Prototype National Research Platform

Frank Wuerthwein
Interim Director
San Diego Supercomputer Center

Long Term Vision

- Create an Open National Cyberinfrastructure that allows the federation of CI at all 3,900 accredited, degree granting higher education institutions.
  - Open Science
  - Open Data
  - Open Source
  - Open Infrastructure

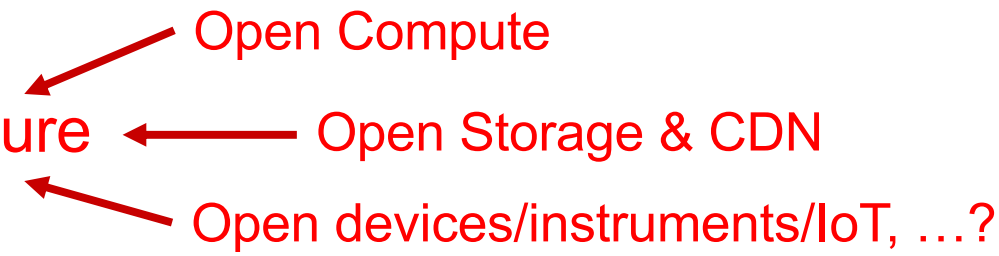

Openness for an Open Society

Multiple Ways to build Open Infrastructure Federation(s)

- At the container orchestration layer, e.g. K8S as implemented in Pacific Research Platform.
- Federating independent container orchestration frameworks, e.g. via SLATE or via Admiralty on K8S.
- Federating storage and compute clusters at the storage and batch system layer., e.g. via OSG.

More cybersecurity control requires more effort to join & operate.

By building compute & data federations
we enable sharing.
Sharing leads to democratization of access.

