



# The Ecosystem for Research Networking (ERN): Democratizing Access to Research Instruments and Data

All Hands Meeting  
March 23 – 24, 2023

# Welcome & Introductions

# Ecosystem for Research Networking (ERN)

## Vision:

Simplify, support, catalyze, and foster multi-campus collaborations and partnerships between academic institutions of all types and sizes across the U.S. that advance the frontiers of research, pedagogy, and innovation.

## Mission:

To achieve the vision through a consortium of academic institutions, research facilities, core service providers, network providers, and industry partners, both public and private, **organized around a shared interest in supporting and enabling collaborative data and computation-enabled science by providing standards, blueprints, policies, and training associated with the design and implementation of an infrastructure to access data and research instruments**, a distributed federated environment designed to simplify, support, and encourage collaborative science, scholarship, and education.

To realize the mission and vision, **ERN will enable collaborations for democratization of access to research instruments, technical expertise, infrastructure, services, and resources** to lower barriers to participation for scientists engaged in collaborative research across institutional and disciplinary boundaries.

# ERN Steering Committee



John Barden



Maureen Dougherty



Forough  
Ghahramani



John  
Goodhue



Jim Griffioen



Vasant Honavar



Florence Hudson



Ron Hutchins



Jim Kyriannis



David Marble



John Moore



Carrie Ramp



Bruce Segee  
Emeritus



Barr von Oehsen



Yifeng Zhu

# Purpose of this AHM

# Purpose of this AHM

To connect with the ERN community to:

- Discuss current and future ERN projects, collaborations and partnerships
- Give updates on the ERN Working Groups' efforts and activities
- Discuss next steps on democratizing access to research instruments
- Share workflow and remote access initiatives
- Discuss scope of the ERN
- Discuss future funding opportunities, future workshops, and other topics of interest

# The ERN All Hands Meeting

## Topics of Interest:

- Security
- Network Infrastructure
- Organization - Governance, Bylaws, Committees
- Research Instrument and Data Abstractions for Accelerating Discovery
- Machine-actionable Data Management and Curation
- Leveraging the Current CI Ecosystem
- Translational Computer Science
- Education, Training, and Workforce Development
- Diversity, Equity, and Inclusion

# Day 1 Agenda

1:00 - 1:15 Welcome & Introductions

1:15 - 1:30 Purpose of this AHM

1:30 - 2:15 Keynote - Democratization of Research  
Speaker Alan Blatecky - The Missing Millions

2:15 - 3:30 State of the ERN

- Architecture and Federation
- Broadening the Reach
- Structural Biology, Materials Discovery, CS, Policy and Next Steps

3:30 - 4:00 Break

4:00 - 4:30 End of Day 1 Wrap-up

5:30 - 7:00 Reception



# Day 2 Agenda

8:00 - 9:00 Meet and Greet Breakfast

9:00 - 9:15 Welcome and set tone for the rest of the morning

9:15 - 10:15 Breakout Groups - Discussion Topics:

10:15 - 10:45 Break

10:45 - 11:30 Open Discussion - Summary of Breakout Groups

11:30 - Noon Action items and Final Thoughts

# State of the ERN

# Eastern Regional Network 2019

# Vision & Mission: Original Version

- **Vision:** To simplify multi-campus collaborations and partnerships that advance the frontiers of research and innovation.
- **Mission:** Through a partnership of educational institutions, research facilities, regional network providers, and Internet2, the ERN is committed to providing layered and transparent access to shared data and computing facilities for research projects located at and between partner sites to address the growing need for a regional research platform designed to support a diverse set of science drivers and education needs. The resulting layered approach offers the research community across the region access to a broad range of services and resources that are not available on any one campus alone.

# Vision & Mission: Current Version

## Vision:

Simplify, support, catalyze, and foster multi-campus collaborations and partnerships between academic institutions of all types and sizes across the U.S. that advance the frontiers of research, pedagogy, and innovation.

## Mission:

To achieve the vision through a consortium of academic institutions, research facilities, core service providers, network providers, and industry partners, both public and private, **organized around a shared interest in supporting and enabling collaborative data and computation-enabled science by providing standards, blueprints, policies, and training associated with the design and implementation of an infrastructure to access data and research instruments**, a distributed federated environment designed to simplify, support, and encourage collaborative science, scholarship, and education.

