SC25 Network Research Exhibition: Demonstration Publishable Abstract

IRI Compute Job Portability

Justas Balcas/Xi Yang, ESnet, jbalcas@es.net, xiyang@es.net

Abstract

The Integrated Research Infrastructure (IRI) initiative, led by the U.S. Department of Energy (DOE), seeks to establish a seamless, secure, and programmable scientific ecosystem connecting experimental instruments, high-performance computing (HPC) centers, storage platforms, and network infrastructures. The overarching goal is to enable integrated, cross-facility scientific workflows that reduce time-to-insight and enhance collaborative discovery.

The IRI architecture plans to support multi-site experimentation and data-intensive workflows by standardizing APIs and protocols for compute, storage, data movement, orchestration, and identity management. Core efforts include the Pathfinder projects, which demonstrate practical integrations across DOE light sources, fusion experiments, climate simulations, and the HPC and data centers at ALCF, OLCF, and NERSC.

To broaden impact and accelerate technology adoption, the IRI design and testing strategy allows for integration with advanced academic infrastructure such as the National Research Platform (NRP) and the FABRIC testbed. These platforms offer flexible, distributed, and high-performance compute and networking resources that mirror DOE facility use cases, enabling early prototyping and experimentation at scale.

This demonstration highlights how IRI services—when extended to platforms ASCR computing facilities (NERSC), and testbed like NRP and FABRIC—can enable end-to-end, portable, and orchestrated workflows. It showcases cross-domain resource federation, real-time network provisioning, and data mobility—all orchestrated through a unified framework. This work represents a critical milestone toward a resilient and flexible research cyberinfrastructure that can scale across multiple agencies and institutions

Goals

This PoC aims to demonstrate a "compute portability" workflow across distributed ASCR's computing facilities, like NERSC, and testbed resources, like FABRIC, NRP, and SC25 Scinet Computing Cluster and data streaming with EJFAT (ESnet/JLab FPGA Accelerated Transport). It will simulate an application workflow where compute jobs can move between facilities depending on resource availability, with integrated data access and transfers.

The goal is to show portable compute jobs running across different sites and facilities, while each might use different technologies (HTCondor, Slurm, Kubernetes), starting with basic job execution and data placement and movement using Janus tools, while providing network guarantees via SENSE. By integrating EJFAT with compute portability mechanisms, this PoC will show how data can be streamed directly into running jobs, bypassing traditional storage bottlenecks, and supporting iterative "compute-to-compute" patterns. The workflow will include automated job placement, network-aware data movement, and dynamic orchestration across heterogeneous environments, demonstrating the feasibility of end-to-end, data-driven scientific workflows that respond in real time based on available conditions.

- Enable tightly integrated workflows across experimental facilities, HPC systems, and high-speed networks.
- Establish unified interfaces and APIs for allocations, data and job movement, scheduling.
- Support data-intensive and time-sensitive applications, like experiment steering and real-time analytics, together with EJFAT.
- Foster modularity and portability, allowing seamless transitions of data and computation jobs across different facilities.

Impacts

This work establishes a blueprint for resilient, flexible, and portable scientific workflows that can span multiple DOE and academic facilities, enabling experiments and simulations previously limited by infrastructure boundaries and limitations. By demonstrating real-time compute combined with high-performance data streaming through EJFAT, researchers can more efficiently leverage distributed resources, reducing time-to-insight for data-intensive applications results. SCinet's infrastructure together with SENSE and National Research Platform are critical for these demonstrations, providing a close to production-quality environment that mirrors operational conditions at other facilities. The lessons learned and tools developed here will guide future IRI strategies, promoting modularity, interoperability, and rapid technology adoption across the broader research community.

Resources

- PoC will interconnect NERSC computer resources on Perlmutter (CPU/GPU), and Caltech/Fabric/NRP for additional CPU-only resources.
- SCinet Cluster Slurm scheduler and NRP Kubernetes control. (Minimum resource requirement: 100

Cores/200GB RAM, 10TB CephFS attached storage, Nodes connected at 100Gbps. Jobs will require 1Core/2GB Ram and partition resources dynamically - can span multiple nodes)

- High-performance networking control via ESnet to link DOE labs and user facilities, like NRP and Fabric. (see NRI106).
- One of the 400G connections to the SCinet NOC to include a dedicated port-mapped, or QoS protected, path to an ESnet Production Router (to support EJFAT demo, may need to go to EJFAT L3VPN) (See Tom Lehman's NRE proposal)
- Public IPv6 /60 Range and DNS Control.

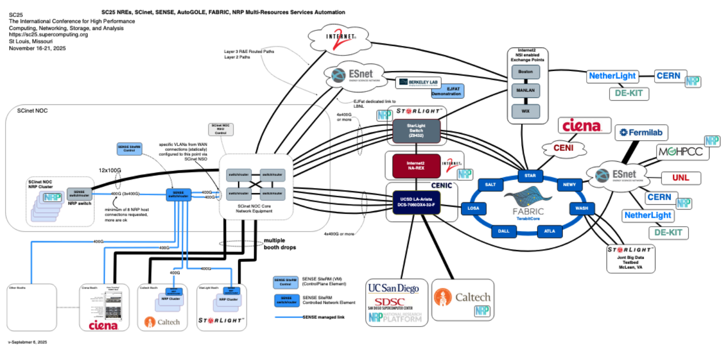

Involved Parties

- Justas Balcas, ESnet, jbalcas@es.net
- John MacAuley, ESnet, macauley@es.net
- Tom Lehman, ESnet, tlehman@es.net
- Xi Yang, ESnet, xiyang@es.net
- Yatish Kumar, ESnet, yak@es.net
- Hari Krishnan, AMCRD, hkrishnan@lbl.gov
- Abdelilah Essiari aessiari@lbl.gov
- Alberto Jimenez ajimenez@es.net
- Ezra Kissel kissel@es.net
- Nicholas Tyler tylern@lbl.gov
- Eli Dart dart@es.net