LHC Networking And SDN/SDX Services

Joe Mambretti, Director, (j-mambretti@northwestern.edu)
International Center for Advanced Internet Research (www.icair.org)
Northwestern University
Director, Metropolitan Research and Education Network (www.mren.org)
Director, StarLight, Co-PI Chameleon (www.startap.net/starlight)
PI IRNC: RXP: StarLight SDX

LHCOPN-LHCONE Meeting
Helsinki, Finland
September 19-20, 2016

iCAIR

STARLIGHT SDX

Macro Network Science Themes

- Transition From Legacy Networks To Networks That Take Full Advantage of IT Architecture and Technology
- Extremely Large Capacity (Multi-Tbps Streams)
- High Degrees of Communication Services Customization
- Highly Programmable Networks
- Network Facilities As Enabling Platforms for Any Type of Service
- Network Virtualization
- Highly Distributed Processes
- SDN/SDX/SDI/OCX/SDC/SDE
- Key Attribute For SDXs = ***Open Services, Architecture, Connectivity***

iCAIR

STARLIGHT SDX


GriPhyN
Data Intensive Science
NEPTUNE
A Fiber-Optic Telescope to Inner Space
NEES
GENOMES to LIFE
BIOLOGICAL SOLUTIONS FOR ENERGY CHALLENGES
U.S. DEPARTMENT OF ENERGY
NG Digital Sky Survey
NEON
National Ecological Observatory Network
BIRN
BIRN: Biomedical Informatics Research Network
www.nbirn.net
CAMERA metagenomics
camera.calit2.net
Carbon Tracker
www.esrl.noaa.gov/gmd/ccgg/carbontrack
CineGrid
CineGrid
www.cinegrid.org
LHC
LHCONE
www.lhcone.net
ALMA
ALMA: Atacama Large Millimeter Array
ANDRILL: Antarctic Geological Drilling
www.andrill.org
OOI OCEAN OBSERVATORIES INITIATIVE
CYBERINFRASTRUCTURE
Providing a link between ocean research and discovery
OOI-CI
ci.oceanobservatories.org
GEON
GEON: Geosciences Network
www.geongrid.org
GLEON
GLEON: Global Lake Ecological Observatory Network
ISS: International Space Station
www.nasa.gov/station
CLASS
Comprehensive Large-Array Stewardship System
www.class.noaa.gov
DØ (DZero)
www-d0.fnal.gov
Two neutron stars orbit each other
LIGO
www.ligo.org
WLCG
Worldwide LHC Computing Grid
WLCG
lcg.web.cern.ch/LCG/public/
PRAGMA
Pacific Rim Applications and Grid Middleware Assembly
www.pragma-grid.net
TeraGrid
TeraGrid
www.teragrid.org
IVOA
IVOA: International Virtual Observatory
www.ivoa.net
Open Science Grid
OSG
www.opensciencegrid.org
the globus alliance
Globus Alliance
www.globus.org
SKA
SQUARE KILOMETRE ARRAY
SKA
www.skatelescope.org
Sloan Digital Sky Survey
www.sdss.org
XSEDE
XSEDE
www.xsede.org
LIGO
STARLIGHT
Compilation By Maxine Brown

LHCONE – LHC Open Network Environment

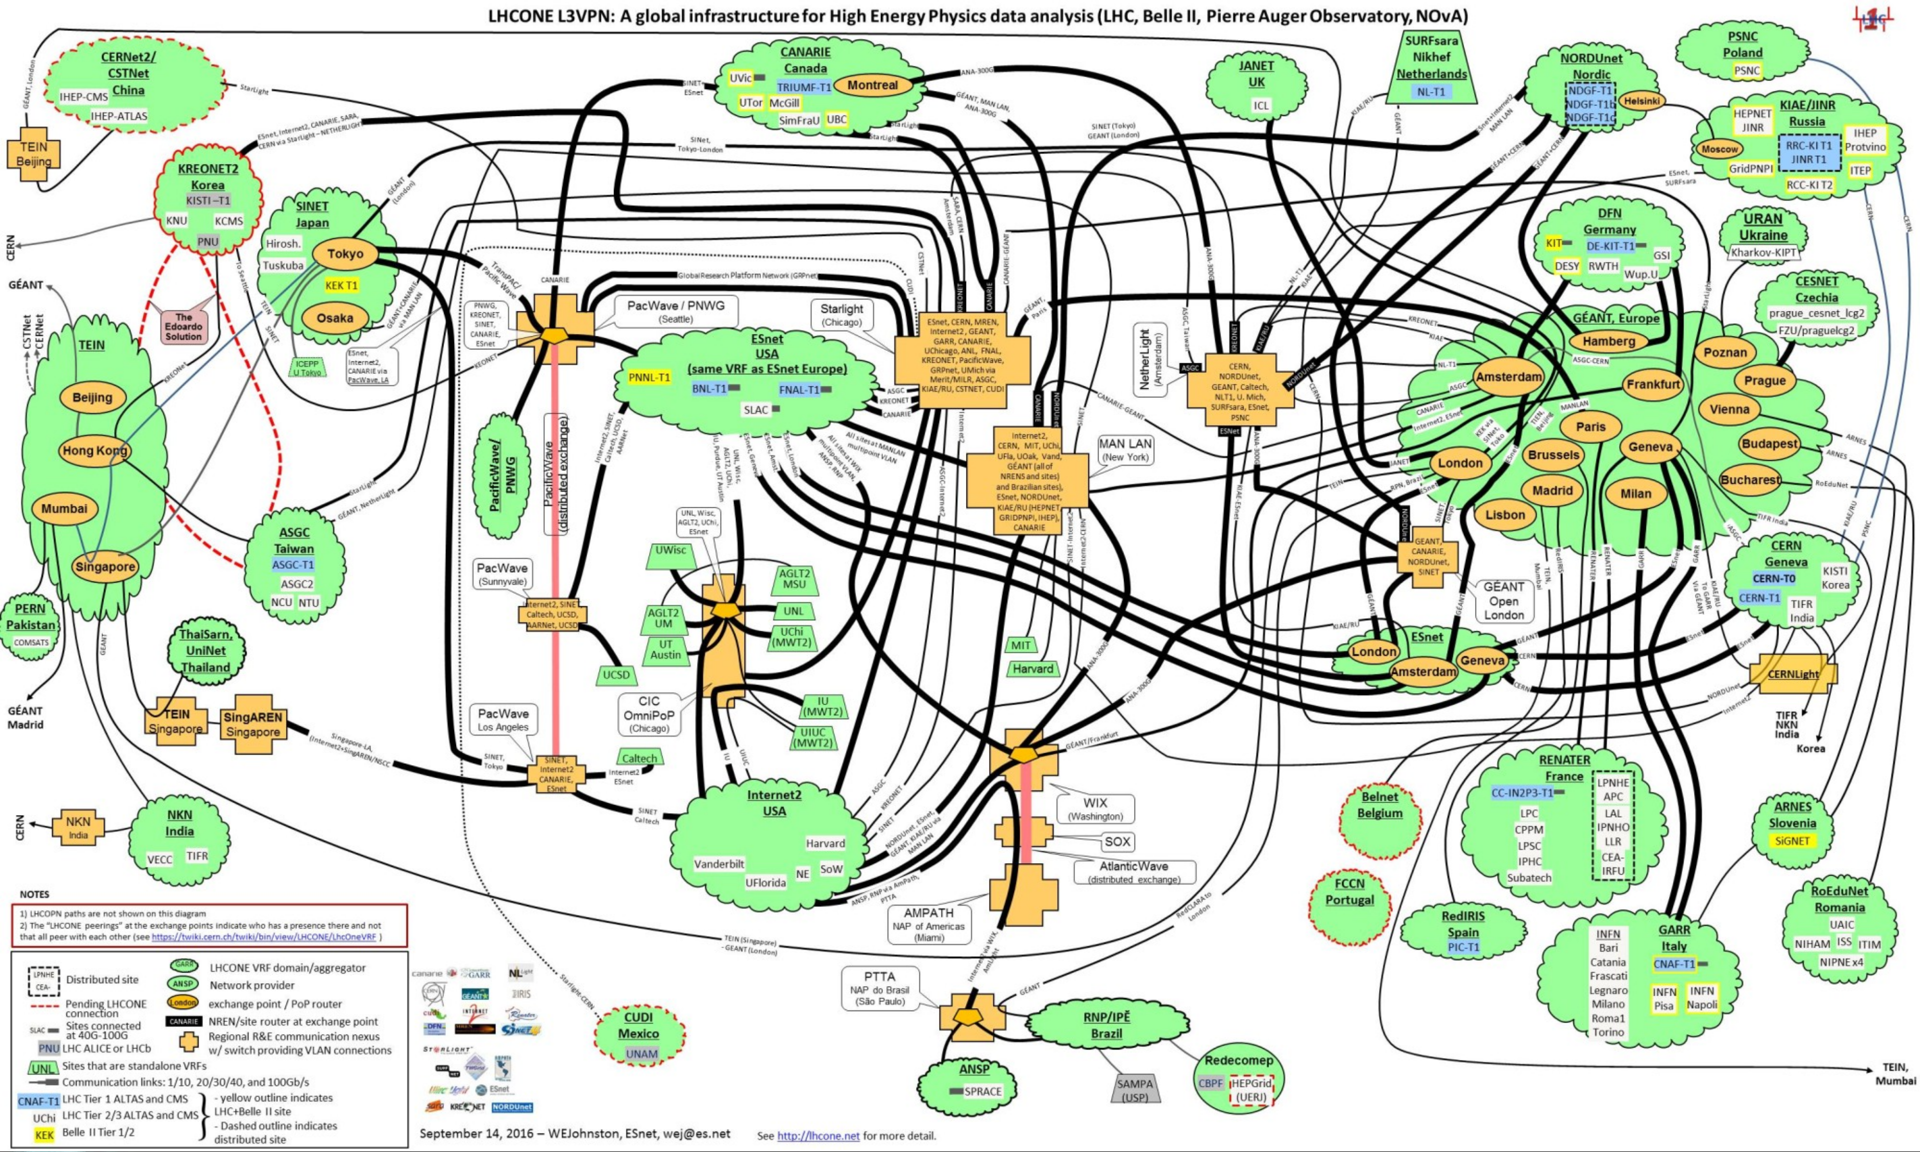

STARLIGHT SDX

StarLight International/National Communications Exchange Facility– "By Researchers For Researchers"

StarLight Is an Innovation Platform For Advanced Communications Services Architecture and Technologies, Including Experimental Testbeds Optimized For High-Performance Data Intensive Applications

Multiple 10GE+100 Gbps StarWave Multiple 10GEs Over Optics – World's "Largest" 10G/100G Exchange First of a Kind Enabling Interoperability At L1, L2, L3

View from StarLight

Abbott Hall, Northwestern University's Chicago Campus

iCAIR

STARLIGHT SDX

Global Network Science
The Global Lambda Integrated Facility As the Basis For Distributed SDN/SDX Testbeds & Prototypes

