FABRIC Introduction and ERN Collaboration

Paul Ruth, RENCI

ERN All-hands Meeting, March 23, 2023

Pegasus - Blast from the Past

- From my SC11 Demo
- Solar Fuels Workflow
- ExoGENI-to-NERSC (Hopper)
- Very early example of our efforts toward FABRIC-like federation capabilities

Thanks Mats!

I love that you still use this image.

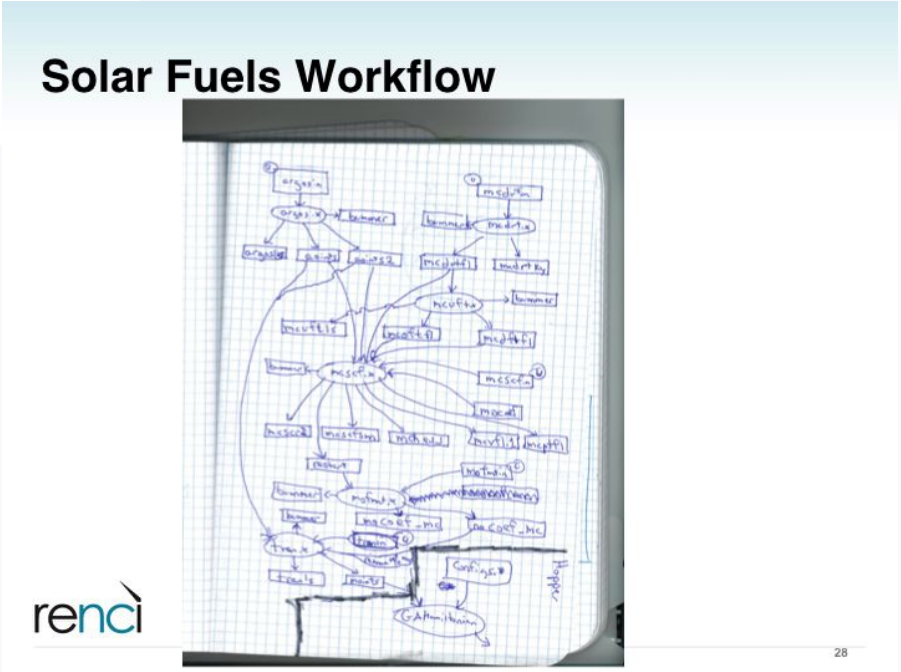

FABRIC

FAB
FABRIC Across Borders

What is FABRIC?

- NSF-funded MSRI-1 (Mid-scale Research Infrastructure)
- Led by RENCI/UNC Chapel Hill
- 5 core team members: University of Kentucky, University of Illinois at Urbana-Champaign, Clemson University and ESnet
- Many other facility partners, including Rutgers
- $20M budget for construction, separately-funded operation phase expected
- Started in 2019, expected completion 09/2023
- FABRIC Across Borders (FAB) - IRNC International Extension

FABRIC

FAB
FABRIC Across Borders

What is FABRIC?

FABRIC is a scientific instrument for studying
distributed applications, Internet protocols, and services:

- A nation-wide programmable network testbed with significant compute and storage at each node, allowing users to run computationally intensive programs and applications and protocols to maintain a lot of information in the network.
- Provides GPUs, FPGAs, and network processors (NICs) inside the network.
- Supports quality of service (QoS) using dedicated optical 100G links or dedicated capacity
- Interconnects national facilities: HPC centers, cloud & wireless testbeds, commercial clouds, the Internet, and edge nodes at universities and labs.
- Allows you to design and test applications, protocols and services that run at any node in the network, not just the edge or cloud.

FABRIC

FAB
FABRIC Across Borders

Why FABRIC?

- The mantra of the last 20 years – ‘Internet is showing its age.’
  - Applications designed around discrete points in the solution space
  - Inability to program the core of the network
- What changed?
  - Cheap compute/storage that can be put directly in the network
  - Multiple established methods of programmability (OpenFlow, P4, eBPF, DPDK, BGP flowspec)
  - Advances in Machine Learning/AI
  - Emergence of 5G, IoT, various flavors of cloud technologies
- Opportunity for the community to push the boundaries of distributed, stateful, ‘everywhere’ programmable infrastructure
  - More control or dataplane state, or some combination? Multiple architectures (co)exist in this space.
  - Network as a big-data instrument? Autonomous network control?
  - New protocols and applications that program the network?
  - Security as an integral component

FABRIC

FAB
FABRIC Across Borders

FABRIC + FAB: Final Topology (Fall 2023)

