMAND at HFIR or TOPAZ at SNS, we will collect live data from the instrument, stream it using an upgraded version of the ADARA protocol to Data Processing Units (Bluefield DPUs), and integrate both with EPICS instrument control.

Leveraging ACE. We are combining cutting-edge AI technology, Frontier and ACE computing resources at OLCF to develop a virtual instrument to help optimize the design of the new IMAGINE-X instrument, accelerate real-time data collection, reduce data redundancy, and enable autonomous online experiment steering. This co-design approach will simultaneously develop the new IMAGINE-X instrument, the virtual instrument, and the computing infrastructure needed to connect HFIR, OLCF and edge computing at the experiment, improving experiment steering with a computational workflow (Figure 3).

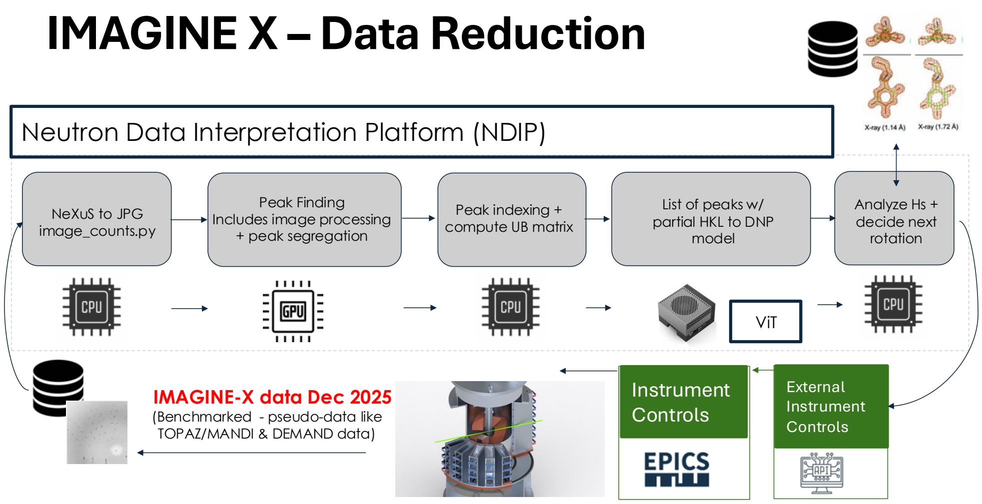

Figure 3. Integrated workflow from instrument to edge to exascale that enables in-situ data processing and automated instrument control for enhanced neutron scattering from biological systems. IMAGINE-X will feature AI-enabled diffraction data analysis software and use GPUs at the edge for structure refinement and phasing. ***Courtesy of SNS.***

New Capability. This project will deliver AI-driven software capabilities that integrate HPC, edge compute and data analysis resources to aid the spin-controlled neutron diffraction experiments.

Next Steps. IMAGINE-X is in the second year of its three-year funding period, and we plan to commission the full instrument in December 2025, deliver the software and integrate the data processing at the HFIR beamline CG4D.

2.1.3 Center for Nanophase Materials Sciences (CNMS)

External Collaborators: Narasinga Rao Miniskar, Aaron Young, Rama K. Vasudevan, Oak Ridge National Laboratory

Pilot Description. The Center for Nanophase Materials Sciences (CNMS) is an instrumentation facility providing access to tools such as atomic force microscopes (AFMs). AFMs are used to analyze the surface and understand the underlying atomic properties of materials by scanning the surface with signal tips (probes) and creating a 3D topographic image based on changes in probe height. The current methodology requires trial and error to identify interesting locations for probing the material. Unfortunately, due to the expensive data analysis, identifying uninteresting locations can take hours to reconcile.

Challenges. The primary challenge encountered during this project was interfacing between the microscope, DAQ hardware, and the IRIS runtime framework. This interface had to be developed to achieve the necessary communication between the microscope and remote server. The insights gained from the development of IRIS-AFM could be leveraged for other projects requiring integration of scientific instruments with computational resources.

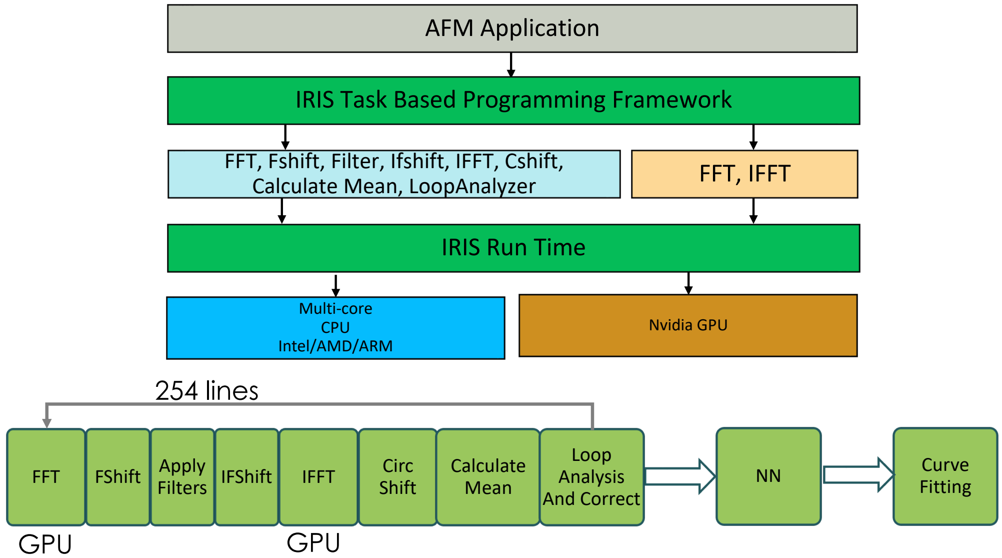

Figure 4. IRIS-AFM Software Architecture: The IRIS framework [3] makes it easier to write code to automatically use a GPU if one is installed.

Leveraging ACE. An ongoing collaboration between OLCF and CNMS is focused on developing a time-sensitive workflow that significantly speeds up the identification of suboptimal placements for AFM users. The first output of this collaboration is a software package called IRIS-AFM (Figure 4). IRIS-AFM collects raw data from the microscope and moves that data into a GPU-enabled server to process the data through filtration methods and extract polarization states. The processed data can then be saved for later analysis as well as sent to a custom real-time visualization program. The visualization program generates butterfly plots indicating the strength of the polarization which can be displayed on the microscope's main workstation and provide the user with immediate feedback about the experiment setup and progress (Figure 5).

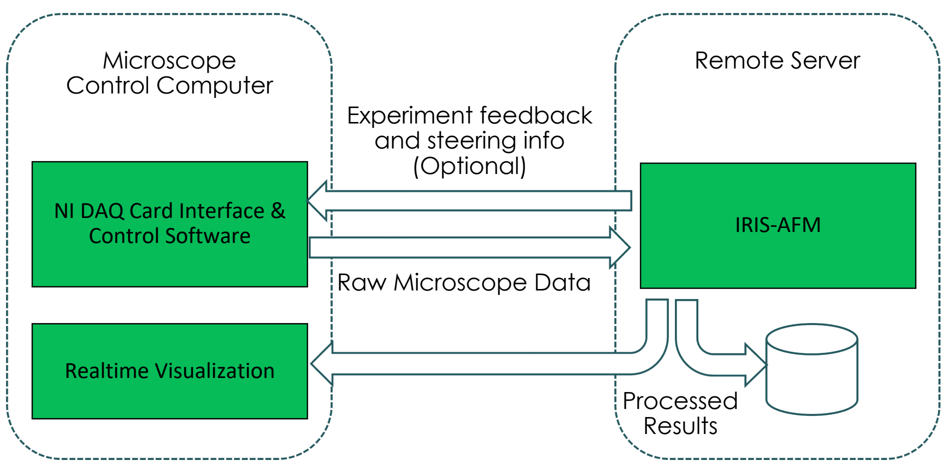

Figure 5. IRIS-AFM performs the bulk of its calculations on a remote server.

The IRIS-AFM project was fully developed and successfully demonstrated to the CNMS team in July 2024, receiving positive feedback. Principal scientist Rama Vasudevan remarked,

> "The IRIS-AFM project enables us to conduct complex experiments on the microscope that, until now, lacked any real-time visualization. The development of a real-time analysis and visualization

framework within IRIS-AFM allows users to quickly visualize results from so-called ‘g-mode’ spectroscopy experiments, enabling us to immediately change relevant experimental parameters and yield insights on tens of thousands of individual ferroelectric switching spectra in a matter of minutes."

New Capability. The ongoing project has produced the IRIS-AFM software package for enabling real-time analysis and visualization of results and mitigates the expensive, time-intensive nature of traditional AFM processing methods.

Next Steps. The current efforts utilize GPU-enabled servers but require relatively low component utilization to generate the necessary feedback information. Our next steps with the CNMS collaboration will involve investigating and developing a machine-learning model that will enable quick scans of the material and identify areas of potential interest for more detailed examination with the goal of reducing the amount of time users spend examining relatively uninteresting ares of the material.

2.1.4 Neutron Scattering Facilities

External Collaborators: Steven E. Hahn, Philip W. Fackler, William F. Godoy, Zachary Morgan, Andrei T. Savici, Christina M. Hoffmann, Pedro Valero-Lara, Jeffrey S. Vetter, Oak Ridge National Laboratory

Pilot Description. The goal of this work is to establish a bridge between advanced computing resources and state-of-the-art neutron science facilities, enabling the efficient processing and analysis of large-scale experimental data. By developing a performance-portable ecosystem that leverages both CPU and GPU architectures, we seek to optimize workflows for complex neutron scattering experiments [4]. This allows for the extraction of detailed material properties with high accuracy and efficiency. The proposed framework not only enhances the computational capacity to handle vast data volumes but also enables near-real-time data processing.

Challenges. The integration of computational and experimental platforms presents significant challenges, particularly in maintaining performance, interoperability, and extensibility as data complexity grows. The evolving nature of next-generation instruments requires a re-engineering of existing algorithms and data workflows to manage the increasing scale and complexity of data. Ensuring seamless communication between disparate systems, managing large data transfers, and optimizing computational efficiency are critical to achieving a truly integrated research infrastructure. These challenges must be addressed to maintain the pace of scientific advancements and to fully exploit the capabilities of modern experimental facilities.

Leveraging ACE. ACE plays an essential role in enabling this application by providing a dynamic and flexible environment that supports the deployment, testing, and optimization of heterogeneous computing resources. ACE’s diverse computing platforms, which include cutting-edge CPU and GPU architectures, allow for rigorous performance evaluation and tuning of the proposed workflows. This capability ensures that our workfklow can fully utilize the computational power available, leading to significant improvements in processing efficiency and speed.

The workflow begins with the collection of raw data from neutron scattering experiments conducted at SNS (Figure 6). This data is captured using advanced instruments, which produce large volumes of complex information. The workflow then involves a data reduction process where the raw experimental data is transformed into a more interpretable format through specialized algorithms. These algorithms, implemented in a performance-portable ecosystem, leverage both CPU and GPU resources to efficiently process and analyze the data. The workflow also includes steps for normalizing and binning the data, which are essential for calculating neutron scattering cross-sections. This computationally intensive process is streamlined through the use of proxy applications that simulate the workload on different computing architectures, ensuring

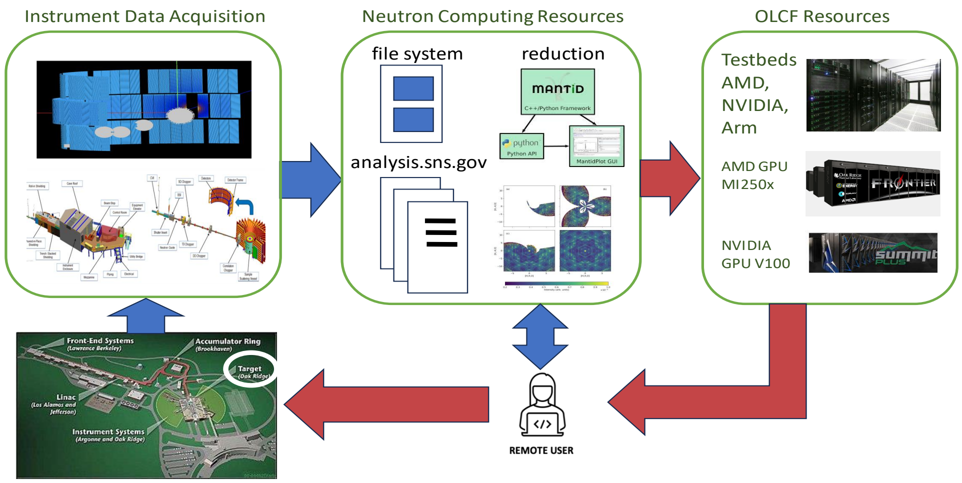

Figure 6. Representation of current (blue) and potential (red) integrated facility workflows for performance-portable codes [4]. ***Courtesy of SNS.***

scalability and performance. The workflow culminates in the generation of high-quality, scientifically meaningful outputs that can be used to study material properties.

New Capability. The project has developed a streamlined workflow that integrates heterogeneous computational architectures with SNS experiments delivering detailed material properties for a broad range of scientific domains.

Next Steps. Moving forward, we will focus on refining its performance-portable ecosystem to further enhance computational models and integrate advanced AI-driven methodologies. Expanding support to additional computing architectures and optimizing the management of data movement and storage will be critical to scaling the framework capabilities. Additionally, efforts will be made to improve the adaptability of the system to accommodate the ever-growing demands of future scientific experiments.

2.1.5 NW-BRaVE

Pilot Description. The NW-BRaVE project, a collaboration between OLCF and the Pacific Northwest National Laboratory (PNNL), aims to enhance biopreparedness by understanding the molecular mechanisms driving pathogenesis and disease transmission. This initiative focuses on creating a robust platform to study pathogen-host interactions, particularly within the context of cyanobacteria and their viral pathogens, cyanophages.

Challenges. The project's complexity involves integrating diverse datasets (structural, genomics, proteomics, and more) across multiple facilities, addressing the challenges of data federation, accessibility, and interoperability. The need to manage and harmonize metadata, ensure data provenance, and comply with varying institutional data management policies further complicates the workflow.

Leveraging ACE. ACE plays a pivotal role by providing the infrastructure to test and refine DataFed, a federated scientific data management tool, across multiple institutions. It supports the development of AI-driven applications and complex workflows. ACE's environment allows for secure, scalable testing of data management and computational workflows, ensuring that the tools meet security and operational requirements before deployment on larger HPC systems.

In FY24, the project successfully set up and began integrating DataFed with ACE's infrastructure, including configuring Data Transfer Nodes and Globus endpoints. The team has begun testing the platform's capabilities, allowing IT administrators to try and refine DataFed in a controlled environment, ensuring readiness for broader deployment.

New Capability. The scientific data management tool, DataFed, has been deployed within the ACE testbed. This capability will facilitate integrating scientific data management tools between OLCF and PNNL to foster improved collaboration.

Next Steps. The project focus will shift to fully operationalizing DataFed with 500 TB of storage, enabling real-world testing by IT administrators and scientific users at PNNL. Additionally, the project will continue to refine AI algorithms and workflows on ACE's NVIDIA hardware.

2.2 Biological and Environmental Research (BER)

2.2.1 Earth System Grid Federation (ESGF)

Pilot Description. The Earth System Grid Federation is a distributed collection of earth and climate science data. The goal of ESGF is to facilitate the advancement of earth system science (such as climate change) through deploying software infrastructure for management, dissemination, and analysis of model output and observational data. Our team is creating a powerful suite of workflow tooling that enables the seamless management of these data.

Challenges. ESGF's IRI challenges include accelerating discovery and democratizing access, and we envision addressing these challenges by leveraging nearly every component of ACE. Foremost, new ESGF applications can benefit from the diverse compute infrastructure that support potential machine learning workflows on these data, and the available 'scale up' capacity made possible by the testbed being adjacent to Frontier. Given that ESGF is a 7 PB data set, the large storage space made available by the Polis storage solution in ACE and the ease of data transfer capabilities provided by Globus enable seamless data management for application users. Finally, the Facility API, which is currently present on Defiant, seamlessly provides secure remote access for users (e.g., via Jupyter notebook) and applications, thus allowing computational scientists to more efficiently gain value from ESGF's data assets.