To realize the mission and vision, **ERN will enable collaborations for democratization of access to research instruments, technical expertise, infrastructure, services, and resources** to lower barriers to participation for scientists engaged in collaborative research across institutional and disciplinary boundaries.

# ERN Origins and Evolution

- **Germ of an idea:** at 2017 National Research Platform meetings
  - Can the regional research platform idea work in the Northeast?
  - Rutgers, OSHEAN, KINBER
- **Coalition of the Willing:** January 2018 gathering at Rutgers
  - Can we find something of common interest to work on?
  - KINBER, Rutgers, OSHEAN, MGHPCC, Internet2, NYSERnet
- **Resource Federation Proof of Concept:**
  - Can we work together to prototype a potentially beneficial resource federation idea?
  - Rutgers, MGHPCC, Syracuse, NJ Edge, U Maine, Google, Internet2
  - “Early if not elegant” approach; emphasis on getting people to work together

# Origins and Evolution

- **Northeast perfSONAR Mesh:**

- Can we work together toward a regional science DMZ?
- Federation sites plus
  - BU, Brown, Colby, CEN, KINBER, Network Maine, OSHEAN, UNH
- Less-specialized technology allows expanded participation

- **Growing interest:**

- Maybe we really can do this?
- [in progress] expanded steering committee; per-project working groups
- All of the above plus
  - U Delaware, NJIT, U Buffalo, Bucknell, UMass Amherst, Yale, Princeton, BTAA, U Kentucky, Case Western Reserve

# Resource Federation Proof of Concept

The screenshot shows a Slurm web interface window titled 'Sview (1)' displaying a table of job status. Below it is a terminal window showing the output of 'squeue' and 'sacct' commands.