Prototype National Research Platform

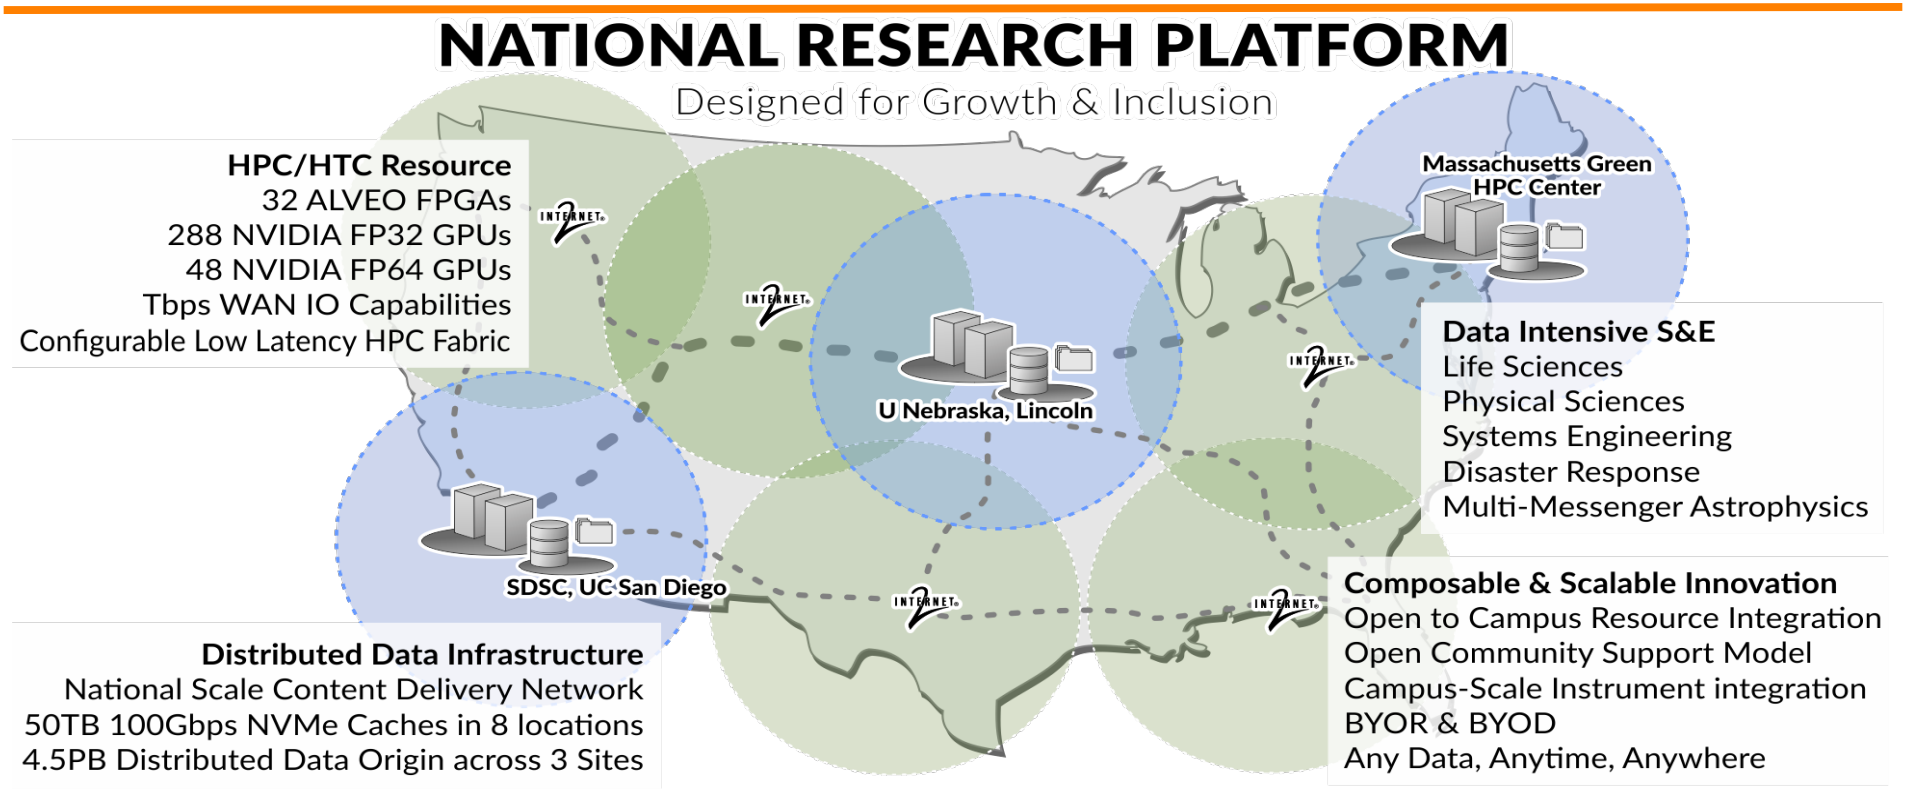

Nationally Distributed Category-II System (NSF)

3 year prototyping phase followed by 2 year allocation phase.

SDSC SAN DIEGO SUPERCOMPUTER CENTER

Funded as NSF 2112167

Be Open To All

- During 3 year prototyping phase, we will try to explore all NSF supported allocation mechanisms.
  - XRAC via Access
  - NSF credit via PATh
  - Testbed allocations as for Chameleon, CloudLab, FABRIC, …
- Will work with NSF on what mechanisms fit best for NRP during year 4-5 allocation phase.

Innovation Thrusts

- I1: Innovative composable network fabric that allows “rack” of hardware to behave like a single “node” connected via PCIe.
- I2: Innovative application libraries to expose FPGAs hardware to science apps at language constructs scientists understand (C, C++ rather than firmware)
- I3: A “Bring Your Own Resource” model that allows campuses nationwide to join their resources to the system.
- I4: Innovative scheduling to support urgent computing, including interactive via Jupyter.
- I5: Innovative Data Infrastructure, including national scale Content Delivery System like YouTube for science.

5 Thrusts of Innovation

I3: Bring Your Own Resource

- Offer campuses to join their hardware to our K8S cluster such that we provide OS & system software support
  - "Co-owners" provide only hardware maintenance
- Support K8S federation via Admiralty.
  - Federate with PRP, CHASE-CI, Expanse, OSG, …
- NRP accessible as an "OSG site" and "OSG Access Point".
- We also provide user support
  - Experiment with scalable user support via staff monitored slack channels.
  - Create community that supports itself
    - Think social networking and open source community as examples.

We strive to understand scalability of growth in resources and growth in user support.

I5: Innovative Data Infrastructure

- Build on both the OSG Data Federation (see Derek Weitzel's talk) and the PRP Regional Ceph clusters.
  - NRP includes 3 Ceph cluster & 8 Data Federation Caches
- Automated replication between Ceph clusters
  - User specify what part of their namespace should be replicated where.
- Experiment with allowing BYOR storage as part of this infrastructure, including the replication.

Any BYOR compute location in continental USA is within ~500 miles of an NRP data cache.