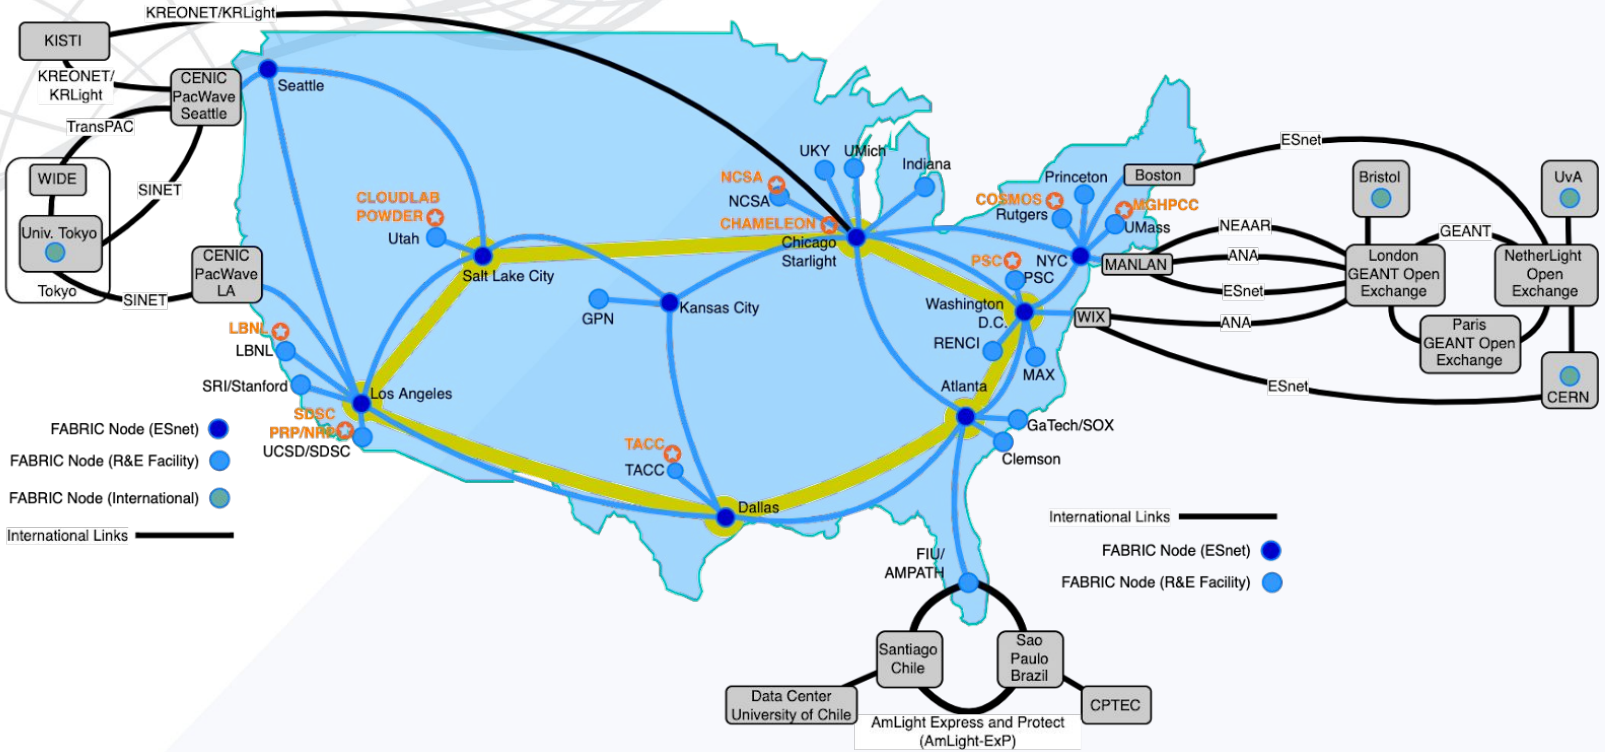

FABRIC

FAB
FABRIC Across Borders

FABRIC Topology - with FAB Sites

FABRIC: Current Production Deployment

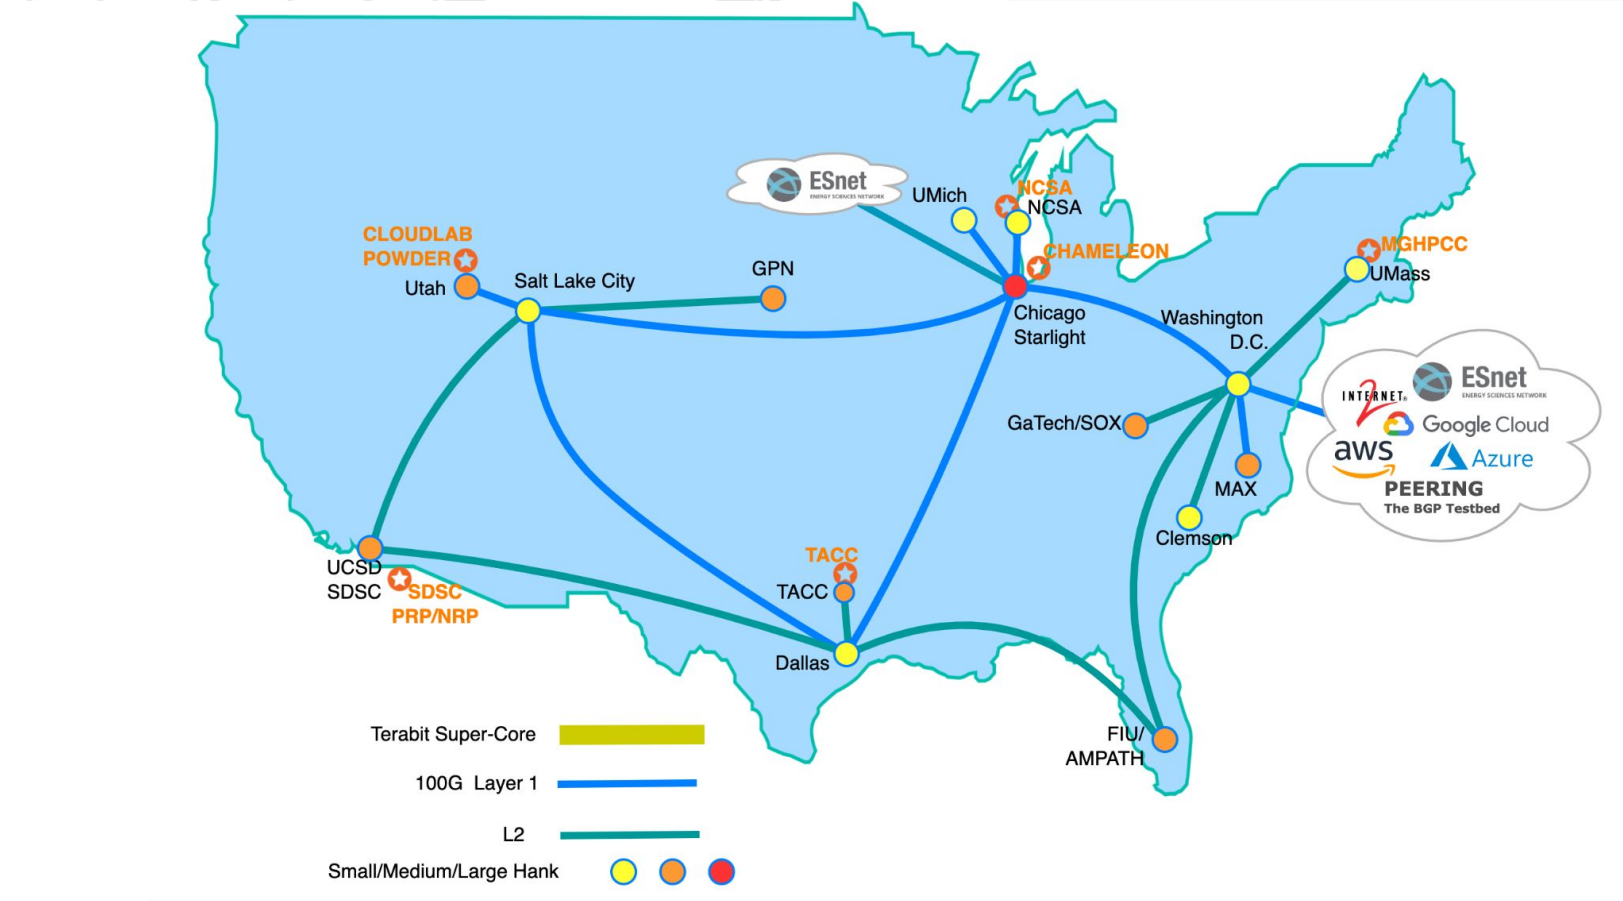

Key Features

- Network as part of computing continuum
  - ‘Everywhere-programmable’ using different abstractions (P4, OpenFlow, others)
  - Diverse compute, storage capabilities in places where routers typically reside today
  - Dedicated 100G optical links between many sites
  - Support new paradigms in network aware applications and protocols
  - Ability to peer with Internet at IPv4 and IPv6
- Network as a scientific instrument
  - Pervasive measurement collection capabilities in- and outside the slice available to researchers
  - GPS-disciplined PTP clock sources at every site
- Serve a broad range of scientific domains and applications
  - Concerned with data transport for big-data science, cyber-security, terrestrial and 5G hybrid network architectures, federated ML/AI, Internet measurements and many more

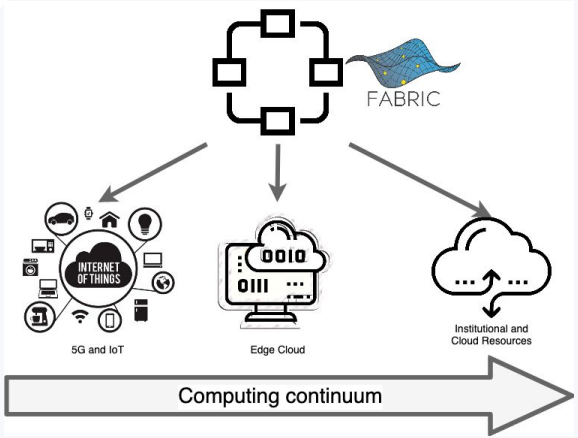

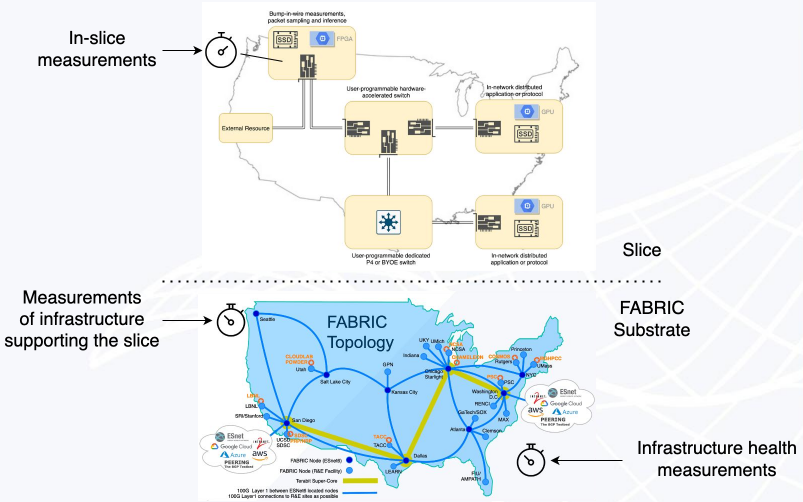

FABRIC

FABRIC “Hank”