JobID	Partition	UserID	Name	State	Time Running	Node Count	NodeList	ClusterName
67111034	ern	babbott	test.slurm	RUNNING	00:00:20	1	edge	edge
67111035	ern	babbott	test.slurm	RUNNING	00:00:19	1	gcloud-compute1	gcloud
67111036	ern	babbott	test.slurm	RUNNING	00:00:19	1	gcloud-compute2	gcloud
67111038	ern	babbott	test.slurm	RUNNING	00:00:18	1	its-nefnode-3	otto
67111032	ern	babbott	test.slurm	RUNNING	00:00:20	1	mace1	mace
67111033	ern	babbott	test.slurm	RUNNING	00:00:20	1	mace2	mace
67111039	ern	babbott	test.slurm	RUNNING	00:00:14	1	node1	acg
67111037	ern	babbott	test.slurm	RUNNING	00:00:15	1	node3	acg

```
ern* mace up 3-00:00:00 2 idle mace1-2]
ern* edge up 3-00:00:00 1 idle edge
ern* otto up 3-00:00:00 3 idle its-nefnode-[1-3]
ern* mfedhpc up 1-00:00:00 3 idle mfedhpc[1-3]
ern* gcloud up 3-00:00:00 2 idle gcloud-compute[1-2]
ern* acg up 3-00:00:00 3 idle node[1-3]

Submitted batch job 67111032
Submitted batch job 67111033
Submitted batch job 67111034
Submitted batch job 67111035
Submitted batch job 67111036
Submitted batch job 67111037
Submitted batch job 67111038
Submitted batch job 67111039

JOBID PARTITION NAME USER ST TIME NODES NODELIST(REASON)
67111037 ern test.slu babbott R 0:15 1 node3
67111039 ern test.slu babbott R 0:14 1 node1
67111035 ern test.slu babbott R 0:19 1 gcloud-compute1
67111036 ern test.slu babbott R 0:19 1 gcloud-compute2
67111038 ern test.slu babbott R 0:18 1 its-nefnode-3
67111032 ern test.slu babbott R 0:20 1 mace1
67111033 ern test.slu babbott R 0:20 1 mace2
67111034 ern test.slu babbott R 0:20 1 edge

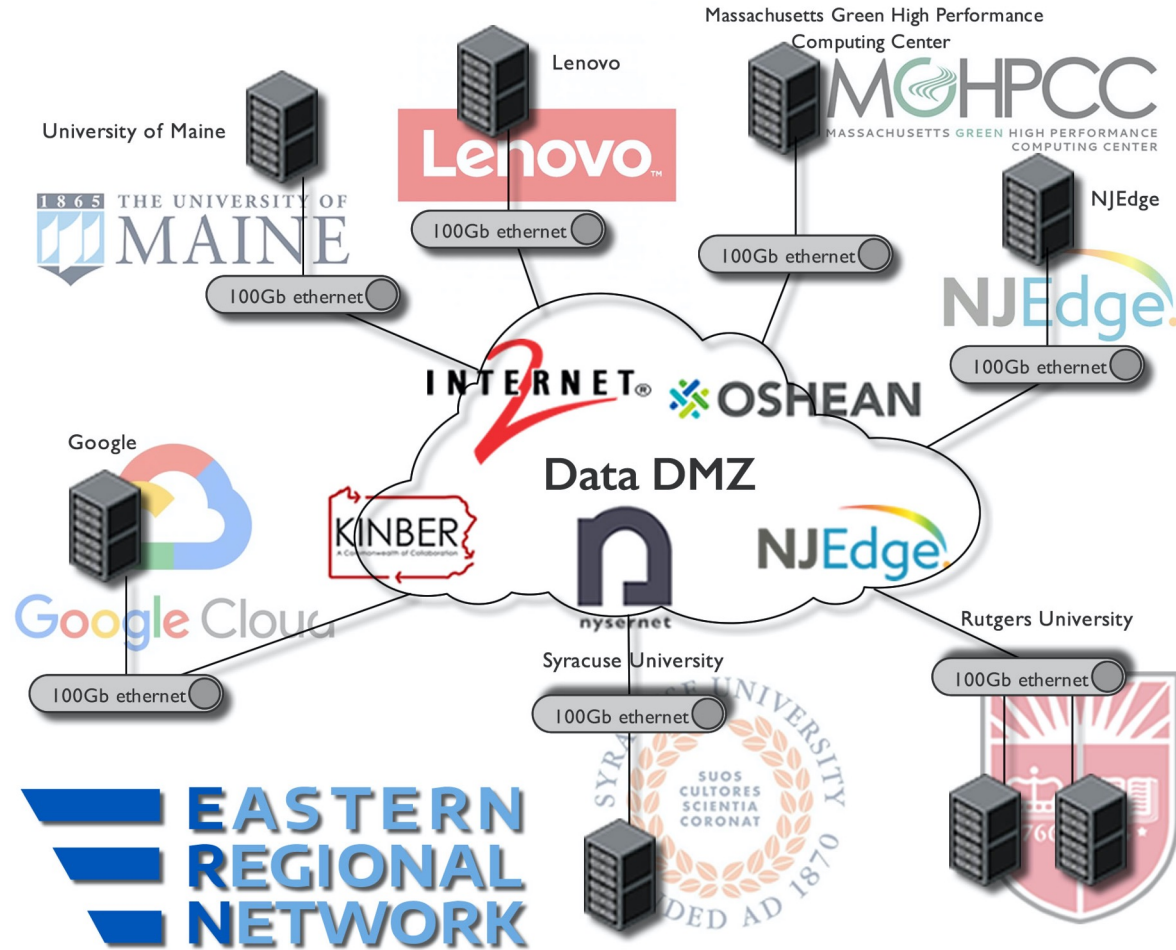
PARTITION CLUSTER AVAIL TIMELIMIT NODES STATE NODELIST
ern* mace up 3-00:00:00 2 alloc mace1-2]
ern* edge up 3-00:00:00 1 alloc edge
ern* otto up 3-00:00:00 1 alloc its-nefnode-3
ern* gcloud up 3-00:00:00 2 alloc gcloud-compute[1-2]
ern* acg up 3-00:00:00 2 alloc node[1,3]
ern* otto up 3-00:00:00 2 idle its-nefnode-[1-2]
ern* mfedhpc up 1-00:00:00 3 idle mfedhpc[1-3]
ern* acg up 3-00:00:00 1 idle node2
```

- Six sites, five states
- Ability to launch containerized (Singularity) jobs as well as traditional HPC jobs, from any site to any other site including commercial cloud, via SLURM scheduler.
- Working with Google, Cisco, Internet2, SchedMD (SLURM), and OSG on enhancements
- InCommon authentication and authorization
- Policy based federation of computing resources
- Reserving slices of the network
- Cloud bursting
- OSG federation with containerized SLURM (slurmd)
- Testing different data sharing approaches as well as filesystems

