



# FABRIC and FAB

Anita Nikolich
University of Illinois

Ilya Baldin
RENCI




# FABRIC/FAB Leadership Team

| | | | |
|---|---|---|---|
| Ilya Baldin (RENCI)  | Anita Nikolich (UIUC)  | Inder Monga (ESnet)  | Jim Griffioen (UKY)  |
| KC Wang (Clemson) 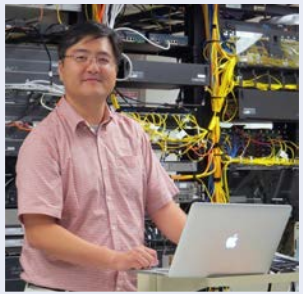 | Tom Lehman (Virnao)  | Paul Ruth (RENCI)  | Zongming Fei (UKY)  |

# Why FABRIC?

- The Internet is one big workaround. Protocols are awful. Security is an add-on.
- Make the future Internet more stateful
  - Cheaper compute - GPUs, FPGAs - enables this approach
  - Routers wouldn't look the way they do today if we could start over!
- ML/AI
  - Network as a 'big-data' instrument: real-time measurements + inferencing control loop
  - Security/anomaly detection, provisioning, etc
  - "Self-driving network" – Juniper marketing buzz
- IoT + 5G - high-speed intelligent network edge
- Continuum of computing
  - Not just "edge" or "cloud" but more seamless structure
  - ***Network as part of computing substrate - computing, fusing, processing data on the fly***

FABRIC

# FABRIC for everyone

**FABRIC Enables New Internet and Science Applications**

- Stateful network architectures, distributed applications that directly program the network

**FABRIC Advances Cybersecurity**

- At-scale realistic research facilitated by peering with production networks

**FABRIC Integrates HPC, Wireless, and IoT**

- A diverse environment connecting PAWR testbeds, NSF Clouds, HPC centers and instruments

