



A Utah-Wyoming Cyberinfrastructure
Water Modeling Collaboration



CI Facilities at the University of Utah and UEN

Team 1 – Steve Corbató
Interim CIO
Adjunct Faculty, School of Computing
University of Utah

AAAS External Review
University of Wyoming – Laramie WY
November 17, 2014



A Utah-Wyoming Cyberinfrastructure
Water Modeling Collaboration



Team 1 CI objectives

- Provide coordinated, high-performance information technology resources and services to CI-WATER team and other UT and WY EPSCoR researchers
- Support research data analysis, management, and curation; modeling; and simulation needs
- Tools: computing cycles, data storage, advanced networking, visualization environments, middleware, software libraries, software development, data centers



Related CI investments

- New data center in downtown SLC (2.4 MW) with dedicated research computing space – University of Utah
- New optical network in northern Utah – BONFIRE – with Utah Education Network (UEN)
 - NSF EPSCoR RII C2 award – Corbato, PI
 - UEN's BTOP award (NTIA/Dept. of Commerce)

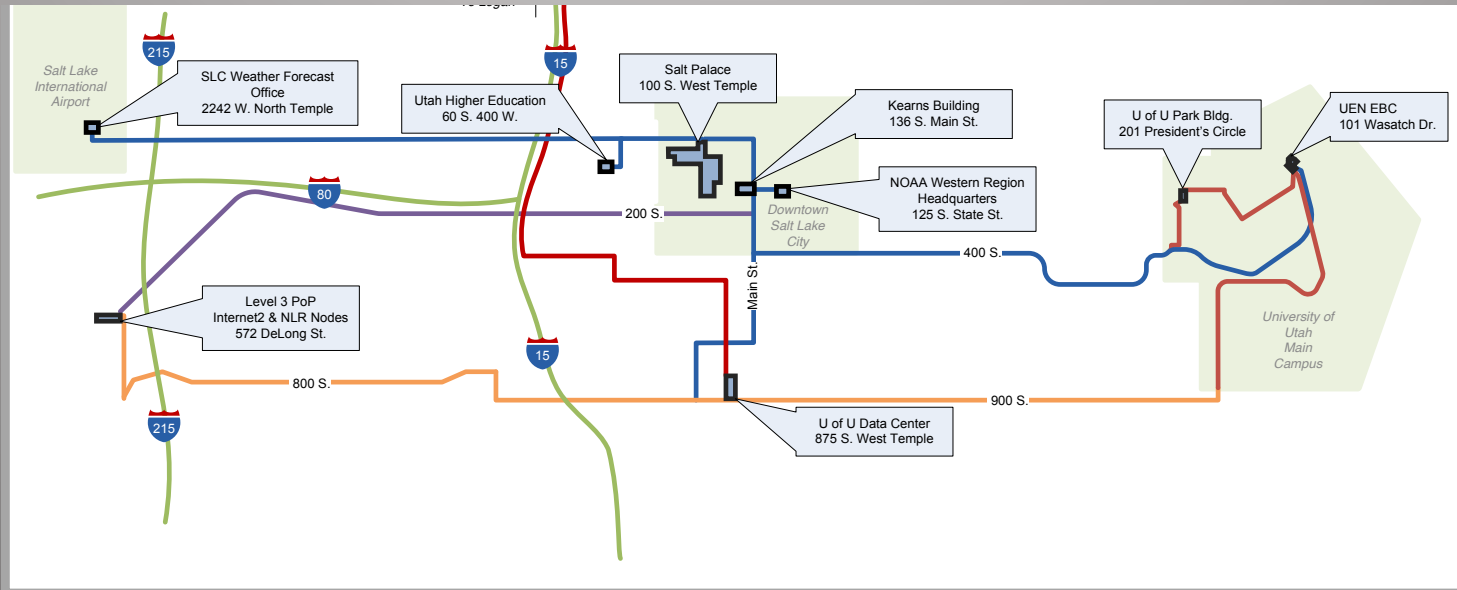
SLC Downtown Data Center



- 74,000+ sq ft² former industrial building south of downtown SLC (~4 miles off-campus)
- Designed for enterprise & HPC (2.4 MW)
- Co-location by research groups & partners
- Low industrial electric power rates in Utah
- High desert climate: energy efficiency
- In production early 2012; CHPC co-located