Recipe for joining:  
<https://github.com/rutgers-oarc/ern-poc>



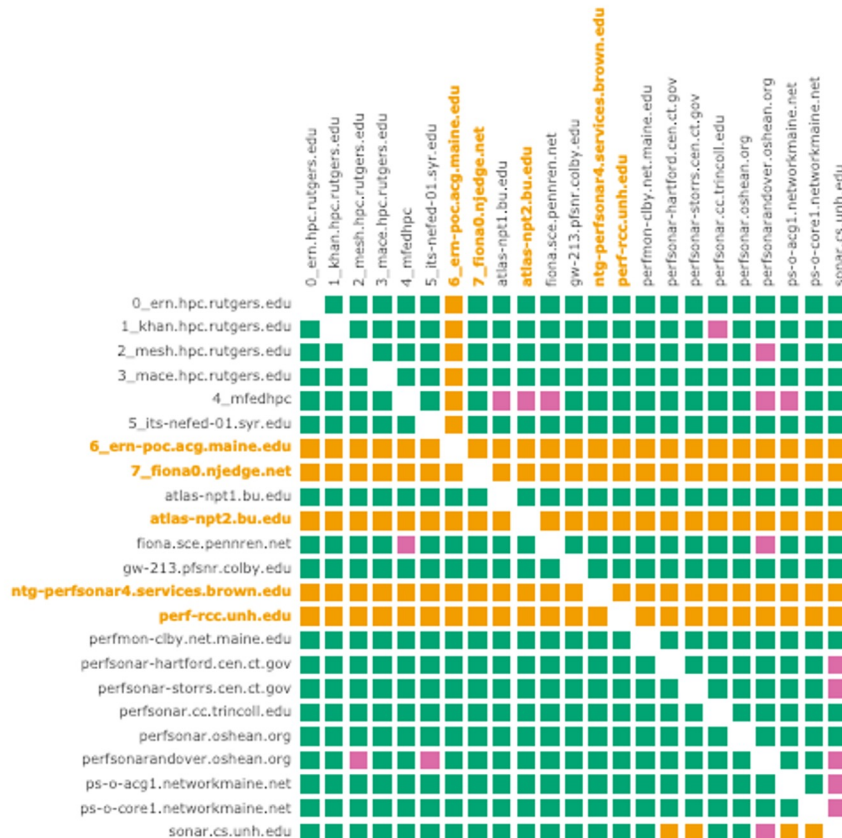
# Resource Federation Proof of Concept



# Federation with Commercial Cloud

- Partnership with Google Cloud and SchedMD (SLURM) to develop federated hybrid environment
  - Allows federation across resources (including cloud services)
  - Allows us to offer one stop shop for users
  - Allows on demand elastic computing
- Initial Proof of Concept
  - Set up federated environment between Rutgers and Google Cloud
  - Google Cloud instance is set up similar to local resource, includes SLURM
  - Connection through Internet2 between both environments
  - Able to submit jobs directly from Rutgers HPC to Google Cloud using SLURM submit script
- Elasticity/bursting capability now being tested
- Next steps: Extend to Jetstream, OSG

# ERN perfSONAR Mesh



- 23 nodes, 14 sites, six states
- New sites welcome
- Still in teething phase
- Improves visibility on connectivity/bandwidth/delay
- Starting point for connectivity improvements (where they matter)
- Prerequisite for a regional data DMZ

Site:

<http://mesh.hpc.rutgers.edu/maddash-webui/>

Recipe for joining:

<https://github.com/rutgers-oarc/ern-poc>

# ERN's Evolution - How we got to where we are now

# NSF CC\* CRIA OAC-2018927



- Received notice of funding during AHM last year (June)
- Project Team: Goodhue, Honavar, Pitt, Segee, and von Oehsen
- Funding period July 1, 2020, to June 30, 2023
- Working Groups
  - Materials Discovery (Shashank Priya, Chair)
  - Structural Biology (Stephen Burley, Chair)
  - Architecture/Federation/Computer Science (Maureen Dougherty, Michael Zink, Co-Chairs)
  - Policies (Ron Hutchins, Chair)
  - Broadening the Reach (Forough Ghahramani, John Hicks, Co-Chairs)
- Collect information to be used for future initiatives and funding opportunities