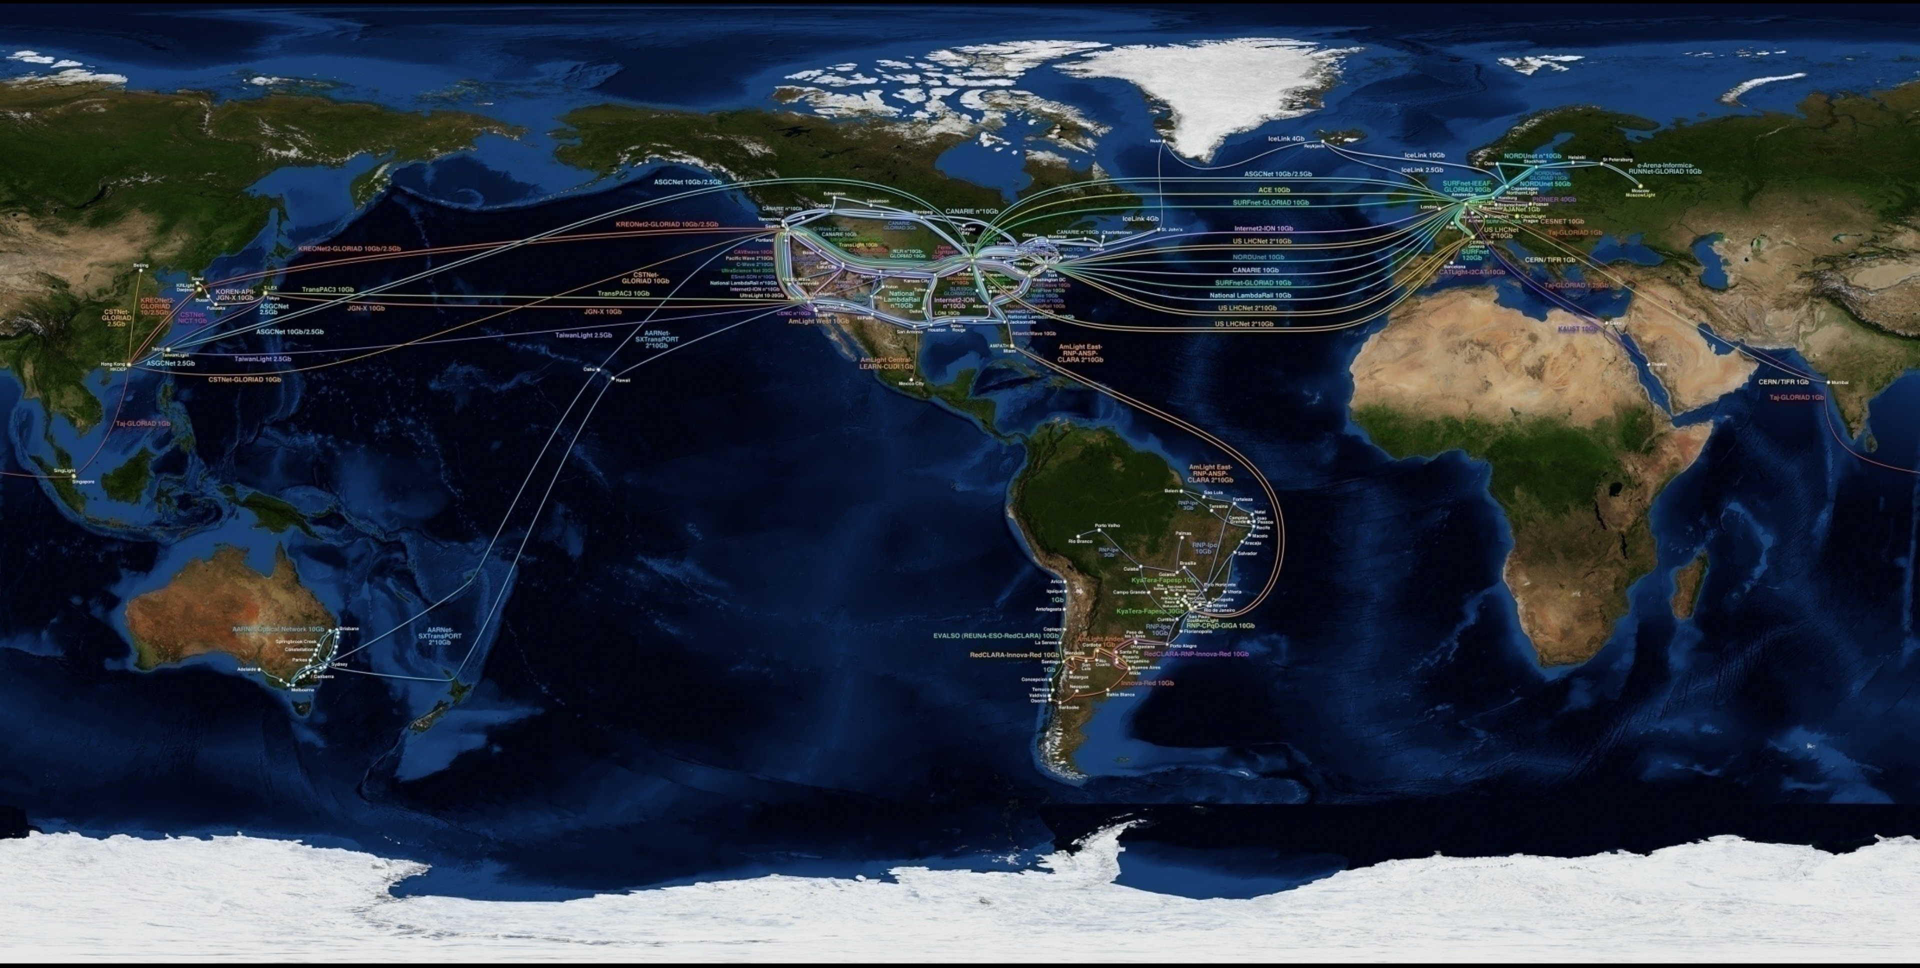

Automated GOLE Fabric

GLIF AutoGOLE Initiative

Source: John MacAuley

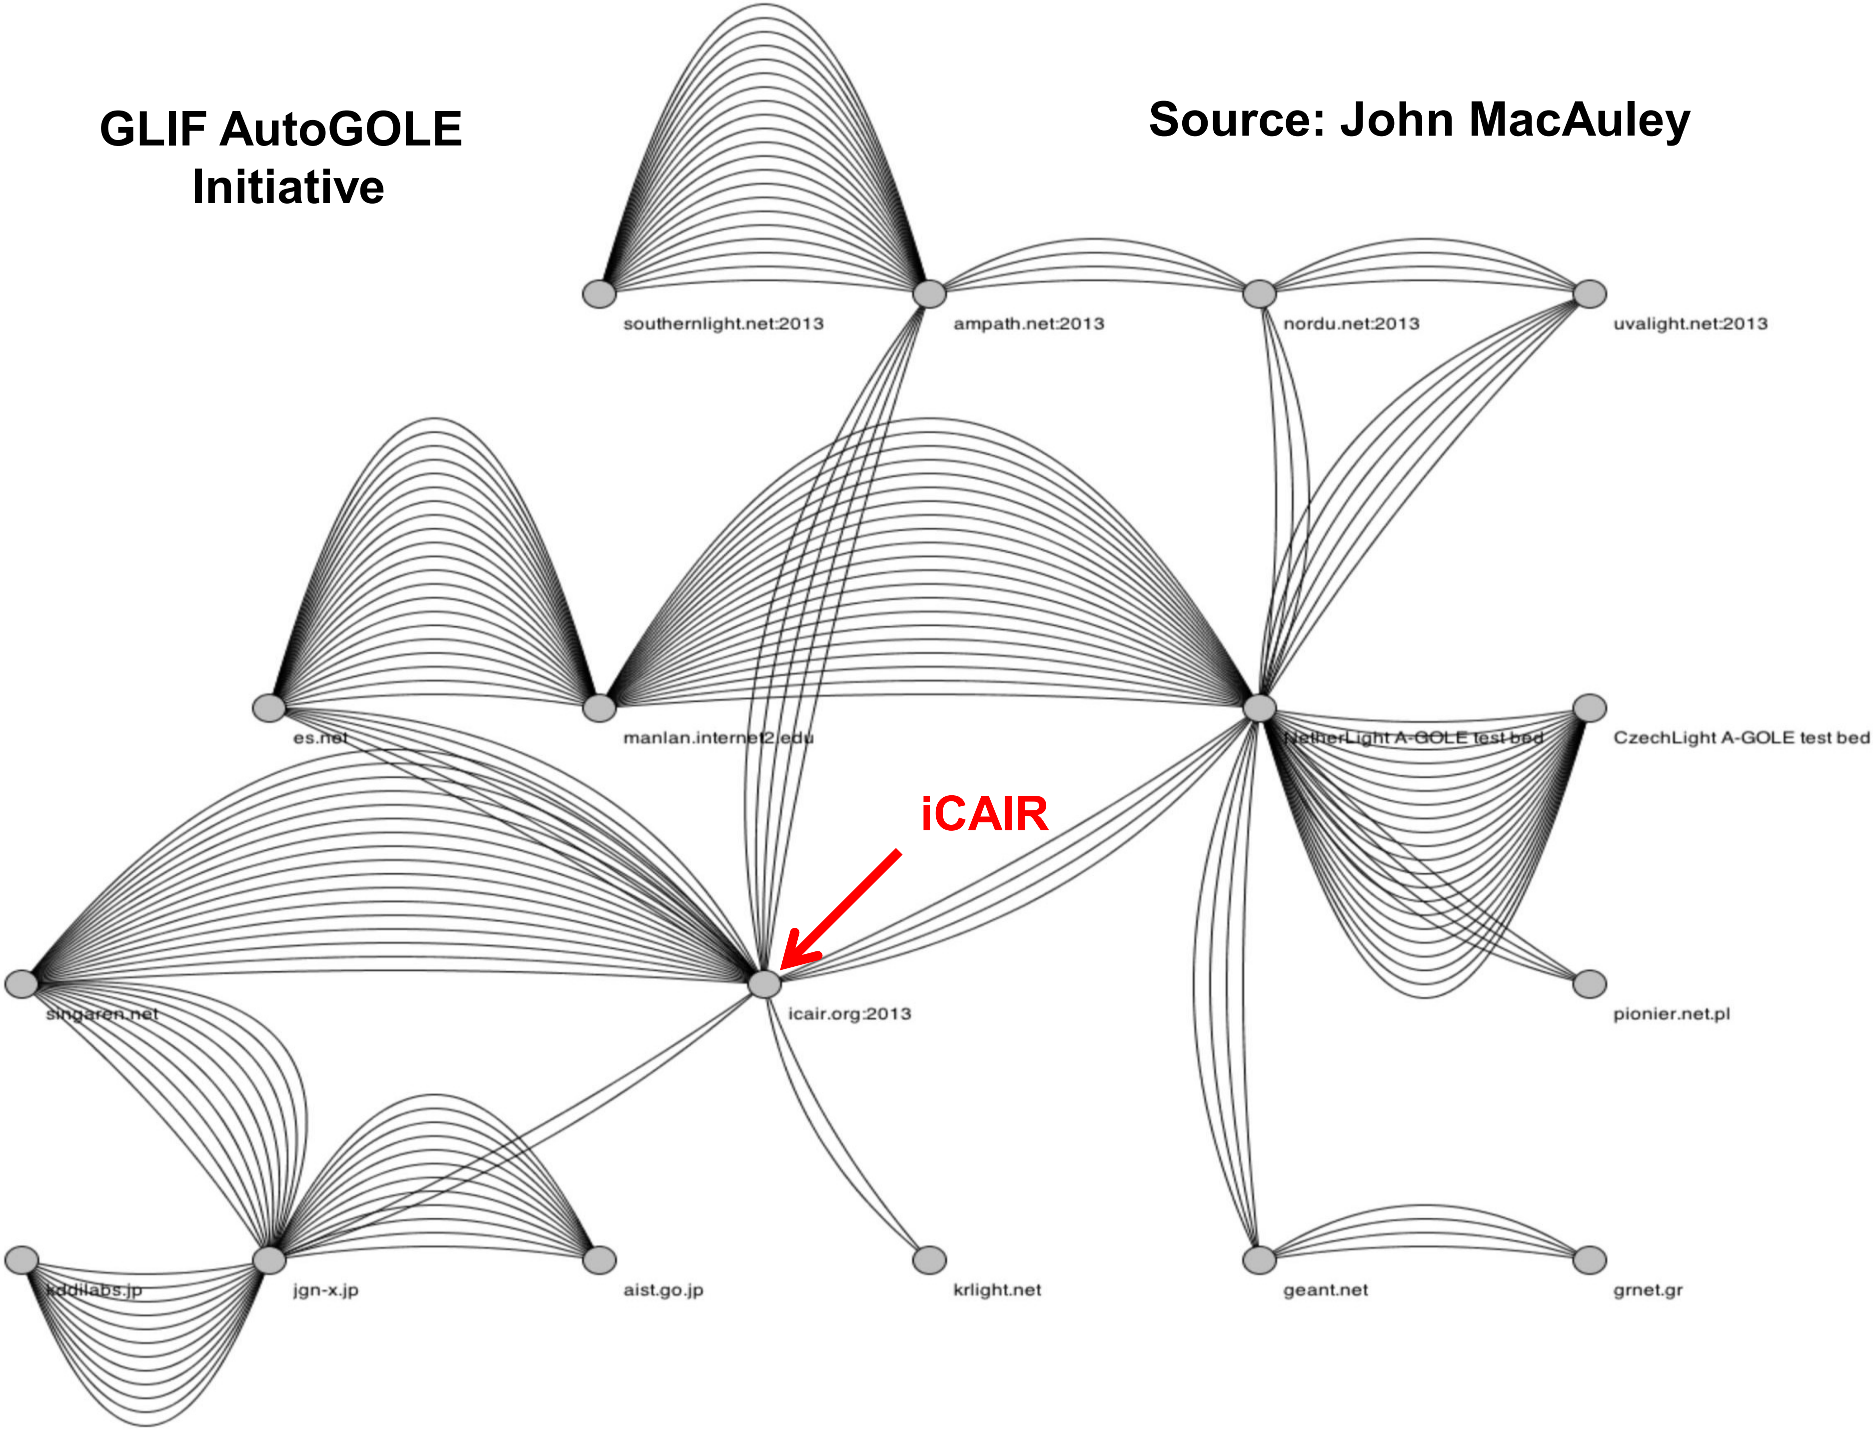

AutoGOLE Dashboard

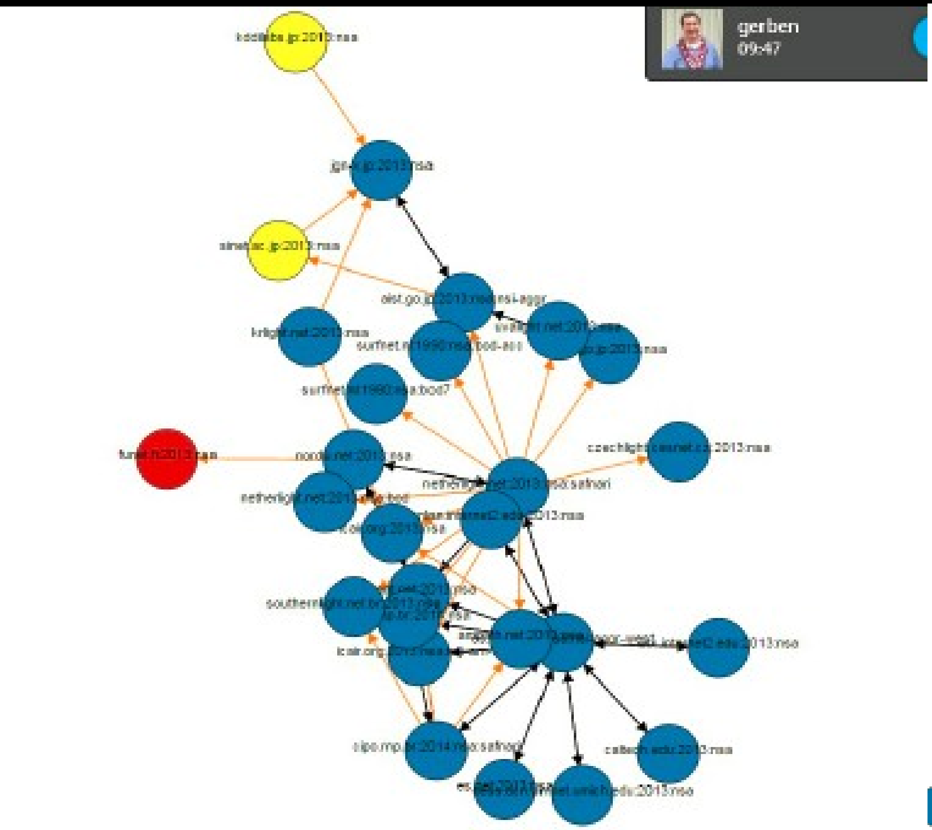

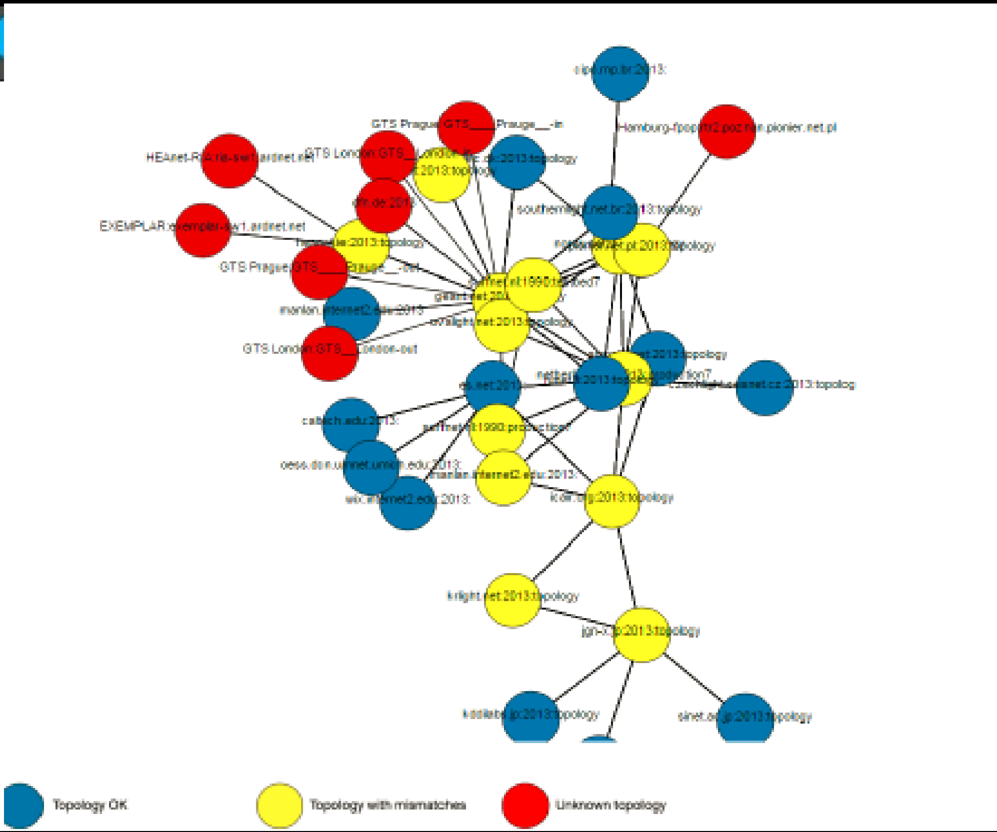

Control Plane

Data Plane

International SDN/OpenFlow Testbed

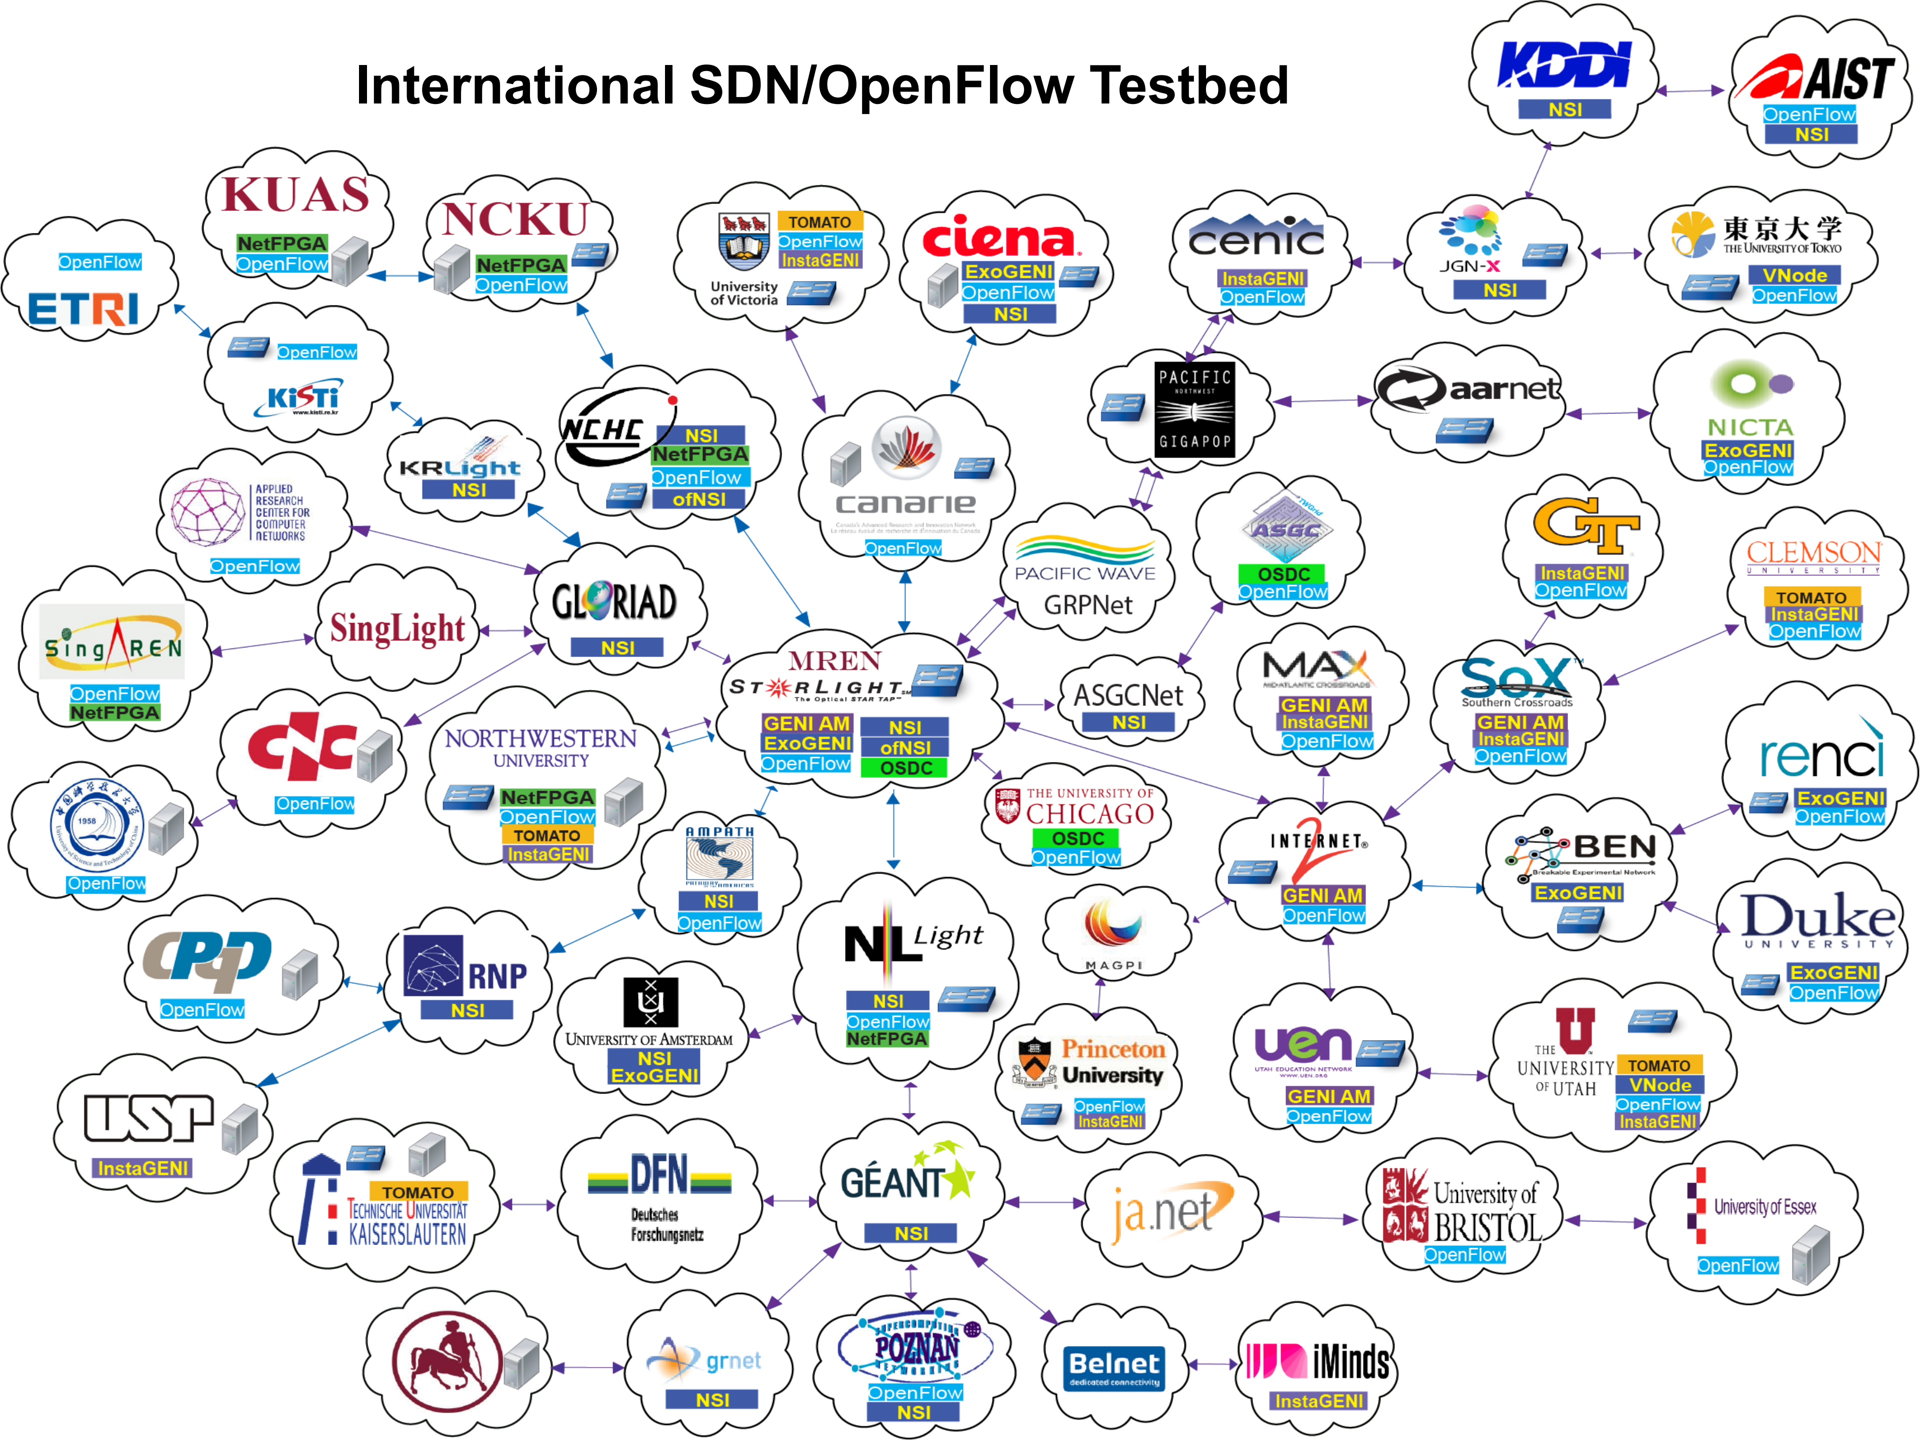

Forthcoming GENI Book: September 2016

Springer

springer.com

Chapter:
Creating a Worldwide Network For The Global Environment for Network Innovations (GENI) and Related Experimental Environments

1st ed. 2016, XVIII, 655 p. 216 illus., 183 illus. in color.

Printed book

R. **McGeer**, M. **Berman**, C. **Elliott**, R. **Ricci** (Eds.)

The GENI Book

- Provides a foundational overview of GENI's core architectural concepts
- Presents a detailed discussion of architecture and implementation
- Includes 24 chapters, divided into five sections, which outline GENI from precursors to architecture, development, applications, and then world federation
- Offers an extensive bibliography

This book, edited by four of the leaders of the National Science Foundation's Global Environment and Network Innovations (GENI) project, gives the reader a tour of the history, architecture, future, and applications of GENI. Built over the past decade by hundreds of leading computer scientists and engineers, GENI is a nationwide network used daily by thousands of computer scientists to explore the next Cloud and Internet and the applications and services they enable, which will transform our communities and our lives. Since by design it runs on existing computing and networking equipment and over the standard commodity Internet, it is poised for explosive growth and transformational impact over the next five years.