**FABRIC Integrates Machine Learning & Artificial Intelligence**

- Support for in-network GPU-accelerated data analysis and control

**FABRIC helps train the next generation of computer science researchers**

# U.S.-Based Testbeds: a short history

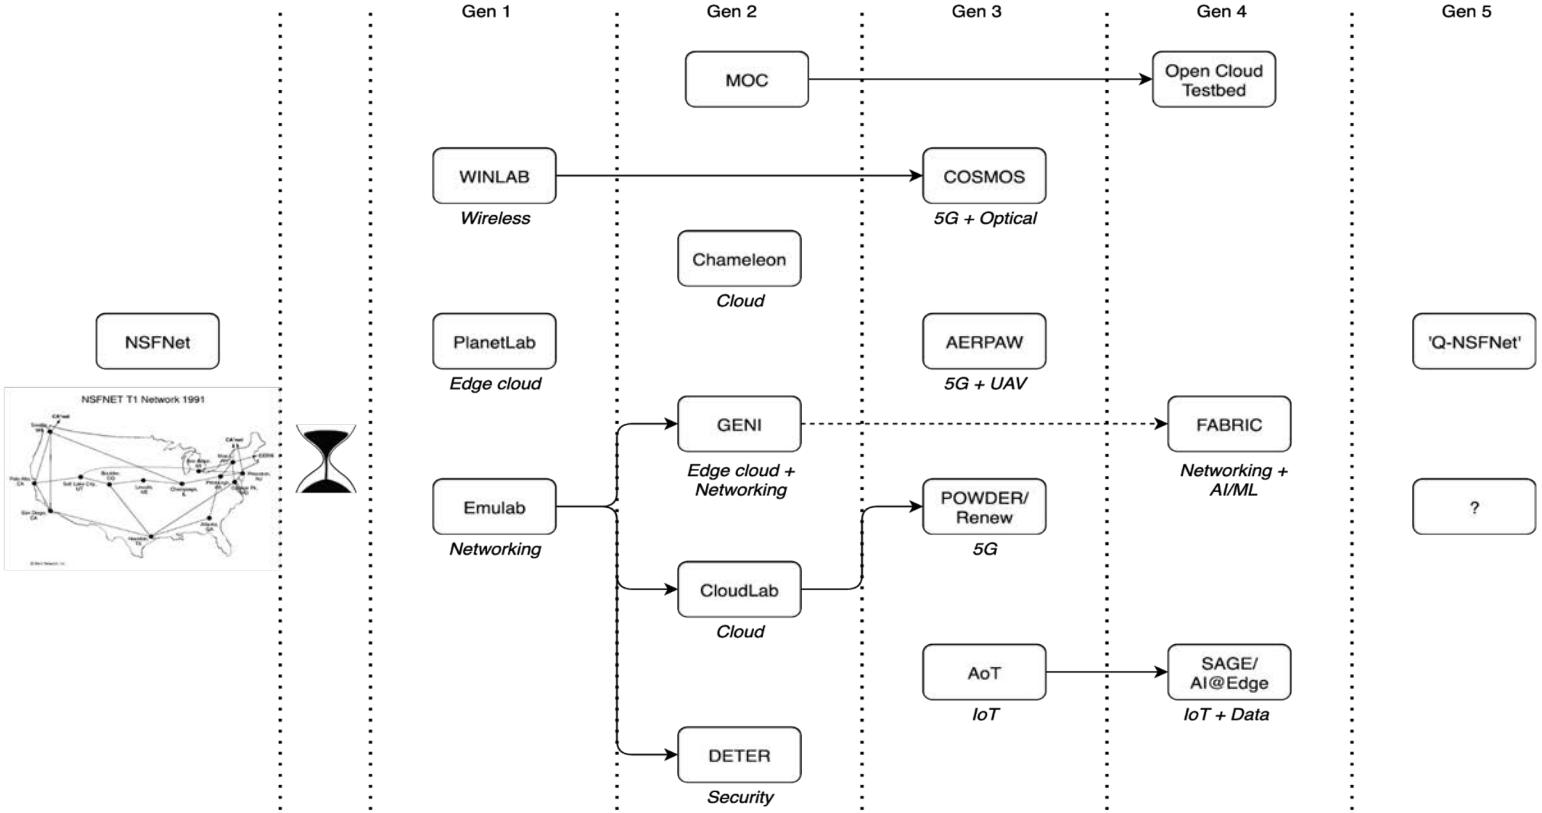

# Past Testbeds were Sandboxes

Did not connect experiment topologies to the Internet or production facilities

FABRIC

# What is FABRIC? An NSF Mid-Scale Project

~~**A testbed that**~~ **enables a *new paradigm for distributed applications and Internet protocols:***

- A nation-wide programmable network with compute and storage at each node. Run computationally intensive programs & maintain information in the network.
- GPUs, FPGAs, and network processors (NICs) inside the network
- Supports Quality of service (QoS) - dedicated optical 100Gb
- Interconnects national facilities: HPC, cloud & wireless testbeds, commercial clouds, Internet, edge nodes, science instruments, universities and labs
- Design and test applications, protocols and services that run at any node in the network not just edge or cloud

FABRIC

# FABRIC Nodes (“hanks”)

- Puts compute and storage in the path of fast packet flows
- Rack of high-performance servers (Dell 7525):
  - 2x32-core AMD 7532 with 512G RAM
  - GPUs (RTX 6000 and T4), FPGA network/compute accelerators
  - Storage - 1TB NVMe drives in servers & pool of ~250TB rotating storage per site.
  - Connect to 100G+ programmable switch
- Reconfigurable Network Interface Cards
  - FPGAs (with P4 support)
  - Speeds : 25G, 100G, 200G+(future)
- Kernel Bypass/Hardware Offload
  - VM/Containers to support full-rate DPDK

FABRIC

Conceptual FABRIC Node ‘Hank’ Overview

a.k.a. ‘A disaggregated router’

