U.S. DEPARTMENT OF ENERGY | Office of Science

THE DOE OFFICE OF SCIENCE

Federated IRI Science Testbed (FIRST): A Concept Note

December 1, 2023

ESnet | OAK RIDGE National Laboratory | LEADERSHIP COMPUTING FACILITY | Argonne NATIONAL LABORATORY | NERSC

Federated IRI Science Testbed (FIRST): A Concept Note

December 1, 2023

Corresponding Authors:
Inder Monga and Chin Guok, Energy Sciences Network (ESnet)
Arjun Shankar, Oak Ridge Leadership Computing Facility (OLCF)

Contributors
Argonne Leadership Computing Facility: Michael E. Papka, Jini Ramprakash, Tom Uram, William Allcock, Ian Foster
Energy Sciences Network: Damian Hazen, Eli Dart
Oak Ridge Leadership Computing Facility: Georgia Tourassi, Sarp Oral
National Energy Research Scientific Computing Center: Sudip Dosanjh, Deborah Bard

DOI: https://doi.org/10.2172/2205149
LBNL-2001553
24-SN-35056

Table of Contents

Concept Description

The Department of Energy's (DOE's) vision for an Integrated Research Infrastructure (IRI) [1] is to empower researchers to smoothly and securely meld the DOE's world-class user facilities and research infrastructure in novel ways in order to radically accelerate discovery and innovation. Performant IRI arises through the continuous interoperability of research workflows with compute, storage, and networking infrastructure, fulfilling researchers' quests to gain insight from observational and experimental data. Decades of successful research, pilot projects, and demonstrations point to the extraordinary promise of IRI but also indicate the intertwined technological, policy, and sociological hurdles it presents.

Creating, developing, and stewarding the conditions for seamless interoperability of DOE research infrastructure, with clear value propositions to stakeholders to opt into an IRI ecosystem, will be the next big step. Governance, funding, and resource allocation are beyond the scope of this document: it seeks to provide a high-level view of potential benefits, focus areas, and the working groups whose formation would further define the testbed's design, activities, and goals.

The Case for a Testbed

Researching, developing, prototyping, and testing new complex workflows involving scientific and compute infrastructure across multiple administrative domains remains challenging. Exploring novel or innovative approaches in the production computing and networking infrastructure is the norm today, but potentially disruptive approaches — such as the application of AI to improve operations, resilient workflows, use of edge computing, in-network processing with FPGAs, etc. — are often constrained and sometimes impossible because of the potential risk to operational and cybersecurity integrity. When multiple user facilities are involved, the task is even harder: practitioners must coordinate time windows, security exceptions, and failure modes, as well as different mission requirements, to allow such workflows to be prototyped and tested at scale.

A specific challenge in the DOE complex is that, out of necessity for their science missions, multiple laboratories and facilities have created capabilities and approaches that overlap but serve slightly different purposes. The difficulty in such an environment is to align, extract, and abstract the best components and elevate them to broader production service status across the complex. We need an environment that enables laboratory teams to bring their unique tools and capabilities, iron out interoperability mismatches, and establish a path forward to align and work together toward a common goal.

A best practice in systems design is to create an at-scale test environment in which to build, assess, and improve tool functionality before transitioning technology and systems to production. Such a common test environment, or testbed, can help the broader complex move swiftly to realize an IRI built on a collectively strong set of tools and functionality already shown effective for science.

The idea of building a test environment is not new or unique to DOE or its Advanced Scientific Computing Research (ASCR) program office. The proposed Federated IRI Science Testbed (FIRST) enables a progressive design-experiment-test-refine cycle to establish a shared environment for IRI developers and pilot application users to come together and advance the overall vision by experimenting with the design patterns and addressing the gaps identified in the IRI Architecture Blueprint Activity (ABA) Final Report [2]. It is "federated" in that it is owned and operated by DOE facilities coming together in partnership and experimentation, without a single centralized owner. The distributed nature will allow facilities to coordinate on interoperable components, test innovative infrastructure ideas, and refine facility-specific extensions if necessary.