Benefits of SDN To Science

- SDN Not Only Allows Network Designers To Create a Much Wider Range of Services and Capabilities Than Can Be Provided With Traditional Networks, But They Also Enable:
  - a) More Comprehensive, Highly Granulated Views Into Network Capabilities and Resources Including Individual Data Flows
  - b) Many More Options For Control Over Those Resources, Including Distributed Control y Edge Applications services processes individuals
  - c) Dynamic Provisioning and Adjustment Options, Including Those That Are Automatic and Implemented In Real Time
  - d) Faster Implementations of many New and Enhanced Services
  - e) Enabling Applications, Edge Processes and ***Even Individuals*** To Directly Control Core Resources;
  - f) Substantially Improved Options For Creating Customizable Networks
  - g) Enhanced Operational Efficiency and Effectiveness.
  - h) Etc …

STARLIGHT SDX

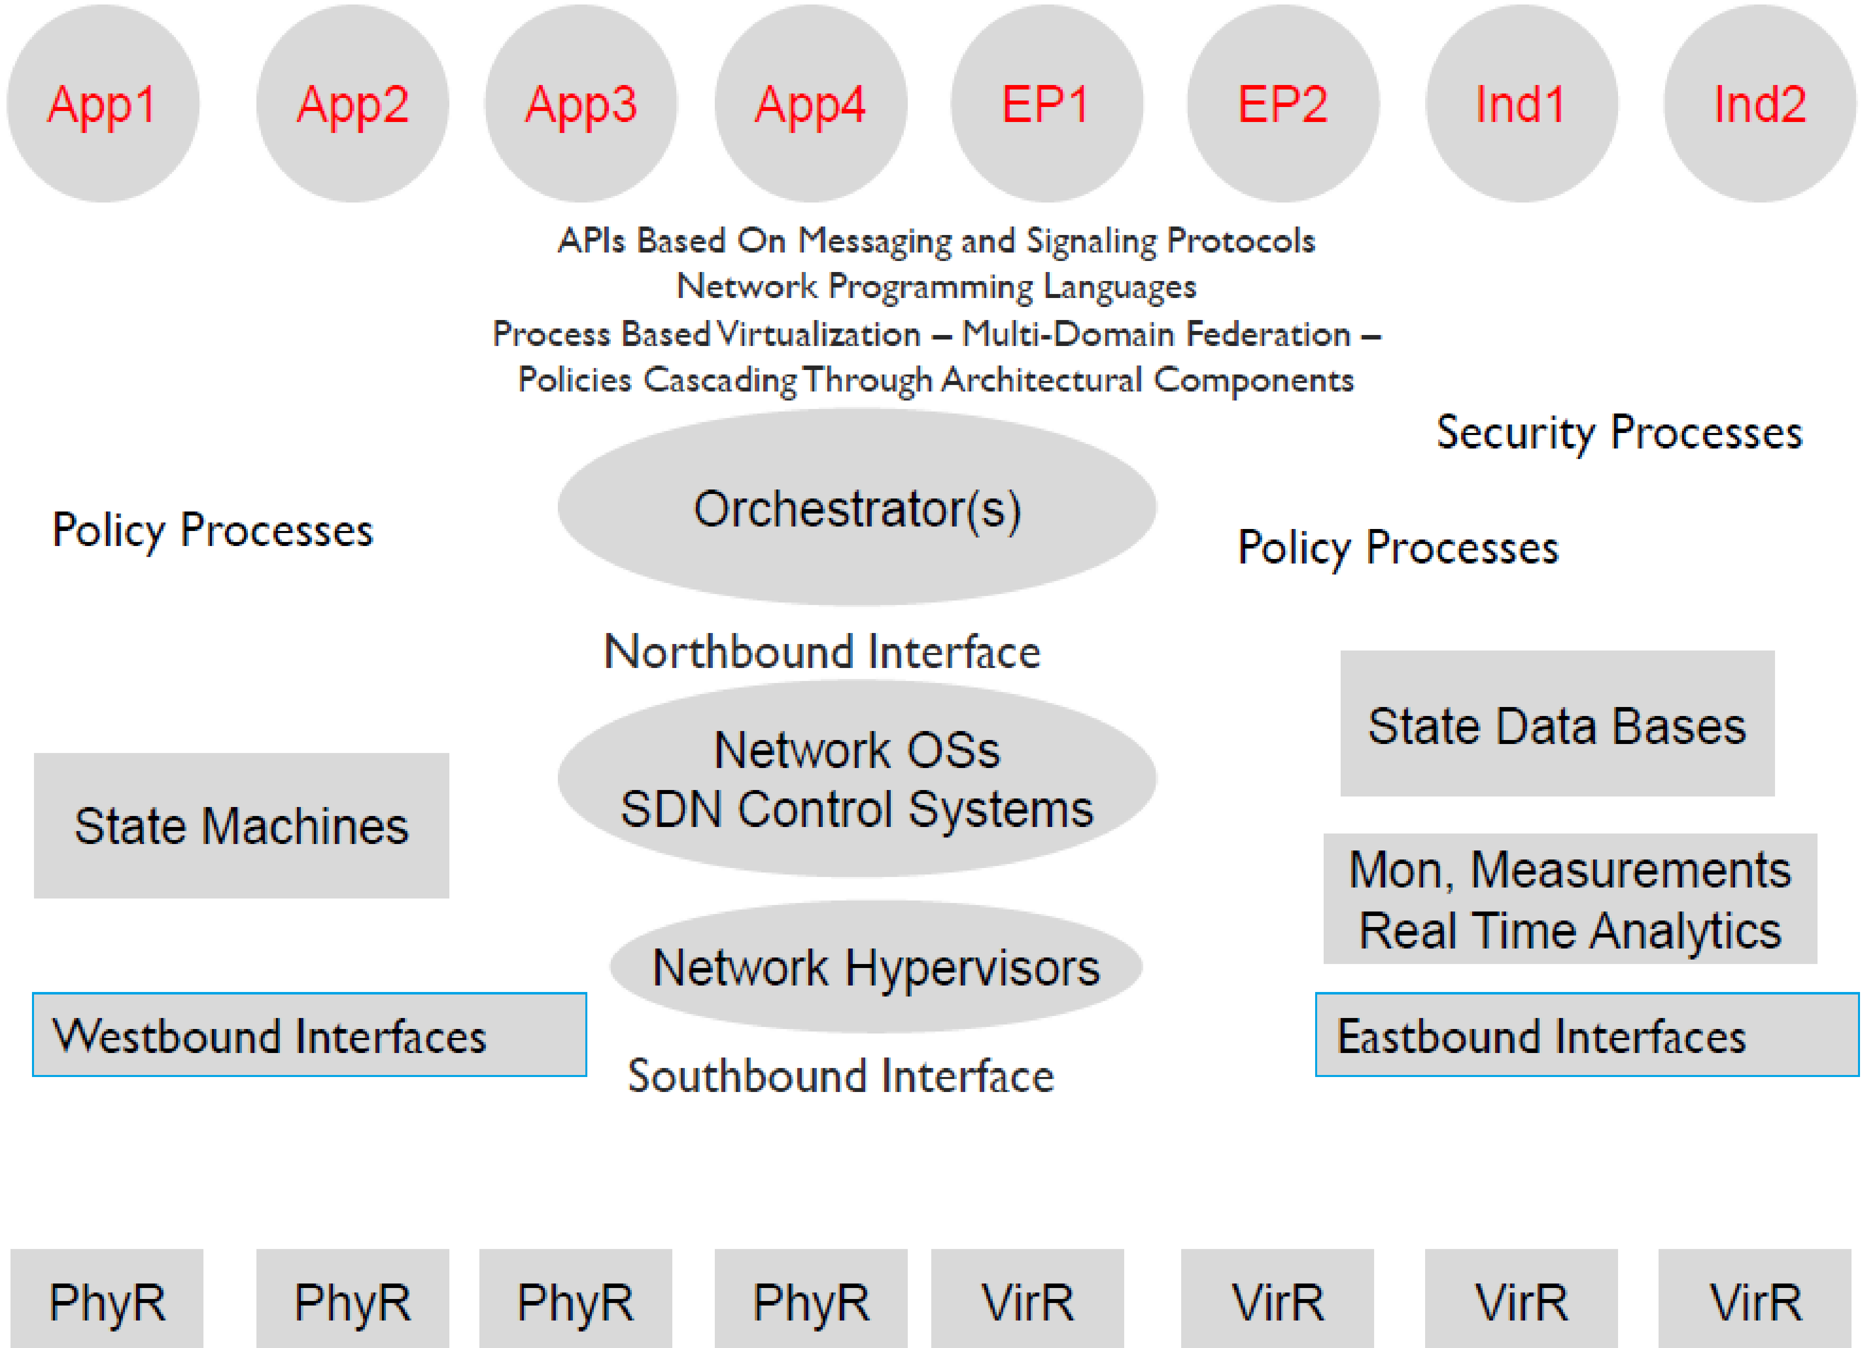
App1
App2
App3
App4
EP1
EP2
Ind1
Ind2
APIs Based On Messaging and Signaling Protocols
Network Programming Languages
Process Based Virtualization – Multi-Domain Federation –
Policies Cascading Through Architectural Components
Security Processes
Policy Processes
Orchestrator(s)
Policy Processes
Northbound Interface
State Data Bases
Network OSs
SDN Control Systems
State Machines
Mon, Measurements
Real Time Analytics
Network Hypervisors
Westbound Interfaces
Eastbound Interfaces
Southbound Interface
PhyR
PhyR
PhyR
PhyR
VirR
VirR
VirR
VirR