Hank: a measured unit of coiled or wrapped yarn or twine

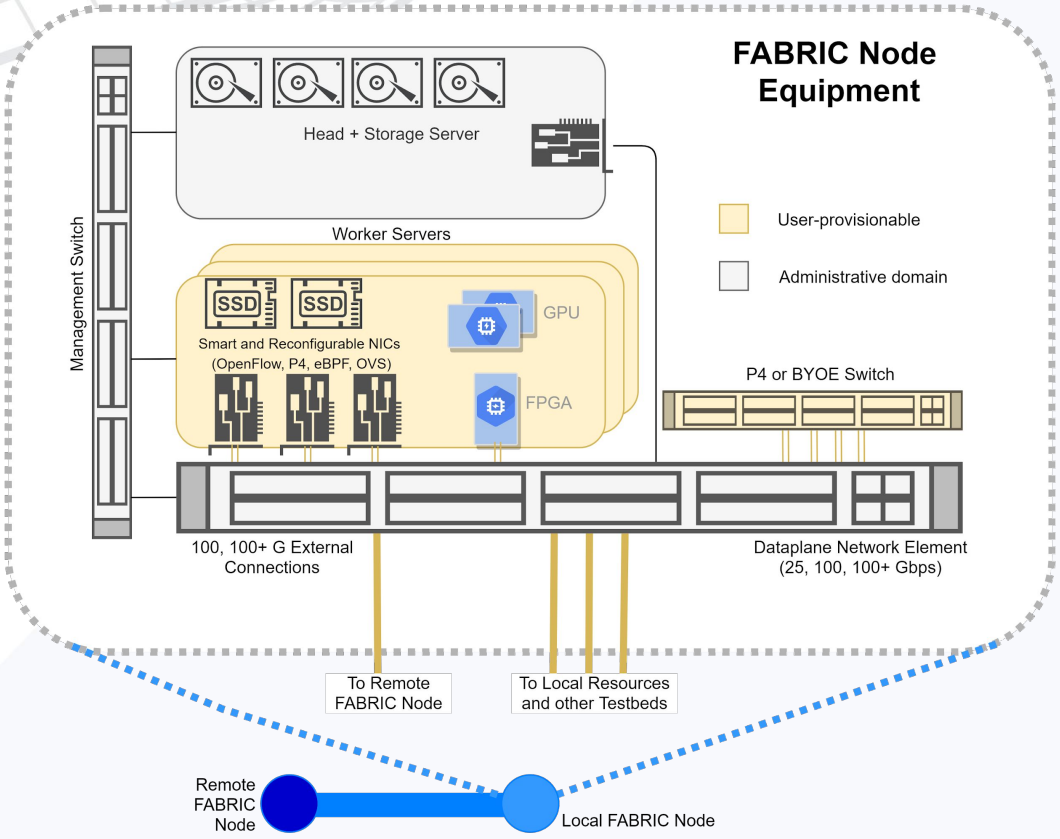

FABRIC

FAB
FABRIC Across Borders

FABRIC Nodes - Network, Compute, Storage

- Interpose compute and storage into the path of fast packet flows
- Rack of high-performance servers (Dell 7525) with:
  - 2x32-core AMD Rome and Milan with 512G RAM
  - GPUs (NVIDIA RTX 6000, T4, A30), FPGA network/compute accelerators
  - Storage - experimenter provisionable 1TB NVMe drives in servers and a pool of ~250TB rotating storage at each site.
  - Network ports connect to a 100G+ switch, programmable through control software
  - Tofino-based P4 switches (4 sites)
- Reconfigurable Network Interface Cards
  - FPGAs (U280 XILINX with P4 support)
  - Mellanox ConnectX-5 and ConnectX-6 with hardware off-load
  - Multiple interface speeds (25G, 100G, 200G+(future)
- Kernel Bypass/Hardware Offload
  - VMs sized to support full-rate DPDK for access to Programmable NICs, FPGA, and GPU resources via PCI pass-through

FABRIC

FABRIC Network Services

- L2 services (dedicated ethernet)
  - L2Bridge - local to individual site
  - L2STS - two sites, any number of interfaces
  - L2PTP - two sites, two interfaces, QoS guarantees
- L3 services
  - FABNetv4 - FABRIC-routed IPv4 service (RFC1918 addresses)
  - FABNetv6 - FABRIC-routed IPv6 services (FABRIC's IPv6 addreses)
  - Optional external peering with the Internet
  - L3VPN service
- Specialized services
  - Port mirroring - mirrors physical ports to experiment
  - Facility Ports - L2 peering with external facilities
    - Chameleon, Cloudlab, **Rutgers Cryo-EM (new!)**, MOC (coming), more…
    - AWS, Google Cloud, Azure (via Internet2 CloudConnect)

FABRIC

Other Features/Services

- P4 workflows
  - On Xilinx U280 FPGAs, (colab with OCT, Northeastern, and ESnet)
  - Tofino switches
- Storage
  - Virtual disk (VM)
  - NVMe drives (PCI)
  - Persistent volumes (Network)
- GPUs
  - NVIDIA RTX 6000
  - NVIDIA Tesla T4
- Monitoring in-slice measurement framework
  - Grafana, Prometheus, Elk, Kibana
- Experiment Management
  - JupyterHub in Google Kubernetes Engine (GKE), via CloudBank
  - Profile sharing and publishing with Trovi (collab with Chameleon)

FABRIC

FABRIC Risk and Security

- Participating in NSF's Trusted CI Framework Cohort (2022)
  - Framework is the *minimum* standard for cybersecurity programs
- Two new part time FABRIC security engineers