Salt Lake City metro optical network - BONFIRE



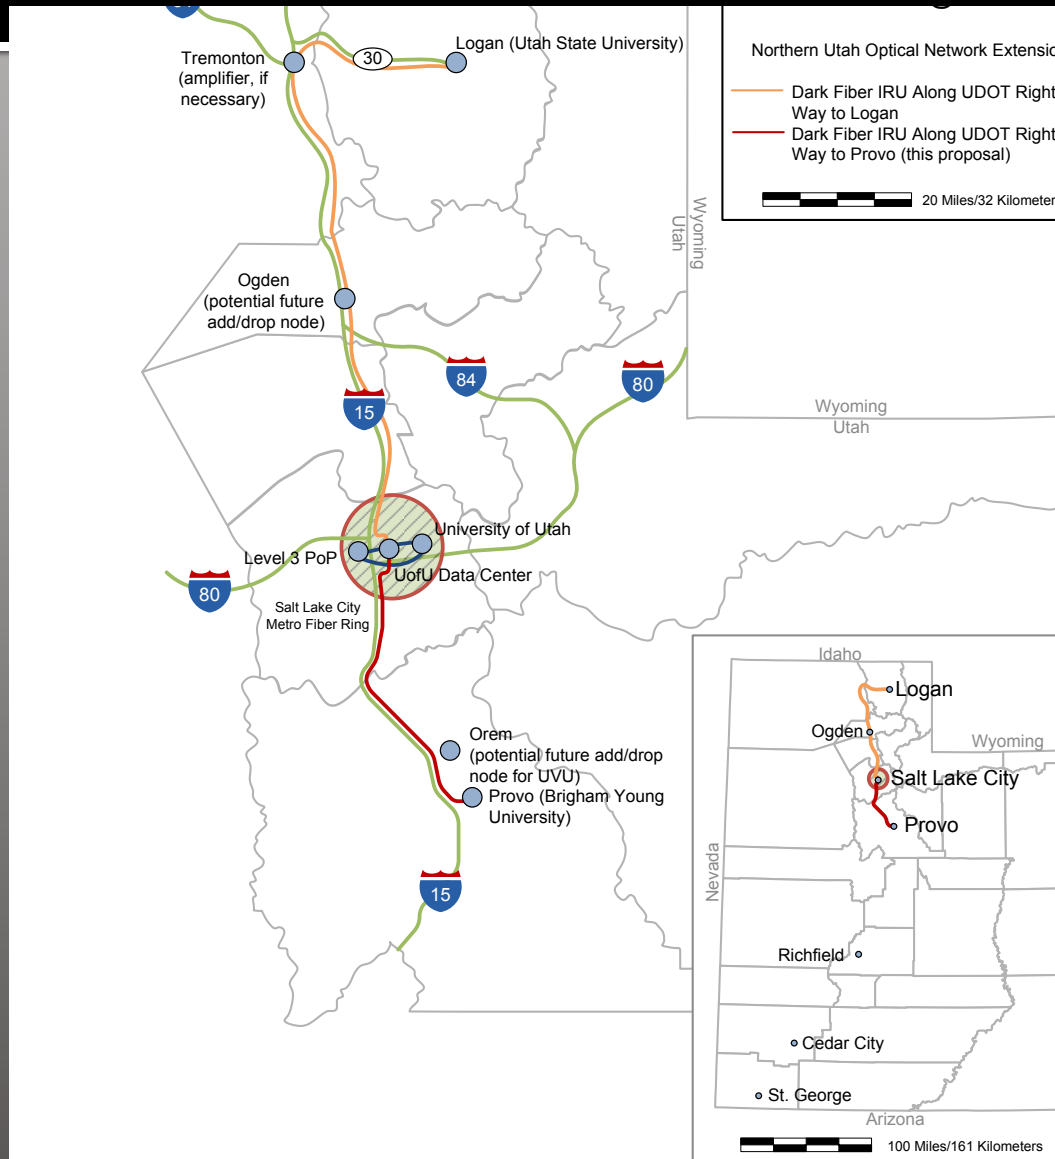
Research@UEN: Salt Lake City Metro Optical Network

