The ERN Cryo-EM Federated Instrument Pilot Project

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

Michael Zink
zink@ecs.umass.edu
University of Massachusetts, Amherst
USA

James Barr von Oehsen
Kenneth Dalenberg
D. Balamurugan
Jason Kaelber
Jeremy Schafer
jbv9@rutgers.edu
kd541@oarc.rutgers.edu
bala.desinghu@rutgers.edu
jason.kaelber@rutgers.edu
jschafer@rutgers.edu
Rutgers University
USA

John Goodhue
jtgoodhue@mghpcc.org
MGHPCC
USA

Wolf Hey
wzh5214@psu.edu
Penn State University
USA

Morgan Ludwig
mludwig@techsquare.com
TechSquare
USA

Boyd Wilson
Cole McKnight
boyd@omnibond.com
cole@omnibond.com
Omnibond
USA

ABSTRACT

Feedback and survey data collected from hundreds of participants of the Ecosystem for Research Networking (formerly Eastern Regional Network) series of NSF (OAC-2018927) funded community outreach meetings and workshops revealed that Structural Biology Instrument driven science is being forced to transition from self-contained islands to federated wide-area internet accessible instruments. This paper discusses phase 1 of the active ERN CryoEM Federated Instrument Pilot project whose goal is to facilitate inter-institutional collaboration at the interface of computing and electron microscopy through the implementation of the ERN Federated OpenCI Lab's Instrument CI Cloudlet design. The conclusion will be a web-based portal leveraging federated access to the instrument, workflows utilizing edge computing in conjunction with cloud computing, along with real-time monitoring for experimental parameter adjustments and decisions. The intention is to foster team science and scientific innovation, with emphasis on under-represented and under-resourced institutions, through the democratization of these scientific instruments.

CCS CONCEPTS

• **Theory of computation** → **Interactive computation**; • **General and reference** → Validation; • **Social and professional topics** → **Computing / technology policy**; • **Networks** → **Network design principles**; • **Security and privacy** → **Security protocols**.

KEYWORDS

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

ACM Reference Format:
Maureen Dougherty, Michael Zink, James Barr 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. In *Practice and Experience in Advanced Research Computing (PEARC '22), July 10–14, 2022, Boston, MA, USA.* ACM, New York, NY, USA, 4 pages. https://doi.org/10.1145/3491418.3535141

PEARC '22, July 10–14, 2022, Boston, MA, USA

ACM ISBN 978-1-4503-9161-0/22/07…$15.00
https://doi.org/10.1145/3491418.3535141

1 INTRODUCTION

Transmission electron microscopes used in cryoEM cost millions of dollars and are increasingly located in national or regional facilities [4]. They can collect hundreds of multi-frame images per hour, with compressed dataset sizes in the terabytes. 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. Currently, insufficient data transfer rates, and significant latency and HPC queue wait times do not permit on-the-fly processing in either public clouds or university clusters. Due to the high cost of acquiring and operating such instruments, typically only scientists at select research centers and (R1) universities 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 feedback from the research community, the Ecosystem for Research Networking (ERN) [2] has partnered with the Rutgers CryoEM & Nanoimaging Facility (RCNF) to launch the CryoEM Federated Instrument Pilot Project. Instruments in the RCNF include a Thermo Fisher Scientific (TFS) Talos Arctica transmission electron microscope (TEM), indirectly attached to the university network via a bridging PC, with limited remote access behind the university's VPN and a secured VNC server. The goal is to simplify collaborative research by removing many of the barriers encountered when accessing these resources. The ERN is creating an easy to use, secure, web-based resource portal (leveraging Open OnDemand) providing remote federated authorized access to the lab's TEM; real-time workflows for experiment parameter adjustments and decisions; local edge computing image pre-processing of raw data; and additional analysis on either Rutgers' private HPC cluster, Amarel, or public cloud resources. This paper discusses phase 1 of this project.

2 CRYO-EM FEDERATED INSTRUMENT PILOT DESIGN

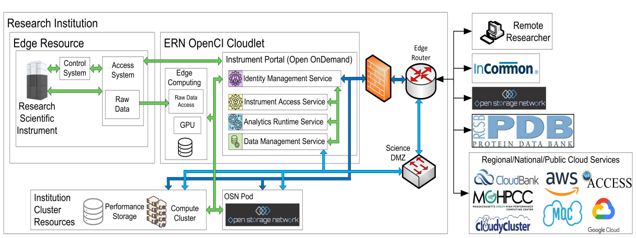

Figure 1: Overview of the ERN Cryo-EM Federated Instrument Pilot Design

The pilot design (see Figure 1) is based on the ERN Federated OpenCI Lab (OCIL) [1] collaborative infrastructure model's Instrument CyberInfrastructure (CI) Cloudlet. Based on the findings of previous ERN research community outreach workshops and pilot projects, it presents a common framework for federated access of research resources. Developed for simple deployment, it utilizes Ansible playbooks and containers for support, maintenance, and portability. The design leverages open-sourced applications, software, tools, vendor interfaces and concepts, with the institution's existing resources and policies, to present a secure user environment. The following OCIL components will be deployed for this pilot: 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. (See PEARC'21 OCIL white paper for more details. [1]).