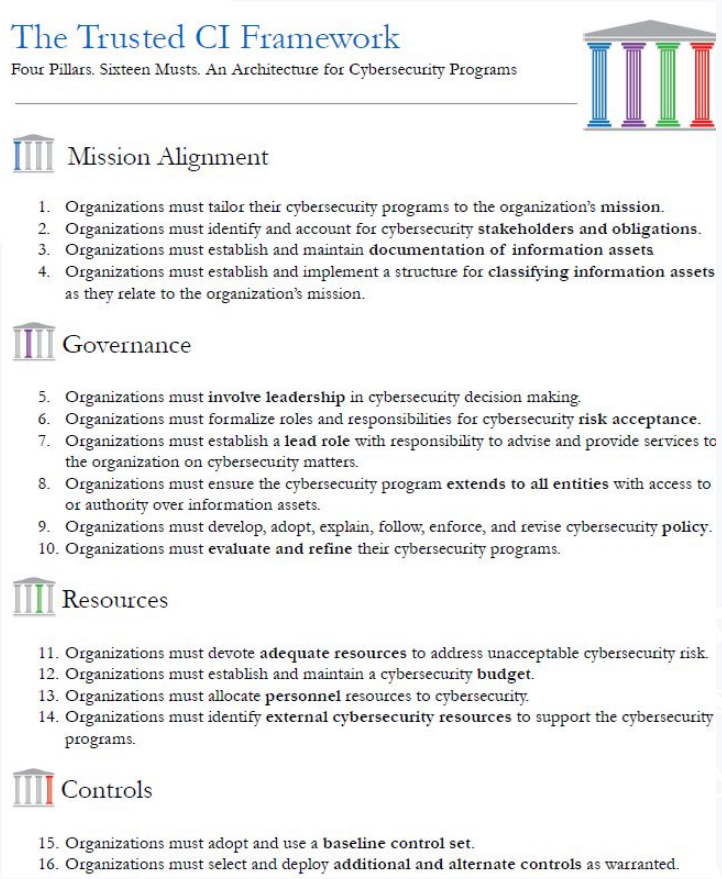

The Trusted CI Framework

Four Pillars. Sixteen Musts. An Architecture for Cybersecurity Programs

Mission Alignment

1. Organizations must tailor their cybersecurity programs to the organization's **mission**.
2. Organizations must identify and account for cybersecurity **stakeholders and obligations**.
3. Organizations must establish and maintain **documentation of information assets**.
4. Organizations must establish and implement a structure for **classifying information assets** as they relate to the organization's mission.

Governance

5. Organizations must **involve leadership** in cybersecurity decision making.
6. Organizations must formalize roles and responsibilities for cybersecurity **risk acceptance**.
7. Organizations must establish a **lead role** with responsibility to advise and provide services to the organization on cybersecurity matters.
8. Organizations must ensure the cybersecurity program **extends to all entities** with access to or authority over information assets.
9. Organizations must develop, adopt, explain, follow, enforce, and revise cybersecurity **policy**.
10. Organizations must **evaluate and refine** their cybersecurity programs.

Resources

11. Organizations must devote **adequate resources** to address unacceptable cybersecurity risk.
12. Organizations must establish and maintain a cybersecurity **budget**.
13. Organizations must allocate **personnel** resources to cybersecurity.
14. Organizations must identify **external cybersecurity resources** to support the cybersecurity programs.

Controls

15. Organizations must adopt and use a **baseline control set**.
16. Organizations must select and deploy **additional and alternate controls** as warranted.

CYBERSECURITY PROGRAM
BUILT WITH THE TRUSTED CI FRAMEWORK

FABRIC

FABRIC Experiments

- FABRIC as a programmable internet core
- Connecting external facilities
- Smart routing, switching, caching, in-network processing, etc. in the core
- Packaged to be repeatable

FABRIC

FABRIC Experiments

- FABRIC as a programmable internet core
- Connecting external facilities
- Smart routing, switching, caching, in-network processing, etc. in the core
- Packaged to be repeatable

FABRIC

Early Experiments

- Users: >250 active (340 registered)
- Projects: >40 (research and education)
- Project topics:
  - P4/SDN
  - Honeypots
  - Named Data Networking (NDN)
  - ServiceX
  - BGP Peering
  - Datacenter protocols
  - Scalable Genome Analysis
  - Fast data transfers
  - Internet Privacy

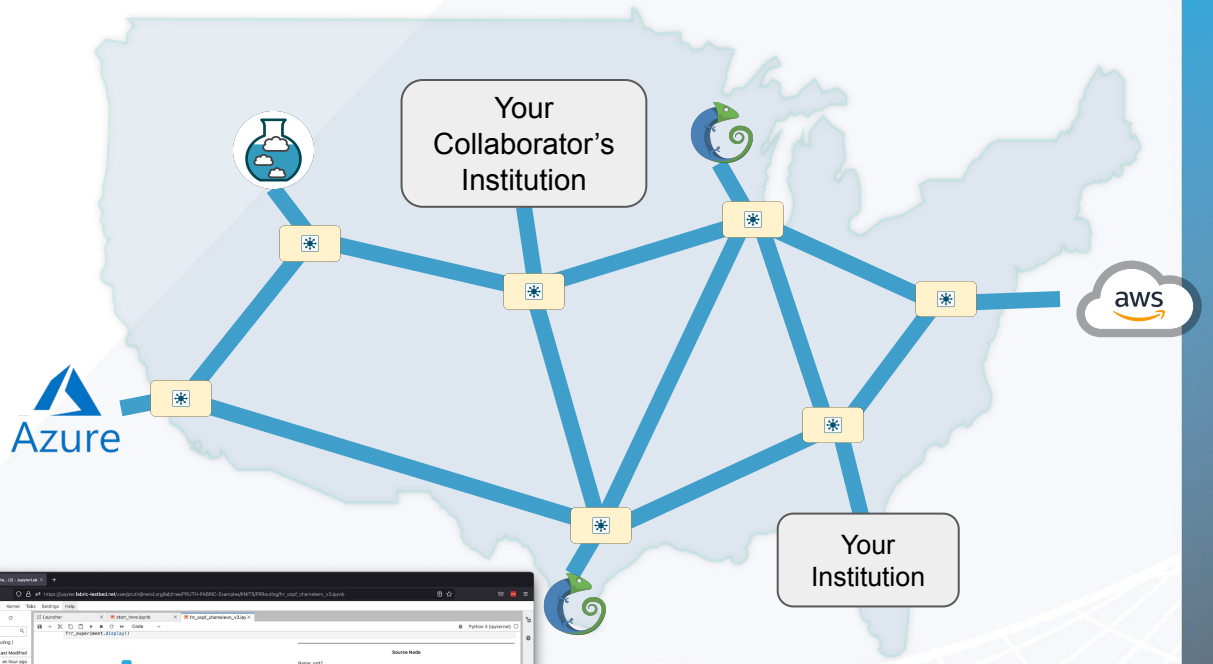

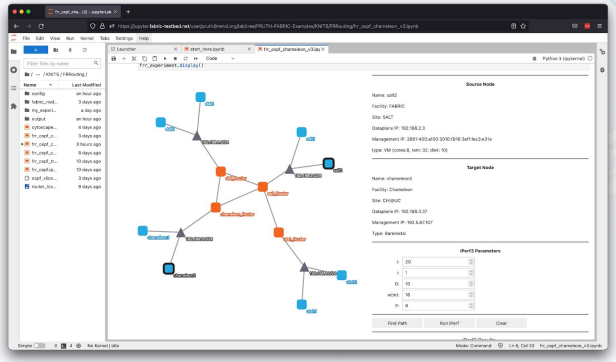

FABRIC

FAB Science Use Cases & Partners

- Astronomy (Vera Rubin Observatory/LSST, Chile)
- Cosmology (CMB-S4)
- High Energy Physics (CERN ATLAS;UChicago) - Rob Gardner
- Weather/Climate (UMiami & CPTEC, Brazil) - Ben Kirtman, Atmospheric Science & Paolo Nobre
- Urban Sensing/IoT/AI at Edge (UBristol) - Dimitra Simeonidou
- 5G across borders, P4/SDN - (UTokyo) Aki Nakao; KISTI (Korea Institute of Science and Technology Information)

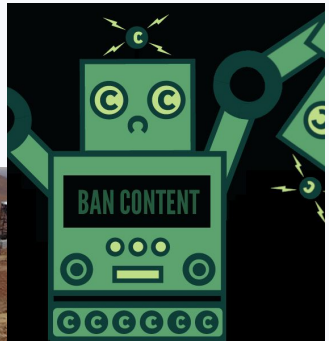

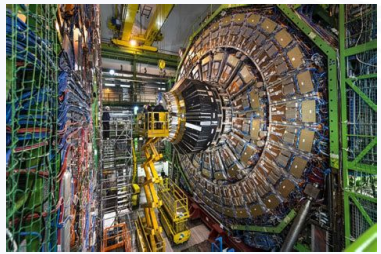

CMB-S4 (Cosmology)