- U of U Campus Fiber
- UTA Light Rail Routes (proposed)
- CENIC/LLC Fiber IRU (through AFS)
- AFS Fiber IRU (proposed)
- Northern Utah Extension (proposed)

1 Mile

Carrier proprietary information included

BONFIRE extended to USU & BYU

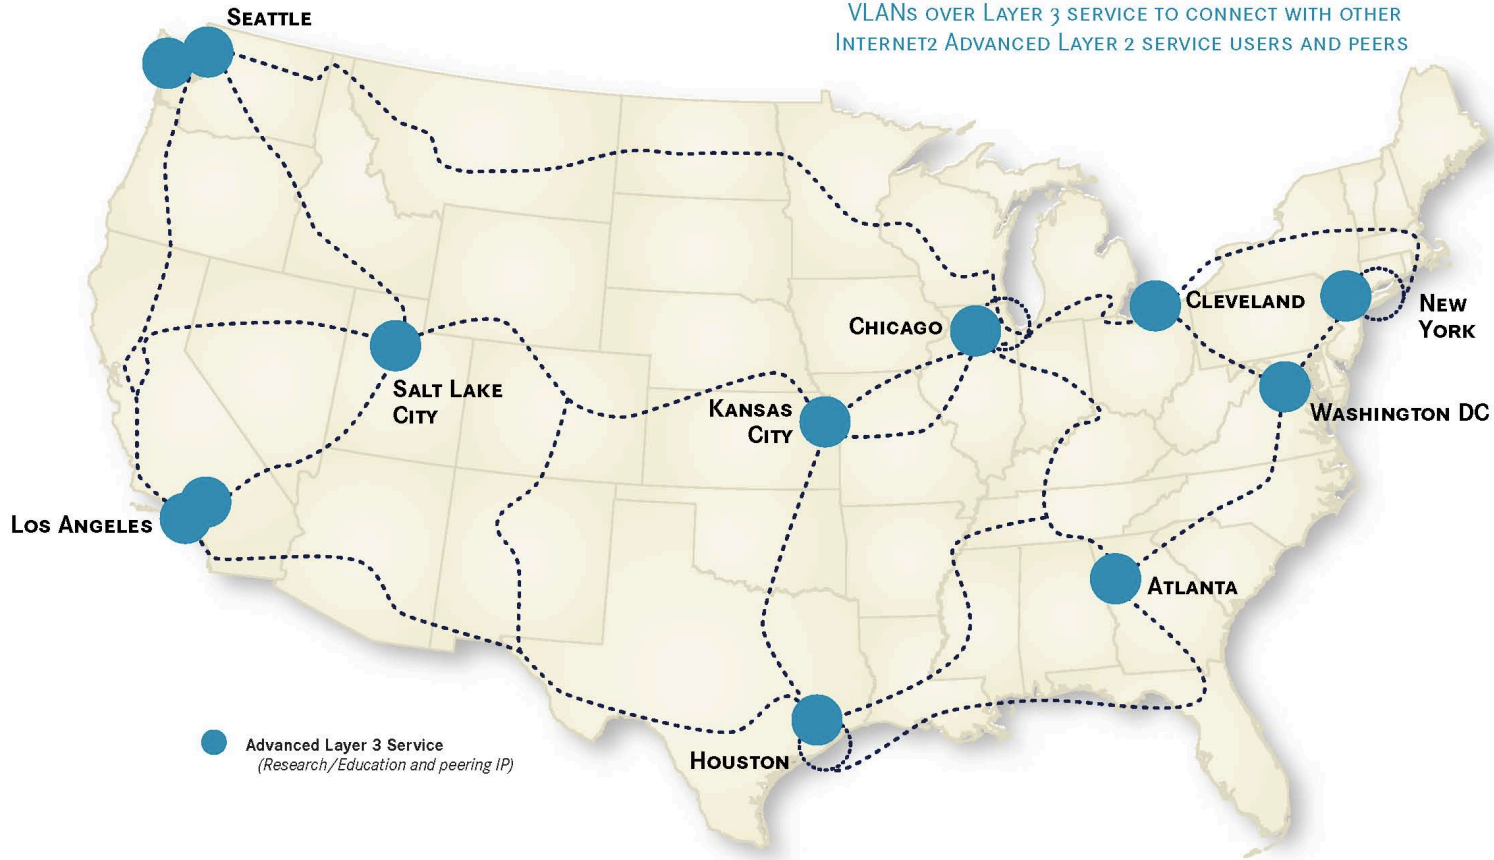




INTERNET2 NETWORK ADVANCED LAYER 3 SERVICE

OCTOBER 2014

VLANS OVER LAYER 3 SERVICE TO CONNECT WITH OTHER INTERNET2 ADVANCED LAYER 2 SERVICE USERS AND PEERS



INTERNET2 NETWORK BY THE NUMBERS

17	JUNIPER MX960 ROUTERS SUPPORTING LAYER 3 SERVICE
34	BROCADE AND JUNIPER SWITCHES SUPPORTING LAYER 2 SERVICE
62	CUSTOM COLLOCATION FACILITIES
250+	AMPLIFICATION RACKS
15,717	MILES OF NEWLY ACQUIRED DARK FIBER
8.8	TBPS OF OPTICAL CAPACITY
100	GBPS OF HYBRID LAYER 2 AND LAYER 3 CAPACITY
300+	CIENA ACTIVEFLEX 6500 NETWORK ELEMENTS
2,400	MILES PARTNERED CAPACITY WITH ZAYO COMMUNICATIONS IN SUPPORT OF THE NORTHERN TIER REGION



