The Evolving ERN Cryo-EM Federated Instrument Pilot Project

Maureen Dougherty
maureen.dougherty@ernrp.org
Ecosystem for Research Networking
Philadelphia, USA

James Barr von Oehsen
barr@psc.edu
Pittsburgh Supercomputing Center
Pittsburgh, USA

Jason Kaelber
jason.kaelber@rutgers.edu
Rutgers University
Piscataway, USA

Jeremy Schafer
jschafer@rutgers.edu
Rutgers University
Piscataway, USA

D. Balamurugan
bala.desinghu@rutgers.edu
Rutgers University
Piscataway, USA

Morgan Ludwig
mludwig@techsquare.com
TechSquare
Cambridge, USA

John Goodhue
jtgoodhue@mghpcc.org
MGHPCC
Holyoke, USA

ABSTRACT

The Ecosystem for Research Networking (ERN) CryoEM Remote Instrument Access Pilot Project addresses the challenges and barriers facing the transition of instrument-driven science self-contained islands to federated wide-area internet accessible instruments. This project's goal is to facilitate cross-institutional infrastructure sharing at the interface of computing and electron microscopy through a federated web portal to access scientific instruments and generate workflows utilizing edge computing paired with advanced computing and real-time monitoring for experimental parameter adjustments and decision making. The intent is to foster team science and scientific innovation, with emphasis on under-represented and under-resourced institutions, through the democratization of scientific instruments. This paper presents a short summary some of the latest challenges encountered with this active project, creating a FABRIC Cloudlet and incorporating Pittsburgh Supercomputing Center's Bridges-2 as an advanced computing option.

CCS CONCEPTS

• **Theory of computation → Interactive computation**; • **General and reference → Validation**; • **Networks → Network design principles**; • **Security and privacy → Security protocols**.

KEYWORDS

Research Computing, Federation, Cloud Services, Edge Computing, Core Facilities

ACM Reference Format:
Maureen Dougherty, James Barr von Oehsen, Jason Kaelber, Jeremy Schafer, D. Balamurugan, Morgan Ludwig, and John Goodhue. 2024. The Evolving ERN Cryo-EM Federated Instrument Pilot Project. In *Practice and Experience in Advanced Research Computing (PEARC '24), July 21–25, 2024, Providence, RI, USA.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3626203.3670592

PEARC '24, July 21–25, 2024, Providence, RI, USA

ACM ISBN 979-8-4007-0419-2/24/07
https://doi.org/10.1145/3626203.3670592

1 INTRODUCTION

Scientific discoveries have almost always been driven by advances in instruments of observation, advances in methods and tools and infrastructure for data analysis and modeling. Today, the experimental instruments have become increasingly sophisticated (and expensive to build and operate), the scientific questions increasingly complex, and the data analysis and modeling increasingly reliant on access to state-of-the-art computing infrastructure, services, and expertise.

Transmission electron microscopes used in CryoEM cost millions of dollars and are increasingly located in national or regional facilities. If scientists can monitor results of their experiments in real-time [5], they can adjust parameters as an experiment goes forward or even optimize instrument utilization because they may know when their target has been achieved—or recognize and terminate a fruitless run. Additionally, due to the high cost of acquiring and operating such instruments, only scientists at select research centers and (R1) universities typically have access to and training on these instruments. The instruments, the computing and data infrastructure and the expertise needed to support and make effective use of them have been largely beyond the reach of faculty and researchers at many under-resourced colleges and universities. In response to these challenges and feedback from a series of NSF (OAC-2018927) outreach activities, the Ecosystem for Research Networking (ERN) [3] has partnered with Dr. Jason Kaelber of the Rutgers CryoEM & Nanoimaging Facility(RCNF) to launched the ERN CryoEM Federated Instrument Pilot Project. The project's goal is to simplify collaborative research by removing many of the barriers encountered when accessing remote research resources, with emphasis on under-represented and under-resourced institutions. This paper is a short summary of the phase 2 development efforts of this active project, creating a FABRIC Cloudlet and incorporating Pittsburgh Supercomputing Center's Bridges-2 into the analytical workflow as an advanced analytics option.