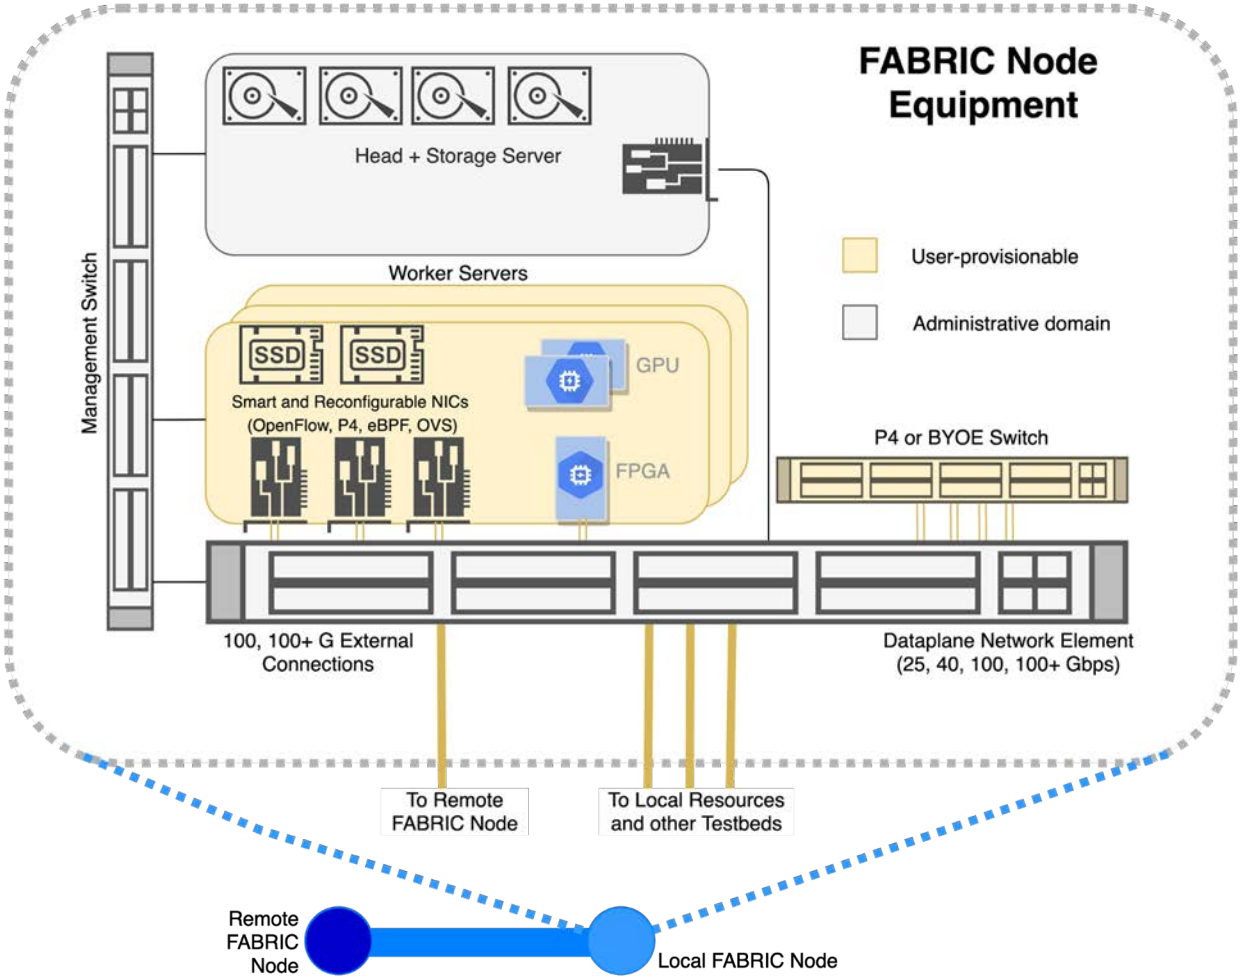

# Value Add: Fine Grained Measurement Capabilities

- GPS clock source using PTP (almost all sites)
- NICs that can do accurate packet sampling/timestamping
- Full packet capture at line speed for designated time
- Smart PDUs to measure power
- Optical layer measurements (where available) & experimentation
- CPU, memory, disk, port/interface utilization and other time-series (software)
- Ability to share measurement data