Leveraging ACE. In FY24, we conducted several demonstrations showcasing the capabilities of our workflows with the CMIP-6 dataset. In the first demonstration, we successfully executed a search, download, and visualization task focused on a near-surface air temperature dataset, utilizing a Jupyter notebook remotely run on Defiant. This workflow, orchestrated by Globus Flows—a system for managing data transfer and compute tasks—was conducted in collaboration with Ryan Chard and Benoit Cote from ALCF (Figure 7). The second demonstration expanded on this by enabling interactive visualization, allowing users to create exploratory graphics that illustrate temporal changes in CMIP-6 data. The third demonstration involved a series of resource-intensive feature engineering workflows that utilized the Facility API to generate features, with Globus facilitating data transfer for validation within the OLCF Jupyter notebook environment. These demonstrations collectively highlight the robust integration of advanced tools and collaborative efforts to enhance data analysis and visualization capabilities.

New Capability. The demonstrations performed in this pilot have showcased advanced workflow tools that support collaborative efforts through improved data management, analysis and visualization capabilities.

Next Steps. In FY25, we will begin our pivot from standalone demonstrations to end-to-end science workflows. Specifically, we will focus on improving support for common workflow tools (e.g., Parsl, Globus, and Dask), the machine learning lifecycle, and common workflow patterns. In particular we plan to focus on patterns present across IRI use cases, including those requiring the use of multiple computers, dynamic

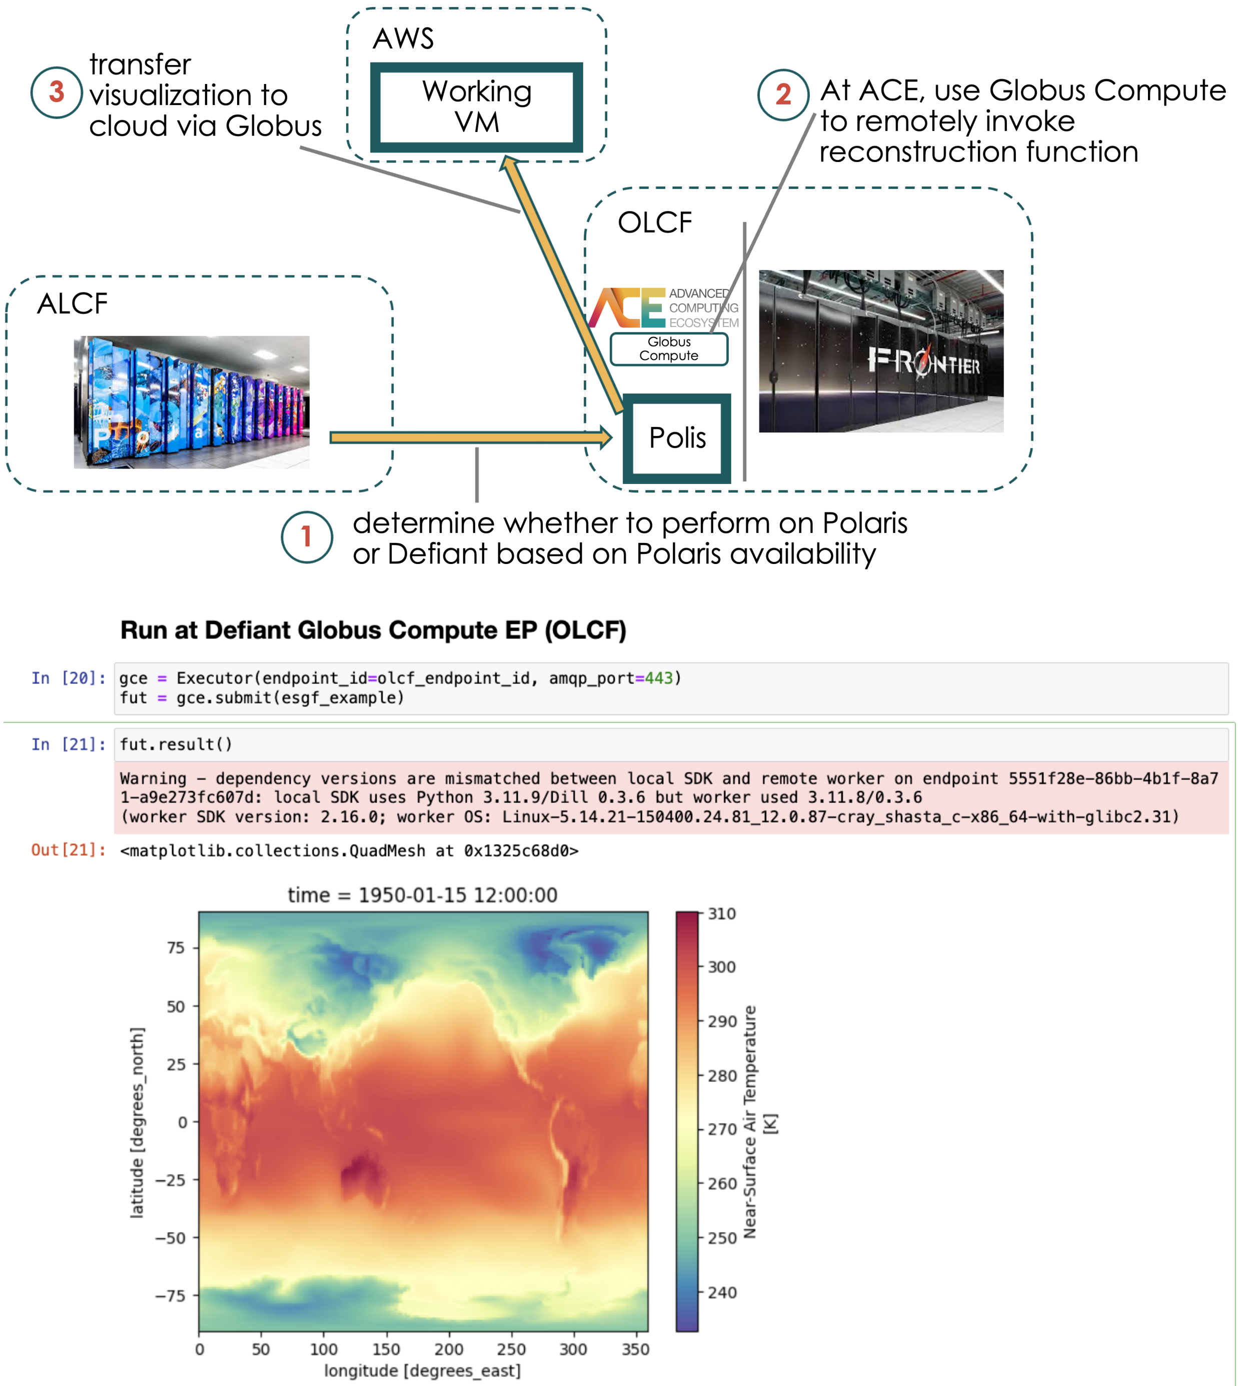

Figure 7. First ESGF demo including search, download, and visualization on Defiant. ***Courtesy of ESGF.***

resource requirements, and human feedback loops [5]. To this end, we plan to align the development efforts of the Zambeze distributed workflow orchestration system [6] with this broad set of goals. Additionally, we would like to provide tighter integration to external tools such as Globus Flows [7] and NERSC's Superfacility API [8].

2.2.2 Joint Genome Institute (JGI)

External Collaborators: Kjiersten Fagnan, Daniela Cassol, Lawrence Berkeley National Laboratory

Pilot Description. The Joint Genome Institute (JGI) is a DOE Office of Science User Facility at Lawrence Berkeley National Laboratory and is part of Berkeley lab's Biosciences area. JGI has developed JAWS (JGI Analysis Workflow Service) as a framework to run users' computational workflows. Its purpose is to improve the re-usability and robustness of workflows in High Performance Computing (HPC), cluster, and cloud environments. JAWS serves as an environment facilitating bridging the gap between user workflows and a target compute infrastructure. Shown in Figure 8 is an architectural view of the various components forming the overall JAWS framework and their interconnections.

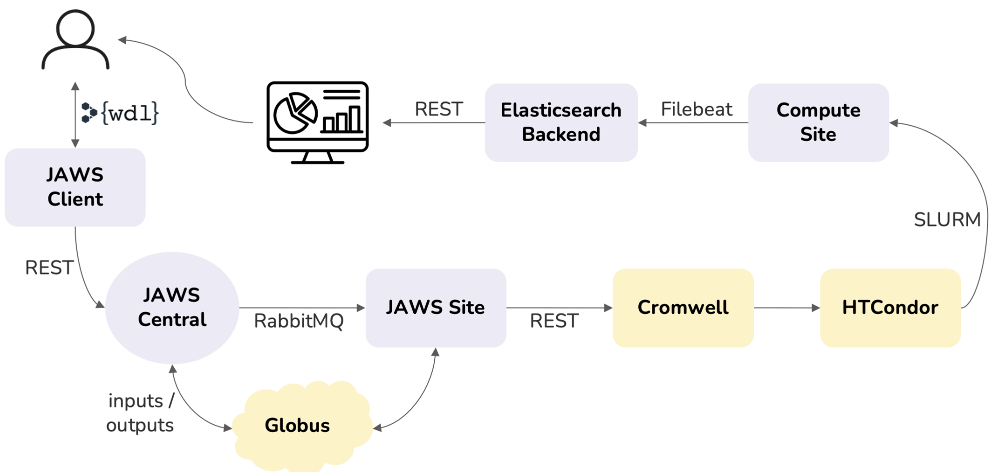

Figure 8. Architecture of the JGI JAWS Framework. ***Courtesy of JGI*****.**

Challenges. As can be seen in Figure 8, the overall architecture involves a significant complexity due to the presence of diverse and heterogeneous tools running in a distributed compute and storage environment. This heterogeneity of components and the distributed nature of the overall framework poses a significant porting challenge for enabling IRI goals for this application. In particular, the installation of various components (e.g., Globus), network configuration, access management for external users from JGI and usage modes with service user setup proved to be significantly challenging.

Leveraging ACE. In April 2024, a collaborative hackathon was hosted by ORNL, bringing together the JGI-JAWS and ORNL-IRI teams with the objective of tackling the identified challenges. The event featured knowledge exchange and hands-on sessions, resulting in significant progress in addressing these issues. A functional JAWS platform was successfully deployed on the ACE testbed's Defiant cluster, with a demonstration conducted during a joint remote session the following week.

New Capability. The JGI JAWS framework has been established on the ACE testbed resulting in several new capabilities including the deployment of complex software for diverse workflows, direct wide area data transfer, accommodation of long running jobs and ability to onboard external users.

Next Steps. Further development of the correct configuration of Globus for seamless remote data transfer, access, and permissions within the workflow remains an ongoing effort of this project. Once completed, this setup will enable JGI users to operate the JAWS platform remotely with full functionality.

2.2.3 Earth Observation

External Collaborators: Takuya Kurihana, Oak Ridge National Laboratory

Pilot Description. Earth observation (EO) satellites, ground-based observation networks, and earth system models are sources of vast, multi-modal datasets that are invaluable for advancing climate and environmental research. These datasets support the development of improved predictive models to better understand complex processes, functions, and feedback across the earth system landscape. The goal of this work is to develop a workflow to automate and enable data management and analysis of these datasets to help drive innovation in earth system sciences.

Challenges. The scale and complexity of the earth system datasets pose significant challenges for processing and analysis, which are distributed across different geographical locations and organizational boundaries. The manual processes currently required for data preprocessing and analysis are time-consuming and computationally intensive, underscoring the need for automated and reliable orchestration capabilities across these advanced research facilities.

Leveraging ACE. ACE enables this application by providing a powerful and flexible testbed that integrates the various computational resources required for the workflow. ACE supports the seamless orchestration of tasks across multiple facilities, ensuring high scalability, efficiency, and reduced processing times for the AI-driven analysis of climate data. The use of Globus software tools within ACE further automates data transfer, computation, and workflow management, making the process more efficient and less prone to human error.

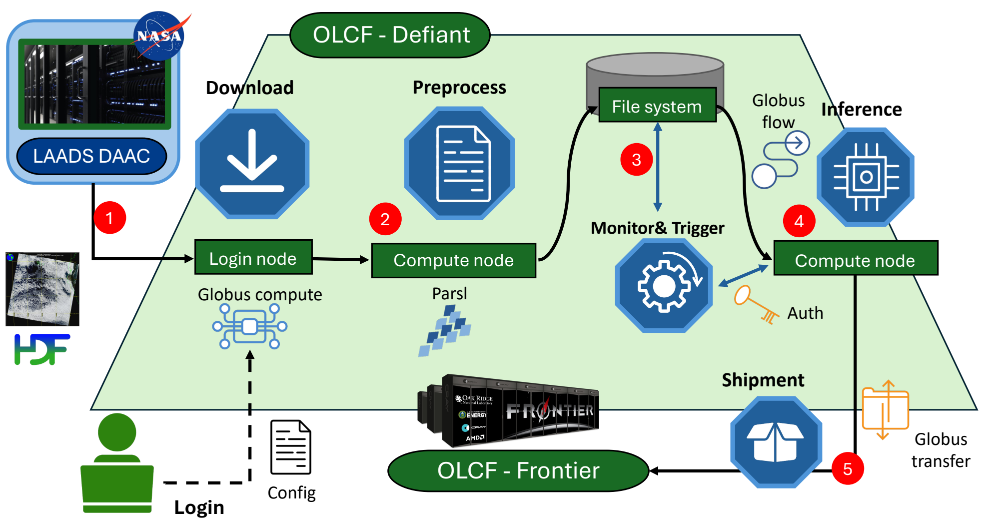

Figure 9. The earth observation ML workflow automates four data processing stages using the Globus software ecosystem. (1) Download: MODIS products are downloaded from NASA DAAC to OLCF's ACE Defiant. (2) Preprocess: MODIS swath images are decomposed into ocean cloud tiles. (3) Monitor & Trigger: A script monitors preprocessing and triggers inference. (4) Inference: Predicts one of 42 AICCA cloud classes. (5) Shipment: Labeled data is transferred to Frontier for further analysis [9].

New Capability. We have developed a multi-faciltiy workflow to automate the collection, preprocessing, AI inference, and data movement tasks for large-scale climate datasets, enabling rapid analysis of petascale data to extract new insights into climate patterns (Figure 9) [9].

Next Steps. The next steps involve extending the current workflow to support more dynamic AI applications, including continual learning, model retraining, and real-time inference on both batch and streaming data. The workflow will evolve to incorporate foundation models that can be adapted for new tasks, improving the accuracy and efficiency of climate predictions. Additionally, there is a focus on integrating advanced provenance tracking and telemetry tools to enhance reproducibility and reliability across different computational environments, as well as developing a federated pipeline-as-a-service platform to streamline the creation and sharing of ML workflows within the scientific community.

2.3 Nuclear Physics (NP)

2.3.1 DELERIA

External Collaborators: Eric Pouyoul, Mario Cromaz, Kiran Vasu, Ezra Kissel, Eli Dart, Lawrence Berkeley National Laboratory

Pilot Description. Deleria (Distributed Event-Level Experiment Readout and Integrated Analysis) is a versatile software toolkit designed to stream and analyze results from physics experiments in real-time, using wide-area networks. It separates data transport and analysis, making it adaptable to various experimental needs. Built on open-source software and standard internet protocols, Deleria operates on off-the-shelf hardware with container technologies. The toolkit comprises components like a forward buffer, storage service, analysis client, event builder, and a schematic orchestrator to configure and coordinate these elements via HTTP-based APIs.

Originally developed to generalize the data pipeline for the Gamma Ray Energy Tracking Array (GRETA) at Lawrence Berkeley National Laboratory, Deleria demonstrates the scalability and applicability of these design patterns. GRETA, a state-of-the-art gamma-ray spectrometer, processes up to 500k events per second, requiring significant computational power for real-time analysis. While initially deployed on a local cluster, Deleria aims to extend this capability over wide-area networks, utilizing remote HPC facilities for more complex analyses.