2 CRYO-EM FEDERATED INSTRUMENT PILOT DESIGN

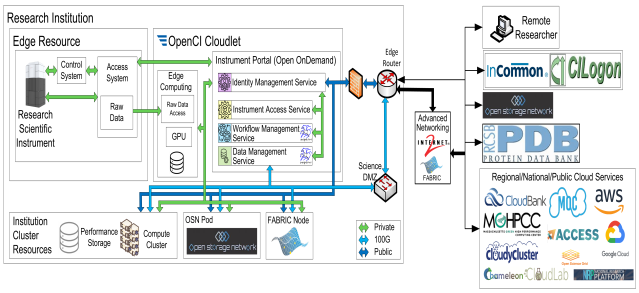

Figure 1: ERN Cryo-EM Federated Instrument Pilot Design

The initial design of the ERN pilot project's Instrument Cloudlet (Cloudlet), (see Fig. 1) hosts a user-friendly, secure, web-based resource portal enabling remote federated authorized access to the lab's transmission electron microscope (TEM); real-time workflows for experiment parameter adjustments and decisions; local edge computing image pre-processing of raw data; and advanced analysis on either Rutgers' private HPC cluster "Amarel," or external analytical compute resources from national centers or public cloud service providers. The design components include the scientific instrument (edge resource), research computing (edge computing), performance storage, advanced networking, instrument access portal, identity management service, instrument access service, analytics runtime service and data management service. While this project targets a particular set of instruments, the goal is to address the broader research community resources and leverage open-sourced applications, software, tools, vendor interfaces and concepts.

Phase 1 [2] established the basic infrastructure and framework for the Cloudlet. A customized, containerized Open OnDemand[4] (OOD) presents the web-based Instrument Portal with CILogon for authentication, an access mapping file for building, and launches a separate OOD user environment restricting access to authorized scientific instruments, and to provide a consistent user environment avoiding data access and transfer issues. The application CryoSPARC is used for monitoring, data management and workflow management of the experiment analysis workflows, leveraging local edge computing resources for pre-image processing with real-time parameter adjustments and Rutgers' Ameral cluster for advanced analytics of 2-D and 3-D protein mapping. Updates to the infrastructure were required to create and extend three secure network pathways/VLANs into the Cloudlet for isolated, secure network paths to access the TEM, the Ameral cluster and the public internet and to cross mount the TEM streaming raw data directory and Ameral project directory. This included firewall rule set restrictions on service ports and IP address ranges. VNC and websockify were installed and paired with OOD's proxy forwarding to facilitate communications to the TEM's controller. The remote lab researcher, through a Rutgers VPN, is able to use their web browser to login to the Instrument Portal using their institution's credentials, connect to the remote scientific instrument, configure their parameters, access the scientific application CryoSPARC and launch successful experimental workflows with real-time parameter adjustments. Dr. Kaelber is actively using the Cloudlet for his research workflows.

3 PHASE 2 - EXTENDING ACCESS

Phase 2 builds off the foundational Cloudlet development of Phase 1, introducing a new FABRIC node Cloudlet deployment, and expands advanced compute analytics resources availability to the Pittsburgh Supercomputing Center's(PSC) Bridges2. Expanding the FABRIC testbed to include the Cloudlet provides multiple options for those institutions looking to share their resources. For institutions with limited resources, a FABRIC Cloudlet can be installed on a remote FABRIC network node providing the broader research community remote access to their various scientific instruments through the Cloudlet's Instrument Portal without the cost and maintenance overhead of a local system. For local, production Cloudlet installations, a FABRIC Cloudlet testbed environment allows for development and enhancements of Cloudlet services without impacting the production Cloudlet and active scientific efforts. Incorporating an external compute resource such as Bridges2 expands advanced compute resource choices, providing more analytical options for the researcher and their scientific workflows.

3.1 FABRIC Cloudlet Implementation

FABRIC[1] is Adaptive ProgrammaBle Research Infrastructure for Computer Science Applications. It is a research infrastructure testbed designed to securely validate and prototype designs, facilitate experiments that closely emulate real-world production environments and enable the evaluation of test applications, protocols and services that run on any node in the network, not just at the edge or in the cloud. There are 33 United States and 4 international sites, distributed across commercial collocation spaces, national labs and campuses interconnected by high speed, dedicated optical links. Site hosts consist of storage, networking and a variety of compute options.