Software Defined Networking Exchanges (SDXs): Motivations

- With the Increasing Deployment of SDN In Production Networks, the Need for an SDN Exchange (SDX) Has Been Recognized.
- Current SDN Architecture/Protocols/Technologies Are Single Domain Centralized Controller Oriented
- SDXs Are Required To Interconnect Increasing Numbers Of SDN “Islands.”
- SDXs Provide Highly Granulated Views Into (and Control Over) All Flows Within the Exchange – i.e., Much Enhanced Traffic Engineering and Optimization.
- Options for Many New Types of Services and Capabilities, e.g., Encyption E2E, Ultra High Resolution Digital Media, Support for Data Intensive Science
- WH Office of Science and Technology Policy – Large Scale Science Instrumentation
- Democratization Of Exchange Facilities – Options for Edge Control Over Exchange Flows

Architectural Components 1

- Hybrid Networking Services (Multi-Service, Multi-Layer, Multi-Domain), Including Integration of OF and Non-OF Paths
- Multi Domain Resource Advertisement/Discovery/Signaling – Including Edge Signaling
- Support for Multi Domain Integrated Federated Path Controllers of Many Types
- Mechanisms For Topology Exchange Services
- Control and Network Resource APIs
- Network Programming Languages (e.g., P4, Frenetic)
- Abstraction Definitions
- AP/Service Signaling and Policy Bundling & Distribution
- New Types Of Primitives
- BGP Extensions and Substitutes
- NDL Schemas
- Orchestration Processes

Other Architectural Components 2

- North/South/East/West Bound Interfaces
- Network OSs
- Network Hypervisors
- State Information Data Bases
- Data Modeling Languages (e.g., YANG)
- Controller Federation Processes
- Hybrid Services/Services Federation/Services Chaining
- Granulated Resource Access (Policy Based), Including Through Edge Processes
- Foundation Resource Programmability
- Programmability for Large Scale Large Capacity Streams
- Data Plane Processes
- Network Function Virtualization (NFV)
- Measurements – Including Real Time Analytics
- Distributed Virtual NOC Operations

StarLight GENI L2 SDX

- Implementation Of Key Software and Hardware Components Of a Layer 2 SDN/OpenFlow Exchange (SDX) Between GENI Layer 2 Network Resources and Other Research Networks.
- Providing Tools For Experimenters To Request And Receive Resources from the Exchange That Are Fully Integrated With GENI Standard Interfaces, Such As the GENI Clearinghouse, the GENI API, GENI Stitching, While Maintains SDN Architectural Attributes
- Integrates GENI Tools & Experimenter Tools From Other Participating Networks
- Supports Open Exchange Policies
- Ensures Resilience/Reliability/Paths
- Ensures Capabilities for Intra and Inter-Domain Interoperability – Without Requiring Intermediaries
- Provides Monitoring and Measuring Capabilities
- Ensures Capabilities and Options For On-Going Enhancement, Expansion, Extendibility

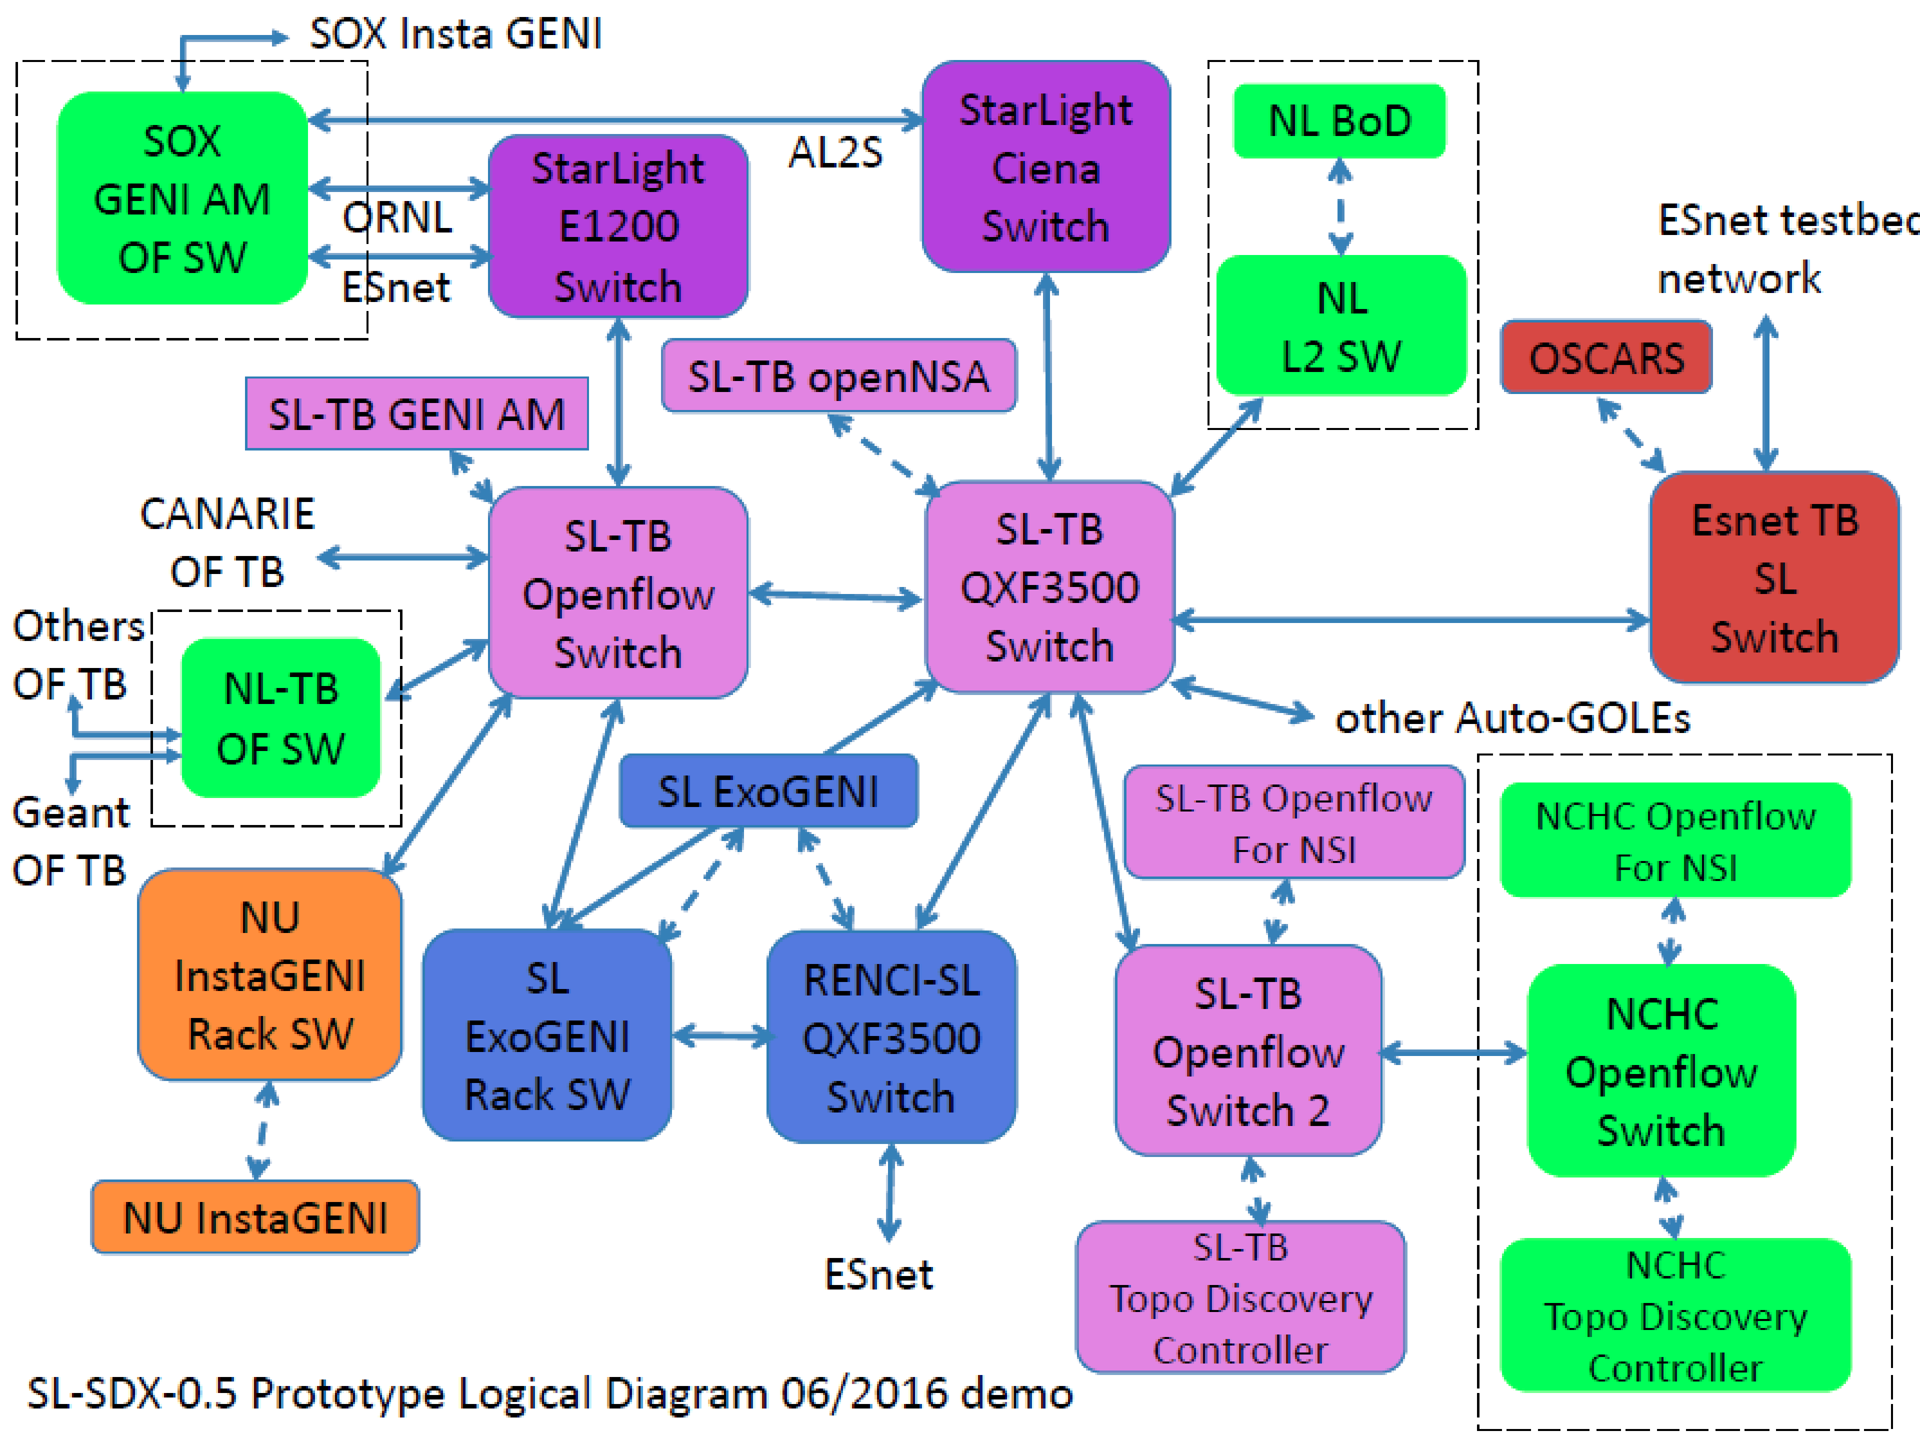
SOX Insta GENI
SOX GENI AM OF SW
ORNL
ESnet
AL2S
StarLight E1200 Switch
StarLight Ciena Switch
NL BoD
NL L2 SW
ESnet testbed network
OSCARS
SL-TB GENI AM
SL-TB openNSA
CANARIE OF TB
SL-TB Openflow Switch
SL-TB QXF3500 Switch
Esnet TB SL Switch
Others OF TB
NL-TB OF SW
Geant OF TB
other Auto-GOLEs
SL ExoGENI
SL-TB Openflow For NSI
NCHC Openflow For NSI
NU InstaGENI Rack SW
SL ExoGENI Rack SW
RENCI-SL QXF3500 Switch
SL-TB Openflow Switch 2
NCHC Openflow Switch
NU InstaGENI
ESnet
SL-TB Topo Discovery Controller
NCHC Topo Discovery Controller
SL-SDX-0.5 Prototype Logical Diagram 06/2016 demo