The Cloudlet represents a hybrid between instrument-local computing and cloud/remote computing whereby the cloudlet performs select, typically high-I/O operations locally while seamlessly integrating with a remote HPC architecture. The Cloudlet is designed to be the bridge device to the scientific instrument through a web-based Instrument Portal. The Instrument Portal and its supporting services, presenting a secure, user friendly, web-based environment for real time analytics, reviewing, tuning, and filtering. The edge computing approach of the Cloudlet is a novel way to address a key infrastructural limitation faced by the field and is synergistic with those advances. Through the Instrument Portal, remote researchers would be able to login via credentials from their home institution, identify the portal's available resources, request access at certain times and/or levels of control, and utilize the instrument and resources in real time on low-latency connections, including cloud computing for analysis. Instrument administrators would register and manage their scientific instruments, review requests and either accept or decline them.

We anticipate significant challenges with identity management (authentication, access and authorization) throughout the environment, particularly with the instrument interface and data management throughout the architecture. Additional challenges include: how to simplify guest researcher accounts; interfacing with a variety of hardware and vendor interfaces potentially using different access methods; concerns about data and metadata ownership, access and management; financial responsibility for compute cycles on institution resources and in the cloud; local institution's security and network policies; secure container management; and making support, maintenance and portability as simple as possible. The planning and development process worked to address some of these challenges, by contacting various entities for their input and feedback on design solutions, including the TEM vendor ThermoFisher Scientific. It was determined a small, phased approach to the deployment of the design would be optimal for achieving our goals, starting with one institution and on-boarding others as we progress.

Open OnDemand[3] with CILogon for authentication was selected as the foundation of the portal. It is a well established web portal for accessing HPC resources, built on open source software, securely hosted behind a web proxy incorporating federated authentication. It has a customizable landing page, whose default installation includes onDemand applications for file management, job management and monitoring, and command line access, aligning with several Cloudlet services. OOD is well supported with training, documentation, and workshops, and encourages the community to customize their native installation to meet the needs of their researchers. There is a strong user community providing their solutions to common needs and wants through open sourced repositories. By implementing OOD for our instrument portal landing page, we can leverage these existing resources, and contribute our efforts to this thriving community.

For container development, Podman, an open-source container management tool using Open Container Initiative (OCI) and developed by RedHat was selected. It provides full image management, container life cycle management, supports groups of containers that share resources and are managed together (pods), has no running manager daemon, and supports multiple container image formats. Podman provides added layers of security with rootless containers, resource isolation of containers and pods, and because it does not have a running manager daemon, should still be responsive as utilization increases.

3 PILOT - PHASE 1 INSTRUMENT ACCESSIBILITY

Phase 1 of the pilot builds the basic framework for accessing and utilizing the scientific instrument through the Instrument Cloudlet for internal researchers, achieved in two steps. Step 1, establishes VNC/noVNC communications between the remote researcher and the scientific instrument. The Cloudlet is installed, and the appropriate hardware, software and configurations are provided to remotely access the TEM through a noVNC browser connection. Remote researchers will launch real-time processing workflows leveraging the Cloudlet's edge computing to generate pre-processed images, and the Amarel cluster for 2D and 3D mapping and structure prediction. Step 2 would provide the Cloudlet's base Instrument Portal interface, a secure, user friendly environment for instrument accessibility (step 1) for remote researchers, in a manner that is easy to maintain and manage, leveraging open source wherever possible. Phase 1 required team discussions with vendors, hardware and software updates and upgrades to systems and network infrastructure to address integration, accessibility, latency and throughput across the workflow pathway.