With the Cloudlet actively running workflows, and components still under design and development, the project team explored options for a testbed environment to avoid impacting the production environment. New components or services would be developed, stabilized, and evaluated in the testbed, then deployed into the production Cloudlet. FABRIC was chosen for our testbed environment as Rutgers was in the process of installing a FABRIC node. As with Phase 1 development, the team engaged the subject matter experts to discuss configuration build options and deployment.

The team reviewed the Cloudlet design and testbed concept with the FABRIC team, discussing infrastructure requirements and configuration challenges. The initial concept was to create a virtual machine (VM) from a slice of a FABRIC node, with compute and storage for edge computing, and the use of a Facility Port to make the connection to the external infrastructure element, the TEM, into the VM. This was the first application of a Facility Port to access a scientific instrument, and would require the lab's private VLAN to be extended to the FABRIC Cloudlet as had been done for the production Cloudlet (ERN Cloudlet). The secure public network would also extended into the FABRIC Cloudlet, necessary for remote access into the Instrument Portal. With the VM configured, an OS and supporting software would be installed and configured. Cloudlet prerequisites (public hostname, public IP address, DNS registration, SSL certifications, CILogon secrets and IDs), would be obtained and the current Cloudlet Github repo would be loaded and installed following the published README information. Due to installation delays of the Rutgers FABRIC node, a remote FABRIC network node was required for implementation. Since this is part of the FABRIC infrastructure, not Rutgers, this would be a more complex build and more likely how most institutions would stand up a FABRIC Cloudlet.

Through the FABRIC portal, a FABRIC account with associated Cloudlet team members was created. Of the three FABRIC network node resource locations assigned to the account (Massachusetts, TACC and Washington DC), Washington DC was chosen based on lowest possible latency possible. With the assistance of a lead FABRIC administrator, a FABRIC slice was configured with dual GPUs and 1TB of storage, and two IPV6 interfaces. The resulting VM ran a bare bones Rocky 8 OS, and two ethernet interfaces (IPV6 by default), accessing the FABRIC internal network, with no access to the public internet though external IPs and ports can be made available to the VM. One interface is for the data plane and the other for the management plane. The FABRIC team was provided an IP address for the RCNF lab's internal secure VLAN and Rutgers' firewall access lists were updated. The FABRIC administrator created the Rutgers Facility Port and connected it to the VM, which was then turned over to the Cloudlet team.

SE Linux was installed with associated kernel updates for security access restrictions and adjustments. Firewall access lists were updated, service ports restricted with nftables blocking the rest, aligning with the ERN Cloudlet configuration. CryoSPARC master instance was installed and then the Cloudlet github repo was downloaded, the parameterized configuration file updated, and the install script run. The install erred out as nat64 was unable to resolve the host name. It was determined that IPV6 was preventing podman from properly network address translating. This was addressed by adding '–net=host' to any command attempting to access the network in the build script. Lack of public internet access at this stage was impacting configuration.

The Cloudlet requires a registered DNS public IP address that can be routed back to it for SSL certificates, Open OnDemand, OIDC, authentication by CILogon, and access to the Instrument Portal. This is not currently supported by FABRIC's internal network so the teams worked together to come up with routing solutions that would maintain an interface into the FABRIC network, and provide the public internet access needed. After numerous attempts, a simple solution came to light. By adding the FABRIC extension FABnet4Ext to the FABRIC account, the data plane interface was converted to IPV4 and became a Layer 3 network and the VM's public interface. Routing was configured with IPV4 port forwarding turned on to enable the data plane for public network traffic flow. The VM's public hostname and a public IP address were then registered with DNS. Due to security pathway concerns, FABRIC does not normally recommend peering the data plane as a public interface to the internet for most sites, but allowed it for this joint effort. The management plane was left at IPV6 for both the internal FABRIC communications and the Layer2 circuit from Washington DC to Rutgers' secure VLAN with the research instrument. This secure ssh tunnel was established by connecting the Facility Port to this VLAN and updating the network routing and firewall access lists with the VM's private IP address. This allowed the TEM's camera streaming output raw data file system to be cross mounted to the FABRIC Cloudlet. Complications with the cross mounting of the Ameral project file system led to the decision of all analysis being local to the FABRIC Cloudlet for this initial deployment. With network routing addressed, the last of the Cloudlet prerequisites were addressed, and CryoSPARC installed. The build was completed, and the FABRIC Cloudlet turned over to the researcher for testing.