I1: Composable Systems

One of two identical systems,
connected via PCIe Interswitch links

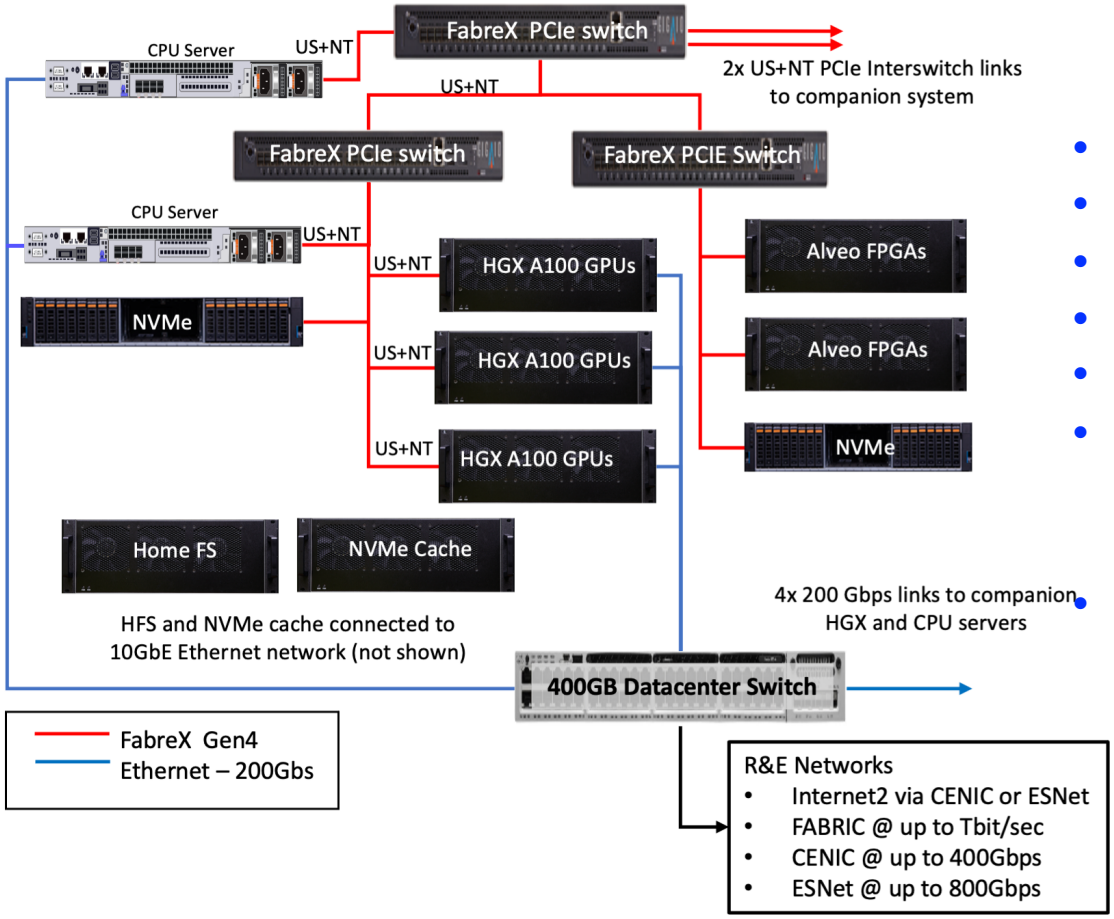

Figure 2.1 High-performance Cluster Network Architecture

- 6 HGX from NVIDIA
  - 48GPUs
  - NVLINK connects 8 a piece
- 350TB of NVMe
- 90x18TB raw /home
- 32 Alveo FPGAs
- 4 additional CPU servers
- Cache for CDN
- Max Connectivity to LAN/WAN:
  - 10x200Gbps via CPU hosts
  - 32x100Gbps via FPGAs
- Possible Expansion: ”Bring Your Own Device”

GigaIO FabreX allows for composing hardware into nodes dynamically.

Science & Engineering Thrusts

- Molecular Dynamics
- Genomics
- Health
- Global Scale Science Instruments
- Campus Scale Science Instruments
- Artificial Intelligence
- Computer Vision and Visualization
- Systems Design & Engineering
- Urgent Computing

Initial Science Drivers from across these thrusts.

Individual collaborators overlap on science drivers and innovation thrusts

Project Details

- NSF Award 2112167
- Award value: $5M Acquisition; $6.25M Operations and Maintenance (expected)
- PI: Frank Wuerthwein (UCSD)
- Co-PIs: Tom DeFanti, Tajana Rosing, Mahidhar Tatineni (UCSD); Derek Weitzel (UNL)
- Partners: UCSD/SDSC, University of Nebraska at Lincoln, Massachusetts Green High Performance Computing Center, Internet2
- Awarded July 1, 2021
- Deployment in H1-2022
- Production Operations: July 1, 2022 – June 30, 2027
- Operated as a resource via XSEDE/ACCESS

Acknowledgements

- NRP is funded via NSF award 2112167