IRNC: RXP: StarLight SDX Key Participants

- PI Joe Mambretti, Director, International Center for Advanced Internet Research
  Northwestern University, Director, Metropolitan Research and Education Network
  Director, StarLight International/National Communications Exchange Facility
- Co-PI Tom DeFanti, Research Scientist, (tdefanti@soe.ucsd.edu)
  California Institute for Telecommunications and Information Technology (Calit2),
- University of California, San Diego, Co-Director, StarLight
- Co-PI Maxine Brown, Director
  Electronic Visualization Laboratory, University of Illinois at Chicago, Co-Director StarLight
- Co-PI Jim Chen, Associate Director, International Center for Advanced Internet Research
- Northwestern University
- Senior Personnel
- John Graham, Senior Development Engineer, Calit2, UCSD
- Phil Papadopoulos, Program Director, UC Computing Systems, San Diego Supercomputer Center, UCSD, Associate Research Professor (Adjunct) Computer Science UCSD
- Tom Hutton, Network Architect, UC San Diego Supercomputing Center, SDSC/Calit2 QI
- Larry Smarr, founding Director of Calit2) a UC San Diego/UC Irvine partnership, Harry E. Gruber Professor in Computer Science and Engineering (CSE) at UCSD's Jacobs School.
- Linda Winkler, Senior Network Engineer, Math and Computer Science Division, Argonne National Laboratory, Senior Network Engineer, StarLight Facility, Technical Director, MREN
- John Hess, Director, Network Engineering and Design, CENIC
- Also, Other Members of the StarLight Consortium, Multi National and International Partners

STARLIGHT SDX

E2E Services Based On Open Architecture For Petascale Sciences (OAPS)

- Integrating Science Workflows With Foundation Infrastructure Workflows (Orchestration)
- Providing Built-In Preconfigured Examples/Templates To Establish Infrastructure Foundation Workflows
- Providing Zero-Touch "Playbooks" For Different Segments of Infrastructure Foundation Workflows After Running The 1st Suite
- Supporting Interactive Control Over Running Workflows
- Providing Portability for Different Infrastructure Foundation Workflows
- Providing Capabilities for Experiment Reproducibility
- Providing Options For Real Time Scientific Visualization
- With Standard Tools

Network as ~~Infrastructure~~ *Instrument*

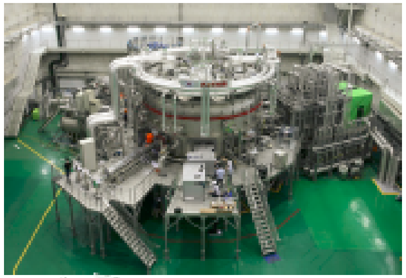

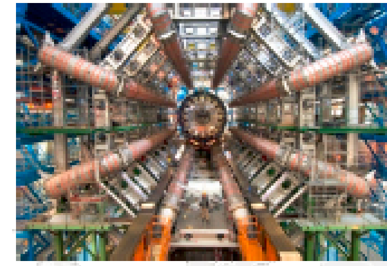

Implementing NSI in production – evaluating SDN

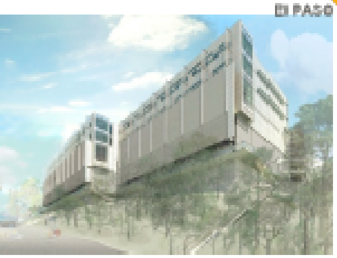

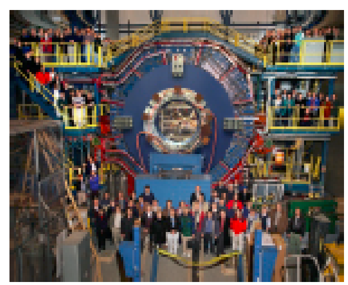

<u>ESnet Vision</u>: Scientific progress will be **completely unconstrained** by the physical location of instruments, people, computational resources, or data.

ESnet

SC15 SDN/SDX/SDI 100 Gbps Demonstrations

- **iCAIR Participated in (Or Led) ~ 20 Sets of 100 Gbps Demonstrations at SC15, Almost All Which Involved Elements of SDN/SDX/SDI**
- ***What's New=> Using Orchestrated SDX Services To Implement and Control WAN "Superchannels," In Part Enabled BY DTNs – Highly Scalable Dynamic Provisioning – A Scalability Not Possible On Today's Networks***

ESnet
ENERGY SCIENCES NETWORK

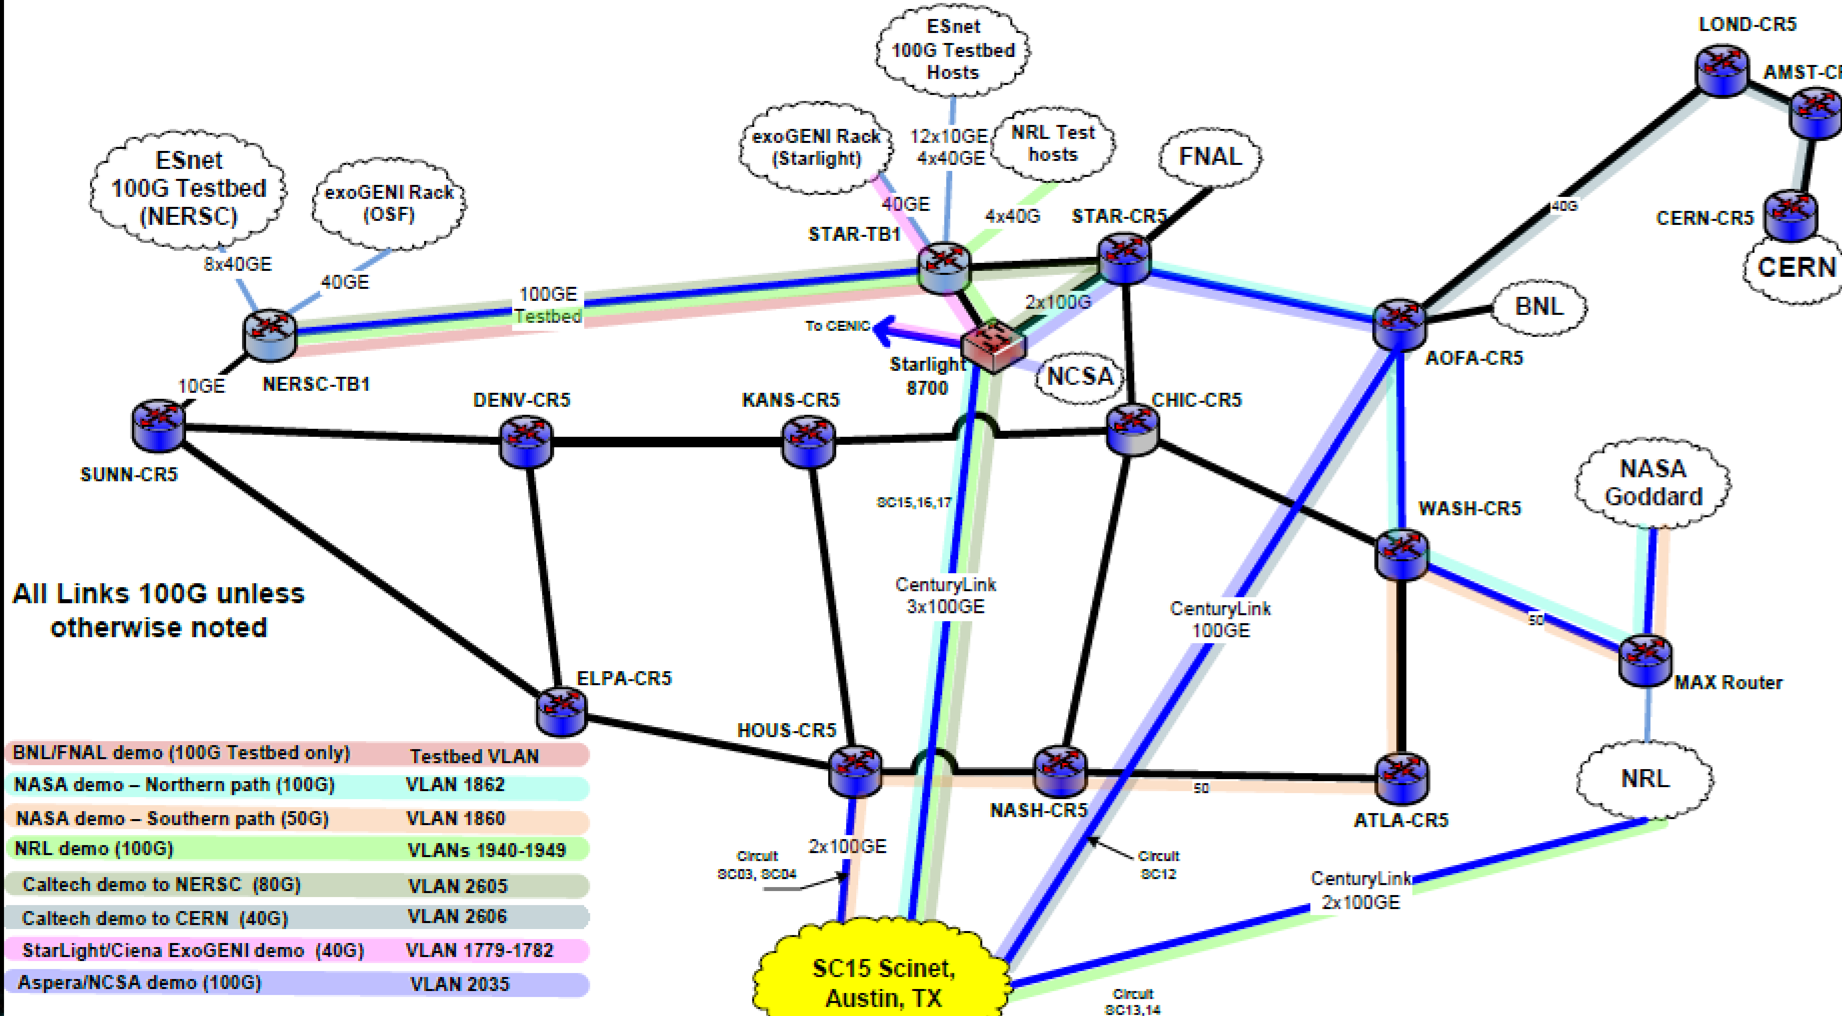

Other Demos 10G or Less:

- ESnet/RENCI demo: NERSC to ANL
- ESnet ENOS Demo: WASH, AMST, CERN
- ANL QoS Demo: DENV, ATLA

SC15 demos – ESnet

Brian Tierney, ESnet 11/6/2015

FILENAME SC15-DEMOS-V9.VSD

ESnet

ENERGY SCIENCES NETWORK

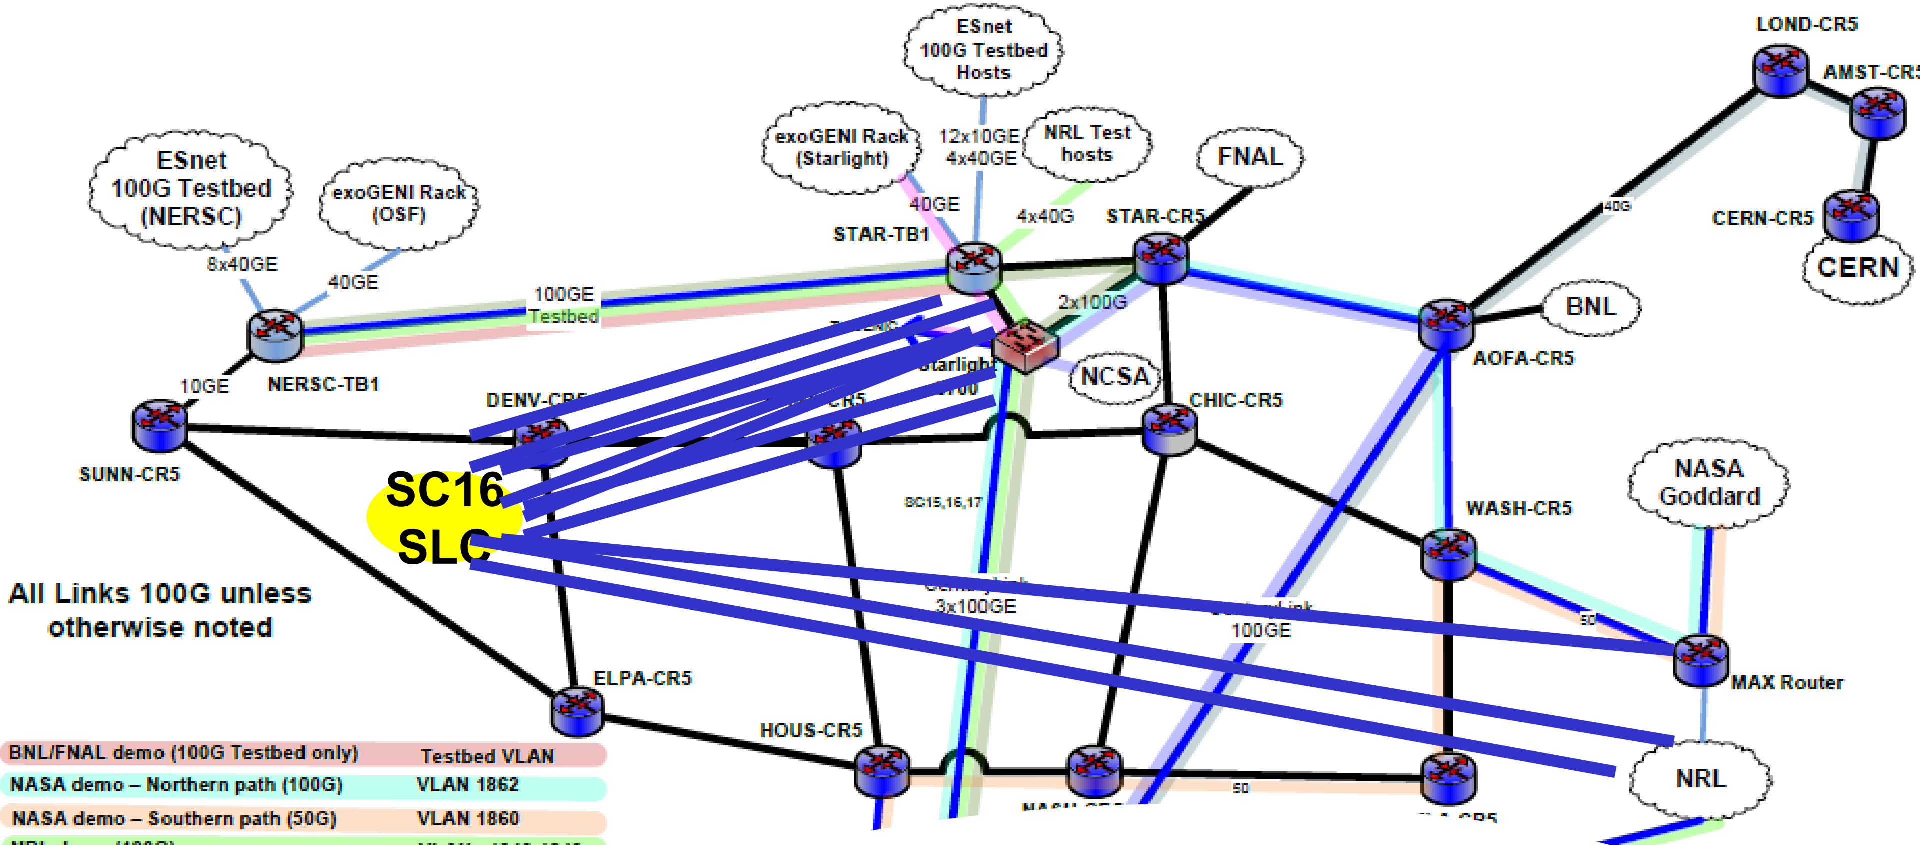

BNL/FNAL demo (100G Testbed only)	Testbed VLAN
NASA demo – Northern path (100G)	VLAN 1862
NASA demo – Southern path (50G)	VLAN 1860
NRL demo (100G)	VLANs 1940-1949
Caltech demo to NERSC (80G)	VLAN 2605
Caltech demo to CERN (40G)	VLAN 2606
StarLight/Ciena ExoGENI demo (40G)	VLAN 1779-1782
Aspera/NCSA demo (100G)	VLAN 2035