Over VPN, our researcher was able to access the FABRIC Cloudlet, launch CryoSPARC to monitor the raw data directory for streaming data, and use local compute and storage resources for image processing and protein structure determination. Through the Instrument Portal the TEM's control panel was reached, the microscope's parameters configured and the test workflow launched with all analysis local. The workflow and test were successful. Access and the workflow pathway went from the remote researcher in New Jersey, to a FABRIC network node in Washington DC, leveraging streaming data from the lab at Rutgers.

The FABRIC Cloudlet design is still under development. The initial VM configuration was created manually and not persistent. If the FABRIC node goes down, the FABRIC slice and VM would have to be reconfigured and the FABRIC Cloudlet re-installed. The FABRIC team recommended a parameterized configuration file and script to be executed by the FABRIC slice initialization for configuring the VM, and the use of a FABRIC Jupyter notebook for portability and reproducability. These recommendations are currently under development.

3.2 Pittsburgh Supercomputing Center Bridges-2 Integration

A recent association with the Pittsburgh Supercomputing Center (PSC) motivated the integration of the ACCESS resource, PSC's Bridges-2 cluster, into the ERN Cloudlet infrastructure as an advanced analytics option. This posed several challenges, most focusing on data management and the CryoSPARC application. Streaming intermediate files from the Cloudlet need to transfer onto Bridges-2, requiring a cross mount via a secure network pathway of either the intermediate output directory onto Bridges-2 or a Bridges-2 directory cross mounted onto the Cloudlet. The UID:GID namespace would need to match across the systems, ensuring proper data ownership and access. The CryoSPARC master instance on the Cloudlet would need to launch jobs onto Bridges-2's queue, and communicate with the CryoSPARC worker instances running on the Bridges-2 compute nodes. Latency and job queuing times were also a concern. To map solutions to these issues, the Cloudlet and Bridges-2 collaborated.

The Cloudlet team applied for and received an ACCESS Explore account, which provided access to storage (PSC OCEAN), GPU service units, and Regular Memory service units on Bridges-2. The group was given full permissions to the project's shared folder on the OCEAN project folder on Bridges-2. A CryoSPARC worker instance was installed into the project's folder. Rutgers' network infrastructure was updated to establish a secure VLAN between the ERN Cloudlet and the Bridges-2 resources. This enabled the cross mount of the project folder on Bridges-2 to the Cloudlet, and communications between the CryoSPARC master instance running on the Cloudlet and the CryoSPARC worker instances running on the Bridges-2 compute nodes.

The initial attempt to mount the project directory from Bridges-2 to the ERN Cloudlet failed. The username and associated UID:GID on Bridges-2 did not match the Cloudlet's UID:GID for the associated user. A new local Cloudlet user account with matching, username, UID:GID was created and the mount was successful. This is not the ideal solution and will require further investigation into a secure, more dynamic solution for user namespace and external storage resources as data security and integrity is a priority.

The Bridges-2 project cross mount is the location for the edge-computing pre-image processing output, which is the input for the protein structure determination advanced analysis on Bridges-2. A new CryoSPARC SLURM job script template was created to identify this data location for monitoring of data and compute job submission to Bridges-2. Upon the detection of new data in this directory, the CryoSPARC master instance ssh's to the Bridges-2 headnode and submits jobs. The job launches the CryoSPARC worker instance on the compute node using the detected data and communicates job status with the CryoSPARC master. Concerns over CryoSPARC database and the storage restrictions of the ERN Cloudlet led to the reuse of the existing CryoSPARC master ports, its database, and the data directory location. The running CryoSPARC master and its associated database, including data, application user names and passwords, was shut down and restarted with the new configuration. We did encounter a CryoSPARC version mismatch between the master and worker instances, requiring an update to match versions.