# Policy Working Group

- Melissa Cragin
- Maureen Dougherty
- Anthony Elam
- Jenni Evans
- Clark Gaylord
- John Goodhue
- Jim Griffioen
- Helen Hill
- John Huffman
- Vasant Honavar
- Ron Hutchins\*
- Barr von Oehsen
- Mike Zink

## Strategy for Culture Change

June 11th, 2019, Brian Nosek

Posted in: [Reproducibility](#), [Open Science](#), [Culture Change](#), [Behavior Change](#)





RUTGERS



PennState

THE GEORGE  
WASHINGTON  
UNIVERSITY

WASHINGTON, DC

UC San Diego




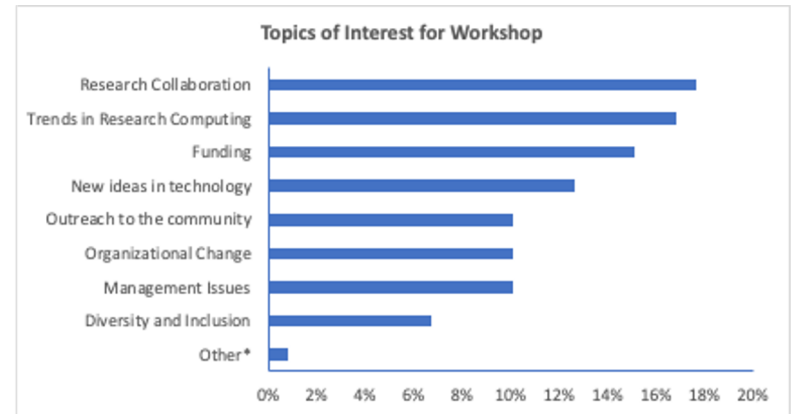
Massachusetts  
Institute of  
Technology



UMass Amherst

# Broadening the Reach Working Group

- Galen Collier
- Scott Frees
- John Goodhue
- Forough Ghahramani\*
- Ventsi Gotov
- John Hicks\*
- John Huffman
- Jim Kyriannis
- Derek Leydig
- Julie Ma
- Dave Marble
- John McNutt
- Linh Ngo
- Todd Price
- Barr von Oehsen
- Jennifer Oxenford
- Andy Sherman
- Scott Valcourt







RUTGERS



PennState

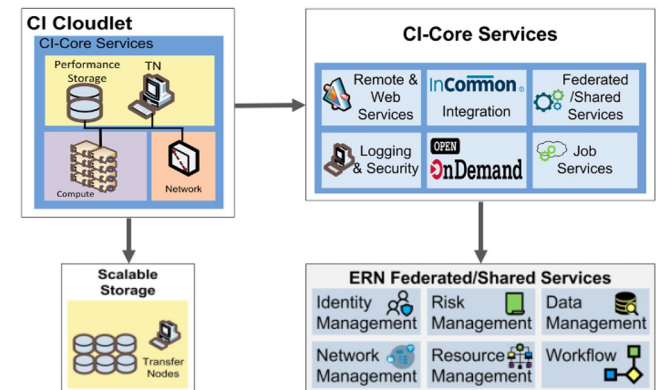


UMassAmherst



# Architecture & Federation Working Group

- Tom Barton
- Steve Cousins
- Eva Deelman
- Heidi Dempsey
- Bala Designhu
- Maureen Dougherty\*
- Adam Focht
- Forough Ghahramani
- Jim Griffioen
- Ventsi Gotov
- John Goodhue
- Vasant Honavar
- Dave Hudak
- Ron Hutchins
- Mahmut Kandemir
- Orran Krieger
- Eric Lyons
- John McNutt
- Barr von Oehsen
- Chuck Pavloski
- Bruce Segee
- Marc Sugarman
- Robert Settlege
- Scott Valcourt
- Ann West
- Boyd Wilson
- Scott Yockel
- Matt Zekaukas
- Mike Zink\*





RUTGERS



PennState



UMassAmherst



USC University of Southern California



HARVARD UNIVERSITY



THE UNIVERSITY OF MAINE

LAFAYETTE COLLEGE



Red Hat




# Structural Biology Working Group

- Stephen Burley\*
- Chris Hill
- Arek Kulczyk
- Chen Xu
- John Goodhue
- Barr von Oehsen
- Ed O'Brien
- Mike Zink
- Bala Desinghu
- Maureen Dougherty