IN SUPPORT OF
U.S.UCAN

NETWORK PARTNERS

ciena



INDIANA UNIVERSITY

infinera

JUNIPER
NETWORKS





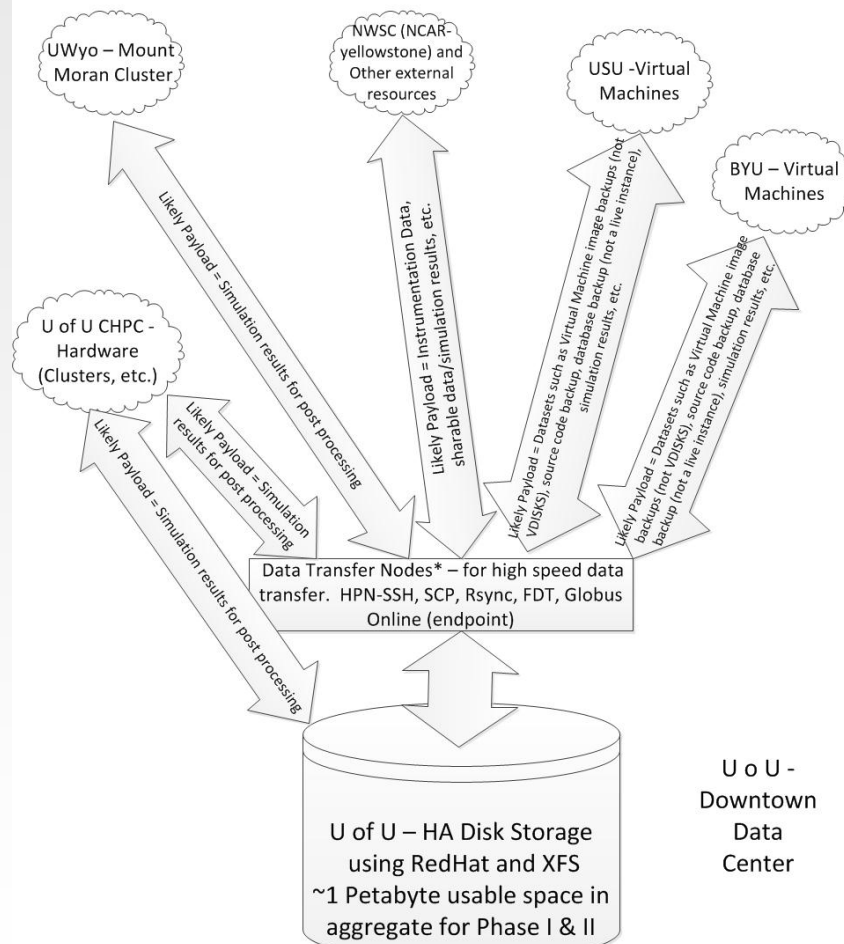
CI-WATER data repository for Big Data

- We moved away from our previous file storage system due to software reliability and performance issues – implementation delayed
- Acquired 1 Petabyte usable space (~\$230/TB – Dell) and made available mid-2013
 - Includes Data Transfer Nodes (DTNs)
 - Archive tape capability added in 2014 – 690 TB
- Critical support for CI-WATER atmospheric modeling work of Court Strong
- Close coordination with another campus Big Data project
 - Partnered with Adam Bolton (Physics & Astronomy) to support Utah's role as the U.S. data management site for Sloan Digital Sky Survey 4 (SDSS-4) – 175 TB



- Shared data repository among CI-WATER and iUTAH institutions
- Supports both fast I/O computation and long-term data archiving
- Hosted at UofU DDC
- Leverages high-speed, secure data transfer nodes (DTNs) as advocated by NERSC and ESnet
- Uses Globus toolset

CI-WATER STORE – EPSCoR Data Repository Data Access Schematic – 2/8/2013



* Data Transfer Nodes are utilized for “put & get” type of operations. The software stack typically offers fast transfer of data.



Three subsequent major NSF awards

- CC-NIE Integration project: *Science Slices*
 - Corbato, PI, University of Utah (ACI-1341034)
- CloudLab open cloud computing platform testbed
 - Rob Ricci, PI, University of Utah (CNS-1419199)
- ACI-REF – Univ. of Utah is one six pilot participants
 - Jim Bottum, PI, Clemson (CNS-1338155)

Utah CC-NIE Integration project: *Science Slices* (NSF #ACI-1341034)

PI: S. Corbató; co-PIs: A. Bolton, T. Cheatham, R. Ricci, K. Van der Merwe; SP: J. Breen, S. Torti

Premise (Rob Ricci): What if we flipped the concept and built our Science DMZ on top of SDN infrastructure, rather than just plugging our SDN testbed into the DMZ?