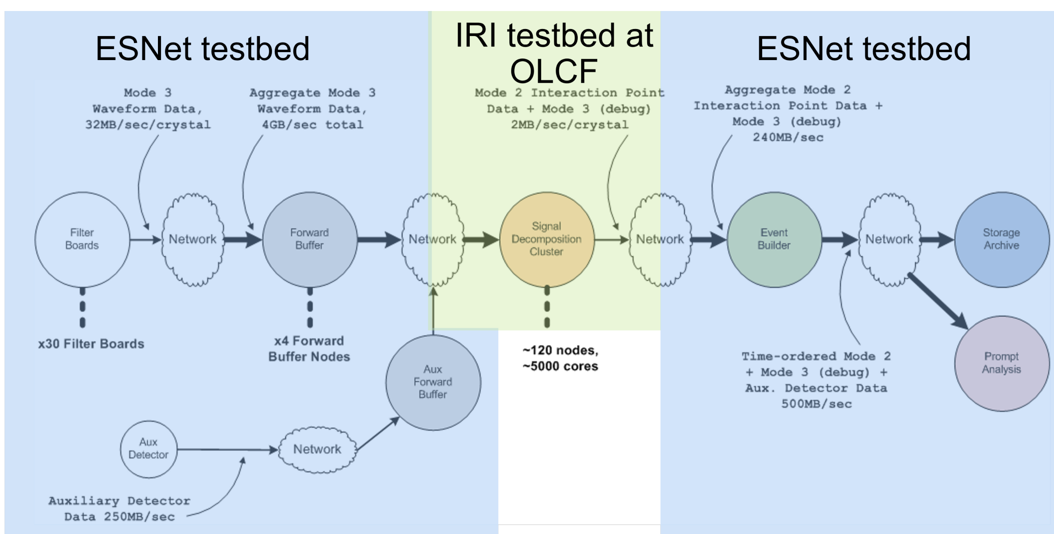

Figure 10. The GRETA/Deleria computing pipeline. ***Courtesy of DELERIA.***

Challenges. Deleria supports a time-sensitive streaming model for detectors, positioning it within the time-sensitive pattern in the IRI context 10. While GRETA operates at a single facility, Deleria envisions a multi-facility setup to enhance data rates and analysis complexity, feeding real-time decisions back to experimenters. This requires: (1) Seamless connection of components over wide-area networks; (2) Resource provisioning across multiple facilities within a single workflow; (3) Guaranteed resource availability during experiments. These challenges necessitate both technical advancements and policy decisions.

Leveraging ACE. Deleria has been deployed across the ESnet and ACE testbeds to demonstrate a distributed experimental computing pipeline. Deleria containers are launched on Defiant through the Slurm API, streaming data between ESnet hosts and Defiant for online analysis. The network setup involves Layer 2 VPN services to ensure secure, efficient data transfer across the testbeds. Deleria's capabilities were demonstrated in three phases. In Phase 1, we emulated a single detector and data stream with analysis running on Defiant, successfully provisioning resources and transferring events between the ESnet and ACE testbeds. Phase 2 involved simulating GRETA's quad detector with four detectors, achieving a data rate of 4 Gb/s, and experimenting with different event sizes and analysis delays to understand performance trade-offs. Finally, in Phase 3, we scaled the setup to 36 simulated detectors, achieving a sustained bi-directional data rate of approximately 35 Gb/s. With further tuning, we expect to hit an aggregate data rate of 40 Gb/s, the current network limit.

New Capability. The project has established a pipeline for real-time streaming of simulated data from GRETA detectors across the ESnet and ACE testbeds to Defiant for analysis.

Next Steps. Deleria will transition from simulation to live data streaming around February 2025, when the GRETA detector modules are completed. This will validate real-time data streaming and analysis over long distances. Additional support for OLCF services like APIs, data streaming, and container automation is planned. The toolkit will also be further generalized to support a broader range of experiments, with enhancements to the forward buffer, analysis client, and orchestration capabilities.

2.3.2 ESnet JLab FPGA Accelerated Transport Load Balancer (EJFAT)

Pilot Description. The EJFAT project started as a collaboration between ESnet and JLAB to improve processing for event data from various DOE accelerator facilities. The idea was to create a load balancer that can take input from Data Acquisition (DAQ) channels and distribute it to multiple compute resources for processing. The implementation has taken the form of multiple custom FPGAs that handle the reception and distribution of UDP packets in real-time. Multiple FPGAs can be used to increase the overall data capacity and the number of FPGAs, DAQ channels and compute nodes can all be adjusted independently. (See [10] for more details about EJFAT's design and implementation.) At the OLCF, we are using the Defiant cluster in the ACE testbed as a compute resource that can accept data from the EJFAT load balancer. This provides the EJFAT developers with a third laboratory environment to experiment with and one whose network architecture and policies are different from the initial two labs.

Challenges. One of the main challenges regarding EJFAT is the fact that compute resources require publicly routable IP addresses. For testing purposes, we added public IPv4 addresses to four of Defiant's compute nodes, however, this is a new paradigm for OLCF. Typically, only the login nodes for our clusters have public IP addresses and the compute nodes use private RFC-1918 IP addresses. Indeed, OLCF does not even have enough IPv4 addresses to assign one to every compute node.

Additionally, since EJFAT streams packets over UDP, there is no error correction or ability to resend a packet. If a packet is lost or damaged in transit, the data is simply lost. This means the all the network links between the accelerator facility, the EJFAT load balancer and the compute facilities must be designed to minimize packet loss.

Leveraging ACE. Initial tests streamed about 700Mbit/sec into Defiant. The next phase of testing will try to increase the transfer rate to 80Gbit/sec - which is the limit of the current network connections - or to whatever maximum speed the four nodes on Defiant can manage. Interestingly, the design of EJFAT does not require the compute resources to all be in the same facility or data center: one of the initial tests used four nodes on Defiant, two nodes on Perlmutter and one node on a system at JLAB to analyze event data from a single stream.

New Capability. The EJFAT project highlights the necessity for a general-purpose data streaming architecture. Despite significant efforts by multiple engineers to configure network settings, the results were less than optimal: only four nodes on Defiant were usable by EJFAT, the network changes were specific to this project and had no broader applicability, and the adjustments were not feasible for implementation on large production systems. For a more detailed discussion, refer to Section 6.

Next Steps. The EJFAT developers have a plan to add the concept of a "data forwarder" or "gateway" to the EJFAT protocol. That would remove the need for compute nodes to have publicly routable IP addresses and would mesh nicely with the data streaming node concept we have been developing at OLCF. Our plan for the coming year is to work with the EJFAT developers to integrate these new features into our data streaming architecture and test them using ACE testbed resources (See section 4.1.1).

3. ACE COMPUTING ENVIRONMENTS

ACE encompasses a diverse array of cutting-edge computing environments designed to propel scientific discovery and innovation. At the heart of ACE is the integration of three key components: the SummitPLUS program, the Frontier exascale system, and the ACE testbed. The SummitPLUS program extends the capabilities of the already powerful Summit supercomputer, offering new allocation opportunities for 2024. Frontier, the world's first exascale supercomputer, offers unparalleled computational power, allowing researchers to tackle the most complex and data-intensive challenges. Complementing these systems, the ACE testbed provides a flexible and open-access platform for testing and validating emerging technologies and innovative solutions in a controlled environment.

3.1 SummitPLUS

In FY24, OLCF has launched the SummitPLUS program, a new allocation initiative aimed at providing continued access to the IBM AC922 Summit supercomputer from January to October 2024. SummitPLUS opens its doors to researchers from academia, government laboratories, federal agencies, and industry, encouraging proposals for computationally ready projects. The program is designed to extend the life and scientific impact of Summit, which debuted in 2018 as the world's most powerful supercomputer and has since been instrumental in advancing research in climate, energy, public health, and national security. With a peak performance of 200 petaflops, Summit remains a vital resource for the scientific community, enabling large-scale simulations, data-intensive computing, and AI-driven research. By offering substantial allocations to diverse projects, SummitPLUS ensures that researchers can continue to leverage Summit's unparalleled capabilities for groundbreaking scientific discoveries.

SummitPLUS plays a pivotal role in advancing the IRI program by providing a robust platform for executing and evaluating complex scientific workflows. SummitPLUS contributes to IRI's vision by facilitating projects that encompass traditional simulations, data science, and AI applications, thereby fostering the convergence of simulation, data, and learning. This initiative supports the IRI program's goals of seamless collaboration and data sharing, enabling researchers to tackle data-intensive challenges with enhanced efficiency and effectiveness.

During the SummitPLUS call, we received 166 proposals and awarded 107 allocations. Among these, 9 projects were specifically aligned with IRI themes. These 9 projects encompass all three IRI patterns: time-sensitive, data integration-intensive, and long-term campaigns. The SummitPLUS program has thus facilitated substantial scientific advancements within the IRI program. Notably, one project involves steering experiments at the SLAC LCLS X-ray laser facility, guided by foundational models trained on Summit, highlighting the program's impact on cutting-edge research.

3.2 Frontier

OLCF's Frontier exascale system represents a monumental leap forward in computational capabilities, providing unprecedented power and speed to tackle the most challenging scientific problems. As the world's first exascale supercomputer, Frontier is designed to perform over a quintillion calculations per second, leveraging advanced architectures that include AMD Epyc CPUs and AMD Instinct GPUs. This powerhouse system facilitates high-fidelity simulations, complex data analyses, and large-scale machine learning applications, making it an indispensable resource for researchers across various scientific disciplines. Frontier's immense computational capacity not only accelerates discovery but also enables the exploration of new scientific frontiers that were previously out of reach.

Frontier's capabilities are integral to the advancement of the IRI program, which aims to create a cohesive and interconnected research environment. By providing a robust platform for executing IRI workflows and patterns, Frontier enhances the ability to process and analyze massive datasets in near real-time, supports the integration of diverse computational resources, and enables seamless collaboration across multiple research domains. The system's advanced architecture allows for the efficient execution of time-sensitive and data-intensive applications, as well as long-term campaigns, driving innovations in areas such as climate modeling, genomic research, and materials science. Moreover, Frontier's role in the IRI program extends to fostering new collaborations and partnerships, democratizing access to cutting-edge computational resources, and setting a new standard for performance and scalability in scientific computing. Through these contributions, Frontier not only advances the goals of the IRI program but also propels the broader scientific community toward groundbreaking discoveries and innovations.

3.3 Testbed

The ACE testbed is a distinctive feature of the OLCF, offering a centralized, sandboxed environment designed to deploy and evaluate heterogeneous computing and data resources. This testbed plays a crucial role in advancing the productization of new HPC technologies in alignment with the OLCF and DOE missions. By providing an open-access environment equipped with HPC production-capable resources, the ACE testbed enables researchers and HPC system architects to assess existing and emerging technologies without the constraints typically associated with a production environment.

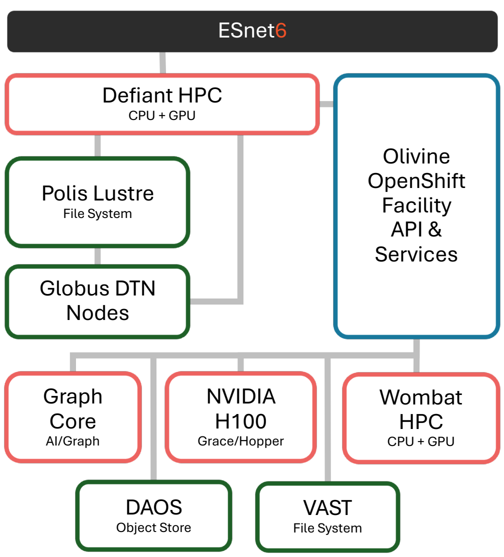

Figure 11. The ACE testbed provides diverse cutting-edge computing resources for HPC workloads, including Defiant with AMD EPYC™ CPUs and MI100 GPUs, Wombat with heterogeneous AArch64 CPUs and NVIDIA™ GPUs, GraphCore's AI acceleration with BOWPod$_{16}$, Holly with NVIDIA H100 GPUs, Quokka's general-purpose Intel™ Xeon™ nodes, and Olivine's OpenShift™ cluster for long-running services.

The ACE testbed is structured around several key areas of interest, which are pivotal to the evolution and enhancement of HPC systems:

- IRI Workflows and Patterns – Developing and testing workflows that are both time-sensitive and data-intensive, ensuring they meet the rigorous demands of scientific research.

- Emerging Compute Architectures – Exploring advanced computational techniques and architectures, including evolving processor architectures and compilers, AI appliances, and reconfigurable computing systems.
- Emerging Storage Architectures – Investigating novel storage solutions, including object storage, to manage the growing data needs of HPC applications.
- Emerging Network Architectures – Evaluating innovative network technologies, such as Data Processing Units (DPUs), to enhance data throughput and reduce latency.
- Cloudification of Traditional HPC Architectures – Implementing cloud-based solutions for HPC, including multi-tenancy and preemptible queues, to increase flexibility and resource utilization.

The ACE testbed is poised to continue its pivotal role in advancing HPC technologies and methodologies. Upcoming efforts will focus on further enhancing the testbed's capabilities, expanding its range of applications, and fostering collaboration with a broader scientific community. By addressing key challenges and exploring innovative solutions, the ACE testbed aims to drive the next generation of scientific discovery and computational excellence.

The ACE testbed offers a variety of cutting-edge computing resources designed to support diverse HPC workloads and facilitate comprehensive evaluations. It includes Defiant, a 36-node cluster with AMD EPYC™ CPUs and AMD MI100 GPUs (previously part of the Frontier early access system). Wombat, a heterogeneous AArch64 cluster, featuring CPUs like Fujitsu A64FX, Ampere™ Computing Altra™, NVIDIA™ Grace™, and NVIDIA GPUs with InfiniBand networking. In addition, two stand-alone nodes provide dedicated AI acceleration with a BOWPod_{16} cluster (GraphCore) and a Supermicro server containing 8 NVIDIA H100 GPUs connected via HDR InfiniBand (Holly). Another cluster, Quokka, comprises 16 general-purpose nodes with Intel™ Xeon™ CPUs connected via NDR200 InfiniBand. Olivine serves as an OpenShift™ Kubernetes platform cluster for long-running or internet-facing services.

The ACE testbed is poised to continue its pivotal role in advancing HPC technologies and methodologies. Upcoming efforts will focus on further enhancing the testbed's capabilities, expanding its range of applications, and fostering collaboration with a broader scientific community. Future hardware additions include the InspireSemiconductor Thunderbird "supercomputer on a chip" packaging 1,536 64-bit custom RISC-V CPU cores per chip. The Wombat cluster, established in 2018, continues to evolve, supporting specialized research fields and utilizing NVIDIA BlueField™ DPUs. Table 2 presents detailed descriptions of the ACE resources.

The ACE testbed, in collaboration with industry vendors, incorporates advanced storage solutions to support high-performance data operations. The Polis storage system, based on Lustre, offers approximately 1.6 PB of capacity, primarily using spinning disks with some flash storage, and is connected to Defiant for high-speed data access. VastData, an NFS-over-RDMA storage appliance, provides around 600 TB of flash storage integrated with the InfiniBand (IB) fabric. Additionally, the DAOS object storage system includes eight servers, each with about 30 TB of flash storage, connected via dual NDR200 IB connections, providing a total aggregate bandwidth of 3.2 Tbps.

The ACE testbed has been successfully integrated into the expansive IRI testbed ecosystem through the Energy Sciences Network (ESnet). This integration marks a significant enhancement of the IRI testbed, which will now encompass advanced computing testbeds from the Argonne Leadership Computing Facility (ALCF), the National Energy Research Scientific Computing Center (NERSC), and Jefferson Lab (JLab). Each of these testbeds, interconnected via the high-speed ESnet, forms a robust and collaborative research infrastructure designed to push the boundaries of computational science and data-driven research. The network architecture, illustrated in Figure 12, showcases the intricate connections and the strategic deployment of resources across the ASCR facilities.

Resource	Description
Defiant	• 36 nodes with AMD Epyc CPUs and 4 AMD MI100 GPUs per node. • Slingshot 10 networking. • Previously served as the Frontier early access system.
Wombat	• Heterogeneous AArch64 cluster • Different CPUs: Fujitsu A64fx, Amere Compting Altra, Nvidia Grace • Nvidia Ampere and Hopper GPUs in the Ampere and Grace nodes • InfiniBand networking - HDR & NDR
GraphCore	• Dedicated AI acceleration appliance • BOWPod$_{16}$ configuration: `https://www.graphcore.ai/products/bow-pod16`
Holly	• Single Supermicro server equipped with 8 NVIDIA H100 GPUs • 8 HDR IB network connections (1.6tbps aggregate bandwidth)
Quokka	• 16 general-purpose nodes with Intel Xeon CPUs • NDR200 networking
Olivine	• Openshift (Kubernetes) cluster • OVN-Kubernetes advanced networking • 1808 CPU cores, 3.1 TB memory • 1 NVIDIA A100 GPU • Useful for long-running jobs or intenet-facing services • Some nodes have high-speed network connections to Defiant and some nodes have high speed IB connections to the Polis filesystem
Future Hardware	• InspireSemiconductor Thunderbird - HPC accelerator based on the RISC-V architecture - expected Q3 2024 • Defiant compute technology refresh - expected Q3 2024

Table 2. Summary of ACE resources description.