In fact, this cyberinfrastructure can be the foundation of a flagship R&D platform for IRI design initiatives at the confluence of cutting-edge user facilities and evolving science workflows, with a direct line of sight to production.

This R&D platform will have different expectations from a production facility, both from the users and maintainers of the testbed, thus lowering the bar for trying out new ideas and deploying innovative equipment, tools, and technologies. By encouraging a rapid try-test-fail cycle, FIRST should generate viable, sustainable, and scalable implementations that can be transitioned to upcoming production infrastructures.

Ultimately, the goal of FIRST is to provide new capabilities, aligned with IRI ABA Final Report, to scientists who need to:

- Integrate and analyze large data sets;
- Move large data sets seamlessly from experiments to compute facilities over performant networks;
- Adjust computational and observational science instruments for time-sensitive computing needs (e.g., through near-real-time feedback);
- Support sustained long-term campaigns requiring multi-facility and co-scheduled allocations;
- Employ new algorithms and methods, including AI and ML.

The IRI testbed will tie together experimental and observational instruments, ASCR compute facilities for large-scale analysis, and edge computing for data reduction and filtering using Energy Sciences Network (ESnet), the high performance network and DOE user facility. The testbed will provide pre-production capabilities that are beyond a demonstration of technology.

Anticipated Benefits

The benefits of an IRI Testbed are multifold:

- It provides a safe space to test new capabilities and exercise novel workflows within a well-defined environment with calculated collateral risk and impact footprints.
- It enables multiple DOE assets such as instruments, user facilities, and high performance computing centers to create interoperable workflows leveraging their unique capabilities that can be eventually deployed in production.
- It is highly controlled, implying experiments are repeatable and reproducible, and can facilitate root-cause analysis.

The above characteristics help drive the purpose of FIRST: to enable results-driven changes to be made in the production workflows and/or infrastructure before broad deployments of IRI workflows occur, thus improving science outcomes.

Initiation and Activities

In order to meet the benefits articulated, FIRST will be designed with certain principles in mind.

- The testbed should allow for real workflows to be executed over actual (e.g., instrument/beamlines) and/or representative infrastructure (e.g., dedicated test network, storage, compute).
- At a bare minimum, the testbed should imitate operational infrastructure and/or include potential designs and configurations that could be prototyped, scaled up, and eventually adopted. It should allow for novel hardware, software, and policy approaches to be included for experimentation.

- The testbed should have functionality that allows it to be isolated physically or logically from the production space, with potentially distinct operational (e.g., breakable, time-to-repair), usage (e.g., allocation), and/or security (e.g., testbed users, different AuthN/AuthZ) policies.
- The testbed design and deployment choices will be prioritized by interoperability considerations that support the IRI ABA principles and guidance.
- The testbed will aim to be open to the DOE national laboratories and collaborators and will operate in a manner that facilitates rapid development by prioritizing rough consensus with working demonstrations of functionality. (The "Request for Comments" model of the Internet Engineering Task Force[1] is a successful model to emulate.)

FIRST should, as much as possible, leverage elements from existing partnerships, accomplishments, and learnings from various projects under the umbrella of current ASCR facility projects, including Superfacility[2], INTERSECT[3], Nexus, and other initiatives launched by DOE SC facilities. The design and deployment of FIRST will be enabled through a set of working groups described below.

Addressing Application Drivers

Application exemplars will be selected by ASCR in consultation with other program offices within the Office of Science, most likely based on specified criteria. We suggest the following considerations be weighed when choosing and prioritizing applications on the testbed:

- Scientific impact;
- Need for greater integration with computing and network facilities;
- Range of computational/network drivers;
- Applicability to other science domains;
- Feasibility given resource constraints.

An important caveat to note: FIRST is deemed an experimental "instrument" that will allow us to build knowledge on how to support IRI-driven workflows. Understanding and developing the IRI testbed attributes — such as capabilities, governance, access policies, and management — will be part of the learning experience in jointly building and operating this testbed. It is to be expected that as we build and run the testbed, our understanding of the various attributes will change, and we will have to evolve our practices accordingly.