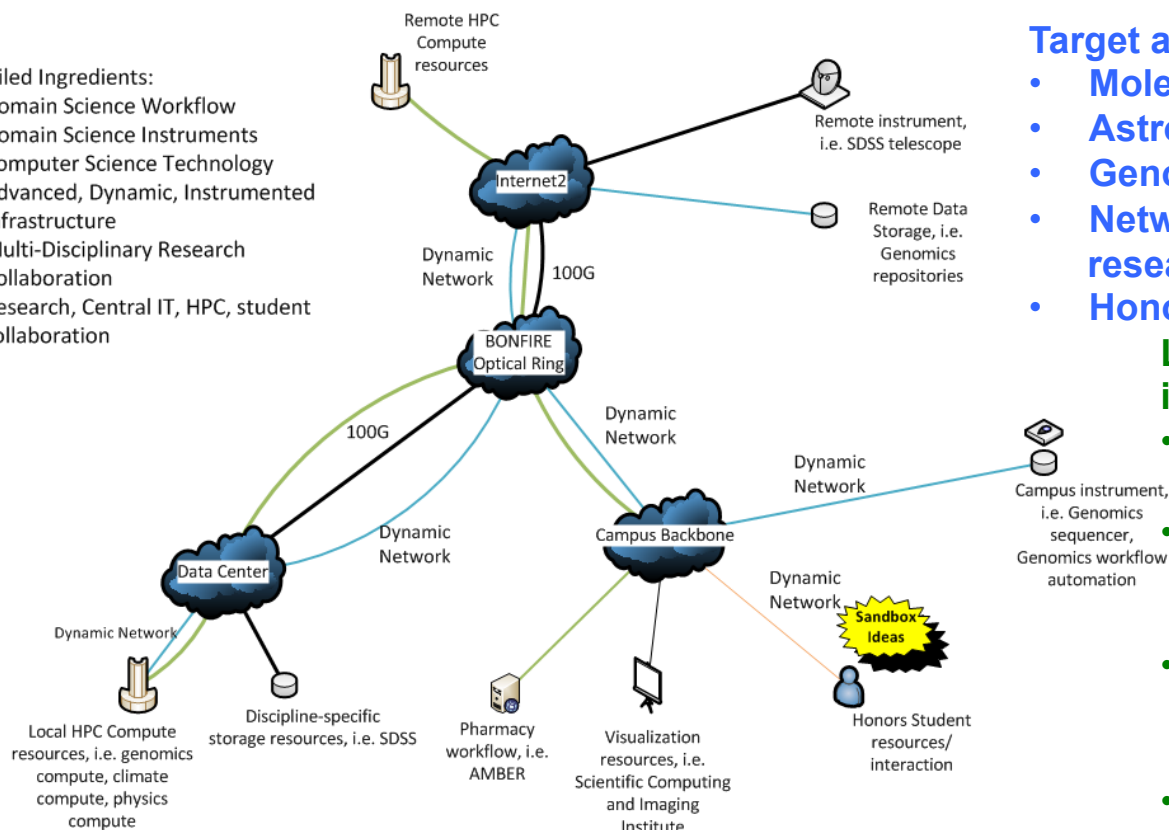
1) Building a dynamic Science DMZ on top of an SDN-based framework (GENI)

Detailed Ingredients:

- Domain Science Workflow
- Domain Science Instruments
- Computer Science Technology
- Advanced, Dynamic, Instrumented Infrastructure
- Multi-Disciplinary Research Collaboration
- Research, Central IT, HPC, student collaboration

2) Extending slices to key campus labs, HPC center, and the Honors residential community

3) Working closely with central IT, campus IT Governance, and Utah Education Network



Target areas:

- Molecular dynamics
- Astrophysics data
- Genomics
- Network/systems research
- Honors students

Leverages new infrastructure:

- Downtown Data Center
- Utah Optical Network (BONFIRE)
- NSF MRI for novel cluster (Apt)
- Campus Net Upgrade