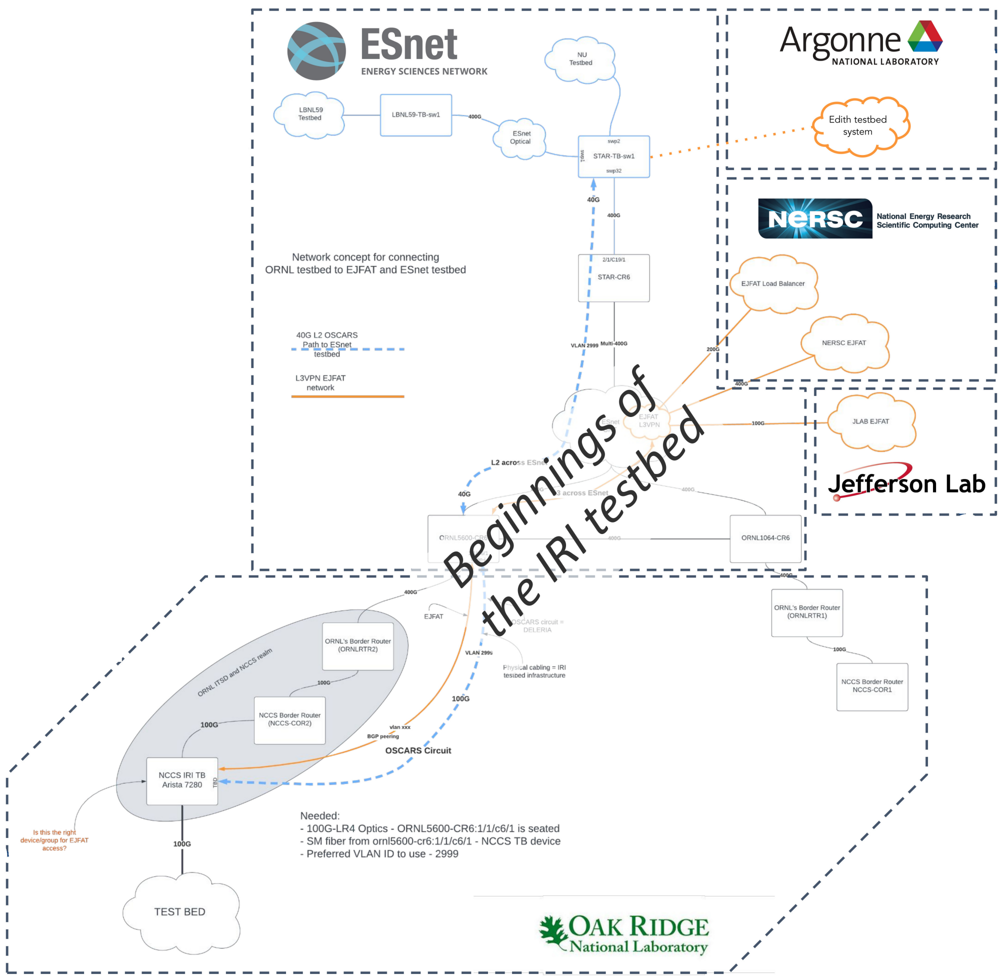

Figure 12. Network diagram of the early ESnet testbed showcasing interconnectivity between ACE and testbeds at partner ASCR facilities [11].

4. RESEARCH AND DEVELOPMENT OF FOUNDATIONAL TECHNOLOGIES

4.1 Data Movement

4.1.1 Data Streaming

Cutting-edge science is increasingly data-driven due to the emergence of scientific machine learning models that can guide scientists toward fruitful areas of exploration. Experimental science facilities such as light and neutron sources, particle colliders, and radio astronomy telescopes are also producing raw measurement data at rates that exceed available data storage and computing capacity at those facilities. As a result, scientific workflows are being developed that concurrently couple experimental science facilities with HPC facilities to enable analysis of observational data while the experiment is ongoing, and where analysis results are potentially fed back to the experiment in a time-sensitive manner for use in control decisions such as instrument tuning or to steer the experiment. These scientific workflows require foundational technologies that enable bidirectional streaming of data between the experimental science facilities and HPC systems located within computational facilities. To meet the time-sensitive demands, such technologies should emphasize memory-to-memory data movement (i.e., moving data directly from the memory of producer applications to the memory of consumer applications) and avoid requiring data transfers between file systems as the basis for moving data. Although file-based data movement is already well-supported by the DOE ASCR user facilities, it has been shown to be a bottleneck for certain use cases [12, 13]. We refer to technologies enabling such a bidirectional data streaming capability as *Data Streaming to HPC*, or DS2HPC for short.

The ASCR IRI Task Force's Architecture Blueprint Activity (IRI ABA) produced a final report [14] that identifies three integrated science patterns important to the success of DOE integrated scientific research. Of these patterns, two encompass potential use cases for DS2HPC. In the "Time-sensitive" pattern, an experimental feedback loop may require low or near real-time latency for data streaming into and out from the HPC system. In the "Data Integration-intensive" pattern, analysis performed on HPC systems is focused on integration of large quantities of data from disparate sources. Use cases matching this pattern may involve multiple independent data stream sources, such as widely distributed sensors or a set of unique instruments/detectors. A recent report [15] from ESNet analyzes the prevalence of IRI blueprint patterns across the breadth of DOE-SC programs, including Biological and Environmental Research (BER), Basic Energy Sciences (BES), Fusion Energy Sciences (FES) High Energy Physics (HEP), and Nuclear Physics (NP). The report's findings indicate that for all the use cases studied by ESNet, a mere 1% do not include a time-sensitive or data integration-intensive component in their scientific workflows, and 39% exclusively use one or both of these patterns. Since data streaming is identified as a key gap for both patterns in the IRI ABA report, our expectation is that a DS2HPC capability will have a broad user base across DOE's science portfolio.

Our recent report [16] examines several science use cases that could benefit from a DS2HPC capability provided by OLCF. We identified eight data streaming workload characteristics that help demonstrate the diversity of streaming needs, as reproduced below in Table 3.

Our analysis of these characteristics for the science use cases suggests that a single technological solution for data streaming cannot be a one-size-fits-all solution that meets the needs of all use cases. Instead, we proposed an architectural approach that is flexible and supports the deployment of many technological solutions, thereby providing the best opportunity for utility to many current and future use cases. We also introduced the eight core requirements of this architectural approach for DS2HPC, building upon prior work by SciStream [17]. These core requirements are briefly summarized as:

- R1: Data streaming from the memory of data producers to the memory of data consumers.

Use Case Characteristic	Description	Values or Metric
Experiment Concurrency with HPC	When the data analysis or processing on HPC occurs in relation to the scientific experiment that produces the data.	*Concurrent* - Analysis or processing occurs during the runtime of the experiment. *Sequential* - Analysis or processing occurs after the experiment has completed.
Workflow Pattern	How the results of the data analysis or processing on HPC are used in the scientific workflow.	*Monitoring* - Results are used to monitor the progress of an ongoing experiment. *Steering* - Results are used to modify the configuration or actions of an ongoing experiment. *Design* - Results are used for design or configuration of subsequent experiments.
Data Production Periodicity	Whether experiment data is streamed to the HPC system continuously or in periodic bursts.	*Continuous* - Data is produced at a relatively constant rate across the lifetime of the experiment. *Bursty* - Data is produced by the experiment in bursts separated by periods of no data production.
Data Consumption Semantics	How data consumers on the HPC system receive the data.	*Push-to-HPC* - Data is streamed to the HPC consumer(s) immediately upon production. When there are multiple consumers, this method requires a pre-determined data distribution among consumers (e.g., round-robin). *Pull-from-HPC* - The HPC analysis is notified when data is available on the stream and consumers choose when and how to pull the data.
Data Stream Elements	The type of data element that is produced and consumed via the data stream.	*Events* - Each element is a unique event. The data contained within an event must provide the required information for establishing uniqueness from all other events. *Files* - Each element corresponds to a file. *Messages* - Each element is a message containing arbitrary data.
Data Stream Persistence	Whether data from the stream should be persisted, and for how long.	*None* - Stream data is available until consumed. *Space-limited* - Stream data is buffered within a durable storage area of a specified size where it remains available until being overwritten by more recent data. *Time-limited* - Stream data is buffered in a durable storage area where it remains available for a specified duration after it is first consumed. *Persistent* - Stream data is persisted to a durable storage area where it remains available until it is explicitly deleted.
Data Stream Bandwidth	The desired data throughput for a data stream. The stream will deliver at least the desired throughput when data is produced at an equivalent or higher rate.	megabytes/second (MB/s) gigabytes/second (GB/s)
Data Stream Latency	The desired maximum latency for end-to-end delivery of a data element from a producer to a consumer. The individual data elements are expected to be small (e.g., less than a kilobyte).	seconds (s)

Table 3. Data Streaming Characteristic Definitions

- R2: Secure data streams that span multiple domains, each with a possibly distinct security context.
- R3: Separation of data streaming control and data planes.
- R4: Adaptability to both advanced reservation and on-demand allocation of data streaming resources.
- R5: Data stream endpoints may be both a data producer and a data consumer.
- R6: Data streams may be utilized by multiple concurrent endpoints.
- R7: Flexibility to deploy alternative technologies for the streaming data plane solution.
- R8: Streaming data plane paths may traverse multiple networks within a single security domain.

The proposed architectural approach to support *Data Streaming to HPC* for the OLCF is shown below in Figure 13. This architecture focuses on providing the necessary capabilities for streaming workflows to establish secure, bidirectional, memory-to-memory data streams between application processes at external science facilities and application processes running within an OLCF HPC system. *Data Streaming Nodes* (DSN) that serve as gateway hosts bridging the public WAN and internal OLCF networks using high-speed network adapters are a key physical infrastructure component of the architecture. DSNs are user-allocatable resources managed by a container orchestration system. Multiple DSNs may be allocated to a single workflow to meet specific network bandwidth, message processing rate, or resilience needs of a deployed streaming service. DSNs are independent from OLCF Data Transfer Nodes (DTNs) to avoid interference from file-based data transfers.

Science workflows interact with the OLCF Facility API (see Section 4.2.1) to authenticate and deploy a *Data Streaming Service* (DSS) on DSNs in concert with a corresponding compute job on an OLCF HPC system. A DSS is a set of one or more coordinating processes that provide the data plane for streaming data between external and internal *Application Stream Endpoints*. The data streaming functionalities, abstractions, and semantics provided by a DSS may vary across technological solutions.

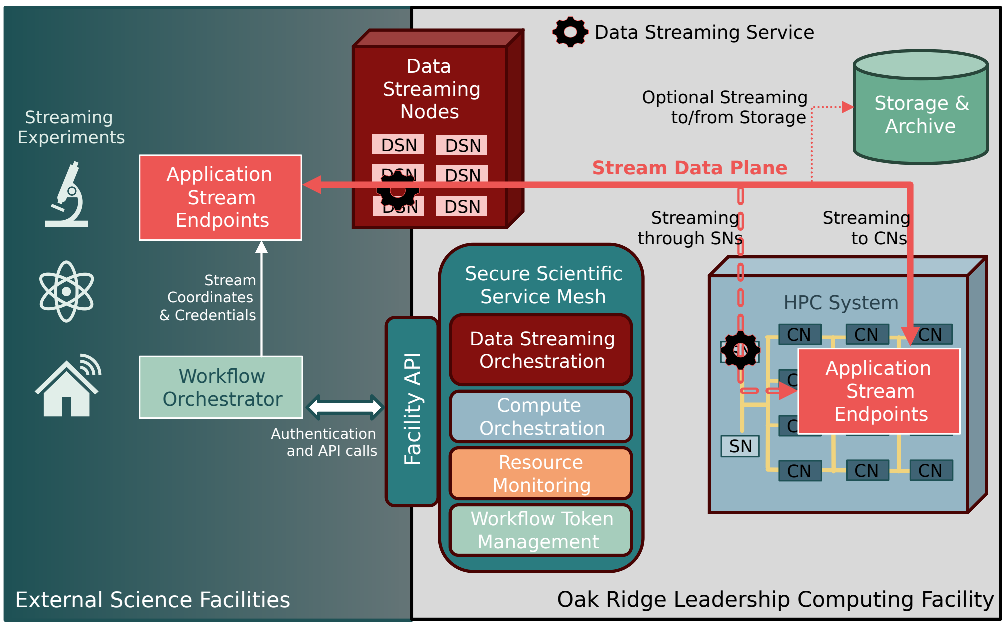

Figure 13. OLCF Architecture for Data Streaming to HPC.

We are currently working to demonstrate a full end-to-end realization of the DS2HPC capability on the ACE Testbed using a commodity streaming framework (i.e., RabbitMQ). Four data streaming nodes with 100 Gbps network adapters have been installed and configured within the Olivine OpenShift cluster. A prototype implementation of the Data Streaming Orchestration service within the OLCF Facility API has been completed that enables users to request a RabbitMQ deployment across one or more DSNs. A benchmark application simulating external science data producers and consumers and internal HPC application stream endpoints is under development to evaluate both latency-sensitive and throughput-oriented streaming.

4.1.2 SciStream Deployment

One of the goals of enabling DS2HPC capability on OLCF's ACE testbed is to deploy existing architectures and toolkits that address the challenges of memory-to-memory data streaming from scientific instruments to remote HPC environments. Deploying and exploring these existing toolkits would enable us to tackle data streaming challenges from scientific workflows to testbed compute resources, particularly from the perspective of external connectivity. Additionally, it would enable us to evaluate complementary technologies like Globus, which can be integrated with these streaming frameworks for enhanced functionality.

SciStream [17] is one such tool that tackles the infrastructural challenges necessary to enable these memory-to-memory data transfers between instruments and HPC. SciStream addresses three primary challenges: handling data transfers across security domains between data producers (e.g., scientific instruments) and consumers (e.g., HPC systems); supporting delegated authentication and authorization within broader scientific workflows; and decoupling sophisticated identity and access management from applications, minimizing changes and reusing existing security architectures.

An initial discussion with the SciStream team welcomed a collaborative effort to deploy it on the ACE testbed. SciStream's architecture utilizes gateway nodes that act as intermediaries, connecting the internal instrument/HPC network with the external wide area network (WAN). The Olivine cluster in the ACE testbed includes four high-bandwidth Data Streaming Nodes (DSNs) with both internal and external 100Gbps connectivity. These DSNs are ideally suited to serve as gateway nodes for SciStream. The tool supports various setups, allowing for one or more DSNs depending on specific use cases and throughput or data load balancing requirements.

SciStream is composed of three software components: SciStream User Client (S2UC), SciStream Data Server (S2DS), and SciStream Control Server (S2CS). S2UC orchestrates end-to-end data streaming on behalf of an end user, S2DS operates on gateway nodes to act as a proxy between the internal network (LAN or HPC interconnect) and the external WAN, and S2CS manages resources on a gateway node, including initiating and terminating S2DS processes. All of SciStream's components - S2UC, S2DS, and S2CS - can be deployed as Docker containers within the Olivine OpenShift Kubernetes platform cluster. Additionally, SciStream integrates with existing authentication and authorization systems, such as Globus Auth, ensuring secure communication between participating facilities.

4.2 Interfaces

4.2.1 Secure Scientific Service Mesh - S3M

The Integrated Research Infrastructure (IRI) initiative depends on supplemental interfaces to extend access to operational facility-provided HPC resources. This dependency stems from the requirement to enable different modes of access to systems than what has been enabled historically. It is less about enabling access to HPC systems for individuals and more about enabling HPC system access on the backend of externally managed software systems. These externally managed software systems may be domain-specific science portals, experimental workflow systems hosted at instrument facilities that need to offload compute, or a data analysis workflow system that is integrated with another data-oriented DOE user facility. In all of these scenarios the desire is to integrate and utilize the HPC facilities as a backend computational resource. One way to think about it is with the HPC facility as an accelerator, analogous to a GPU in an HPC system, within a larger scientific workflow.

The challenge is an operational one. To enable these new modes of integration, HPC facilities must provide supplemental and operationally trusted interfaces. Developing these new supplemental interfaces will require significant software engineering efforts on behalf of the operational staff. Additionally, HPC facilities must develop clear security policies for how these supplemental interfaces can be integrated into externally manage software systems. In summary, the operational challenge of IRI is extending facility-provided HPC resources to external systems in a trusted and secure way. Traditionally, HPC facilties have provided their resources to users who are vetted during the application process. These users then access the facility environment directly with their user credentials. In the IRI case it is not about vetting users for individual access but integrating software systems that users drive scientific workflows through.