Transitioning from Testbeds to Production

Each site will provide resources that are representative of, but safely separated from, its production resources. In many cases, logical separation of production resources will be the mechanism to ensure workflow development is done in a representative environment. Hardware architectures, schedulers, accounting, and the rest of the software stack should be as similar as possible. Being separated from production allows some flexibility in terms of breaking changes to configuration, but security must remain at the forefront — both for all the traditional reasons, but also because there are security challenges that are key to the success of IRI. This similarity of architecture will ease the transition to our production machines, while the greater flexibility will enable rapid iteration in the solution space around IRI.

1 https://www.ietf.org/about/introduction/ : mainly, "**Rough consensus and running code**: We make standards based on the combined engineering judgement of our participants and our real-world experience in implementing and deploying our specifications."
2 https://www.osti.gov/biblio/1875256
3 https://www.ornl.gov/intersect

The FIRST effort will require framing and engagement, active participation, and support from user facilities and program managers across the DOE's Office of Science (SC) programs to ensure that we connect technologists and scientists through interoperable infrastructure that truly supports the generation of new IRI workflows with the potential to accelerate science. This effort will be co-designed through partnerships between the ASCR facilities and select SC experimental and observational facilities and will build on a wealth of experience and previously stated requirements [1-11]. The testbed activities will also require regular conversations and involvement of the production operations to get early feedback on their requirements and operational needs, concerns, and/or gaps as IRI workflows are implemented and run on the testbed.

Participation and Governance

The intent of the IRI testbed is to encourage open participation and span the user facilities and other assets across Office of Science activities. These may include national labs and/or assets used by science collaborations in universities. Funding and resource allocation for the broad implementation and maintenance of this testbed is not covered in this document. The ASCR program office, in coordination with other Offices of Science program leadership, will establish a governance structure, similar to an Integrated Project Team, that will allow for complex policy and participation decisions.

While an IRI testbed promotes experimentation, it is also future looking. There will be a strong component of research participation needed for the infrastructure selection, software and API development, and overall workflow innovation expected to be built over this testbed. The complexities of transitioning a new, innovative workflow showcased on the IRI testbed to production will be an important step to get benefit from the research, but the responsibility of transition to production does not lie within the domain of the people building and operating this testbed and will be handled at the programmatic levels. Transition to production and operation will be carefully planned as part of an established program. We expect that the IRI Testbed explorations will help determine the maturity level of tools and capabilities so that they may be transitioned to production IRI systems.

Governance will be set up such that program managers in DOE can coordinate on policy and prioritization topics while field operations at the laboratories can coordinate prioritization and work on technical matters. The activities of the testbed will take place through a set of working groups with representatives from across the ASCR facilities with expertise to establish the underlying infrastructure and interested domain-science pilot application users who understand that this is a design playground. The deployment of resources and activities will be decided by the governance process factoring in the prioritization of goals and the availability of resources.

Focus Areas for Working Groups

The following provides a high-level view of potential focus areas used to further the various activities of the testbed. Working groups would be formed within one or more areas such as hardware infrastructure, software, services, policy and governance, outreach, and engagement and empowered to make decisions and help define the scope. These working groups will convert this concept note into an actual design and eventual deployment.

The focus areas include the following:

- **Hardware Infrastructure** – identifying the hardware infrastructure necessary to build the physical testbed. Activities would include establishing the various physical resources of the IRI testbed and how they contribute to the overall testbed design.
- **Testbed System Software** – designing, developing, and integrating software needed to manage and use the testbed environment. This includes management and control software needed to reserve and provision resources, manage access, and monitor status. Interface definition and documentation will be a part of this activity.
- **Federated ID Management** — recognizing the security posture and policies of the participating facilities and determining appropriate AuthN/AuthZ and access control mechanisms.
- **Services** — defining the services that the testbed will provide and assisting in the interoperability between the workflow and the testbed.
- **Policy and Governance** — defining the policies and governance models for the testbed, to include metrics and resource allocation for appropriate use of the testbed as well as methods of coordination and communication between facilities and DOE to ensure alignment of program offices' interests.
- **Outreach and Engagement** – establishing user-facing activities to help understand the scientific use cases and provide any necessary assistance to users. The expectation is that users will gain access based on the end-to-end IRI application use-case. Users will be brought in with a realistic expectation of science outcomes until a transition to operations and production is programmatically supported.