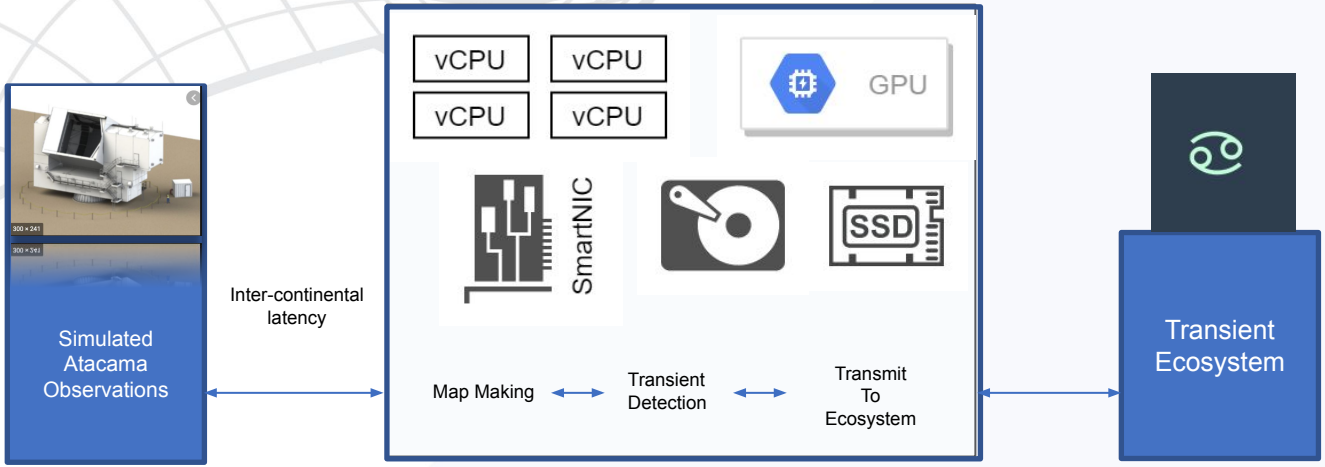

- Don Petravick, Gregory Daues (UIUC/NCSA)
- Designed/deployed CMB-S4 experiment(s) on FABRIC
- Simulated observatory source at FIU (projected actual path)
- In-network data processing
- Implemented a shell on top of FABlib to control their experiment

FABRIC

FAB
FABRIC Across Borders

User Interface

- FABRIC Portal, Jupyter Hub, User Knowledge base all operational
- Integrated with CI Logon for federated identity
- Accessible and used by early experimenters

FABRIC Knowledge Base with User Forums

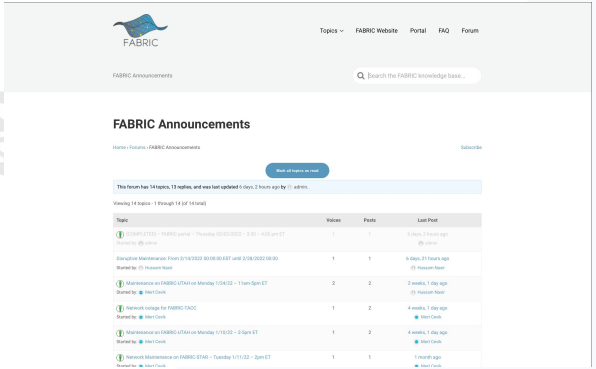

FABRIC Portal Home Page

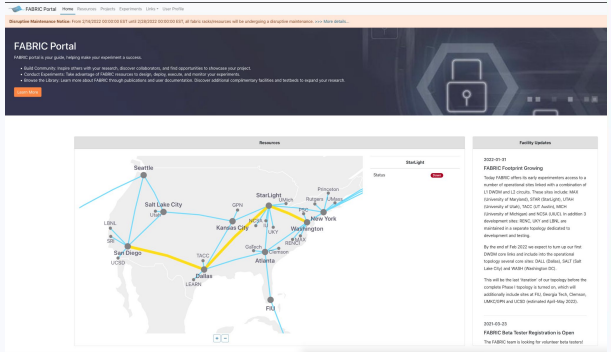

FABRIC Portal Experiments page

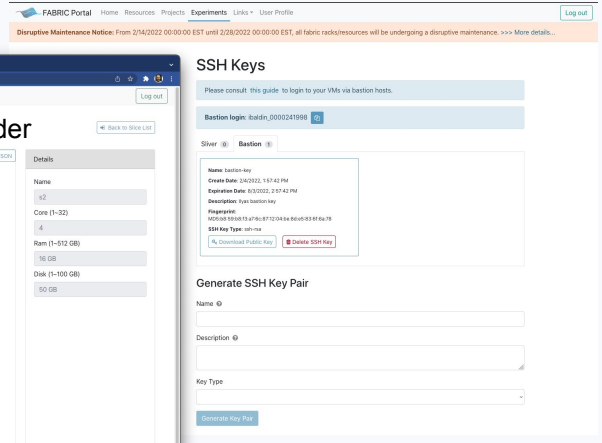

FABRIC Jupyter Hub

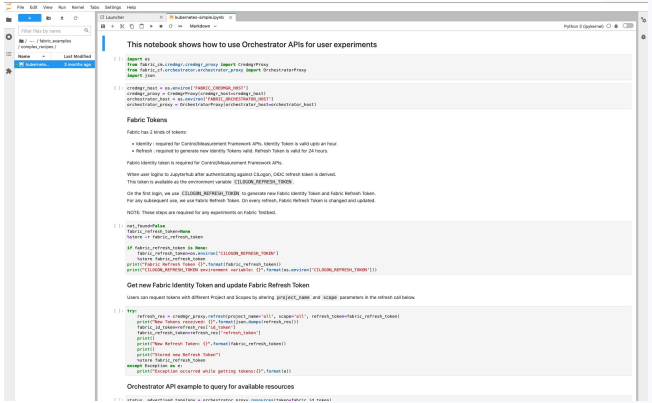

FABRIC Portal Slice Viewer/Builder

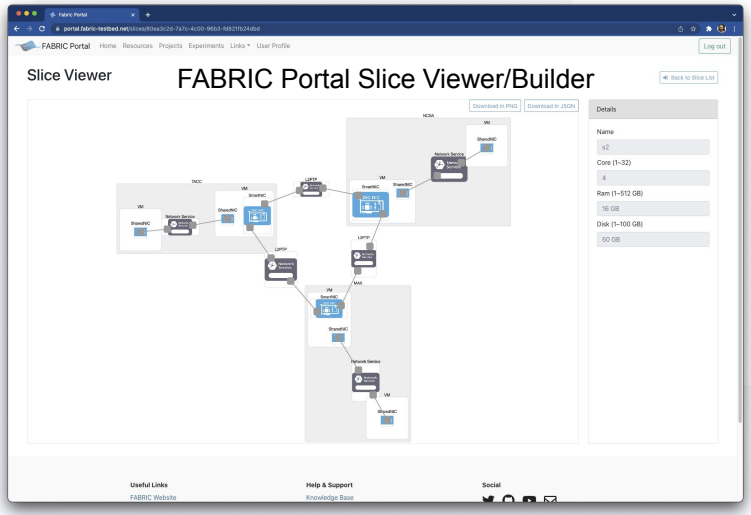

Experimenter Workflow

- Portal: https://portal.fabric-testbed.net/
  - Manage projects, keys, tokens
- Learn: https://learn.fabric-testbed.net/
  - Read docs
- Discuss: https://learn.fabric-testbed.net/forums/
  - Post topics/replies with the FABRIC community
- Jupyter: https://jupyter.fabric-testbed.net
  - Design, develop, deploy, and run experiments
- Publish (coming soon!)
  - Share packaged experiments in the port