The OLCF's approach to these challenges is to develop a Secure Scientific Service Mesh (S3M). S3M will enable a flexible and secure framework for workflow integration. Within the S3M framework the OLCF

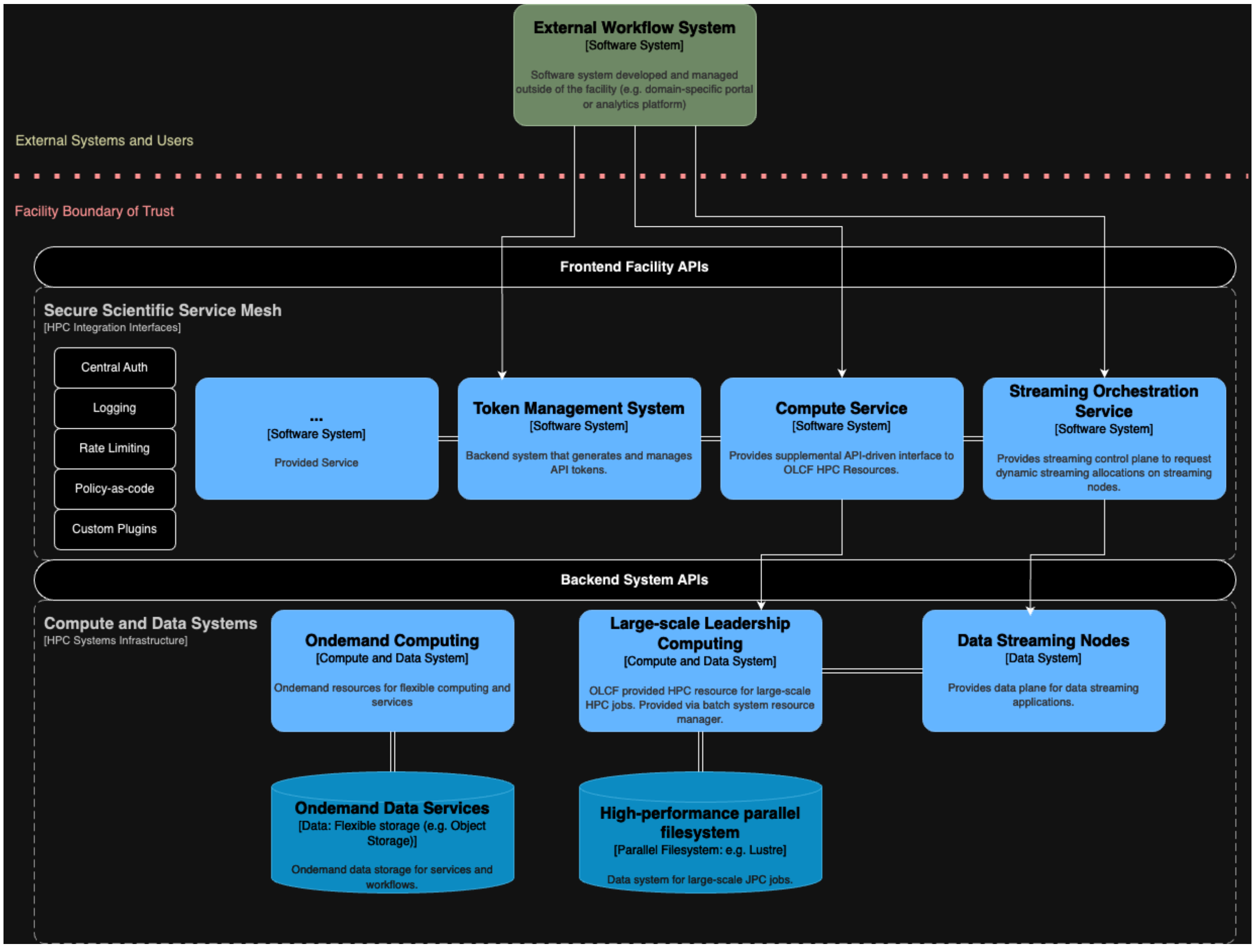

Figure 14. S3M integration with backend HPC systems and external workflow systems

will develop and integrate services that expose supplemental interfaces for the IRI. These services will enable flexible and secure interfaces to OLCF-provided resources, through a Facility API and other potential mechanisms, for new modes modes of application trust and access in support of IRI. At the core, S3M will provide a firewall for scientific web-based workflow integration. This includes the ability to define policy-as-code, implement rate limiting, handle centralized auth, and perform log analysis.

From a user experience and integration perspective, S3M extends compute and data system access in new ways. This doesn't replace traditional modes of access (e.g. SSH) but provides a controlled bridge to other systems. By embracing this, it empowers OLCF staff to develop and provide new frontend capabilities but it also allows for clear integration patterns between OLCF and systems outside of OLCF. We can start to see OLCF as an ecosystem that can be utilized similar to a cloud environment.

4.3 Scheduling

4.3.1 Preemptive Scheduling

Defiant has the ability to support an on-demand queue for high priority workflows through the use of preemption as part of the Slurm workload manager configuration. Through this configuration, lower priority jobs are subject to an early, yet graceful, termination to free up system resources required for more urgent, high priority jobs. Some projects, depending on the time sensitivity of recurring workflows, can make great use of this capability.

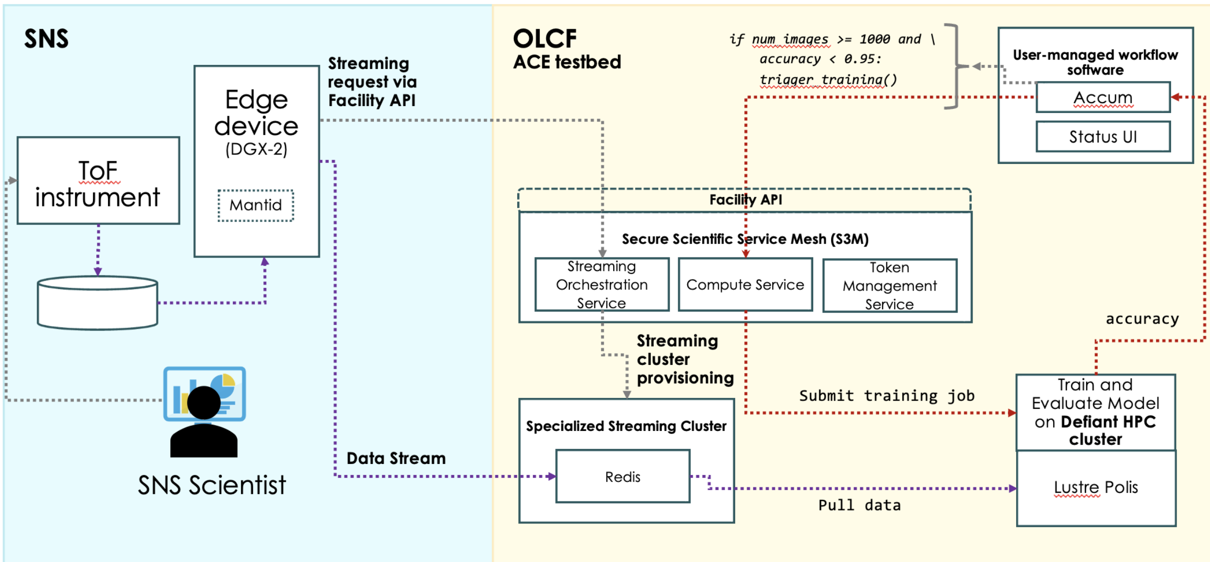

Figure 15. Edge To ExaAI example workflow with streaming and compute

Slurm provides a Quality of Service (QoS) feature that allows for certain effects to be applied to an associated job which will have an impact on the scheduling process. At present time, multiple QoS's have been created on Defiant, each providing varying degrees of priority and preemptive capabilities. A higher priority QoS will grant a submitted job a higher priority in the queue and will have the ability to preempt jobs associated with a lower priority QoS if sufficient resources are not available for the highest priority workload. QoS assignments are granted to Slurm accounts, which reflect the organization project membership. Users then inherit QoS permissions through the accounts that they are members of. Using this strategy, we can grant higher priority on the system on a per project allocation basis.

QoS priorities are not applied to job submissions by default on Defiant. Rather, the user must specify which QoS they wish to use with their submission. The QoS's available to a given user will be the ones that are granted to their project allocations. In the event that the system ever experiences a full load, i.e. there is a wait for submitted jobs to run, and multiple jobs are submitted with the same QoS, normal behavior returns with a fight for priority. To prevent this from happening and negating the preemption capabilities, recurring reservations can be put in place which guarantee resources are available for specific projects that require them at certain times. This customized node availability must be requested and coordinated between users and system administrators, but it provides the benefit of configuring the system in such a way that can guarantee high priority workloads are able to run at specific times.

In order to validate and verify the scheme of Slurm workload manager, we set up and carried out a simulation of its behavior. We created 8 dummy users and automated a stochastic job submission scenario from these user's accounts where the users submitted a random number of jobs at random intervals. Each job was in turn generated with randomly generated parameters such as walltimes, QoS requirements and node counts. The jobs were setup to run an mpi based program across the allocated nodes with nominal IO requirements. We observed the scheduler behavior in the next 24 hours and collected the jobs data. We plotted the data to observe how the jobs ended up for each of these 8 dummy users. The plot shown in Figure 16 shows how each user's jobs fared. In addition to giving insights on scheduler performance, and more importantly, this exercise helped us create instrumentation that could be used to evaluate a given Slurm policy configuration.

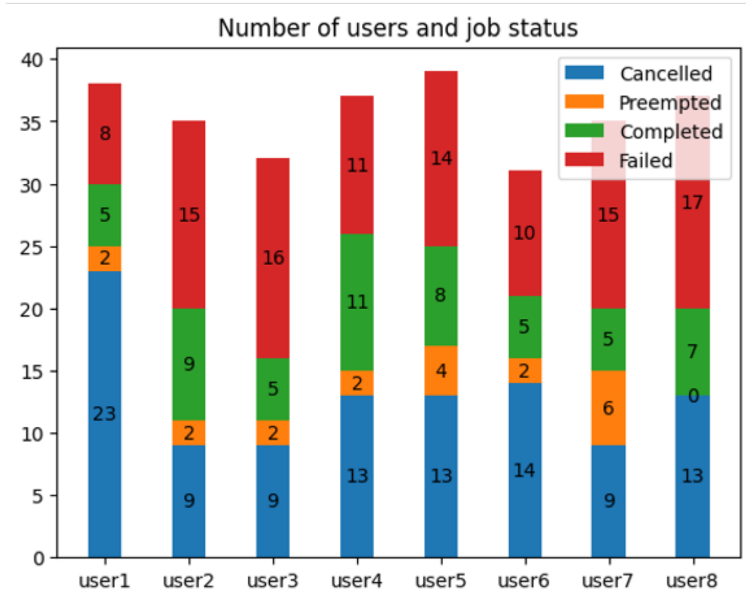

Figure 16. A Depiction of Scheduler Performance for Stochastically Submitted jobs on behalf of 8 Dummy users on the ACE / IRI testbed.

4.4 Software Deployment and Portability

4.4.1 Containers

The increasing complexity of scientific applications and the diversity of computing environments pose significant challenges for software deployment and portability. Ensuring that applications can run efficiently across different systems without extensive modifications is crucial for maximizing resource utilization and accelerating scientific discovery. Containers offer a solution by encapsulating applications and their dependencies, allowing them to be executed consistently across various platforms.

In the context of the IRI, the primary challenge is ensuring portability across diverse computing environments while maintaining compatibility with specialized hardware like GPUs and high-performance interconnects. Achieving this without sacrificing the seamless user experience is crucial. Additionally, managing the complexity of containerized environments, along with ensuring security and compliance with institutional policies, further complicates the deployment process. Performance parity with native executions remains important but is secondary to the overarching issue of portability.

The ACE testbed addresses these challenges through the use of Apptainer [18] (formerly Singularity) as the container runtime. Apptainer was chosen because it allows users to build and run container images without requiring additional privileges, enhancing security. Its use of single-file container images, known as SIF (Singularity Image Format) files, simplifies the management of storage, distribution, and execution of scientific applications. Unlike Docker's multilayered images, SIF files are easier to handle on parallel filesystems and can still be pushed to OCI-compatible registries like DockerHub for broader distribution.

Apptainer is deployed on Defiant with detailed documentation available for users. This documentation provides instructions on building and running single-node containers as well as multi-node containers requiring MPI. Ensuring MPI compatibility is critical for HPC applications. The solution involves building containers with an MPI implementation compatible with the host's MPI libraries and mounting the host's optimized MPI libraries into the container at runtime. This method ensures that containerized applications can communicate effectively over MPI and leverage the host's high-performance interconnects.

To facilitate GPU utilization within containers, Apptainer supports Nvidia and AMD GPUs through the `--nv` and `--rocm` flags, respectively. These flags enable containerized applications to access GPUs as they would in a native environment, ensuring no performance degradation. For MPI integration, Defiant uses HPE Cray's optimized MPI libraries. The containerized application is built with an MPI implementation compatible with Cray MPICH, and the required libraries are mounted from the host into the container. This allows the container image itself to be built elsewhere and then be run on whichever system hosting a compatible MPI implementation to the one in the container image. This approach ensures high performance and portability across different HPC systems.

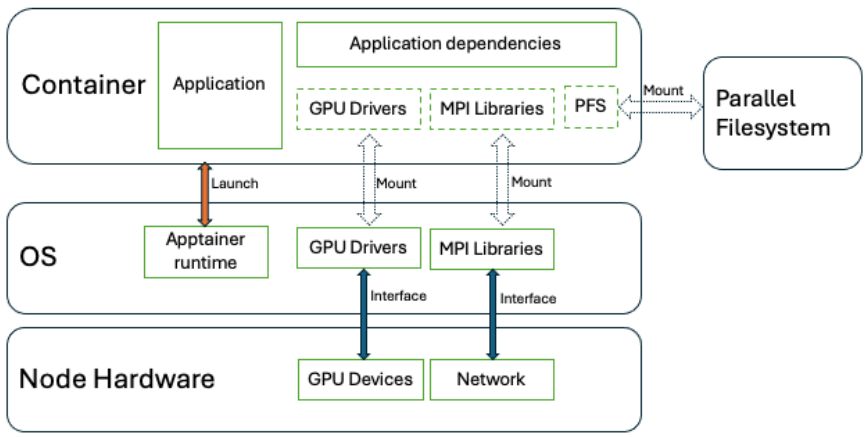

Figure 17. A container encapsulates an application and its dependencies. A container is launched by Apptainer, and it mounts GPU and MPI libraries from the host system, as well as mounting the parallel filesystem.

To ensure the robustness of containerized applications, a comprehensive test suite is being developed. This suite includes tests for MPI, ROCm, miniapps, and sample applications from AMD's Infinity Hub container registry. The test suite helps identify regressions after system or software upgrades and refines the procedures for running containers on the ACE testbed, maintaining optimal performance and reliability.

Looking forward, the full potential of containers in the ACE testbed will be realized as pilot projects progress. The focus will be on further refining containerization practices, enhancing documentation, and expanding the test suite to cover more applications and scenarios. These efforts will help reduce the complexity of maintaining repeatable builds and deploying software on new systems. By encapsulating applications and their dependencies within containers, researchers can achieve greater portability and faster setup times, ultimately enhancing productivity and collaboration in HPC environments.

4.4.2 Multi-Tenancy

IRI use cases will present a significantly more diverse collection of workloads than has been traditionally deployed on OLCF systems, and will likely exhibit different resource requirements and runtime behaviors from our current application set. OLCF system architectures have historically consisted of node designs that incorporate a large number of replicated processing elements (such as multiple GPUs and dozens of CPU cores). These designs were driven by a capability use case that includes workloads capable of saturating the hardware resources of a large number of compute nodes in a system. In contrast, a large portion of IRI workloads are expected to require substantially less than the entirety of a compute node's hardware resources (i.e., a single GPU and a handful of CPU cores). In order to support these workloads efficiently, we are exploring and developing multi-tenancy capabilities for OLCF systems in order to more efficiently leverage available computational resources.

Our intended multi-tenant model is based on a space sharing approach that allocates subsets of a compute node's hardware resources to independent workloads belonging to separate applications, jobs, and users. These resource partitions are then independently managed using hardware virtualization features to make each partition appear as a virtual compute node that is managed, allocated, and accessed identically to other physical nodes in the system. This model differs from other forms of virtualization based multi-tenancy that oversubscribe hardware resources in order to maximize utilization. Our approach does not support over-subscription, and instead explicitly determines resource partition boundaries based on the ability to provide both performance and security isolation between any two partitions.