CloudLab

updated: 10/23/14

www.cloudlab.us

CloudLab

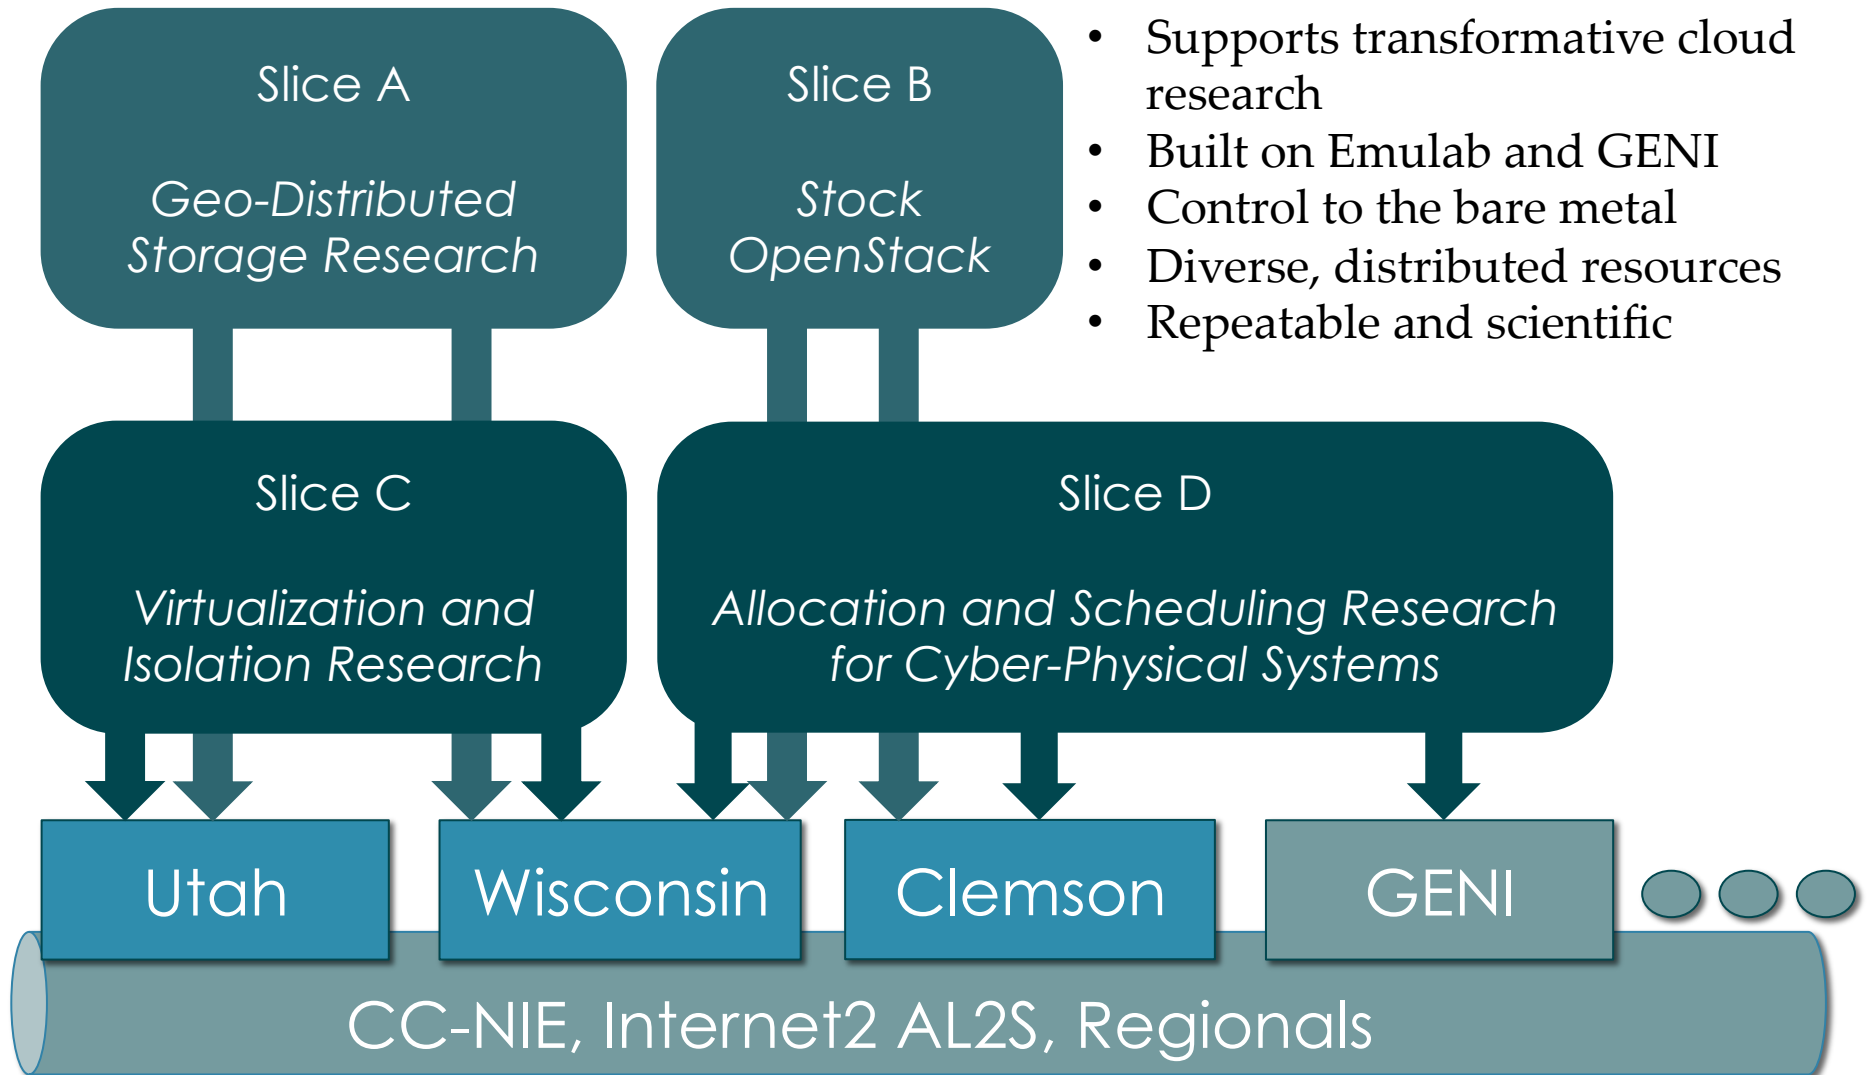


Supported by National Science Foundation under Grant #1419199



What Is CloudLab?

- Supports transformative cloud research
- Built on Emulab and GENI
- Control to the bare metal
- Diverse, distributed resources
- Repeatable and scientific





CloudLab's Hardware

One facility, one account, three locations

- About 5,000 cores each (15,000 total)
- 8-16 cores per node
- Baseline: 4GB RAM / core
- Latest virtualization hardware
- TOR / Core switching design
- 10 Gb to nodes, SDN
- 100 Gb to Internet2 AL2S
- *Partnerships with multiple vendors*