FABRIC

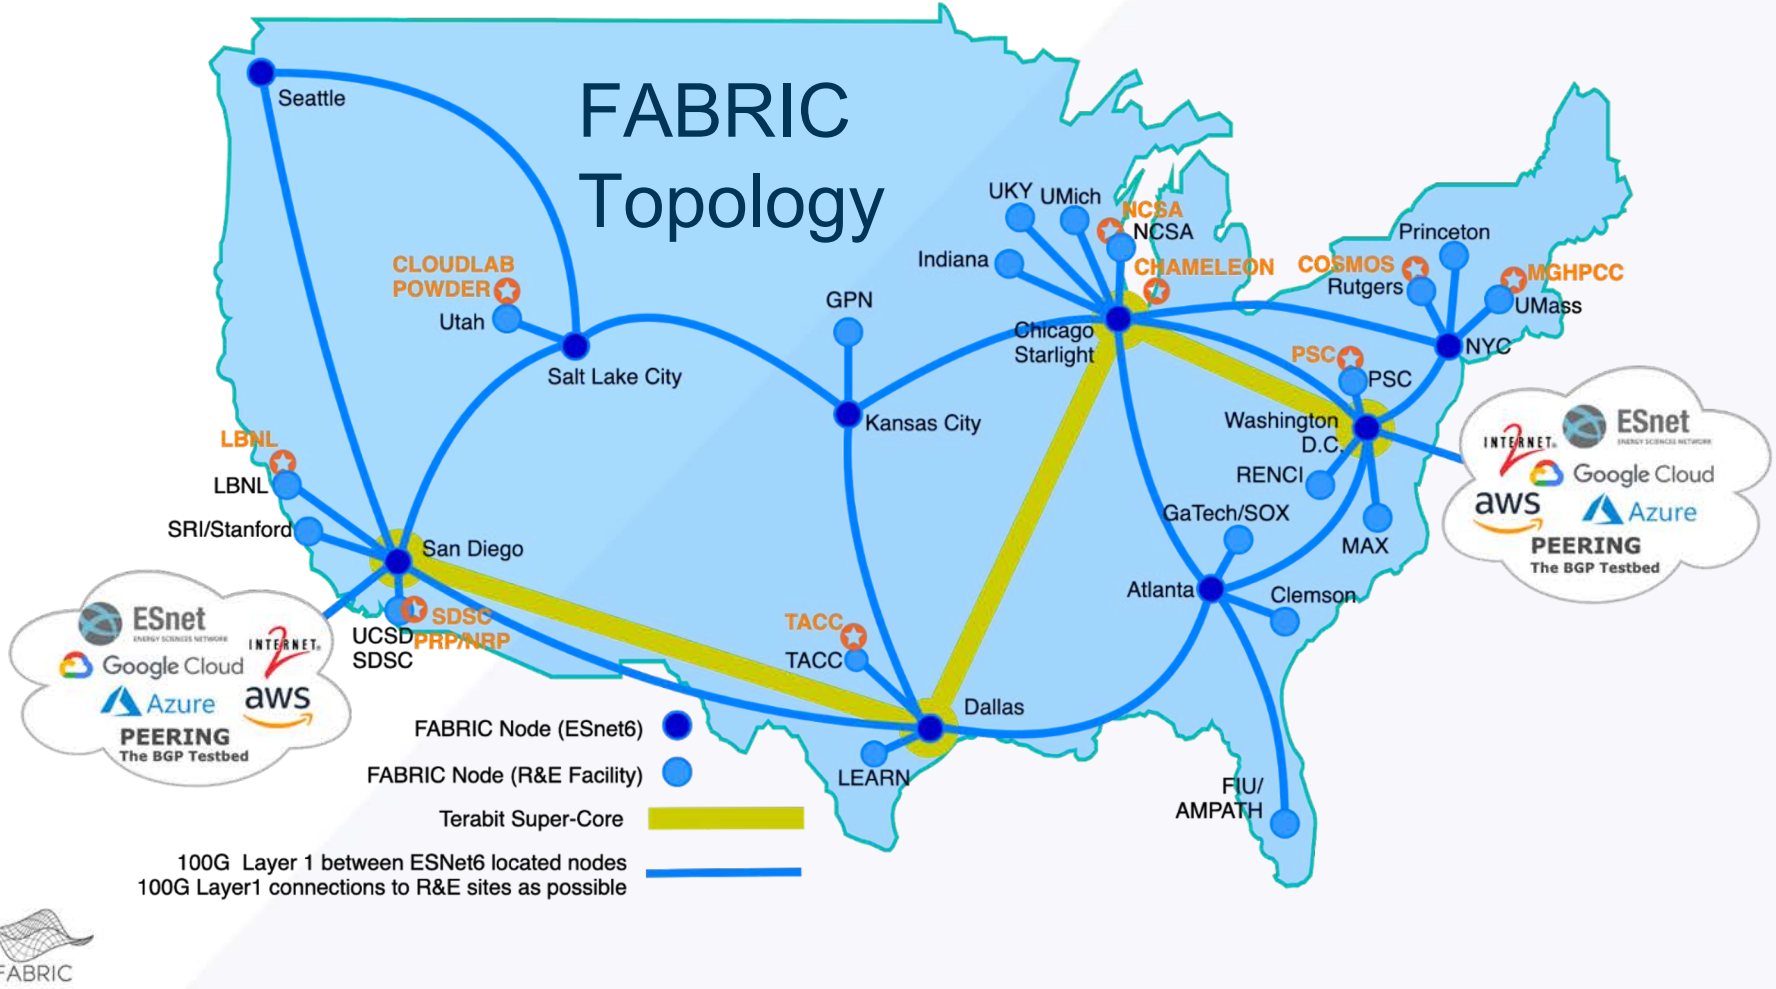
FABRIC Topology
Seattle
CLOUDLAB
POWDER
Utah
Salt Lake City
GPN
Kansas City
UKY
UMich
NCSA
NCSA
Indiana
CHAMELEON
Chicago
Starlight
COSMOS
Rutgers
Princeton
MGHPCC
UMass
NYC
PSC
PSC
Washington
D.C.
RENCI
GaTech/SOX
MAX
Atlanta
Clemson
LBNL
LBNL
SRI/Stanford
San Diego
UCSD
SDSC
SDSC
PRP/NRP
TACC
TACC
Dallas
LEARN
FIU/
AMPATH
ESnet
Google Cloud
Azure
aws
INTERNET2
PEERING
The BGP Testbed
FABRIC Node (ESnet6)
FABRIC Node (R&E Facility)
Terabit Super-Core
100G Layer 1 between ESNet6 located nodes
100G Layer1 connections to R&E sites as possible