5*100 Gbps From StarLight
3* 100 Gbps from Wash DC

Other Demos 10G or Less:

- ESnet/RENCI demo: NERSC to ANL
- ESnet ENOS Demo: WASH, AMST, CERN
- ANL QoS Demo: DENV, ATLA

SC15 demos – ESnet
Brian Tierney, ESnet 11/6/2015
FILENAME SC15-DEMOS-V9.VSD

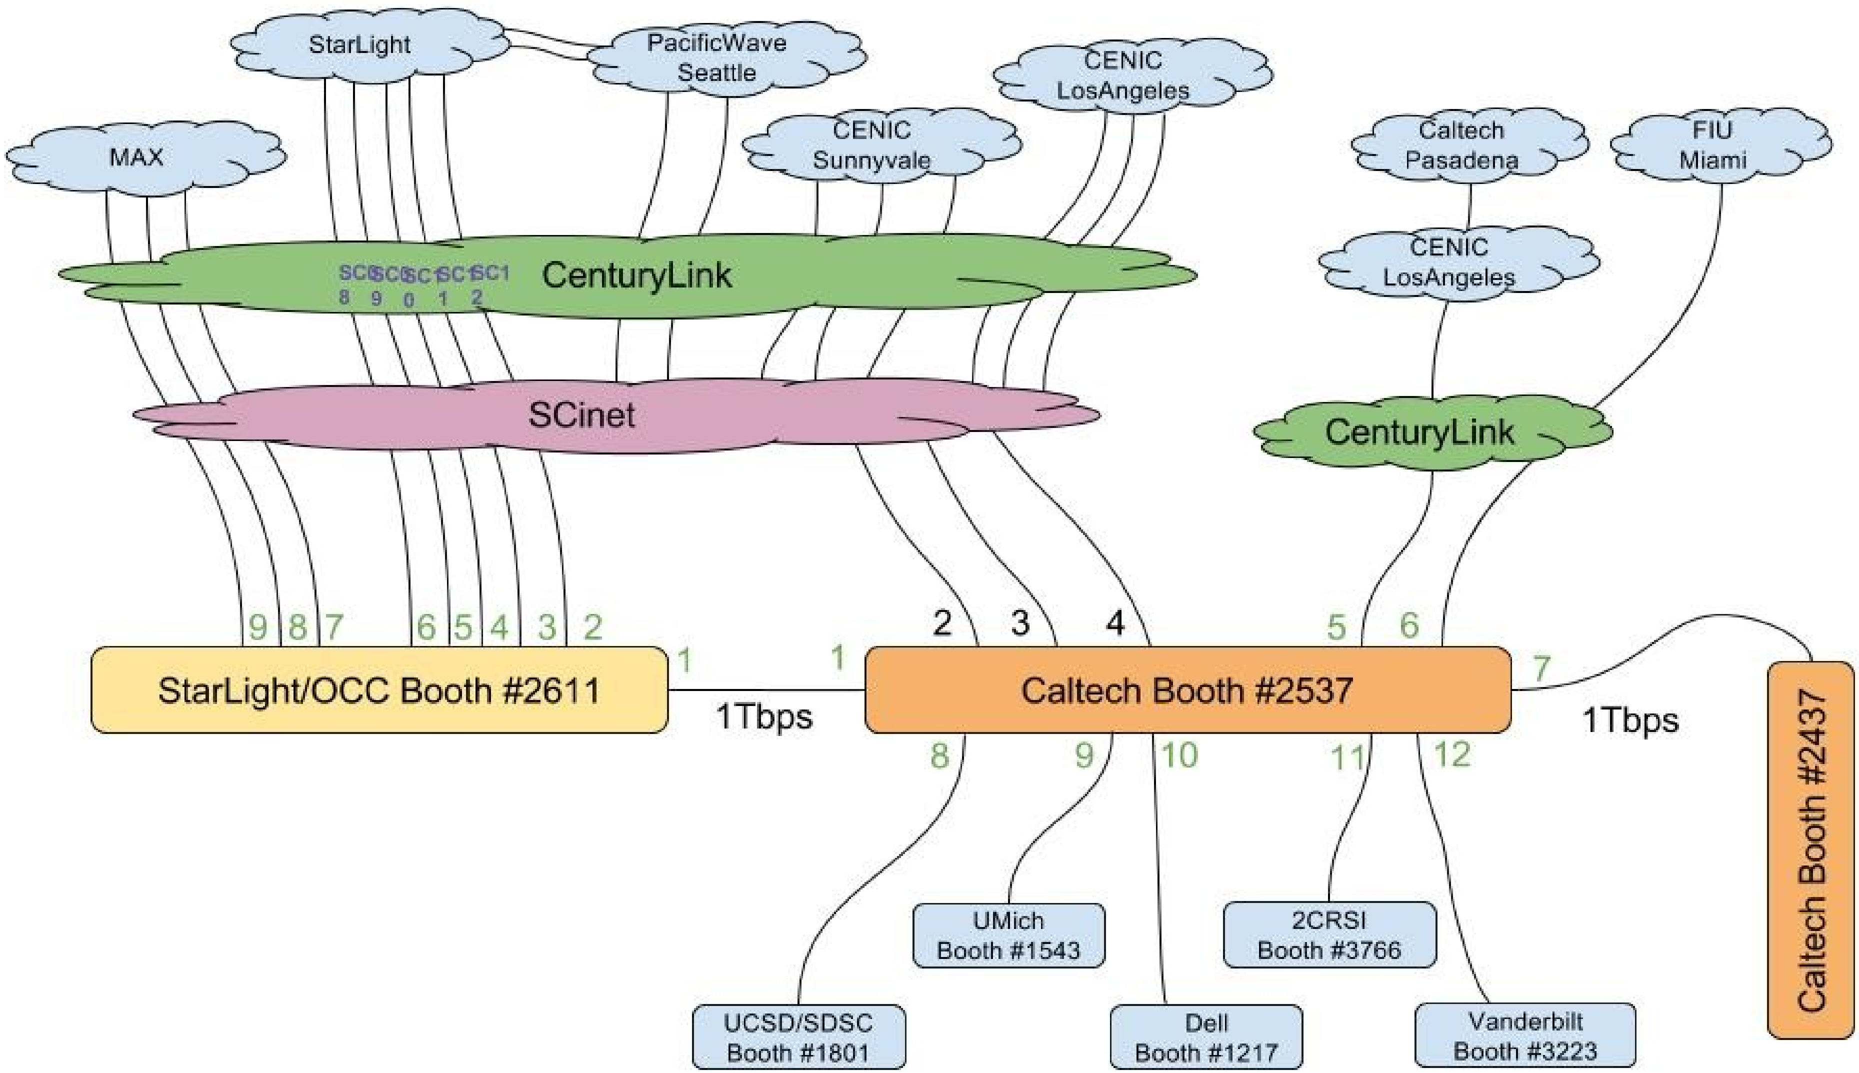

Image Developed By azher Muhgal

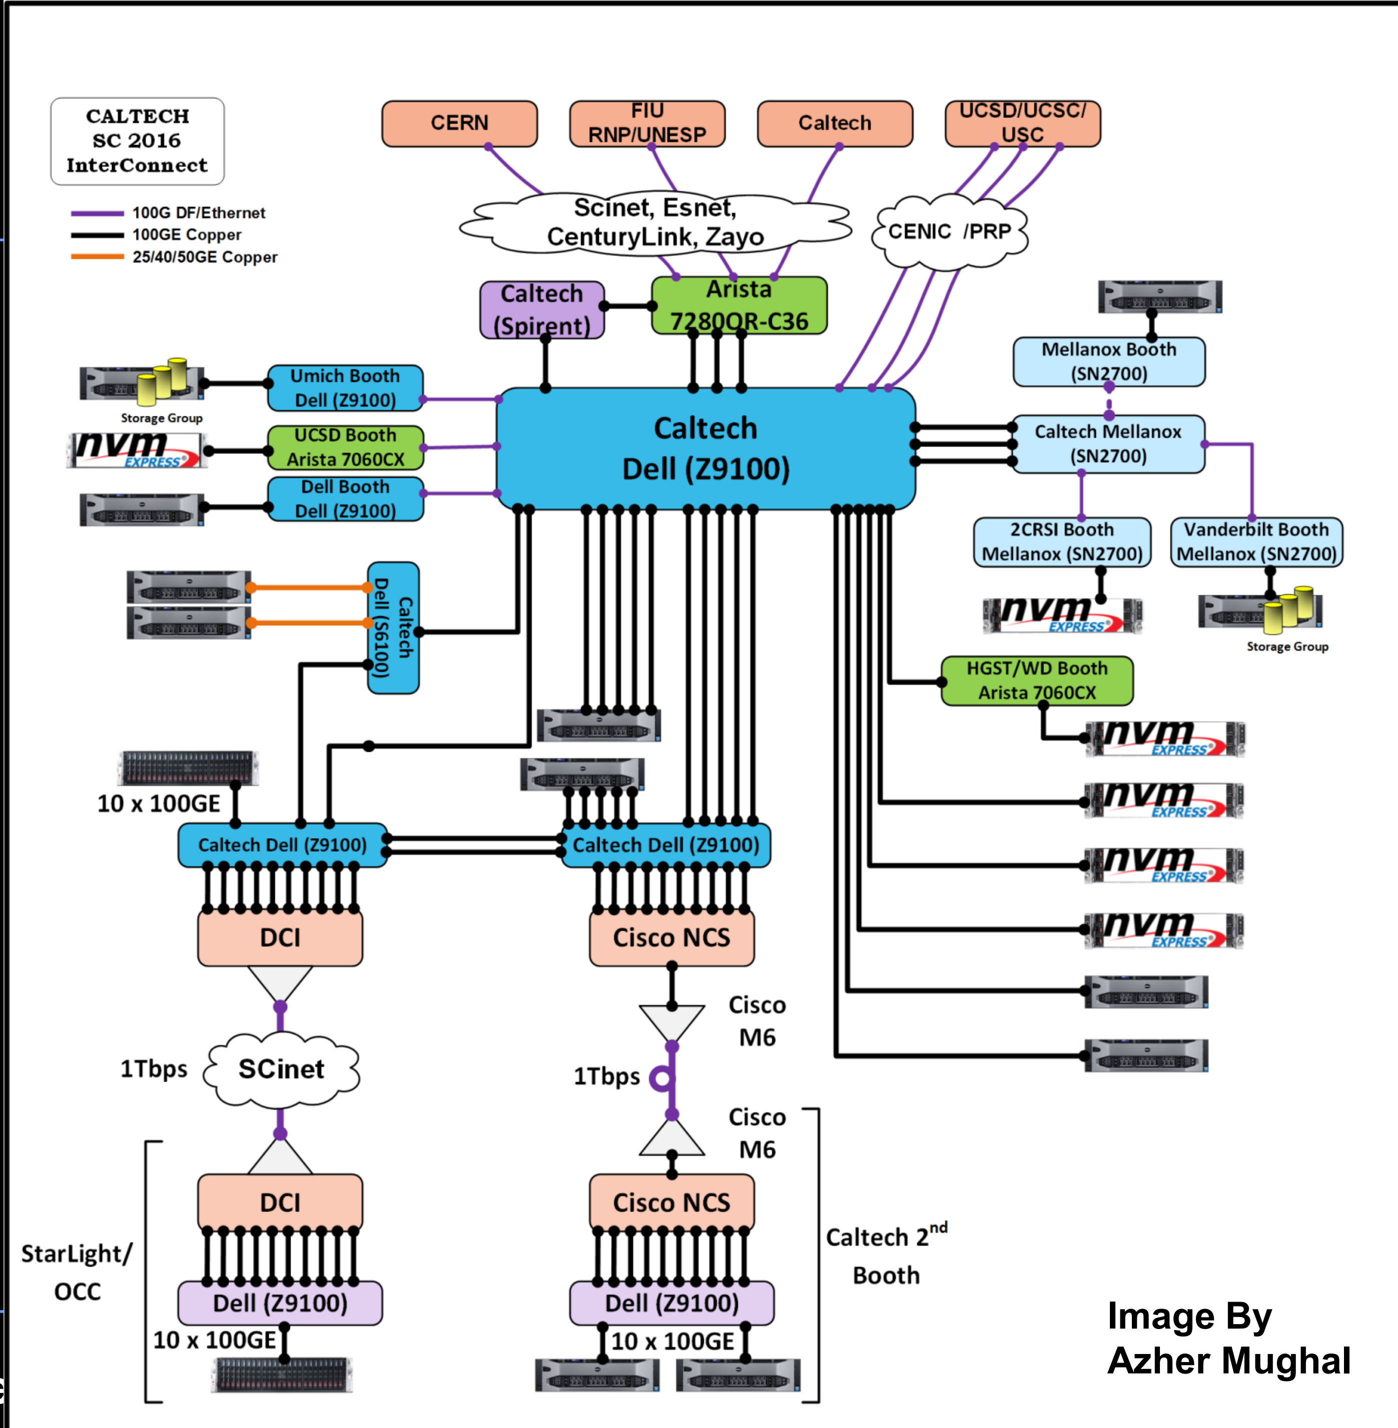

Image By
Azher Mughal

De

PetaTrans: Petascale Science Data Transfer

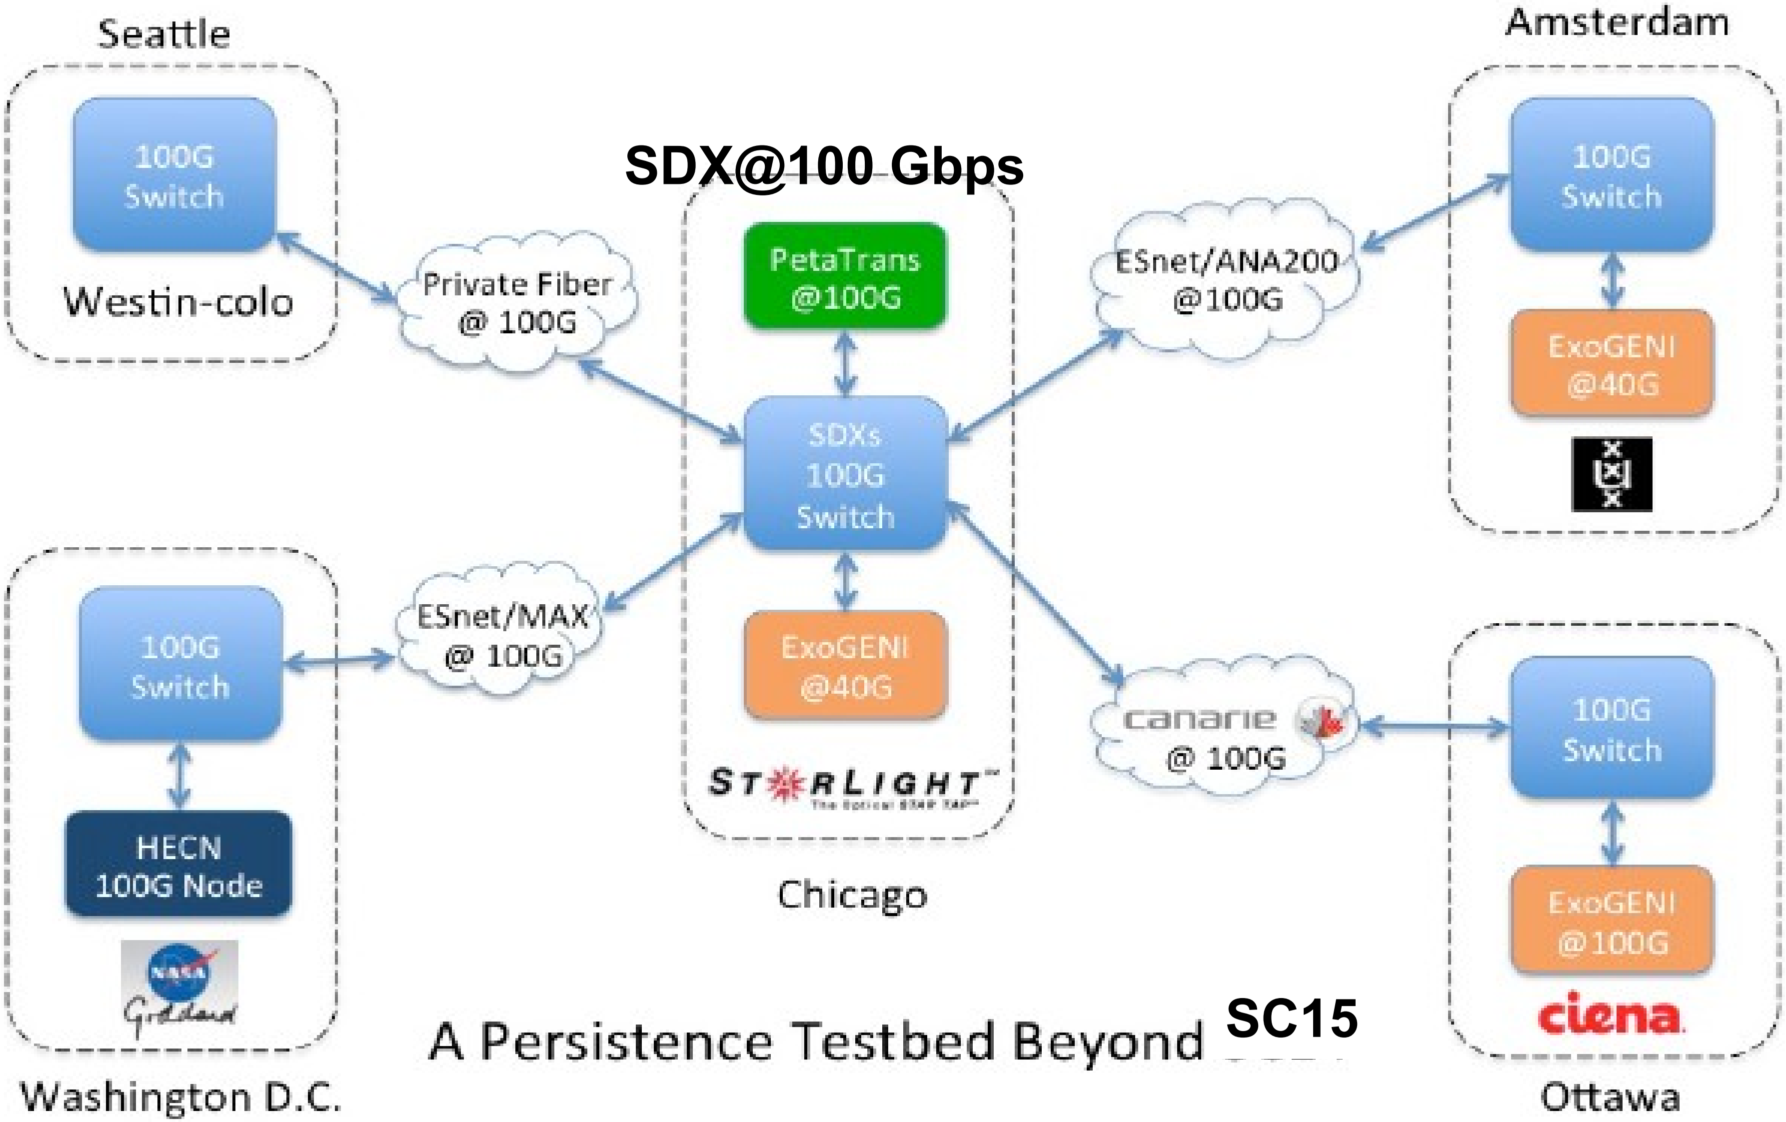