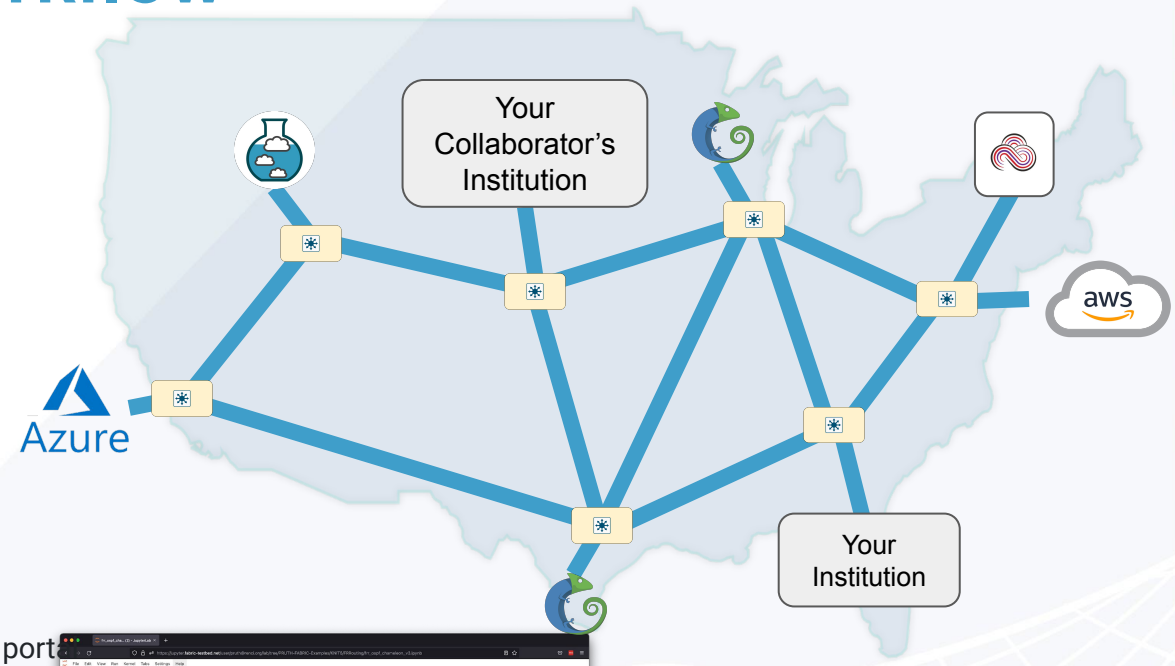

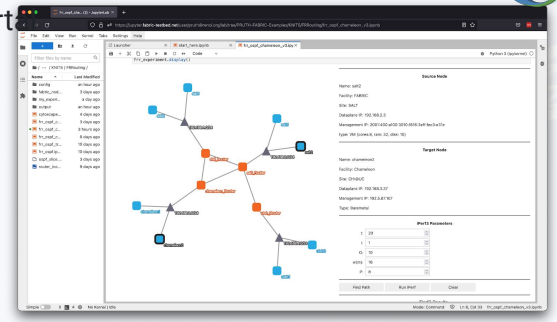

FABRIC

Federated Experiments

- FABRIC as a programmable internet core
- Connecting external facilities
- Smart routing, switching, caching, etc. in the core
- Packaged to be repeatable

FABRIC

Federated Experiments

- FABRIC as a programmable internet core
- Connecting external facilities
- Smart routing, switching, caching, etc. in the core
- Packaged to be repeatable

FABRIC

Federated Experiments

- Packaged for Jupyter
  - Notebooks plus tools
- FRRouters
  - Dockerized
  - 3 sites
  - Gateways for local network
- Protocol: OSPF
  - Many other protocols available
- Local networks
  - FABRIC
  - Chameleon

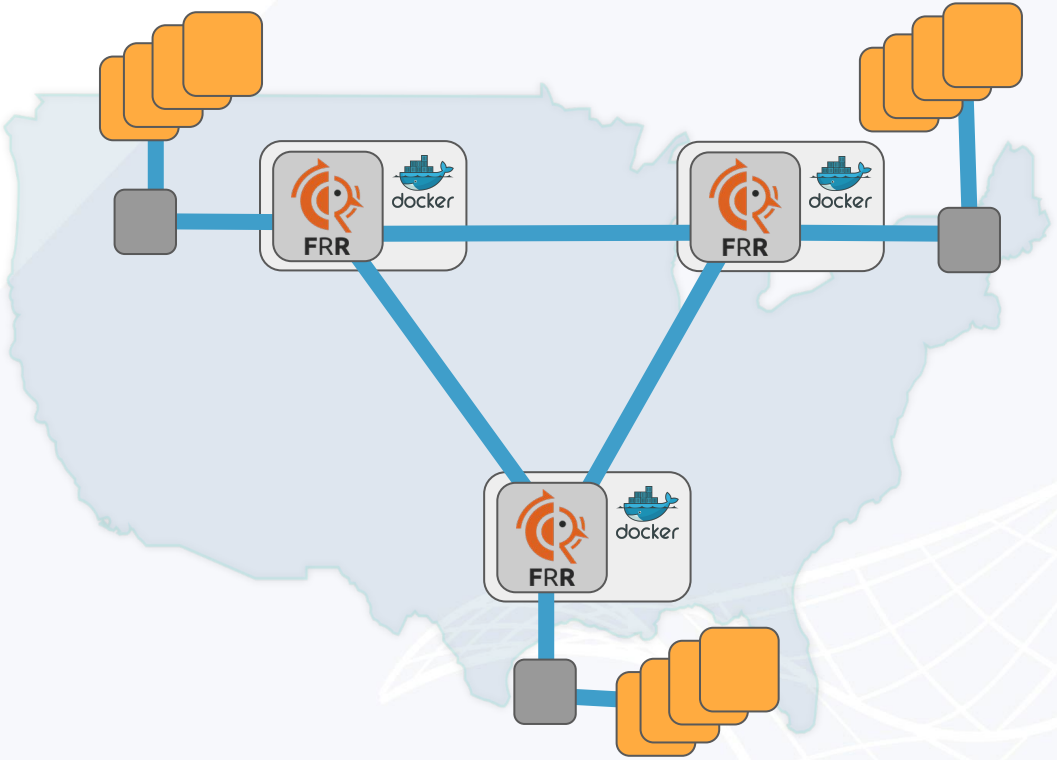

FABRIC

Federated Experiments

- Packaged for Jupyter
  - Notebooks plus tools
- FRRouters
  - Dockerized
  - 3 sites
  - Gateways for local network
- Protocol: OSPF
  - Many other protocols available
- Local networks
  - FABRIC
  - Chameleon
  - MOC
  - Your Institution