Security isolation is a primary design goal of our approach, since IRI use cases will necessitate external users and organizations co-existing with the current OLCF user community. These external users will require secure partition boundaries that prevent interaction and resource contention between different IRI users as well as between IRI and OLCF users. Virtualization not only provides isolation due to lowering the abstraction level to near physical hardware, but also by supporting the integration of advanced confidential/trusted computing features that combine hardware based confidential computing with the virtualization architectural features such as AMD SEV and Intel TDX. By leveraging hardware based virtualization extensions to enforce partition boundaries between tenants, we can effectively co-locate multiple independent software environments on a single compute node while ensuring that each tenant is effectively isolated from other users as well as other OLCF infrastructure integrated with other parts of the system. This in turn enables much greater flexibility in the software environments available to users, and providing the means for users to have a greater ability to specifically configure the software environment to suit their workloads requirements.

Our current efforts in this area are concentrating on integrating virtual machine (VM) management capabilities into current testbed system environments. Specifically, we are deploying the capability to create virtual nodes on the Defiant testbed system and integrate those virtual nodes into the existing system management infrastructure. By doing so, we are able to support much more dynamic system configurations by provisioning VM instances as needed using the existing system management tooling. These VM instances are initialized and then managed as any other node in the system, and support multiple different software personas that map to different workload classes that we expect to be useful for IRI use cases. One such system configuration integrates with the existing Slurm service to allow job dispatch to a virtual node via a special job queue that maps to virtual nodes that are comparatively smaller in scale to the full compute nodes. This in turn allows multiple jobs to be co-scheduled on the same node by allocating two virtual nodes that happen to reside on the same physical node.

Virtualization based multi-tenancy opens up several avenues for increasing resource utilization and maximizing the available resources to support IRI use cases that are not yet able to fully saturate the local hardware resources on OLCF production system nodes. As we move forward, we are investigating the capability of fully supporting the available hardware on a compute node such as GPUs and high performance network interfaces. We are also expanding our evaluation of the approach to different hardware architectures such as ARM64 based systems available in NVidia Grace Hopper equipped testbed systems.

4.5 TRUSTID

The Trusted IRI Designs (TRUSTID) technical activity is one of the three technical subcommittees convened by the IRI leadership council in FY24. The main focus of the subcommittee is on the identification and assessment of emergent IRI design patterns in order to recommend security best practices and sow trust amongst participating facilities. TRUSTID has identified Federated Identity as a key capability required for successful support of IRI activities and has launched a pilot activity to deploy and prototype scientific Single Sign On (SSO) across at least two ASCR facilities.

4.5.1 Federated Identity Pilot

The TRUSTID federated identity pilot assessed the readiness of ASCR facilities to support some level of identity federation. We refer to federated identity management as “a process that allows for the conveyance of identity and authentication information across a set of networked systems" as defined by NIST Special Publication 800-63C [19]. At the beginning of the FY24, NERSC was the only ASCR facility that fully supported federation with several national laboratory identity providers. During this FY, OLCF has deployed infrastructure to allow us to federate with outside identity providers and has selected OneID and ORNL’s UCAMS identity databases as the first ones to integrate with. OLCF’s moderate security enclave still requires two-factor authentication, so OneID will be the only identity provider supported for that environment. In contrast, OLCF’s open enclave already supports authentication with UCAMS usernames and passwords, so there are no security significant blockers to enabling federated single sign on with UCAMS and OneID identity providers.

In addition to user-based identity federation, an additional federation design pattern has emerged related to system-to-system trusts that are established on the behalf of a user. ALCF is adopting this model more strongly than NERSC and OLCF through the use of Globus compute and data transfer services. In a system-to-system federation, users receive security tokens from services that they stand up on computational and data resources and copy those tokens to a centralized orchestration point. This type of workflow is not endorsed explicitly at OLCF or NERSC, although OLCF implicitly allows it.

Because each HPC facility supports JupyterHub and Globus data transfer services, we are working on a prototype demonstration that will feature Jupyter Python notebooks and Globus data transfers performed from within a user’s web browser. Authentication to two JupyterHub instances as well as Globus with a single Identity Provider will be the goal of this initial Federated Identity demonstration. Because the demonstration will use Jupyter notebooks, the notebooks will be published to a Git repository that is widely accessible, enabling users to fork the project and develop their own customized workflows. Where possible, the notebooks will utilize new IRI related interfaces such as Facility APIs, streaming data capabilities, and Globus compute endpoints.

4.5.2 Lingering “Iceberg" Problems

The TRUSTID subcommittee has also been identifying security and operational roadblocks that will persist even after deployment of a fully-federated identity management infrastructure at participating facilities. Scientific SSO will not solve access control issues, eliminate account creation paperwork, or provide a single point of contact for proposals and allocation grants across DOE. We have begun to call these issues the ‘Iceberg Problems’ because users’ desire for a cross-facility scientific single sign on capability is just the ‘tip of the iceberg.’ We envision a future where the following processes are federated for ease of use by scientists:

1. **Federated ‘Authorization’** – A PI can add a new research scientist to their project(s) without having to contact each facility for approval.
2. **Federated ‘Allocation’** – DOE allocates time and resources to a PI for a given Science Project in a single programmatic way that each facility trusts.
3. **Federated ‘Vetting’** – PIs and Scientific Users submit their credentials to one place for vetting and Identity Proofing instead to each identity provider required.
4. **Federated ‘Project Control Screening’** – Export control, IRB, and other reviews of scopes of work are reviewed by a single body instead of at each facility the project is active.

Many policies will need to be changed in order to achieve this future, however. DOE policies, lab-wide policies, facility policies, and the prime contracts between DOE and M&O contractors that run the national laboratories all contain directives that must be satisfied as current processes are changed at DOE User Facilities.

5. OTHER IRI DEVELOPMENT AND EVALUATION TECHNOLOGIES

5.1 Integrating Quantum Computing and HPC

Quantum computing represents a transformative technology poised to revolutionize computational capabilities, particularly for applications requiring immense processing power. The integration of Quantum Processing Units (QPUs) with traditional HPC systems is crucial for realizing the potential of quantum computing within the IRI as a future hybrid use case, possibly crosscutting all there IRI workflow patterns. The OLCF at ORNL is spearheading efforts to integrate quantum computing into its HPC ecosystems, thus addressing DOE's mission needs in energy, materials, and computational sciences.

One of the primary challenges in integrating quantum computing into HPC systems is the nascent state of quantum hardware. Current quantum devices, including superconducting circuits, trapped ions, nuclear and electron spin qubits, and optical qubits, are still evolving and face significant hurdles such as qubit decoherence, error correction, and scalability. These challenges are compounded by the need for high-fidelity qubit control and low-latency communication interfaces between quantum and classical processors. Moreover, managing the physical infrastructure, such as cryogenic temperatures and ultra-high vacuum conditions, adds another layer of complexity.

OLCF is developing a comprehensive framework for integrating quantum devices into its HPC systems (Figure 18) [20]. This involves a detailed evaluation and benchmarking of various quantum hardware technologies to identify the most suitable candidates for integration. The goal is to create a hybrid computing environment where QPUs serve as accelerators for specific computational tasks, analogous to the role of GPUs in current HPC systems. This hybrid model allows for the offloading of certain computationally intensive tasks to the QPUs while leveraging classical HPC resources for other parts of the computation.

The integration strategy includes a multi-phase approach:

1. Mission Need and Alternatives Analysis – Identifying science driver applications that can benefit from quantum computing. These applications span energy, earth sciences, materials, and national security domains.
2. Benchmarks and Requirements Gathering – Developing benchmarks to evaluate the performance of quantum devices and gathering detailed requirements for integration.
3. Procurement and Installation – Based on the analysis, procuring suitable quantum hardware and integrating it into the existing HPC infrastructure.
4. Optimization and Operations – Continuously optimizing the hybrid computing environment to ensure efficient performance and transitioning to full-scale operations once the hardware matures.

Quantum computing holds the promise of substantial advancements in various scientific fields by providing computational capabilities that far exceed those of classical systems. For the IRI program, quantum computing could offer breakthroughs in modeling complex systems, optimizing large-scale operations, and solving intractable problems in materials science and quantum chemistry. The successful integration of quantum and classical computing resources is expected to enhance the overall efficiency and effectiveness of the IRI, driving innovation and enabling new scientific discoveries.

The next steps involve rigorous testing and validation of the integrated quantum-HPC framework. This includes deploying prototype systems, conducting extensive user feedback sessions, and refining the software ecosystem to support a wide range of applications. Additionally, OLCF plans to expand its collaboration with quantum hardware vendors and the broader scientific community to ensure the seamless integration and scalability of quantum computing technologies within the IRI framework. This ongoing effort will be critical in positioning quantum computing as a cornerstone of future scientific research infrastructure.

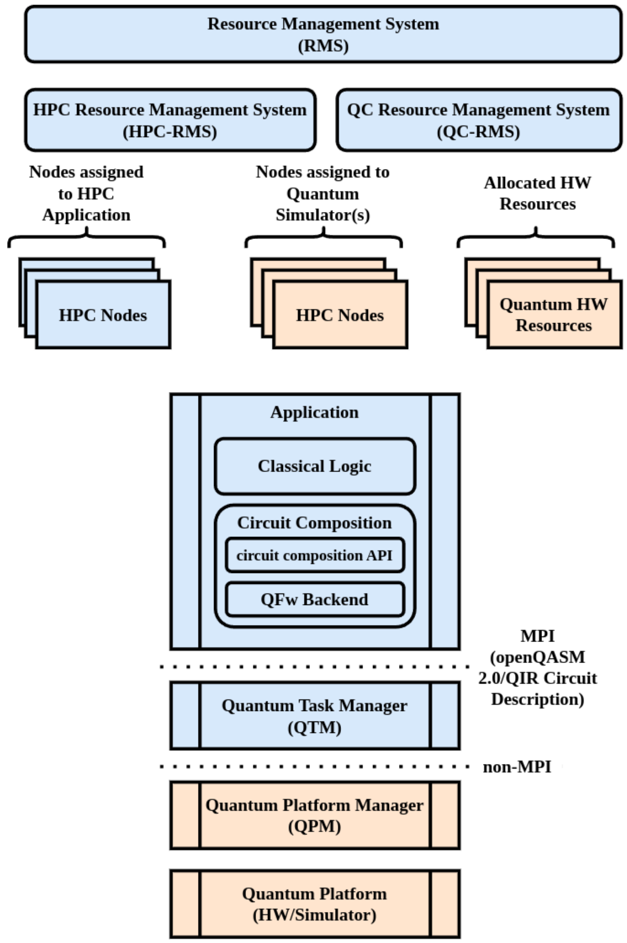

Figure 18. The QC/HPC Integration Framework uses a layered approach, with a quantum-aware resource management system reserving resources. Applications run on these resources and communicate quantum operations via MPI through the Quantum Task Manager, which can modify tasks, such as by circuit cutting. The Quantum Platform Manager then executes the prepared tasks on the quantum platform. Blue boxes represent classical resources, while orange boxes denote quantum resources [20].

5.2 Multi-Facility Workflow Orchestration

In modern scientific research, the complexity and scale of experiments necessitate the integration of diverse computational and storage resources spread across multiple facilities. This demand has led to the development of multi-facility workflow orchestration, which enables seamless management of end-to-end workflows across geographically distributed resources. To address this challenge, we have developed Zambeze [21], an automated and distributed orchestrated framework. By leveraging principles of swarm intelligence, Zambeze orchestrates complex scientific workflows through the management of distributed autonomous agents, which provide essential services such as computing, storage, and data management.

Zambeze addresses the challenges of cross-facility workflow orchestration by providing a comprehensive solution that includes a user-friendly interface, a robust compute fabric, and an efficient data fabric. The user interface allows scientists to define and manage their workflows at an abstract level using standard Python, making it accessible and intuitive. The compute fabric consists of autonomous agents equipped with plugins that enable them to execute tasks on various resources, while the data fabric uses a lazy-transfer model to efficiently manage data transfers between facilities. This approach ensures that workflows can be executed seamlessly across different environments, optimizing resource utilization and reducing manual intervention.

A key component of Zambeze's architecture is its ability to handle scheduling and execution of tasks across heterogeneous resources (see Figure 19). Each agent in the system operates using an asynchronous pull worker model, ensuring that tasks are executed only when the necessary data and computational resources are available. This model enhances the flexibility and scalability of the system, allowing it to accommodate varying workloads and optimize performance. Additionally, the system's modular design allows users to register custom plugins, further extending its capabilities and adaptability to specific research needs.

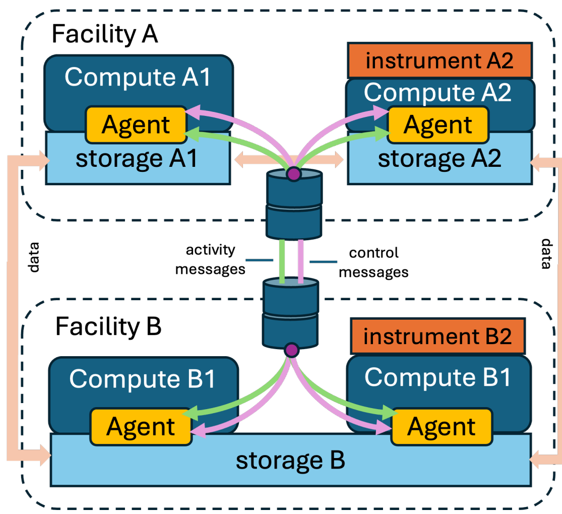

Figure 19. Architecture of a distributed workflow orchestration system on two facilities using Zambeze. Each instrument and computational resource has an agent. Agents can autonomously initiate direct data movement between storage resources [21].

The use of a distributed computing model complicates the comparison of workflow efficiency and accuracy. To address this challenge, we plan to enhance Zambeze by integrating the powerful and flexible FlowCept provenance collection and analysis framework [22]. FlowCept offers plug-ins that capture data at various levels of detail, from GPU temperature to hyperparameter configurations that optimize model performance. This data can be fed back into Zambeze, enabling it to schedule tasks on agents that are not only available but also capable of executing tasks more quickly, with greater energy efficiency, and other considerations.

Zambeze's practical application is showcased through its deployment in an electron microscopy use case, where it effectively orchestrates a complex, multi-facility workflow. This workflow begins with the capture of raw imagery data, often in gigabyte scales, using an edge device connected to an electron microscope. The data is then transferred to a leadership-class HPC system (OLCF's Frontier) where deep learning models are trained using Dask, AtomAI, and PyTorch. These models, developed for tasks like atomic species identification and defect tracking, are further refined and validated on the ACE's Defiant cluster before being deployed back to the microscope for near real-time analysis. Zambeze automates the coordination of these distributed tasks, managing data transfers, resource allocation, and workflow execution across the involved facilities.

5.3 Interconnected Science Ecosystem (INTERSECT)

Advanced research instruments in laboratories and at user facilities produce data at ever-increasing rates to deliver high-impact science. Improvements to individual instruments provide greater quantity and quality of data enabling more sophisticated studies. Automation and robotics allow measuring large systematic datasets without constant human attention and intervention. Combining diverse research tools makes ambitious multimode measurements possible that result in complex datasets. High performance computing facilitates the simulation and modeling of large systems with a fidelity previously thought impossible. Combining these advances to automate entire workflows—instrument setup and tuning, sample synthesis and processing, measurements, data analysis and model-driven data interpretation— and controlling them and making them "smart" with AI/ML will bring about revolutionary efficiencies and research outcomes. This kind of autonomous control of processes, experiments and laboratories will fundamentally change the way scientists work, allowing us to explore high-dimensional problems previously considered impossible and discover subtle correlations invisible until now.

The INTERSECT initiative at ORNL is building interconnected "Smart Labs of the Future" to enable (i) self-driving autonomous experiments that leverage advanced compute systems, scientific instruments and facilities; (ii) multi-domain and/or multi-modal experiments that alter the approach to science; (iii) real-time data analysis and feedback optimizations using AI/ML and edge computing. "Smart Labs of the Future" as envisioned here consist of a human-AI-machine interface (HAMI), experiment management software, autonomous science applications and "Big Data" analysis and management as well as data-driven feedback. INTERSECT will provide the currently missing common, seamless and secure infrastructure required to scale across facilities and science domains. Figure 20 depicts the programmatic structure of INTERSECT, which comprises multiple complementary projects which explore this space.