Global Software-Defined Dynamic Circuits for Data Intensive Science
(PhEDEx – ANSE – PANDA - OpenDayLight)

SC15 Demonstrations of 100 Gbps Disk-to-Disk WAN File Transfer Performance via SDX and 100G FW

An SC15 Collaborative Initiative Among NASA and Several Partners

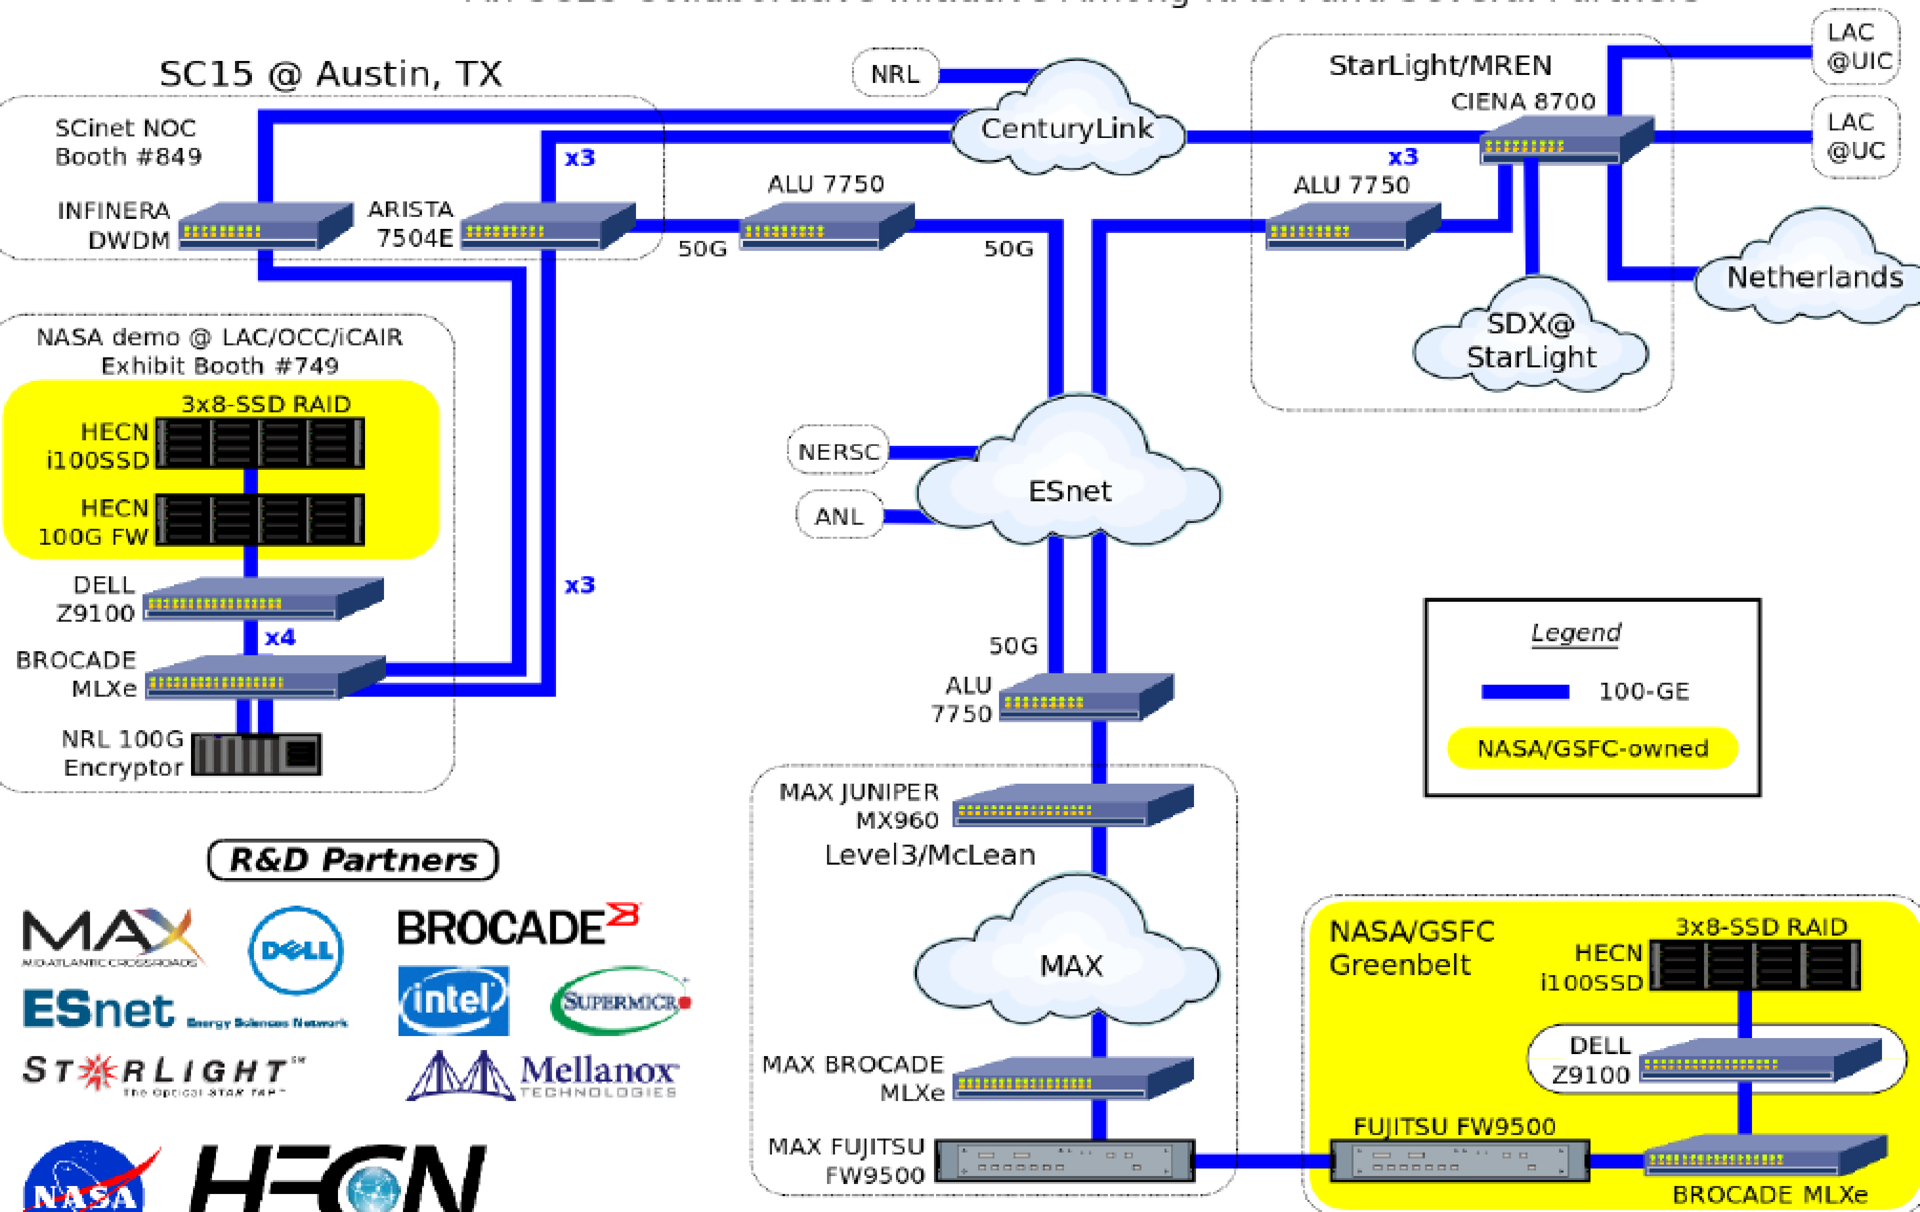

NASA/GSFC High End Computer Networking (HECN) Team
Diagram by Bill Fink / Paul Lang - 11/4/2015

SC16

Demonstrations of 200 Gbps Disk-to-Disk WAN File Transfers using Parallelism across NVMe Drives