FABRIC

# FAB (FABRIC Across Borders) – International

- Japan (University of Tokyo)
- UK (University of Bristol)
- EU (University of Amsterdam)
- CERN

- New Use-cases & Partners
  - Astronomy/Cosmology (CMB-S4, LSST) – test new data processing solutions to work on alert streams and “data in transit”
  - Weather (UMiami) – new workflows for high-res weather data between US & Brazil; Efficiently distributing data on weather events
  - High-Energy Physics (Co-PI Rob Gardner)
  - Urban Sensing/IoT/AI at Edge (UBristol)
  - Computer Science: 5G across borders, P4/SDN, Cybersecurity/Censorship Evasion

FABRIC

# FAB EU Connectivity

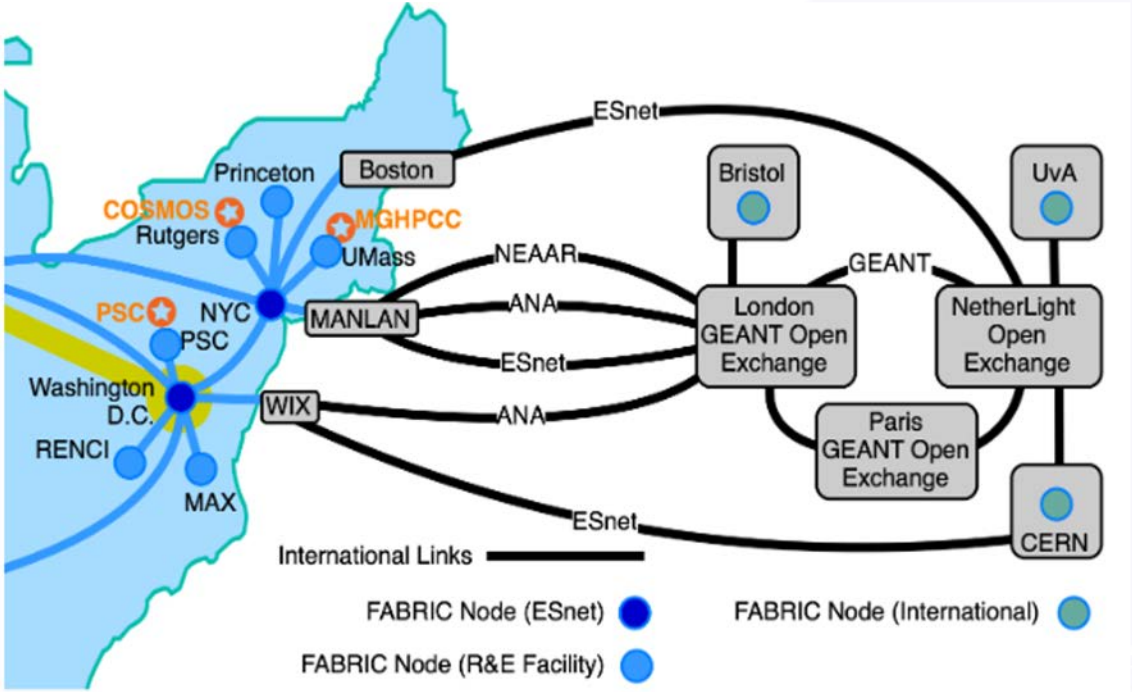


- Ring of three interconnected OpenExchange points (GEANT London, GEANT Paris, NetherLight) facilitates traffic in the region and across international links. Amsterdam Exchange point is operated by SURFnet.
- High degree of flexibility for Layer 2 and/or Layer 3 flows to other EU locations, or back to US.

FABRIC

# Example:Testing HEP data analysis approaches

- Real-time filtering and accelerated HEP data delivery
- Develop and test ML algorithms - inferencing within FABRIC nodes for real-time data processing