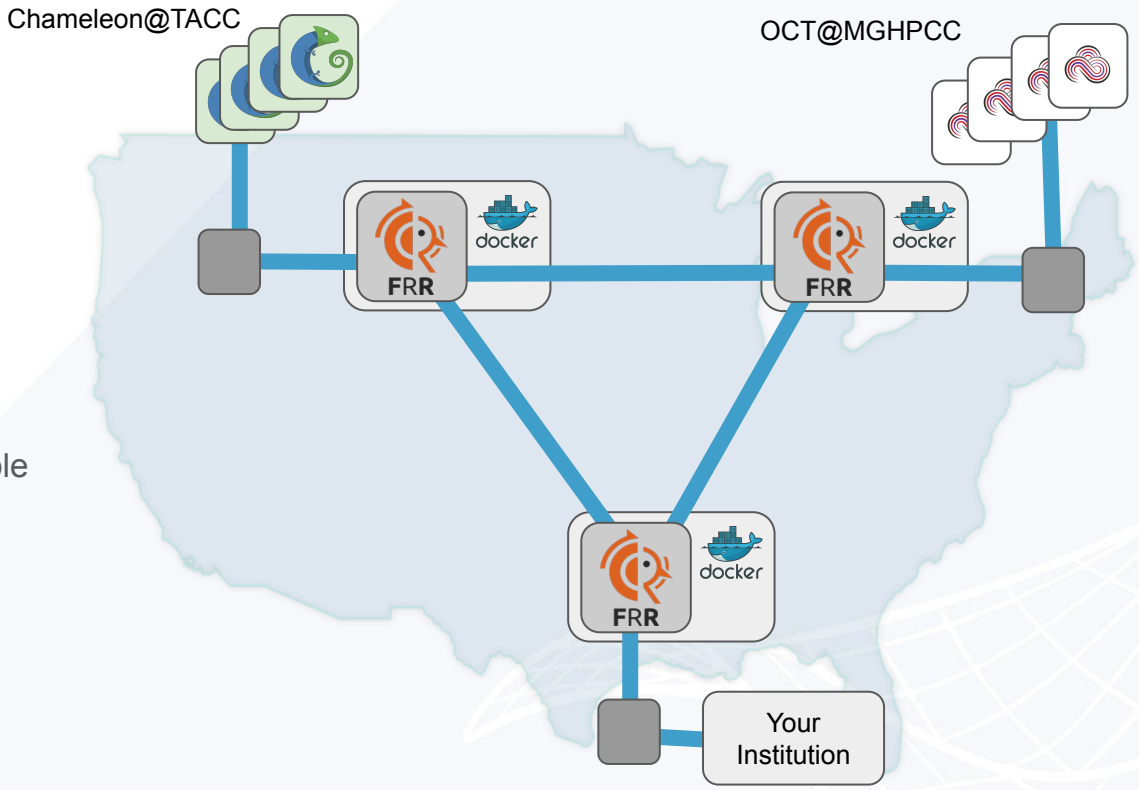

FABRIC

Federated Experiments: Cryo-EM

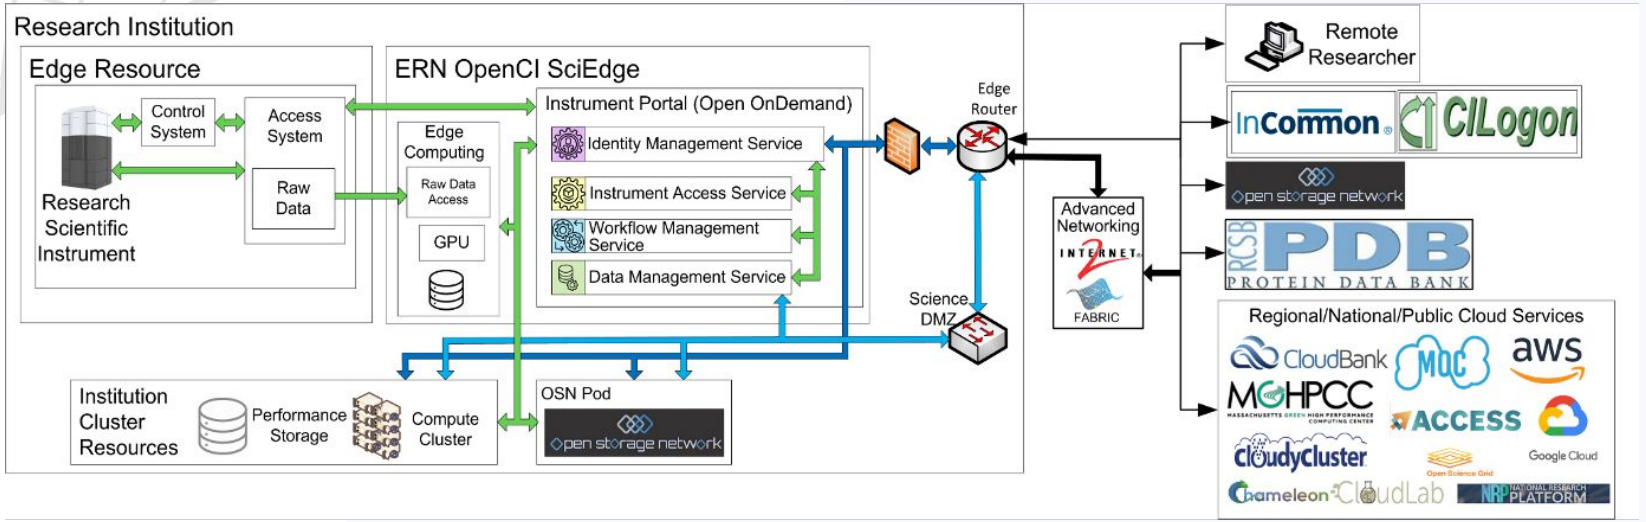

FABRIC

FAB
FABRIC Across Borders

Federated Experiments: Cryo-EM

- Packaged for Jupyter
  - Notebooks plus tools
- FRRouters
  - Dockerized
  - Three sites
  - Gateways for local network
- Protocol: OSPF
  - Many other protocols available
- Local networks
  - Cryo-EM @ Rutgers
  - MGHPCC
  - Chameleon