To determine the optimal method to access the TEM's Instrument Access System and Controls and routing requirements, the team had several discussions with TFS regarding their existing and developing toolsets. The TEM currently utilizes TFS' Remote Access Program for Interactive Diagnostics (RAPID) application with VNC for communications on a proprietary Instrument Access system connected to the lab's private network, but VNC has a known polling latency issue. A review of various TFS projects and potential options, identified TeamViewer under RAPID2 as a potential option. Unfortunately it was not viable due to its inability to route through the proprietary server, necessary for OOD implementation. Ultimately the VNC polling latency issue was ameliorated by polling from a noVNC instance on the Cloudlet. VNC also provides a more generalized solution, for the broader scientific community and instrument deployments. To address phase 1 routing from the institutions edge to the TEM , the firewall whitelist was updated to reflect the necessary VLANS, limited IP address ranges, and appropriate service ports to provide secure access to the Cloudlet and it's access into the lab's private network.

3.1 Step 1 - Basic Remote Accessibility

Step 1 (See Figure 2) establishes remote access of the TEM through a noVNC browser connection, and runs a real-time workflow with the Cloudlet's edge computing. The Instrument Cloudlet was provisioned to host the web-based portal, provide edge-computing resources, and a data transfer node for further analysis either locally or at remote public and private clouds. Installed in a data center as close to the edge resource as possible, it contains dual

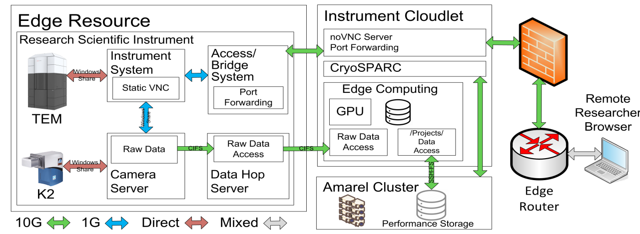

Figure 2: Step 1 basic remote accessibility diagram

AMD EPYC 7452 32-Core processors, 512GB memory, dual Nvidia A100 GPUs, and 2.2TB SSD storage, running Rocky Linux 8.4 with Docker, Apache, iperf3, CryoSPARC and an institution supported instance of LDAP. Switches, interfaces and cables were procured and installed to provide 10G and 100G, after step 1, interfaces to the Cloudlet to support live data stream processing transfer rate requirements. The configuration has been ansiblized for ease of management and maintenance. The TEM's camera output raw data file system can not be directly mounted onto the Cloudlet without additional hardware, so a data hop server is used to cross mount the file system to the Cloudlet. Using 10G connections and mounting read-only CIFS, the file system is first mounted on the data hop server, then onto the Cloudlet. Multiple hops is inefficient, introducing additional latency impacting performance. Under step 2, the raw data folder will be a direct mount on the Cloudlet, eliminating the performance hit. A direct mount of the Rutgers' Amarel cluster projects directories on the Cloudlet would also achieve optimal performance for the workflow. This GFPS filesystem requires licensing and network adjustments for a natively mounted directory, not available at this time. SSHFS provides a low-performance but generalizable solution for directory access. Normally the performance of SSHFS is insufficient to support real-time analysis of cryoEM datasets, however, because the most data-intensive operations are performed on the Cloudlet, the required network I/O to and from the remote cluster is reduced by 1-2 orders of magnitude, so SSHFS is viable for this pilot study.

In November 2021, from an off-institution remote system, a Rutgers' VPN session was established, and from the remote system's local browser, a noVNC address of the Access/Bridge System submitted. Network routing directed the connection to the Cloudlet's noVNC server and port forwarding passed the connection to the Access/Bridge System, which port forwarded to the Instrument System's static VNC process. After a successful login, a workflow was launched. The Cloudlet's CryoSPARC application monitored the cross mounted raw data directory and ran datasets through pre-process imaging analysis with the edge computing GPUs, writing output to the cross mounted projects Amarel file system. CryoSPARC then submitted jobs to Amarel based on the data sets sent to the projects directory, with real-time processing of the pre-processed image files, for 2D alignment and 3D structure refinement. We collected 320 images per hour of a novel complex of the transmembrane protein TolC over a two-day period, processing 2.5TB. Computation output was less than three minutes behind actual data acquisition with no increase over the course of the experiment. Graphically monitoring the output and analytics, the quality of the incoming data was assessed and experimental decisions made. Bandwidth measurements confirm that the requisite data transfer rate from the instrument exceeded 1Gbps, but the data transfer rate between the cloudlet and cluster was low enough to be accommodated even by a 1G connection. Assessment of the workflow determined that data transfer was the main bottleneck, particularly regarding the sshfs mount. Once the data hop server is eliminated with the direct mount, this test will be rerun. Since the pre-process output was significantly smaller than the raw data, data transfers for next step analysis will take less time than transferring the full raw data set. Step 1, with its associated network, software and hardware installations and modifications, successfully achieved its goals. Lab researchers are actively, remotely, utilizing the Cloudlet's edge computing and workflow process for their research efforts.