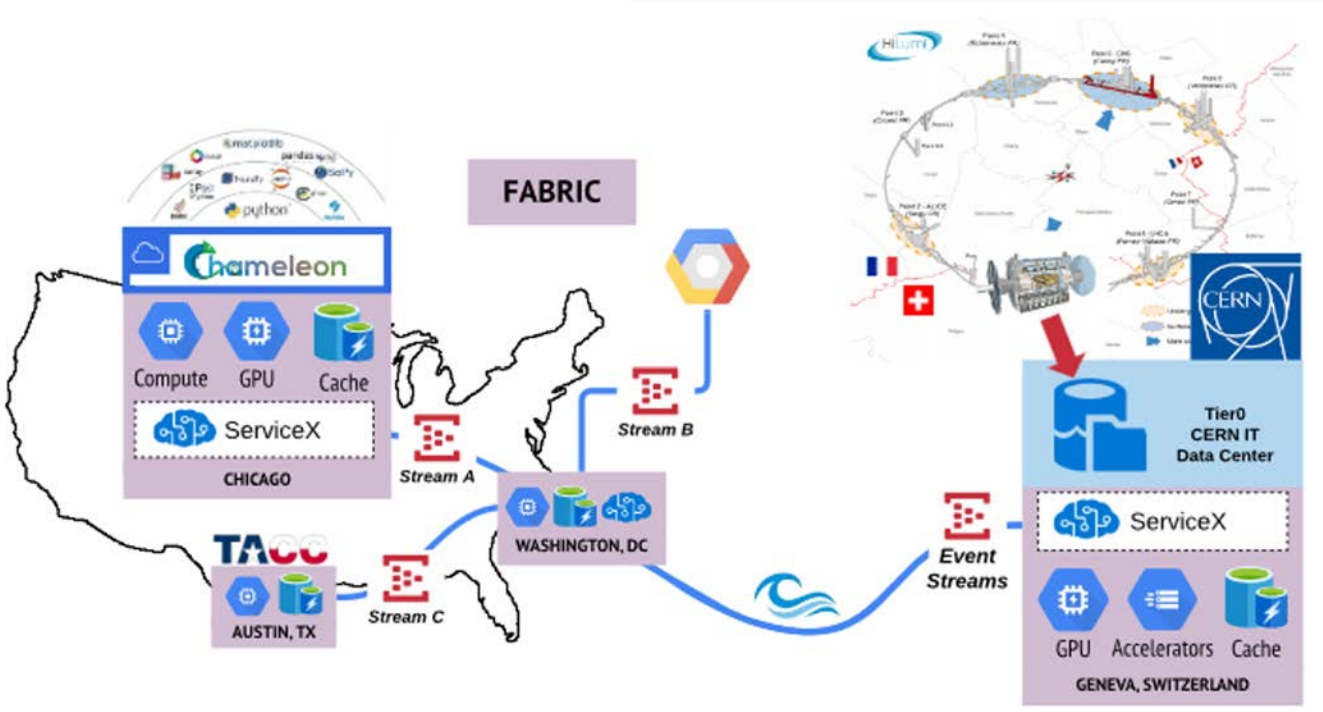


FABRIC

# Timeline

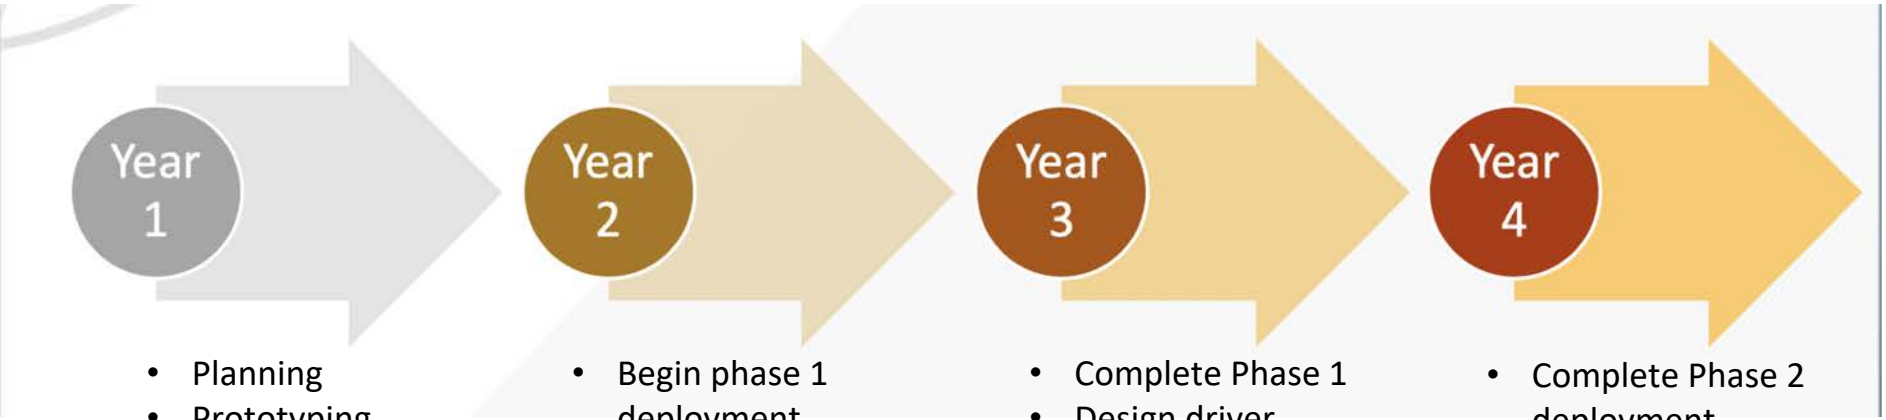


- Planning
- Prototyping
- Software development
- Community building

- Begin phase 1 deployment
- Testing, commissioning
- Initial Experiment on-boarding

- Complete Phase 1
- Design driver experiments and early users
- Begin Phase 2 deployment

- Complete Phase 2 deployment
- Prepare for operations

# Help Us!

- Host a Hank Node
- Join a Community Working Group: Data/Storage, Security, Measurement
- Feedback on topology
- Cool experiments?

FABRIC

# Thank You!

Questions?