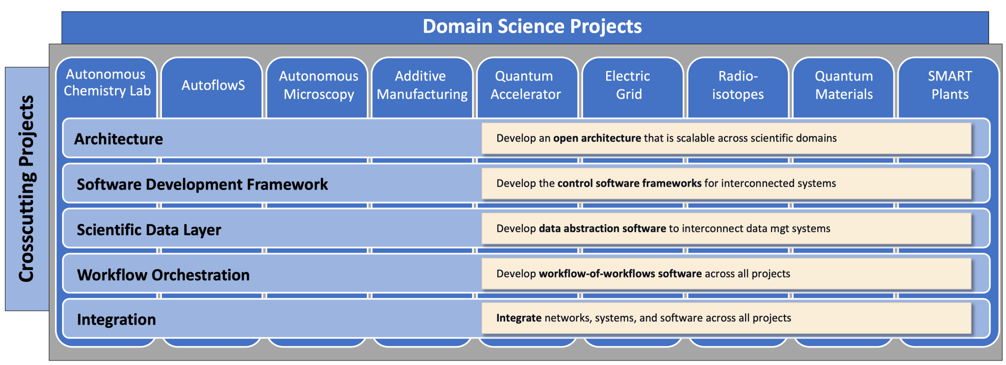

Figure 20. Programmatic structure of the INTERSECT initiative.

The INTERSECT initiative delivers a common computational infrastructure enabling autonomous workflows across multiple disciplines at ORNL. Computer scientists, data scientists, and domain scientists collaborate closely in the spirit of co-design to develop the experimental and computational tools to achieve breakthroughs in diverse science areas. The INTERSECT "system-of-systems" includes data management software, tools enabling data analysis workflows, experiment management software, as well as the integration of AI/ML capabilities. The infrastructure components being developed through these projects are designed to be transferable and form the building blocks of INTERSECT. This approach allows the expansion of automation and autonomy beyond the initial science projects through INTERSECT, minimizes duplication of effort and through the collaboration of computer and domain scientists leads to higher performance and more powerful tools.

Several of the science pilot applications hosted by ACE will be greatly enhanced by being remotely controllable, participating in distributed data movement and management, and adapting to a range of endpoint instruments and sensors. ACE and INTERSECT are complementary efforts directed at these common goals. INTERSECT applications will be able to dynamically and flexibly compose microservices to control a variety of instruments. Interfaces with IRI will provide the ability to remotely discover, enumerate, and control these dynamic INTERSECT systems-of-systems. The systems and resources of ACE are well-positioned to support the design and development of such IRI-facing interfaces.

5.4 Data Lifecycle Management

ACE provides an excellent environment for exploring the data management issues that will accompany IRI applications. ORNL has invested in tools designed to reduce the burden of data management on scientists and application developers.

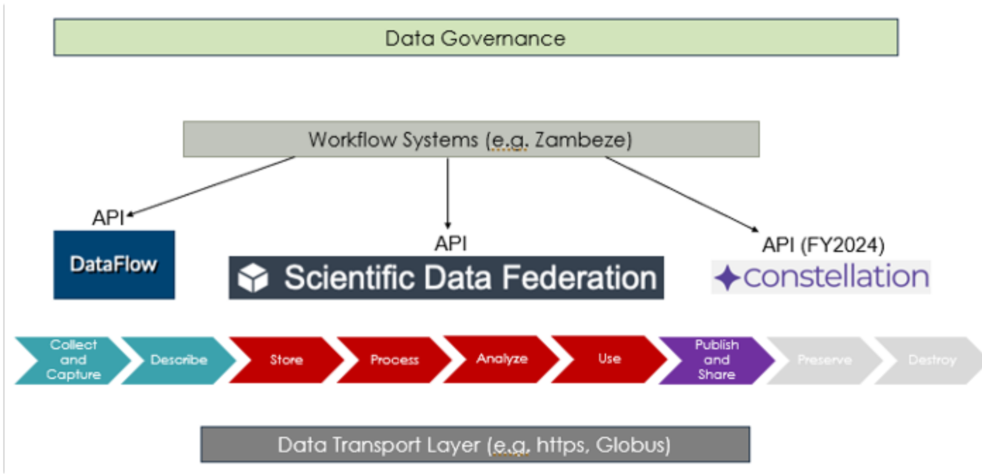

Figure 21. The lifecycle of scientific data managed by ORNL tools.

DataFlow. Data captured at instruments will be used as input by many workflows and applications running on ACE. *DataFlow* is designed to assist end-users in capturing and transferring this data. Typically, an instrument is connected to a commodity computer system where vendor-specific control software is installed, along with a nominal amount of local storage. These computers are frequently air-gapped from networks which have external routes. This is done to preserve a known-good software state on these computers, avoiding complications or delays for experimental campaigns which may result from automatic software or system upgrades to the controlling computer.

DataFlow provides an interface for users to move experiment results from the locally-attached storage at the instrument to much larger, high-performance storage accessible by simulation and analysis codes elsewhere in OLCF. Having DataFlow available on the ACE testbed systems allows users to refine their workflows; this also allows the developers of DataFlow to explore how data movement from instruments will be formalized by IRI.

DataFed. The primary goal of DataFed is to improve scientific data quality by enabling precise early-lifecycle control over data artifacts, with the ability to uniformly share and access data across geographically distributed facilities. DataFed can be thought of as a “tier 2+” distributed data storage system - meaning it is intended for creating and working with data that is of medium- to long-term significance to the owner and/or collaborators. Unlike a tier 1 storage system (i.e. a local file system), DataFed compromises raw data access performance in favor of FAIR data principles.

While DataFed shares many features with tier 3 storage systems (i.e. data archival systems), DataFed allows data and metadata to be modified after it is ingested and specifically includes features for disseminating subsequent changes to downstream data consumers via automatic provenance-based alerts as well as opt-in data subscriptions. DataFed also provides powerful and easy to use collaboration features to encourage “in-band” data-related communication instead of ad hoc and error-prone methods, such as email.

Traditional scientific data management systems (SDMS) are restricted to individual organizations or a small number of organizations connected via a “virtual organization” (VO) configuration. These systems typically support domain-specific and predetermined data workflows that cannot be readily applied to other domains or applications. On the other hand, data cataloging systems are typically single-site and provide access to static datasets. Also catalogs usually only support HTTP data transfers, thus limiting the size of datasets that can be served. Neither SDMSs nor cataloging systems can easily scale-out to accommodate large numbers of users across multiple organizations.

DataFed provides a combination of the features and benefits of both SDMSs and data cataloging services while also supporting big data. For example, DataFed provides storage and access to structured and unstructured heterogeneous raw data with access controls, metadata and provenance capture, and metadata indexing and search; however, DataFed diverges significantly from these systems in a number of ways to better serve the needs of open and collaborative scientific research.

Briefly, DataFed provides the following unique blend of capabilities and benefits:

- Presents a uniform and concise logical view of widely distributed data.
- Supports both general- and domain-specific use cases.
- Manages “living data” throughout critical pre-publication data lifecycle stages.
- Encourages FAIR-principled data practices via user- and community-defined schemas.
- Enhances data awareness with automatic notification of “data events”.
- Scales out and up to enable efficient big data research across organizations/facilities.
- Provides high-quality data management foundation for use by other applications and services.

DataFed provides interfaces that can be used to capture metadata, provenance, and raw data from the creation stage; whereas during analysis, new records may be created and linked to input records or dedicated context records. Pre-publication is supported by providing powerful data organization and data handling capabilities to help ensure that the right data is being published and that it contains proper metadata and provenance information. Note that data publishing systems may have additional metadata requirements that are not available from the data records themselves (i.e. contract numbers, sponsoring organizations, etc.)

Constellation. Constellation is the leadership-class public data repository operated by OLCF. Its purpose is to provide long-term data archive storage for OLCF users and the public at large. Constellation also manages the assignment of Digital Object Identifiers (DOIs) for data artifacts; DOIs are long-lived standards-compliant identifiers for data which may be disseminated by users without fear of staleness. Constellation functions at OLCF as a part of the overall set of data lifecycle tools, but is not involved directly by applications running in the ACE environment.

6. LESSONS LEARNED AND OPEN CHALLENGES

6.1 Policies

Most projects with real-time scheduling needs expect to be able to simultaneously utilize large HPC jobs and high-bandwidth network paths with short notice. Greta/Deleria expect to run using bursty on-demand pattern. While the experiment is on, they will require full bandwidth ($\approx$ 40 Gbit/sec presently) and a running HPC job ($\approx$ 10 nodes presently). On-times happen during a two-week window, scheduled months in advance, and last for a few days. LCLStream expects to run at $\approx$100 Gbit/sec or higher and thousands of GPUs during running experiments, and at a lower, as-available rate for offline data processing at other times.

Supporting these workloads is difficult from a policy perspective, because it requires providing guarantees on available capacity. By creating one partition with nodes that are scheduled either on-demand or pre-emptible, and another, non-preemptible partition, we are able to guarantee that all nodes in the on-demand partition will be available. At the same time, we are also

Working with HPC facility users is an important part of this strategy. Users need to be aware of the scheduling environment's constraints and batch queue policies in order to effectively use the resource. Many users have workloads that either are presently or can be converted to utilize checkpoint/restart methodology. This allows more jobs to be run in the pre-emptible queue, providing overall higher utilization of the center.

6.1.1 Security

The OLCF provides multiple security enclaves to enable projects with distinct security-level needs. As part of a strict Export Control review, the project's scope, data, and software are assessed to determine which controls, if any, are required, and to identify the appropriate EC classification. The OLCF's flagship supercomputer, Frontier, is housed in the "moderate" production enclave and allows Category 1 (Unrestricted) and Category 2 (Export controlled/Proprietary) software and/or data. This enclave supporting the facility's leadership-class computing mission requires multi-factor authentication (MFA) and thorough screenings to ensure that users allowed to access the facility meet DOE guidelines.

In the IRI context, these policies can potentially block several use cases including: automation of workflows, prevent use of shared accounts and robot accounts, and exposing data outside specific projects and even outside the facility.

For example, the facility needs to report utilization and specific characteristics of users using facility resources including but not limited to identity, citizenship, and other demographic data. With non-human accounts that are shared, that is not an option.

6.1.2 Operations

In order to support the full range of patterns explored here, the facility would need to better understand which patterns can co-exist in the flagship system and which ones would be better served utilizing peripheral resources in the ecosystem. The characteristics across patterns have competing and conflicting needs which require different levels and types of support to enable successful scientific campaigns across all of them. Defining how and where each of the patterns can thrive is a requirement in order to understand if additional expertise, support models, or resources are needed. This is an area that we are continuing to explore as part of this effort.

6.1.3 User Experience

6.1.4 Flexibility

Many of the workflows explored in this effort require dynamic changes to the configuration of the system. For example, time-sensitive patterns would require preemption of workloads running on the system, and potentially may require longer walltimes that regularly allowed. Having the flexibility to accommodate batch job requests of competing sizes, runtimes, and priorities is one of the challenges that need to better understood.

6.2 Data Movement Challenges: Networking, Firewalls, and Hardware Configuration

Moving data between storage devices within a network security domain, between network security domains, and between internal and external organizations is a central requirement for computation, occurs nonstop, and requires careful planning and control to carry out securely at high rates. Our data streaming design effort identified data movement needs for each application considered within ACE. Attempting to utilize emerging best practices, we designed a data streaming workflow involving Kubernetes-allocated nodes specialized for high bandwidth and API-initiated streaming workflows. This approach provided all the physical infrastructure to satisfy our planned capacity needs, and was setup within a few months as hardware arrived.

Unfortunately, Kubernetes applies its own paradigms to resource allocation / authorization, network path management, and hardware configuration. Kubernetes APIs were excellent for running and monitoring a variety of data forwarding services that could scale to multiple copies across nodes. However, they were unable to handle high-speed networking. This was due both to network routing assumptions in Kubernetes, and to the fact that Kubernetes as a container framework complicates configuring the network interfaces themselves. We used a pair of high-speed NICs, one internal and one external. Moving all data through a single gateway, however, made it difficult for Kubernetes to distinguish between internal-to-ACE and external-to-internet connections. To account for this, ORNL's traffic routing policies required including their own firewall on the internet-to-Kubernetes path. This wasn't a bottleneck, since the firewall used for this can handle 120 Gbps in each direction, but it illustrates the practical consequence of the mismatch within hybrid cloud/HPC.

6.3 Authentication Patterns

The ACE pilot applications have demonstrated many different authentication patterns – based on their user and data management needs. Greta/Deleria makes use of HPC-provided user-IDs via the Facility or SLURM batch job API, and secures network paths to nodes using custom firewall rules on source and destination addresses. LCLStream uses a combination of mutual TLS certificate-based authentication for its REST-HTTPS APIs, as well as OpenShift tokens for starting and stopping services at OLCF. JGI uses group-level accounts for running services on OLCF that can be utilized by users who are authenticated separately by JGI. Other applications utilize individual user accounts within traditional batch jobs.

A group-level UNIX userid fits most current patterns because the science application teams are tightly integrated and generally employ open-source software and open data. However, larger collaborations like JGI have implemented separate authorization mechanisms to prevent users from data corruption or misuse. These actions can happen accidentally, for example when onboarding new users. Experimental data sources work with multiple different user groups as well, and need to protect raw experimental data from theft or other misuse by competing groups. In addition, the computing facility needs to track resource utilization and system errors at a user-level in order to effectively communicate with user and system teams. Thus, it is important for applications to make use of authentication patterns that forward individual user and group membership to HPC resources.

6.4 Challenges Supporting Microservices

Many of the IRI applications involve running user-requested microservices, for example web-interfaces based coupled to databases that tracking project status. Our OpenShift(TM) Kubernetes platform can run these user-provided API servers. However, there need to be security limitations on what HPC resources can be accessed by applications running on Kubernetes. For example, if a user-managed application serves incoming requests from outside ORNL, there is a heightened security risk in allowing that application to also interact with HPC login nodes or running jobs.

We have used two different approaches to controlling these security risks: facility-based access approval and facility-provided APIs. The primary difficulties encountered for both approaches were centered around user communication. Since only the NCCS network and security teams can see the current access controls applied to a project's microservices, they must work closely with users and NCCS's OpenShift engineering team to setup Kubernetes resources correctly. However, errors in this process can come from the user's application, Kubernetes configuration, or NCCS security and networking settings. Coordinating all four groups is a slow and cumbersome process.

In contrast, facility-provided APIs are developed centrally by NCCS's platforms services group. Hence, it is relatively simple to setup and run these applications. However, strong software engineering practices need to be followed to guarantee security and availability of each service, which takes time. It is also not possible to support custom microservice installation requests from dozens of user groups. In order to provide API-based services generally useful across HPC user projects, user outreach and planning is being done to gather requirements from multiple groups and carry out use studies.

Effectively securing user-managed microservices has several aspects:

- Authenticating and authorizing service configuration and startup/shutdown
- Authenticating and authorizing incoming traffic
- Creating and managing resource access policies
- Authenticating and authorizing requests from the microservice for facility resources

Our current service startup/shutdown authorization is done via OpenShift API, secured by a token provided through a facility-managed web interface. Incoming traffic is allowed for some projects, and only via HTTPS. When allowed, traffic is passed through a facility-managed TLS termination, and sent to the service via either HTTP or a new HTTPS channel. Although the facility provides the originating user to the API via a header, it does not yet allow the service to pass user credentials through to OLCF's internal services. Providing access from services to facility resources is currently allowed only for some projects. When allowed, the facility provides a group-level account with the ability to run commands on a data transfer node and/or access files on some of the high-performance filesystems. The microservice is also able to make network requests directly to OLCF's networked resources, such as API requests to cluster SLURM schedulers.

Granting user-managed microservices the ability to open firewall ports for data channels, or to access files or run commands as individuals (not group service account) is not currently possible. The barrier is technological, since the facility needs to be able to identify the individual connecting to the open port, accessing files, or running commands. Potential solutions are under investigation, involving tokens whose scopes can be narrowed (macaroons[23]), and mutual TLS for authenticating incoming high-speed TCP traffic.

7. OUTREACH AND ENGAGEMENT ACTIVITIES

7.1 Hackathons and Training

7.1.1 OLCF/JGI Hackathon

A 2-day joint hackathon with the JGI team was organized and hosted on-site by ORNL in April 2024. The objective of the hackathon was to port and deploy JGI's JAWS workflows platform to the ACE IRI testbed. Several individuals from the various groups across the board at OLCF participated in the hackathon. A multitude of technical challenges were anticipated in the process. A non-exhaustive list of those challenges are discussed below.