Wisconsin

- **Storage and net.**
- Per node:
 - 128 GB RAM
 - 2x1TB Disk
 - 400 GB SSD
- Clos topology
- *Cisco (OF 1.0)*

Clemson

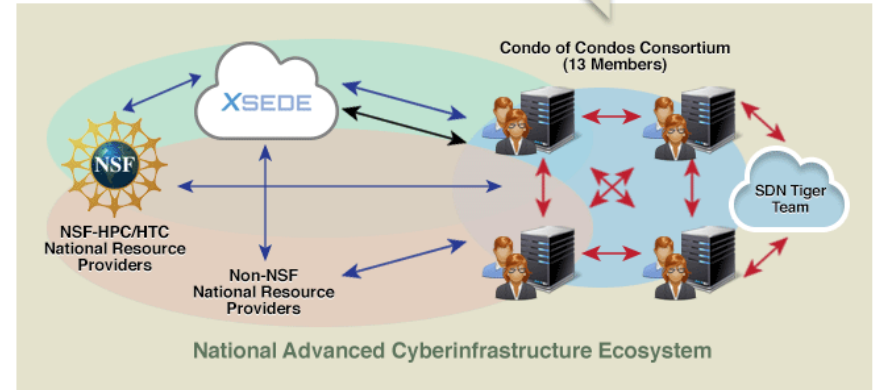


- **High-memory**
- 16 GB RAM / core
- 16 cores / node
- Bulk block store
- Net. up to 40Gb
- High capacity
- *Dell (OF 1.0)*

Utah

- **Power-efficient**
- ARM64 / x86
- Power monitors
- Flash on ARMs
- Disk on x86
- Very dense
- *HP (OF 1.3)*



ACI-REF project: Condo of Condos

NSF #CNS-1338155 – PI: Jim Bottum, Clemson



ACI-REF campus research computing centers are federating to share resources – both human and computational

Also all connect to the Internet2 Layer 2/3 Network at 100G





Next steps and sustainability

- Complete final year expansion of data archive and related systems (with USU)
- Continue HPC partnership with Univ. of Wyoming
- Complete HPC consolidation with Utah State Univ.
- Work with UEN to expand BONFIRE optical network in northern Utah as needed to support demand
- Partner with Marriott Library for campus research data management program
- Leverage NSF CloudLab award for new capabilities to support combined atmospheric/hydrological research
- Leverage NSF ACI-REF approach for federating campus research computing centers for shared user services and resources and more standardized operations