Visit https://whatisfabric.net

Ask info@fabric-testbed.net

Email me: anitan@illinois.edu

FABRIC Software: https://github.com/fabric-testbed


This work is funded by
NSF grant CNS-1935966

FABRIC

# Back Up Slides

# Example FABRIC Use-case Scenarios

Examples of potential uses:

- Bump-in-wire measurements and packet sampling at high bit rates (25, 40, 100, 100+ Gbps)
- Hardware-accelerated switching using Smart NICs, FPGA NICs or P4 switches in individual nodes
- Hosting in-network applications and stateful architectures using a combination of storage and compute resources in individual nodes
- In-network inference, other types of accelerated computing via FPGAs and GPUs
- Connect experiments to external facilities like IoT, 5G, cloud testbeds, public clouds and HPC resources.
- Deploy non-IP protocols on top of wide-area L2 topologies, that may include in-network processing and storage

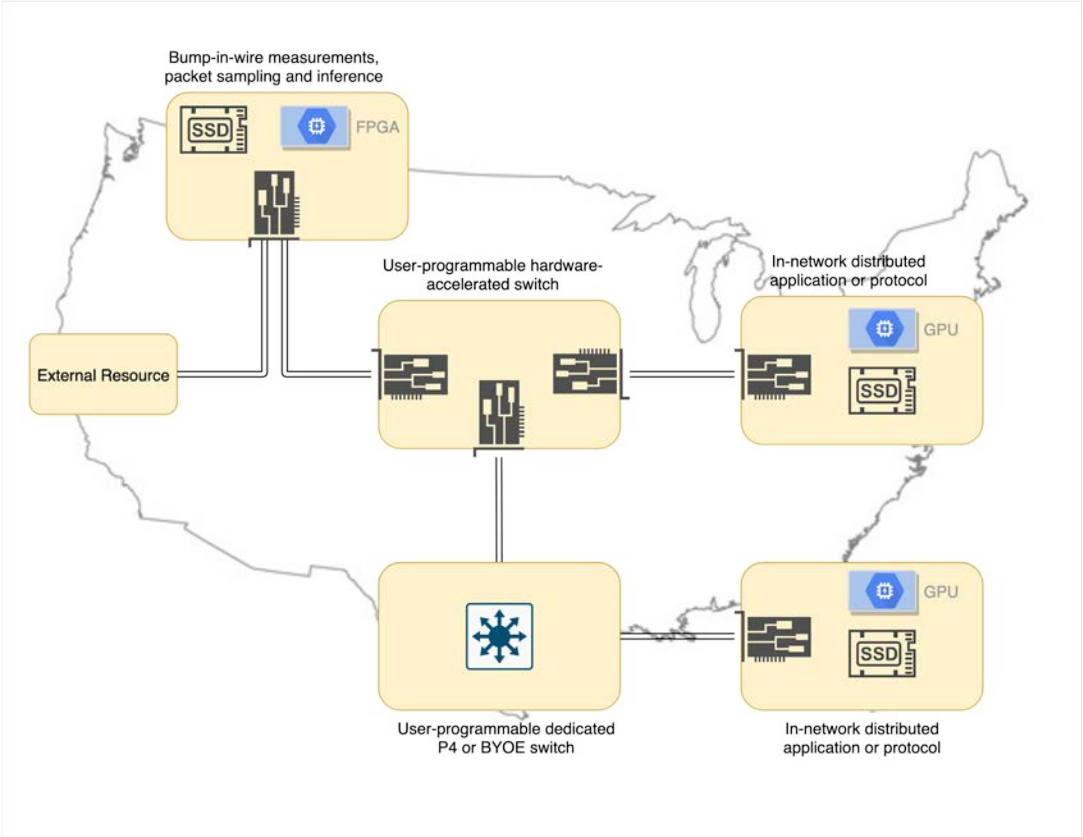

# FABRIC Near-Term use-cases

- FABRIC ‘Science Design Drivers’:
  - SRI - Network Security
  - Georgia Tech - 5G/IoT/Network Resilience
  - University of Virginia - ML/Autonomous Network Management
  - FIU - Named Data Networking and AR/5G
- FAB : domain-oriented use-cases
  - Efficient distribution and in-network fusion of astronomical event data
    - LSST/Vera Rubin and CMB-S4
  - Urban Sensing
    - Connecting COSMOS, SAGE and University of Bristol testbeds
  - Weather science
    - Efficiently distributing data on weather events
  - Computer Science
    - Censorship evasion
    - Private 5G across borders
    - SDX policy negotiation

FABRIC