- Association of the external users with appropriate ACE IRI testbed managed Projects and Accounts was an initial challenge due to the external user's unfamiliarity with the onboarding portal and requirements.
- Accounts, access settings and directory configurations such as collections settings with Globus platform were a prerequisite to ensure smooth data transfer between the JGI and testbed sites. In particular, a service-account was required that would do certain operations on behalf of real-user accounts.
- Integration with Marble/Olivine, an Openshift/Kubernetes cluster and Gitlab runner setup hosted at OLCF in order to automate the orchestration over a remotely operated portal and command line.
- Installation and testing of several software tools such as Globus, Cromwell Workflow Manager, and HTCondor.
- Enabling containers runtime via the Apptainer platform on the testbed.
- Tweaking the jobs scheduling policies and queuing properties in order to best accommodate the JAWS requirements. For instance, updating the maximum walltimes for jobs to 12 hours from the initial 8 hours.
- End-to-end integration with the system scheduler such that the Cromwell workflow management system that is deeply embedded into JAWS is able to create jobscripts and submit them to the testbed's SLURM job scheduler.

Most of the aforementioned challenges and related technical issues were discussed and resolved during the hackathon. The hackathon also served the purpose of bringing the JGI and ORNL teams together fostering fruitful discussions and networking opportunities for a lasting collaboration. As a follow up to the hackathon, weekly virtual meetings continued to refine and fine tune the deployment—in particular the Globus configuration.

7.1.2 Quantum Workshop

ORNL has been collaborating with an Australian company called Quantum Brilliance on topics related to quantum computing for HPC applications. In August of 2024, we hosted a two and a half day workshop involving several of their employees and ORNL staff and discussed a variety of topics including both hardware capabilities and applications/software.

Topics discussed at the workshop included fundamental questions about parallelizing quantum computing algorithms and how typical HPC computational problems such as quantum chemistry could be mapped onto a quantum computer or onto a hybrid classical / quantum computer. Also discussed were practical topics about the use of Quantum Brilliance's emulation software that had recently been installed on the Wombat and Holly computers in the ACE testbed. There were further discussions about the continued collaboration and plans for the future.

7.2 Conferences

7.2.1 Smoky Mountains Conference

The Smoky Mountains Computational Sciences & Engineering Conference organized annually by ORNL has been a long-standing venue for discussing forward-looking technologies and wrestling with their potential impact on the scientific ecosystem. For the past few years, this has included a strong focus on integrating laboratory experiments “into the computational loop” within the scope of ORNL’s INTERSECT initiative. Although that work involves edge compute, data transfer and workflow methodologies, integration with high-performance computing had not been a focus.

This year will host a demo session featuring the combination of both ACE and INTERSECT-developed enabling technologies (Sec. 4) and collaboration-driven science pilot applications (Sec. 2). On the technology side, we tackle core missing components in the HPC ecosystem with talks on i) enabling user-relevant, API-driven workflow systems within OLCF’s secure scientific service mesh concept, ii) the INTERSECT instrument control pattern (and software development kit) for allowing AI-through API driving of laboratory instruments, iii) flexible setup and teardown of API-driven data streaming hardware, and iv) a survey of experimental user facility barriers to utilizing HPC.

On the application side, talks include the high data-rate Greta/Deleria experiment station, the AI-intensive LCLS streaming data analysis and feedback project, HPC-enabled access to quantum computing simulation, and experiment/HPC co-design of the IMAGINE-X dynamic nuclear polarization experiment at HFIR. We will also hear from NSLS-II on their online data analysis and feedback portals. By including video-demos within each talk, the sessions will build a shared picture of what’s possible when these technologies and applications are able to freely intermix. These will fuel design discussions during the meeting’s break-outs. Although the process can be messy at first, common patterns of interaction between people, facilities, data and compute will guide our assumptions as we continue to build and integrate open-source, facility and application software stacks.

7.2.2 Monterey Data Conference

In August 2024, members of the ACE team actively participated in the Monterey Data Conference (MDC), fostering valuable in-person discussions and collaborations. MDC provided an opportunity for ACE team members to engage directly with researchers from ASCR facilities and those involved in IRI science pilots.

Additionally, we presented a poster entitled “Enabling Distributed Research Orchestration Capabilities at ORNL", which showcased key capabilities of the ACE ecosystem. Specifically, the poster highlights the advanced capabilities of the ACE IRI testbed, demonstrated through the integration of the OLCF Facility API and the Zambeze distributed orchestration system. The ACE testbed showcases how these tools can modernize and automate HPC facilities, allowing for seamless orchestration of complex science workflows across a distributed cyberinfrastructure. This enables researchers to focus on scientific innovation while benefiting from streamlined access to compute, storage, and instrument resources.

7.2.3 IRI/HPDF Meeting

ACE team members participated in the DOE Office of Science (SC) IRI/HPDF coordination kickoff meeting held in July 2024. This invitation-only meeting gathered key stakeholders, including Federal Program leaders and community experts, to advance the understanding of community priorities and develop a multi-year roadmap for the IRI program and HPDF project.

8. SUMMARY AND STRATEGIC OUTLOOK

NCCS considers the Advanced Computing Ecosystem (ACE) effort as a strategic investment. ACE provides a flexible vehicle for advancing NCCS' technical and scientific expertise and capabilities for multiple crosscutting projects, and efforts, such as OLCF-6, OLCF-7, INTERSECT, FASST, and IRI.

ACE allows NCCS to evaluate technologies that can have material impact for current and upcoming OLCF projects and SPPs. Current examples include AI appliances, storage technologies and network DPUs. These evaluation efforts enable NCCS staff to keep up-to-date with new technologies.

For IRI, ACE is an invaluable strategic platform for NCCS and OLCF for multiple reasons:

- *Science pilots and collaborations*: ACE allows a flexible and secure environment for science pilots to test out new technologies, perform demonstrations of new methods and capabilities within a permissive semi-production environment. Allowing pilots on the production environment to test and develop is problematic and ACE alleviates this and opens up a path for these pilots to safely develop before applying for production allocations.
- *Foundational technology development*: For integrating experimental and observational user facilities and workflows with leadership computing environments a new set of tools are needed. ACE allows OLCF to develop these tools (e.g., FacilityAPI, streaming services) in a safe production-like environment in close collaboration with the real-world science pilots. This is a strategic capability benefiting for all stake holders.
- *Policy considerations*: Time-sensitive computing requires tools to schedule resources for execution at a more-or-less specific point in time, i.e. the results–to be deemed useful–need to be obtained before a well-determined end time, regardless of potential contention on the resource used. ACE allows the OLCF to measure both the typical lead times required to estimate that critical end time and the amount of resources required to fulfill that requirement. These data will be essential as potential IRI projects move from the testbed environment to the leadership ecosystem. Without this guidance, a finite amount of an exceedingly valuable resource (the leadership platform) will have to be idled for significant amounts of time to ensure utilization is consistent with the time-sensitive demand. Even with the data, significant modifications to the current set of metric for leadership computing will require modification if these use cases are to be properly supported.

The current generation of the ACE infrastructure is mostly left overs from the Frontier project (e.g., Defiant) with some additions we've made within the last year (e.g., GraphCore, H100 and GraceHopper nodes). To keep the ACE testbed viable for evaluation and development, a constant effort is needed. In order to reduce the system administration burden, maintain security and also to increase the value to researchers and developers, we have made plans to do technology refreshes in FY25. These include deploying a new Defiant system with more up-to-date technologies and also a network refresh. NCCS plans to grow ACE with technology refreshes.

Currently the biggest challenge we can see is establishing a healthy "graduation" path for science pilots and developed technologies from ACE to the proper production environment. While all the details are not clear how to enable this effectively and securely, OLCF is working on this problem. This approach is well in line with OLCF-6 project goals, so this exploration is timely. Perhaps one path towards solving this problem is establishing a Leadership ecosystem, resembling the ACE environment for production where the flagship supercomputer is closely associated with a production on-demand computing cluster as depicted in Figure 14 to accommodate bulk of the incoming IRI requests for compute and data capabilities, and surging into the supercomputer exactly and only when those capabilities are exceed. Again, this is an open question for now and a production-grade solution has not been decided upon, yet.

For FY25, OLCF plans to strategically engage with more science pilots where we can demonstrate clear mutual benefit for all stake holders. OLCF also plans to increase its collaboration with other ASCR facilities for tighter integration with their respective testbeds and also integration of their applicable technologies to the OLCF/ACE environment.

REFERENCES

[1] K. B. Antypas, D. Bard, J. P. Blaschke, R. S. Canon, B. Enders, M. A. Shankar, S. Somnath, D. Stansberry, T. D. Uram, and S. R. Wilkinson, "Enabling discovery data science through cross-facility workflows," in *2021 IEEE International Conference on Big Data (Big Data)*. IEEE, 2021, pp. 3671–3680.

[2] B. L. Brown, W. L. Miller, D. Bard, A. Boehnlein, K. Fagnan, C. Guok, E. Lançon, S. J. Ramprakash, M. Shankar, and N. Schwarz, "Integrated research infrastructure architecture blueprint activity (final report 2023)," US Department of Energy (USDOE), Washington, DC (United States). Office of . . . , Tech. Rep., 2023.

[3] J. Kim, S. Lee, B. Johnston, and J. S. Vetter, "Iris: A portable runtime system exploiting multiple heterogeneous programming systems," in *Proceedings of the 25th IEEE High Performance Extreme Computing Conference*, ser. HPEC '21, 2021, pp. 1–8.

[4] S. E. Hahn, P. W. Fackler, W. F. Godoy, K. Maheshwari, Z. Morgan, A. T. Savici, C. M. Hoffmann, P. Valero-Lara, J. S. Vetter, and R. F. da Silva, "Integrating ORNL's HPC and Neutron Facilities with a Performance-Portable CPU/GPU Ecosystem," in *Submitted to the 6th Annual Workshop on Extreme-Scale Experiment-in-the-Loop Computing (XLOOP 2024)*, 2024.

[5] F. Suter, R. F. Da Silva, A. Gainaru, and S. Klasky, "Driving next-generation workflows from the data plane," in *2023 IEEE 19th International Conference on e-Science (e-Science)*. IEEE, 2023, pp. 1–10.

[6] T. J. Skluzacek, R. Souza, M. Coletti, F. Suter, and R. Ferreira da Silva, "Towards cross-facility workflows orchestration through distributed automation," in *Practice and Experience in Advanced Research Computing 2024: Human Powered Computing*, 2024, pp. 1–4.

[7] R. Chard, J. Pruyne, K. McKee, J. Bryan, B. Raumann, R. Ananthakrishnan, K. Chard, and I. T. Foster, "Globus automation services: Research process automation across the space–time continuum," *Future Generation Computer Systems*, vol. 142, pp. 393–409, 2023.

[8] B. Enders, D. Bard, C. Snavely, L. Gerhardt, J. Lee, B. Totzke, K. Antypas, S. Byna, R. Cheema, S. Cholia *et al.*, "Cross-facility science with the superfacility project at lbnl," in *2020 IEEE/ACM 2nd Annual Workshop on Extreme-scale Experiment-in-the-Loop Computing (XLOOP)*. IEEE, 2020, pp. 1–7.

[9] T. Kurihana, T. J. Skluzacek, R. F. da Silva, and V. Anantharaj, "Scalable Multi-Facility Workflows for Artificial Intelligence Applications in Climate Research," in *Submitted to the 6th Annual Workshop on Extreme-Scale Experiment-in-the-Loop Computing (XLOOP 2024)*, 2024.

[10] S. Sheldon, Y. Kumar, M. Goodrich, and G. Heyes, "Ej-fat joint esnet jlab fpga accelerated transport load balancer," 2023. [Online]. Available: https://arxiv.org/abs/2303.16351

[11] B. Brown, H. Finkel, M. Papka, T. Uram, D. Bard, T. Connors, S. Dosanjh, C. Guok, E. Dart, I. Monga, R. Adamson, R. F. da Silva, A. Shankar, A. Boehnlein, G. Hayes, and N. Tran, "Advancing doe's integrated research infrastructure: Update on progress," in *Advanced Scientific Computing Advisory Committee Meeting Presentations, May 29, 2024*. United States Department of Energy, Office of Science, 2024. [Online]. Available: https://science.osti.gov/ascr/ascac/Meetings/202405

[12] J. P. Blaschke, A. S. Brewster, D. W. Paley, D. Mendez, A. Bhowmick, N. K. Sauter, W. Kröger, M. Shankar, B. Enders, and D. Bard, “Real-time xfel data analysis at slac and nersc: A trial run of nascent exascale experimental data analysis,” *Concurrency and Computation: Practice and Experience*, vol. n/a, no. n/a, p. e8019, 2024. [Online]. Available: https://onlinelibrary.wiley.com/doi/abs/10.1002/cpe.8019

[13] S. Veseli, J. Hammonds, S. Henke, H. Parraga, and N. Schwarz, “Streaming data from experimental facilities to supercomputers for real-time data processing,” in *Proceedings of the SC ’23 Workshops of The International Conference on High Performance Computing, Network, Storage, and Analysis*, ser. SC-W ’23. New York, NY, USA: Association for Computing Machinery, 2023, p. 2110–2117. [Online]. Available: https://doi.org/10.1145/3624062.3624610

[14] B. L. Brown, W. L. Miller, D. Bard, A. Boehnlein, K. Fagnan, C. Guok, E. Lançon, S. J. Ramprakash, M. Shankar, and N. Schwarz, “Integrated research infrastructure architecture blueprint activity (final report 2023),” US Department of Energy (USDOE), Washington, DC (United States). Office of Science, Tech. Rep., 2023. [Online]. Available: https://www.osti.gov/biblio/1984466

[15] E. Dart, J. Zurawski, C. Hawk, B. L. Brown, and I. Monga, “Esnet requirements review program through the iri lens: A meta-analysis of workflow patterns across doe office of science programs,” Lawrence Berkeley National Laboratory (LBNL), Berkeley, CA (United States), Tech. Rep., 11 2023. [Online]. Available: https://www.osti.gov/biblio/2205078

[16] M. J. Brim, A. George, A. Shehata, C. Lester, D. Rogers, P. Widener, R. Miller, G. Jansen, R. Ferreira Da Silva, and S. Oral, “A high-level design for bidirectional data streaming to high-performance computing systems from external science facilities,” Oak Ridge National Laboratory (ORNL), Oak Ridge, TN (United States), Tech. Rep., 3 2024. [Online]. Available: https://www.osti.gov/biblio/2338264

[17] J. Chung, W. Zacherek, A. Wisniewski, Z. Liu, T. Bicer, R. Kettimuthu, and I. Foster, “Scistream: Architecture and toolkit for data streaming between federated science instruments,” in *Proceedings of the 31st International Symposium on High-Performance Parallel and Distributed Computing*, ser. HPDC ’22. New York, NY, USA: Association for Computing Machinery, 2022, p. 185–198. [Online]. Available: https://doi.org/10.1145/3502181.3531475

[18] “Apptainer,” 2024. [Online]. Available: https://apptainer.org

[19] “NIST Special Publication 800-63C,” https://pages.nist.gov/800-63-3/sp800-63c.html, 2017.

[20] T. Beck, A. Baroni, R. Bennink, G. Buchs, E. A. Coello Perez, M. Eisenbach, R. Ferreira da Silva, M. Gopalakrishnan Meena, K. Gottiparthi, P. Groszkowski, T. S. Humble, R. Landfield, K. Maheshwari, S. Oral, M. A. Sandoval, A. Shehata, I.-S. Suh, and C. Zimmer, “Integrating quantum computing resources into scientific hpc ecosystems,” *Future Generation Computer Systems*, vol. 161, pp. 11–25, 2024. [Online]. Available: http://dx.doi.org/10.1016/j.future.2024.06.058

[21] T. J. Skluzacek, R. Souza, M. Coletti, F. Suter, and R. Ferreira da Silva, “Towards cross-facility workflows orchestration through distributed automation,” in *PEARC’24: Practice and Experience in Advanced Research Computing*, 2024. [Online]. Available: http://dx.doi.org/10.1145/3626203.3670606

[22] R. Souza, T. J. Skluzacek, S. R. Wilkinson, M. Ziatdinov, and R. F. Da Silva, “Towards lightweight data integration using multi-workflow provenance and data observability,” in *2023 IEEE 19th International Conference on e-Science (e-Science)*. IEEE, 2023, pp. 1–10.

[23] A. Birgisson, J. G. Politz, Úlfar Erlingsson, A. Taly, M. Vrable, and M. Lentczner, “Macaroons: Cookies with contextual caveats for decentralized authorization in the cloud,” in *Network and Distributed System Security Symposium*, 2014. [Online]. Available: https://research.google/pubs/macaroons-cookies-with-contextual-caveats-for-decentralized-authorization-in-the-cloud/