3.2 Step 2 - Instrument Portal Implementation

Step 2 simplifies step 1 by providing VNC access of the TEM through the Instrument Portal and the use of a navbar option, removing manual steps, including the use a unique URL. An Ansiblized OOD would be installed, and modified to authenticate via CILogon and a local institution LDAP instance with VNC access option added to the customized portal. CILogon authenticates through the researcher's own institutions authentication process, returning authorized user credentials. OOD will be containerized and customized via parameterized scripts and persistent configuration files for OS independence, portability and images built dependent on authorized GIDs for new functionality like the TEM access. OOD customization is well documented, but still requires expertise to coordinate and implement these modifications, script the VNC access to the TEM.

The initial step 2 implementation was achieved by the team OOD expert at a remote site, without access to a TEM, used Globus for authentication mapping against a local user file, and used local SSL certificates for Apache communication between OOD and the host apache. Knowing that the Cloudlet implementation at other institutions may be different than our initial install, this was an opportunity to document the migration steps.

To simulate the TEM's static VNC server, a containerized VNC Flask application was launched separately from the OOD instance, running an off-the-shelf VNC server. Upon a successful login to the portal, an OOD image is built with podman and a buildah stack based on the UID and GID from the authentication user mapping, and a rootless context container is launched. The host UID is mapped with the container UID to allow consistent host/container file permissions for mounted folders, while the white-listed GID adds the VNC Server option to the navbar. After successful VNC Server password submission, an HTLM template is called that provides the appropriate ULI which OOD then reverse proxies, redirecting to the VNC server container. A private Github repository was created and populated to provide future contributions and access to community. The user environment presented within the container would be based on the LDAP attributes from the user mapping, and initially would present home directories as a place holder, raw data and /projects/ directories would be mounted as in step 1 and applications would be available via the Cloudlet or OOD navbar. The workflow ran seamlessly in the testbed environment.

The step 2 migration requires a public facing hostname and ip address, SSL certificates, a pathway to the public internet, and the appropriate OOD scripts and configuration files updated to reflect authentication via CILogon, mapping using the local LDAP instance, new SSL certificates and the TEM's VNC server IP address. For security purposes a new VLAN was created to isolate the network traffic and network routing and firewalls updated. Complications with network routing, has delayed progress of the migration. Once resolved, and updates confirmed, step 2 will be verified including workflows similar to step 1. Any issues will be addressed accordingly.

4 CONCLUSION

The initial phase 1 efforts of the ERN Cryo-EM Federated Instrument Pilot remote access to edge scientific instruments by remote internal researchers for analytical workflows using edge computing is feasible and beneficial to their projects, but delays have impacted forward movement. Launched during Covid-19, restrictions impacted experienced staff availability, activities requiring an on-campus staff presence, staff coordination efforts between multiple departments/institutions, and procurement procedures and timelines due to stock availability and transport. As restrictions ease, some of these issues will subside, and the others will be resolved and incorporated into the final pilot design solution and shared with the community through the github repository. The current phase 1 state has been well received by the RCNF lab, who are actively making use of the currently deployed resource. Moving forward, we anticipate the ecosystem being developed through this pilot will break down resource access barriers, foster team science and the democratization of scientific instruments to benefit the research community, with emphasis on under-represented and under-resourced colleges and institutions, who typically do not have access to these resources, nor expertise to take full advantage of them.

ACKNOWLEDGMENTS

The authors wish to thank all the members of the Ecosystem for Research Networking Architecture and Federation and Structural Biology Working Groups for their contributions to the project and the ERN [2] Steering Committee for guidance and support. This work was supported by the National Science Foundation CC* CRIA Eastern Regional Network planning grant (OAC-2018927).

REFERENCES

[1] Maureen Dougherty, Michael Zink, and James von Oehsen. 2021. Identifying Research Collaboration Challenges for the Development of a Federated Infrastructure Response. *ACM PEARC21 conference (PEARC '21)* 40, 1 (2021), 1–4. https://doi.org/10.1145/3437359.3465594
[2] Ecosystem for Research Networking. 2021. The Ecosystem for Research Networking (formerly Easter Regional Network). Retrieved April 12, 2021 from https://ernrp.org
[3] 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
[4] Sriram Subramaniam. 2019. The cryo-EM revolution: fueling the next phase. *IUCrJ* 6, 1 (Jan 2019), 1–2. https://doi.org/10.1107/S2052252519000277
[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