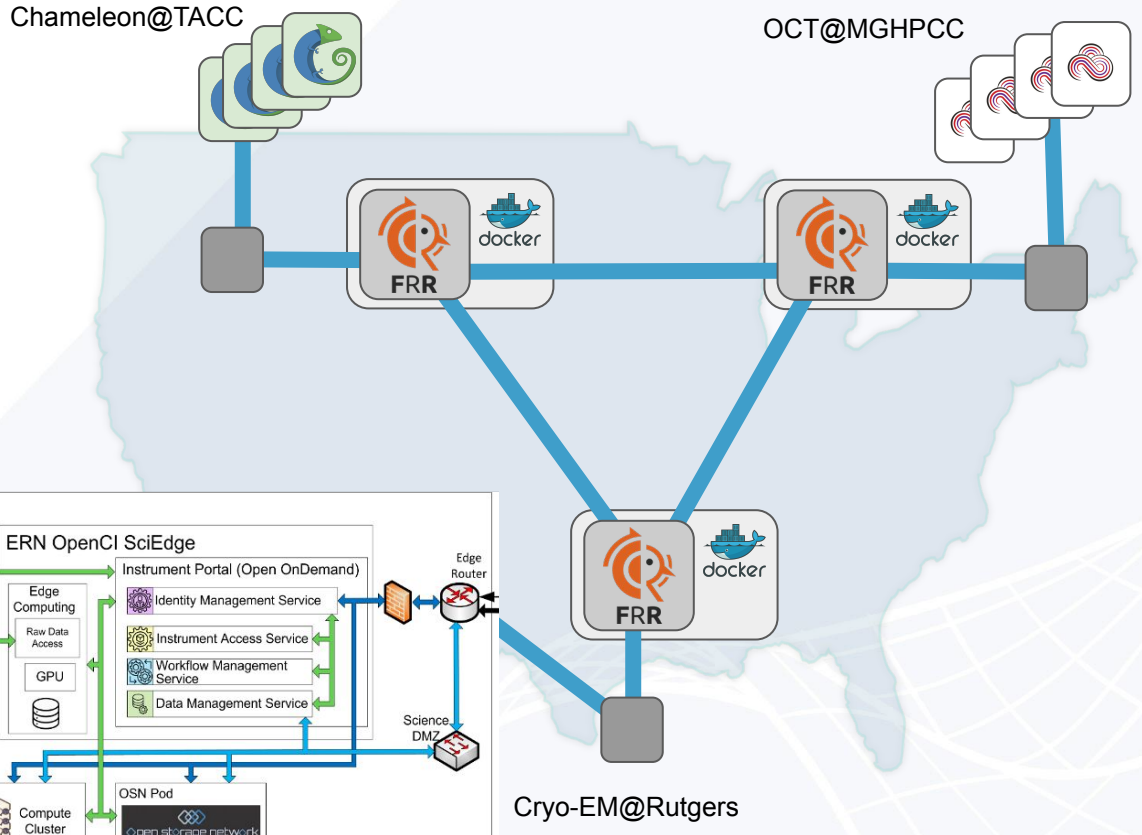

FABRIC

Federated Experiments: Cryo-EM

- Packaged for Jupyter
  - Notebooks plus tools
- FRRouters
  - Dockerized
  - Three sites
  - Gateways for local network
- Protocol: OSPF
  - Many other protocols available
- Local networks
  - Cryo-EM @ Rutgers
  - FABRIC
  - Chameleon

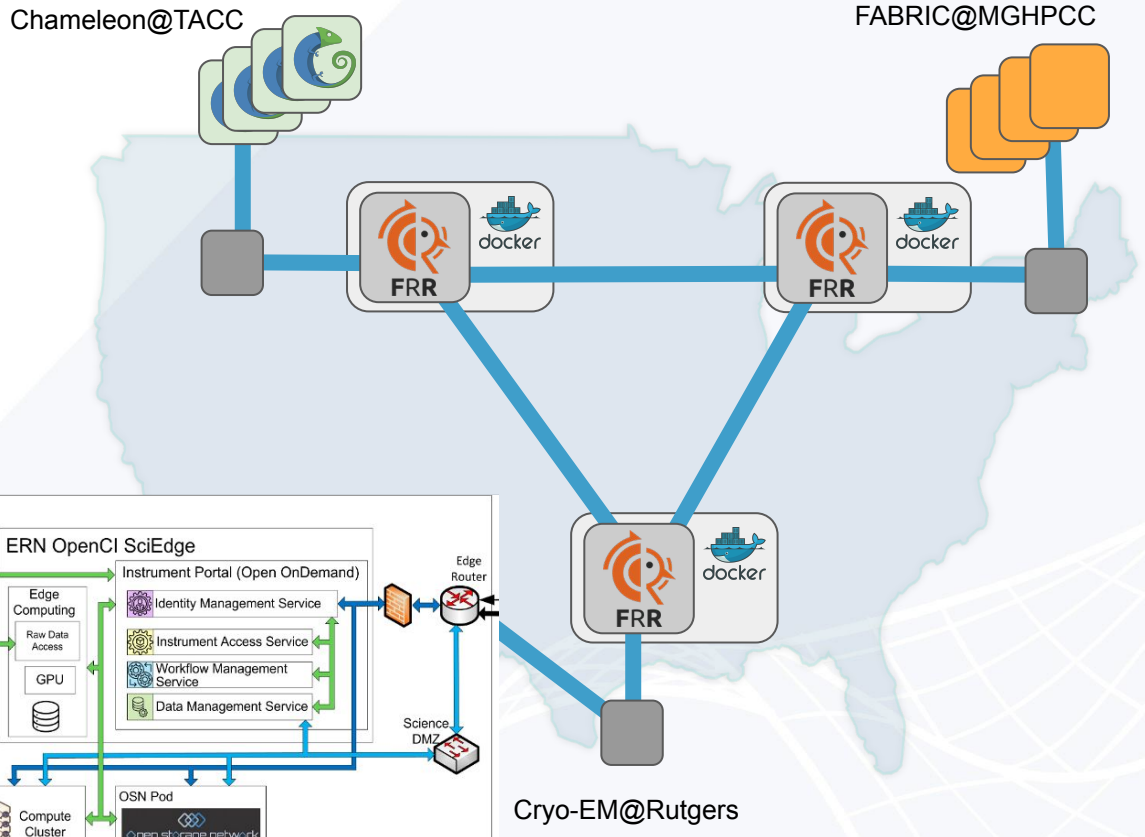

FABRIC

Demo

FABRIC

FAB
FABRIC Across Borders

Thank You!

Questions?

Visit https://whatisfabric.net

Ask info@fabric-testbed.net

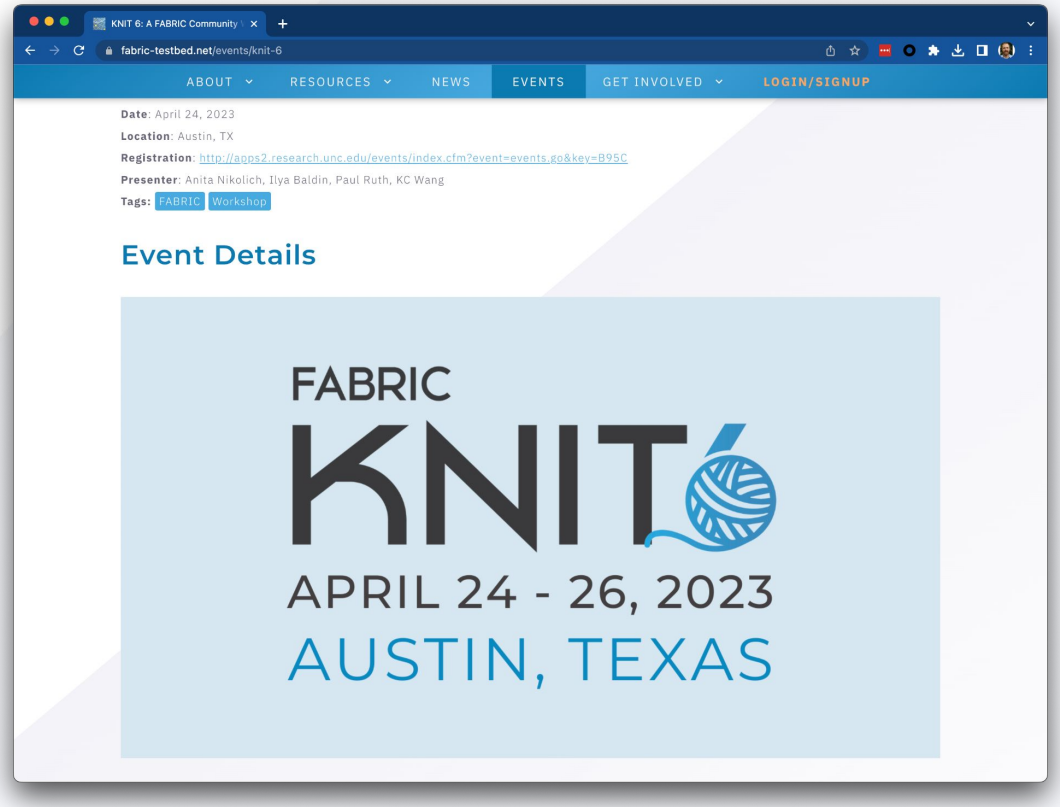

This work is funded by NSF grants
CNS-1935966, CNS-2029261,
CNS-2029235, CNS-2029200,
CNS-2029261, CNS-2029260

FABRIC