# Materials Discovery Working Group



- Liping Yu
- Shashank Priya\*
- Bala Desinghu
- Vasant Honavar
- Alisa Kang
- Emanuela Del Gado
- Barr von Oehsen
- John Goodhue
- Maureen Dougherty



RUTGERS



PennState

UC San Diego



Massachusetts  
Institute of  
Technology

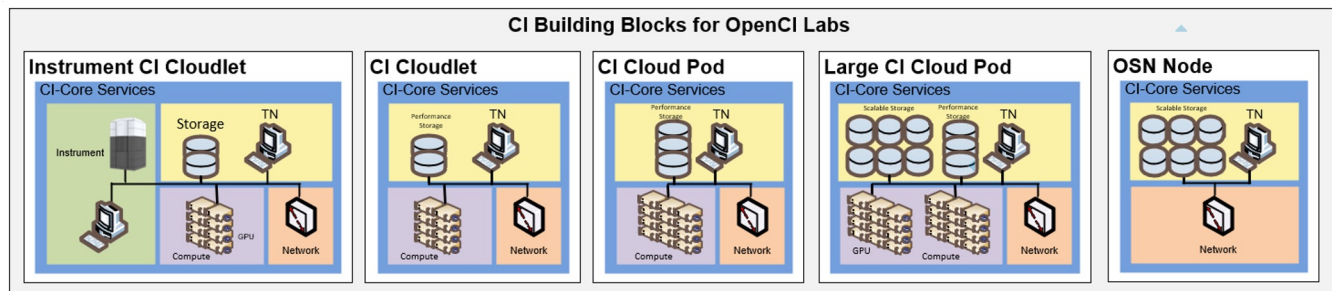
UMassAmherst

GEORGETOWN  
UNIVERSITY

# Planning Grant Workshops and Results

- Three information gathering workshops
  - ERN Architecture and Federation Virtual Workshop, December 2-4, 2020
  - ERN Broadening the Reach Virtual Workshop, December 10-11, 2020
  - Structural Biology The Voice of the Customer Online Workshop, February 11, 2021
- Year One Significant Results
  - NSF Mid-Scale RI-1 solicitation, pre-proposal
  - ERN Federated OpenCI Labs Design
  - ERN Policy Planning
    - Bylaws and charter
  - Four white papers (two publications)
  - Data CI Pilot