Next Steps

The IRI Testbed will be co-designed through partnerships between the ASCR computing and network facilities and key DOE Office of Science experimental and observational facilities. Working groups will be constituted initially with participation by various facility technologists required to build and establish the scaffolding on which the testbed will evolve. Application drivers, including testbed requirements, will be defined by participants from interested SC facilities. Together, the working groups will define an activity cadence and level of effort. SC leadership will set up a progress and reporting milestone structure.

References

[1] "Towards a Seamless Integration of Computing, Experimental, and Observational Science Facilities: A Blueprint to Accelerate Discovery," DOE ASCR IRI Task Force white paper, Mar 8, 2021, https://doi.org/10.2172/1863562

[2] Integrated Research Infrastructure Architecture Blueprint Activity (Final Report 2023), https://doi.org/10.2172/1984466[3]

[3] Report of the DOE Workshop on Management, Analysis, and Visualization of Experimental and Observational data – The Convergence of Data and Computing, https://doi.org/10.2172/1525145

[4] Envisioning Science in 2050, https://doi.org/10.2172/1871683

[5] Office of Science User Facilities Lessons from the COVID Era and Visions for the Future, https://doi.org/10.2172/1785683

[6] E. Lancon, et al., Background and Roadmap for a Distributed Computing and Data Ecosystem. United States: N. p., 2019, https://doi.org/10.2172/1528707

[7] D. Bard, et al., The Superfacility Project Report, https://arxiv.org/abs/2206.11992

[8] R. Vescovi et al., Linking scientific instruments and computation: Patterns, technologies, and experiences, https://doi.org/10.1016/j.patter.2022.100606

[9] S. Somnath et al., Building an Integrated Ecosystem of Computational and Observational Facilities to Accelerate Scientific Discovery, https://doi.org/10.1007/978-3-030-96498-6_4

[10] Exascale Requirements Reviews: https://exascaleage.org; Crosscut report: https://exascaleage.org/crosscut-report/

[11] ESnet Requirements Reviews: https://www.es.net/science-engagement/science-requirements-reviews/requirements-review-reports/

Appendix A: Testbed Design and Establishment

The following are significant activities for an initial IRI Testbed to be constructed with R&D participation from key user facilities with a goal of continuous improvement and evolution. The funding and associated implementation plan are not described in this white paper and will be covered in additional documents or program announcements. This list is representative of how quickly an effort, if resourced appropriately, can be executed:

- **Inventory of deployed testbeds**
  - Physical infrastructure (e.g., hardware specs, topology, etc.)
  - Common Application Deployment Approach (e.g., Docker, etc.)
  - Access and control (e.g., scheduling, APIs, AuthN/AuthZ, etc.)
  - Opportunities to leverage logical separation but colocation with production systems
- **Testbeds integration design**
  - Definition and prioritization of primary test and deployment objectives
  - Physical infrastructure (e.g., how to connect the testbeds, etc.)
  - Testbed architecture (e.g., functional components, domain boundaries, etc.)
  - Security (e.g., security boundaries, interoperable AuthN, etc.)
- **Testbeds integration buildout**
  - Physical infrastructure (e.g., connecting the testbeds, etc.)
  - Coordination functions (e.g., scheduling, monitoring, orchestration, recovery, etc.)
  - Operational policies (e.g., Acceptable Use Policies, Service Level Agreements, etc.)
- **Application mapping and instantiation**
  - Identify representative science use cases that can exploit testbed features
  - Cross-facility workflow porting
  - Application realization and load testing
  - Measurement test-harness operation
  - Lessons tracking and next iteration planning
  - Data management (data publication, DOI minting, etc.)