With all components in place, we conducted a pilot test. As intermediate files populated the monitored directory, Bridges-2 jobs were submitted and launched. The CryoSPARC worker on the compute node failed when attempting to the Cloudlet on what was identified as an unexpected and restricted port. Analysis of the network flow between the CryoSPARC worker and the Cloudlet confirmed that the list of authorized Bridges-2 IP addresses defined was too narrow. Once the firewall access lists and the Cloudlet's nftables were updated, communications between the CryoSPARC master and workers worked and the test workflow completed successfully.

The preliminary workflow design successfully incorporated the external compute resources of PSC Bridges-2 compute resources into the Cloudlet workflow. Additional investigation and evaluation of data security, and matching username, UID:GID solutions; and latency impact on the workflows for broader adoption is required.

4 CONCLUSION

The Phase 2 Cloudlet design facilitates a broader range of both advanced compute resources and deployment and testbed options, particularly for smaller, under resourced institutions. The FABRIC Cloudlet is a versatile component of the FABRIC infrastructure network fostering new scientific discovery. As a testbed, development of modifications and services do not impact the production environment. As an Instrument Portal it offers centralized, remote, secure, shared access to scientific instruments, either locally or for remotely for under resourced institutions looking to share resources. This has the potential to be a new revenue, core service enabling new scientific instrument opportunities for the research community. Traditional HPC activities involve asynchronous processing of data and are temporally separable from the act of raw data acquisition. Real-time computation is traditionally performed at the site of the instrument. Although obtaining a reservation for Bridges-2 for the pilot study was not difficult, its need does point to the philosophical difference in design of edge computation systems for real-time decision making in "wet lab" experiments, where on-demand resource allocation is not a convenience but an integral and essential component of the value added by computation. The magnitude of data reduction provided by the cloudlet will enable even public cloud resources to be brought to bear on this problem, enabling diverse instruments (and users) to benefit from real-time remote control and analytics. By cultivating a flexible, democratic ecosystem for sharing research instruments we are establishing the connective tissue that will bridge a broad range of scientific instruments across scientific domains for a diverse research community establishing a network of resources that will foster team science and collaborations with emphasis on under-represented and under-resourced colleges and institutions. The research community is encouraged to explore these resources, offer feedback and contribute to development through our Github repository, https://github.com/mghpcc/SciEdge.

ACKNOWLEDGMENTS

The authors thank the following for their contributions to the project: all ERN Working Groups; Open OnDemand; FABRIC; and Pittsburgh Supercomputing Bridges-2 teams; and the ERN Steering Committee for guidance and support. This work was supported by the NSF CC* CRIA grant (OAC-2018927) and the NIGMS grant R21 GM140345.

REFERENCES

[1] I. Baldwin, A. Nikolich, J. Griffioen, I. Monga, KC. Wang, T. Lehman, and P. Ruth. 2019. FABRIC: A National-Scale Programmable Experimental Network Infrastructure. *IEEE Internet Computing* 23, 6 (Nov 2019), 38 – 47. https://doi.org/10.1109/MIC.2019.2958545

[2] Maureen Dougherty, Michael Zink, James von Oehsen, Kenneth Dalenberg, D. Balamurugan, Jason Kaelber, Jeremy Schafer, John Goodhue, Wolf Hey, Morgan Ludwig, Boyd Wilson, and Cole McKnight. 2022. The ERN Cryo-EM Federated Instrument Pilot Project. *Proceedings of ACM PEARC22 Conference* 52 (2022), 1–4. https://doi.org/10.1145/3491418.3535141

[3] ERN. 2024. The Ecosystem for Research Networking. Retrieved April 24, 2024 from https://ern.ci

[4] D. Hudak, D. Johnson, A. Chalker, J. Nicklas, E. Franz, T. Dockendorf, and B. L. McMichael. 2018. Open OnDemand: A web-based client portal for HPC centers. *Journal of Open Source Software* 28, 1 (Mar 2018), 235–242. https://doi.org/10.21105/joss.00622

[5] R. F. Thompson, M. G. Iadanza, E. L. Hesketh, S. Rawson, and N. A. Ranson. 2019. Collection, pre-processing and on-the-fly analysis of data for high-resolution, single-particle cryo-electron microscopy. *Nature Protocols* 14, 1 (Jan 2019), 100–118. https://doi.org/10.1038/s41596-018-0084-8