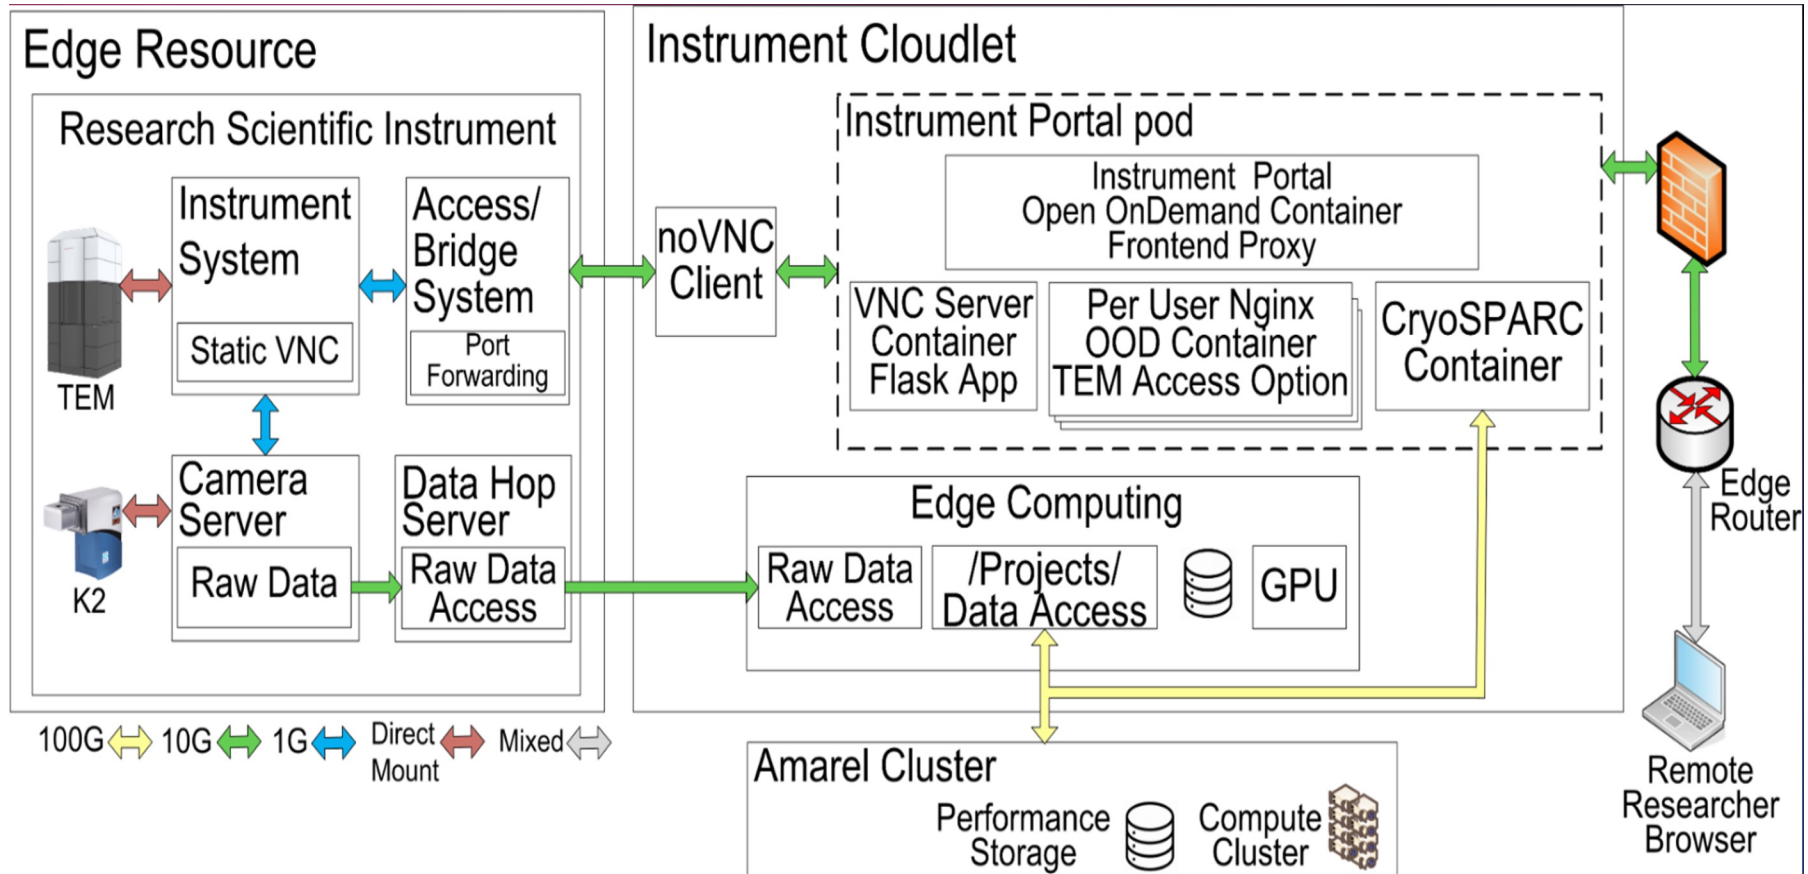
# OpenCI Labs



- Distributed Federated Cloud Architecture
  - Building Block design
- Gateway to instruments on the edge and Core Facilities
  - CryoEM, X-ray diffraction, Scanning electron microscope, Focused ion beam, Transmission electron microscopy, and gene sequencers
- Data access & Distributed Workflows
  - Materials-centric structured databases
  - Data-sharing policies and standardizations
  - Machine learning/Deep Learning for materials discovery across all scales
  - Distributed research computing (virtualization/containers, bare metal)
- Minimal configuration:
  - Open OnDemand, OSN Gateway (or node), and federation



# Step 2 Design





# Publications

- Maureen Dougherty, Michael Zink, Barr von Oehsen, Kenneth Dalenberg, Bala Desinghu, Jason Kaelber, Jeremy Schafer, John Goodhue, Wolf Hey, Morgan Ludwig, and Boyd Wilson. 2022. ``*The ERN Cryo-EM Federated Instrument Pilot Project: Phase 1,*'' 2022 4th Annual Workshop on Extreme-scale Experiment-in-the-Loop Computing (XLOOP), Dallas, TX, USA, 2022, pp. 20-25, doi: 10.1109/XLOOP56614.2022.00009. **BEST PAPER!**
- Maureen Dougherty, Michael Zink, Barr von Oehsen, Kenneth Dalenberg, Bala Desinghu, Jason Kaelber, Jeremy Schafer, John Goodhue, Wolf Hey, Morgan Ludwig, Boyd Wilson, and Cole McKnight. 2022. ``*The ERN CryoEM Federated Instrument Pilot Project.*'' In Practice and Experience in Advanced Research Computing (PEARC '22). Association for Computing Machinery, New York, NY, USA, Article 52, 1–4. <https://doi.org/10.1145/3491418.3535141> **BEST SHORT PAPER!**
- Forough Ghahramani, John Hicks, and Barr von Oehsen. 2022. ``*Broadening the Reach for Access to Advanced Computing: Leveraging the Cloud for Research.*'' In Practice and Experience in Advanced Research Computing (PEARC '22). Association for Computing Machinery, New York, NY, USA, Article 65, 1–5. <https://doi.org/10.1145/3491418.3535143>

# Publications

- Melissa Cragin, Ron Hutchins, Maureen Dougherty, James Barr von Oehsen, Michael Zink, D. Balamurugan, and John Goodhue. 2022. ``*Federating CI Policy in Support of Multi-institutional Research: Lessons from the Ecosystem for Research Networking.*'' In Practice and Experience in Advanced Research Computing (PEARC '22). Association for Computing Machinery, New York, NY, USA, Article 49, 1–4. <https://doi.org/10.1145/3491418.3535167> **BEST SHORT PAPER RUNNER UP!**
- Maureen Dougherty, Michael Zink, and James Barr von Oehsen. 2021. ``*Identifying Research Collaboration Challenges for the Development of a Federated Infrastructure Response.*'' In Practice and Experience in Advanced Research Computing (PEARC '21). Association for Computing Machinery, New York, NY, USA, Article 40, 1–4. <https://doi.org/10.1145/3437359.3465594>
- Forough Ghahramani, John Hicks, and Barr von Oehsen. 2021. ``*Broadening the Reach for Access to Advanced Cyberinfrastructure: Accelerating Research and Education.*'' In Practice and Experience in Advanced Research Computing (PEARC '21). Association for Computing Machinery, New York, NY, USA, Article 49, 1–3. <https://doi.org/10.1145/3437359.3467026>

# Participation

## Academic Institutions:

Alabama State University, Arcadia University, Boston University, Brown University, Case Western Reserve, Clemson University, Columbia University, Cornell University, Delaware State University, Franklin and Marshall College, George Washington University, Georgetown University, Harvard University, Kentucky State University, Lafayette College, MIT, Montclair State University, New Jersey Institute of Technology, Northeastern University, Pace University, Penn State University, Princeton University, Ramapo University, Rutgers University, Southern Connecticut State University, Stevens Institute of Technology, Syracuse University, The College of New Jersey, Trinity College, Tufts University, USMA, University of Arkansas, University of Buffalo, UC Santa Barbara, University of Chicago, University of Delaware, University of Illinois, University of Kentucky, University of Maine-Orono, University of Massachusetts-Amherst, University of Michigan, University of New Hampshire, University of New Haven, University of Utah, University of Virginia, West Chester University of Pennsylvania, and Yale University.

## Organizations/Industry/RENs:

Massachusetts Green High Performance Computing Center (MGHPCC), Northeast Big Data Innovation Hub, Ohio Supercomputer Center, SDSC, Internet2, CAAREN, CEN, KINBER, NetworkMaine, NJEdge, NYSERNET, Great Plains Network, Omnibond, OSHEAN, CILogon, CaRCC, Redhat, SHI, IBM

# What's Next?

# Special Thanks

- Maureen Dougherty
- Joylynn Almeida
- All the ERN working groups
- The ERN Steering Committee
- The National Science Foundation
- University of Delaware





**Thank You!**

Interested in learning more or participating,  
please contact [info@ernrp.org](mailto:info@ernrp.org)

Website: <https://ernrp.org>