An SC16 Collaborative Initiative Among NASA and Several Partners

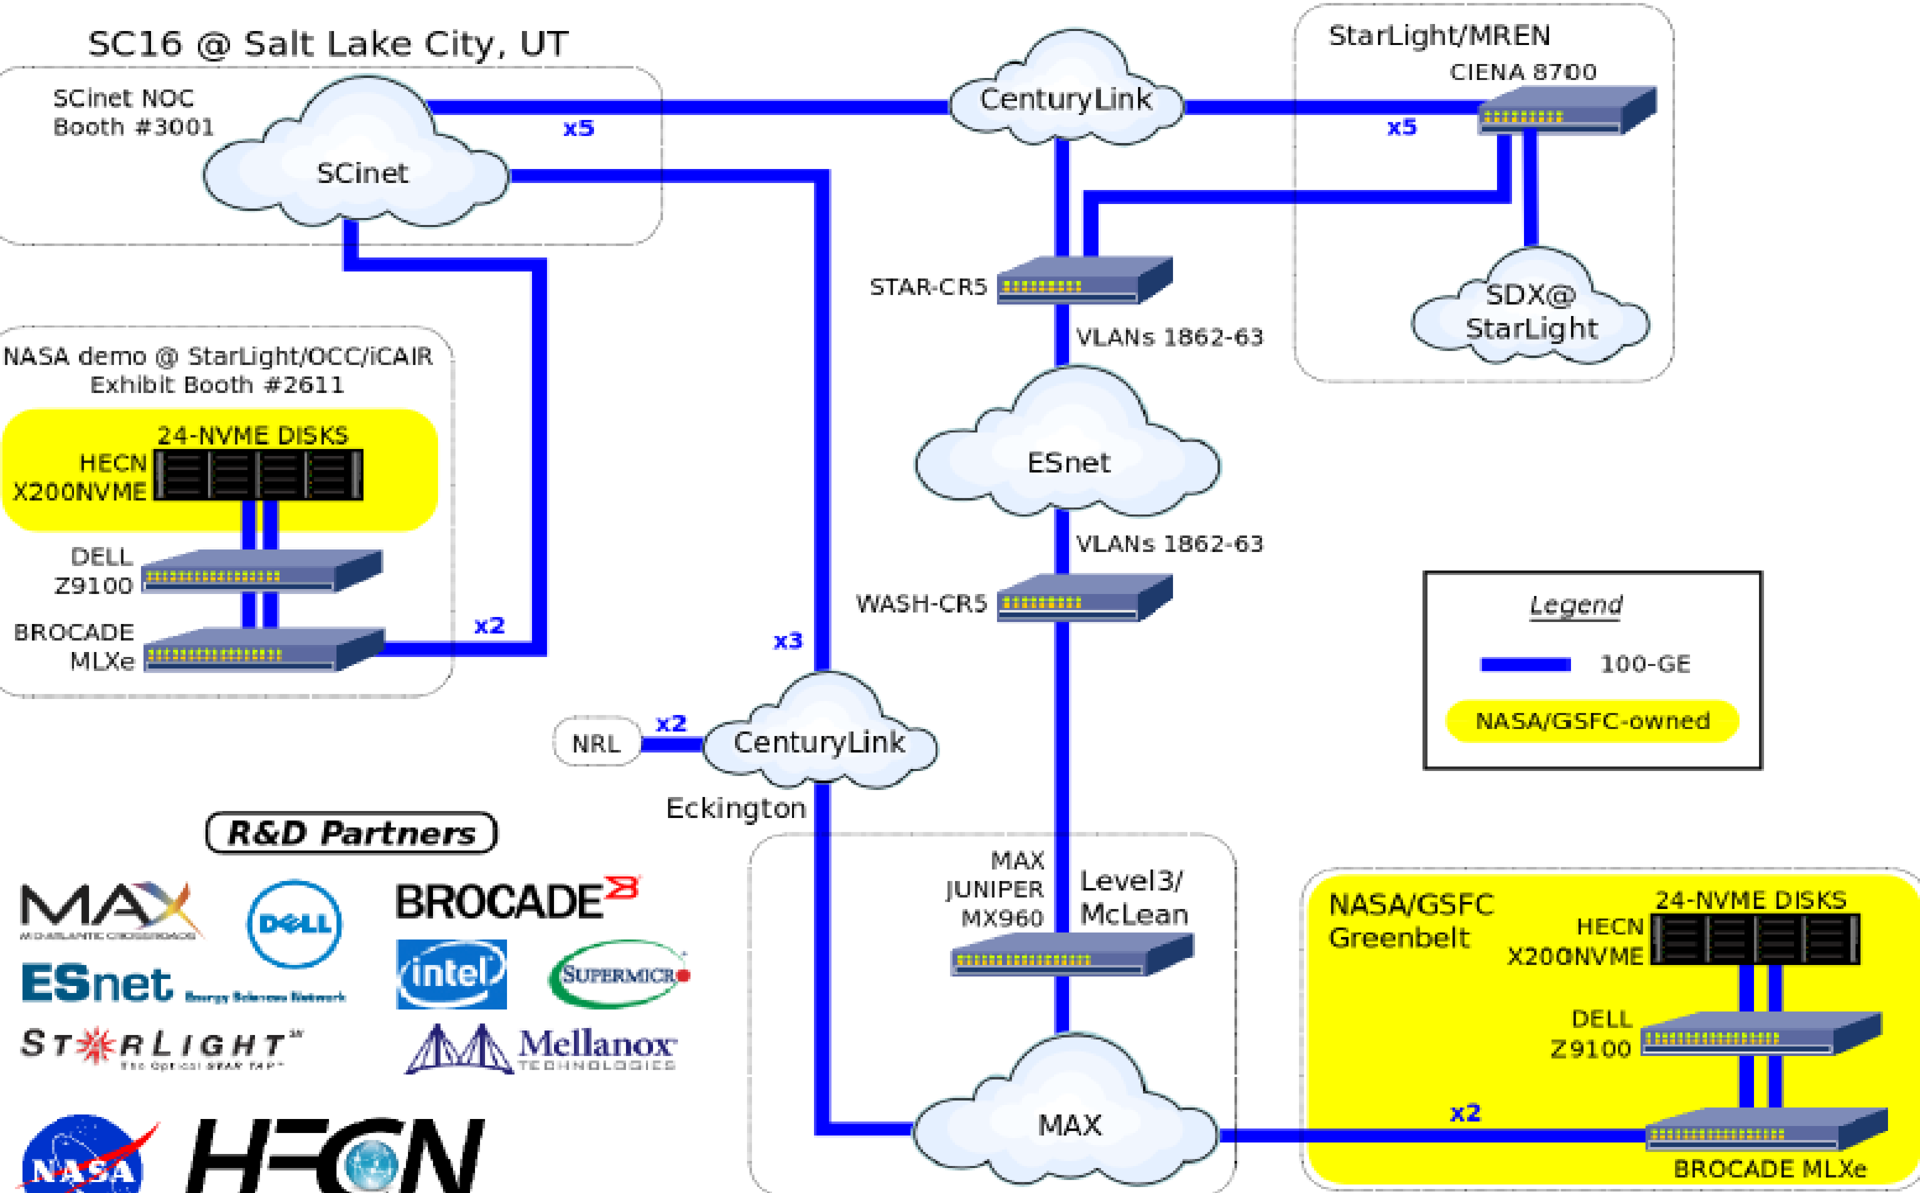

NASA/GSFC High End Computer Networking (HECN) Team
Diagram by Bill Fink / Paul Lang - 8/25/2016

Dynamic Remote I/O Network

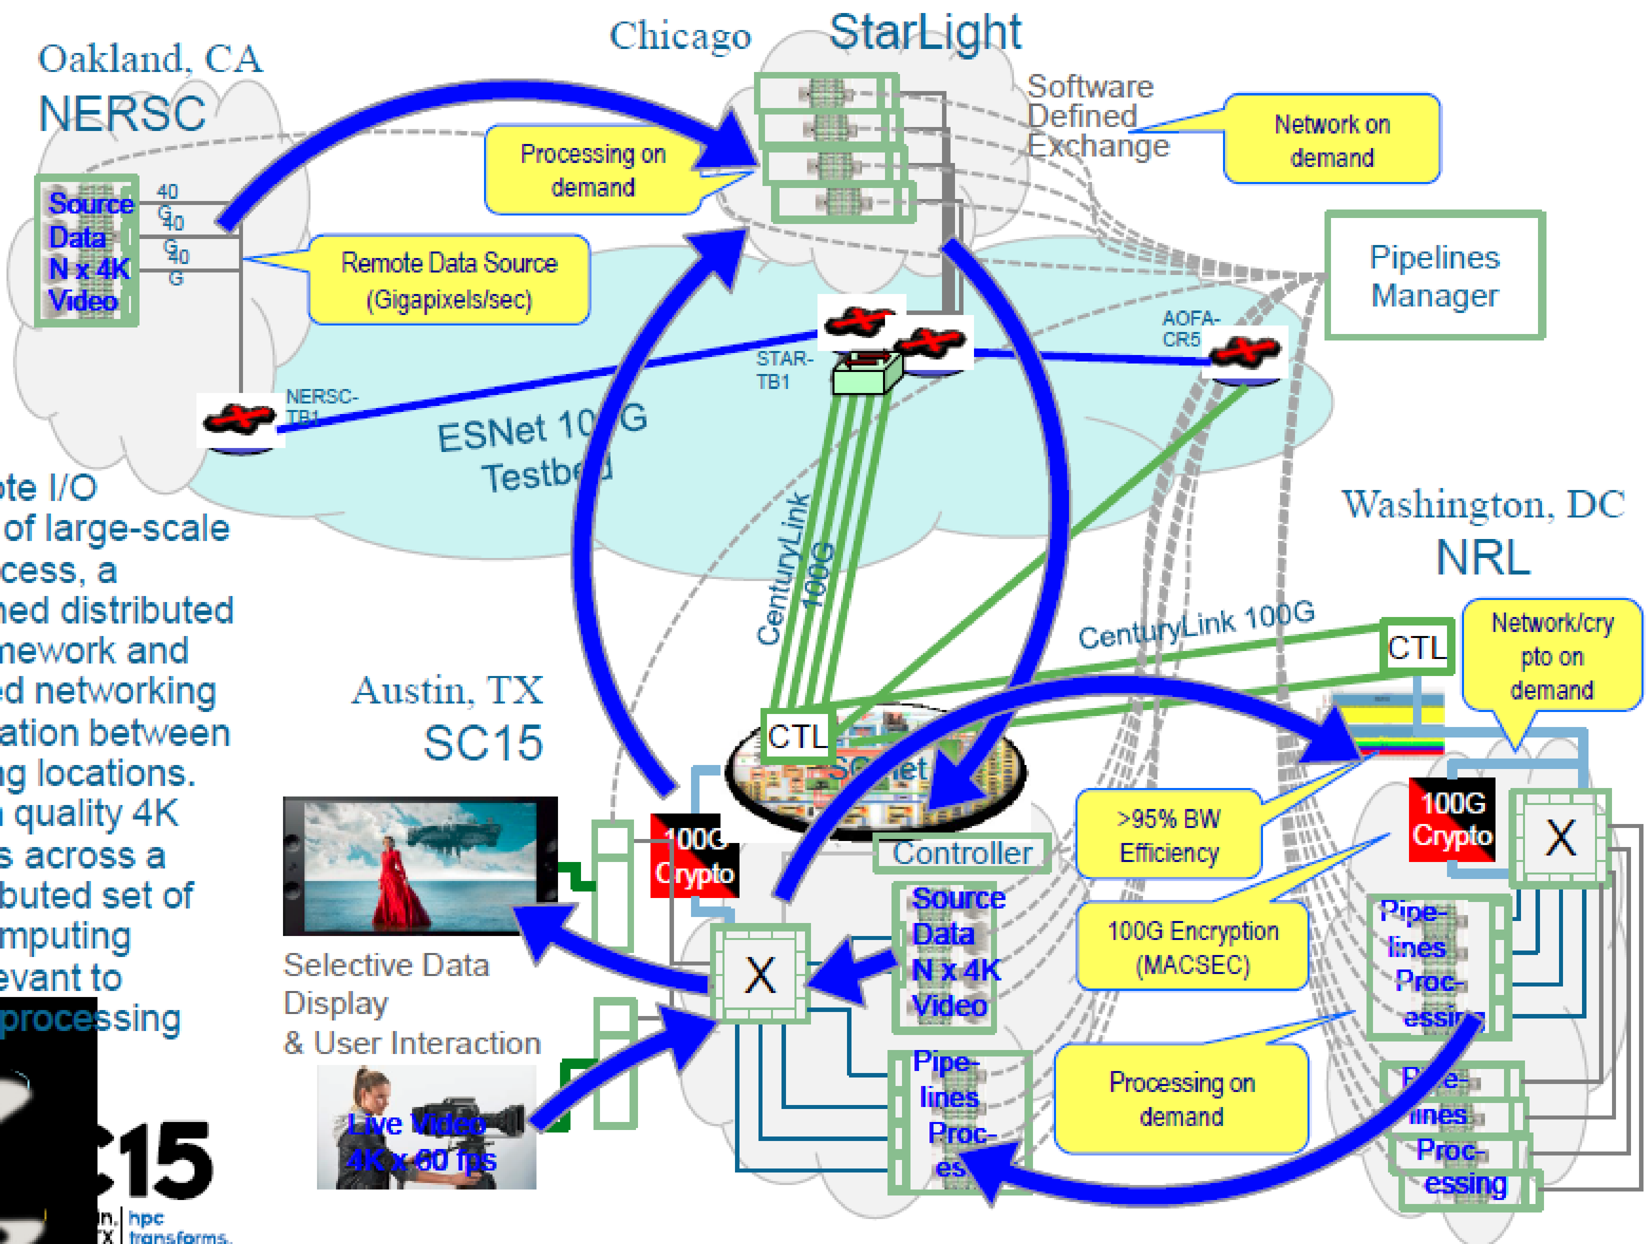

Dynamic Remote I/O Demonstration of large-scale remote data access, a dynamic pipelined distributed processing framework and software defined networking enabled automation between distant operating locations. Live production quality 4K video workflows across a nationally distributed set of storage and computing resources - relevant to emerging data processing challenges.

SC15 hpc transforms.

WAN Demonstration Plans

CALTECH HEP NETWORKING

Remote NSI Sites:

1. Caltech
2. StarLight
3. Univ of Michigan
4. Florida International Univ
5. RNP (Brazil)
6. UNESP (Brazil)

Software used to Provision paths to remote end points:

1. OSCARS
2. OESS
3. NSI
4. OpenFlow

Total of 9 x 100GE links (7 links in Caltech Booth)

NetherLight
Caltech
Fermilab
DE-KIT
GEANT
DFN
StarLight
Vanderbilt
Austin
AmLight
WAN providers
• ESNet
• CENIC
• CenturyLink
• NetherLight
• GEANT
• DFN
• RNP
• ANESP
RNP
UNESP

OpenFlow Traffic Engineering

CALTECH HEP NETWORKING

1. **Install** Flows based on interface statistics or shortest path

2. **Move** Flows dynamically by looking at the interface statistics in the environment and if there is a high bandwidth route available then re-config the flow

3. **Manual** Flow insertion/removal to adapt any requirement

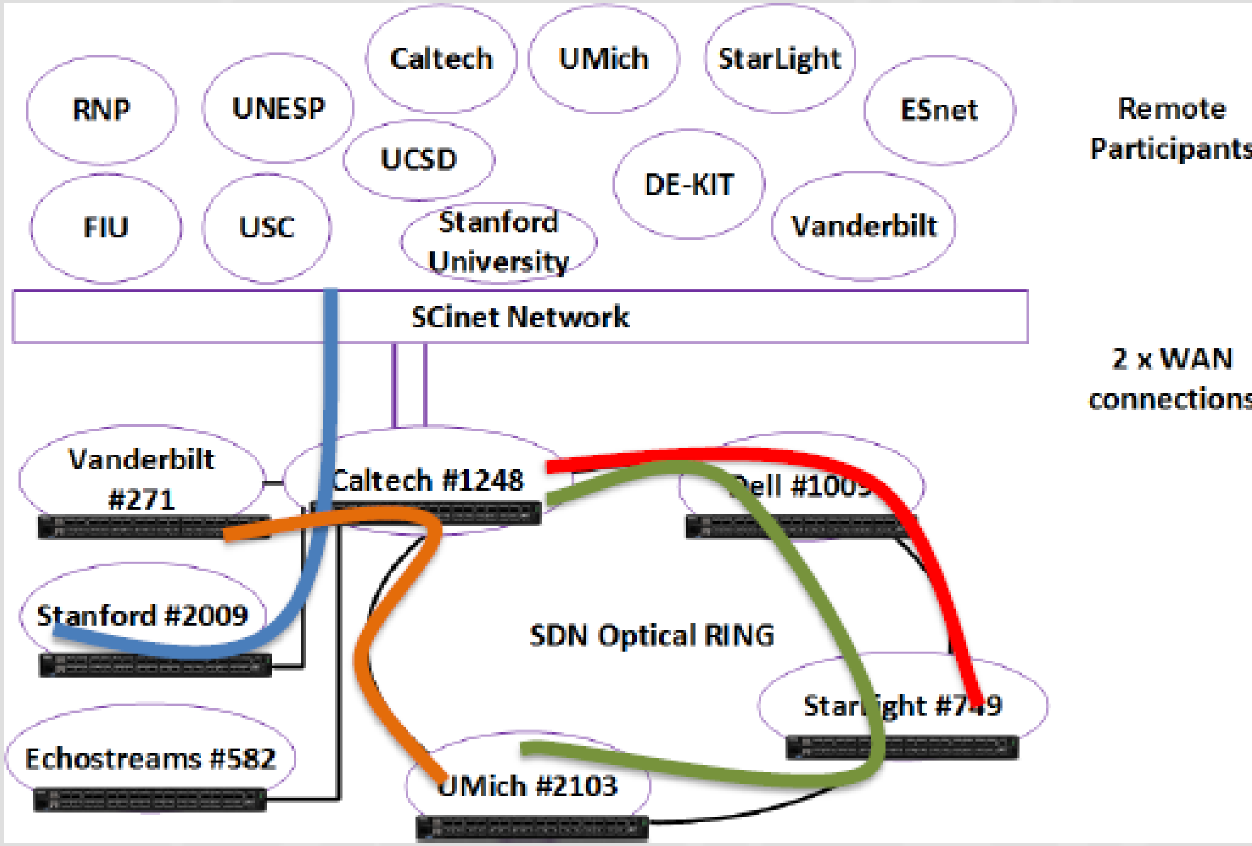

SC16 SDN-WAN Demonstration End-Points
Caltech, UM, Vanderbilt, Dell, Mellanox, 2CRSI, Ciena, Echostreams, StarLight, PRP, FIU, UNESP

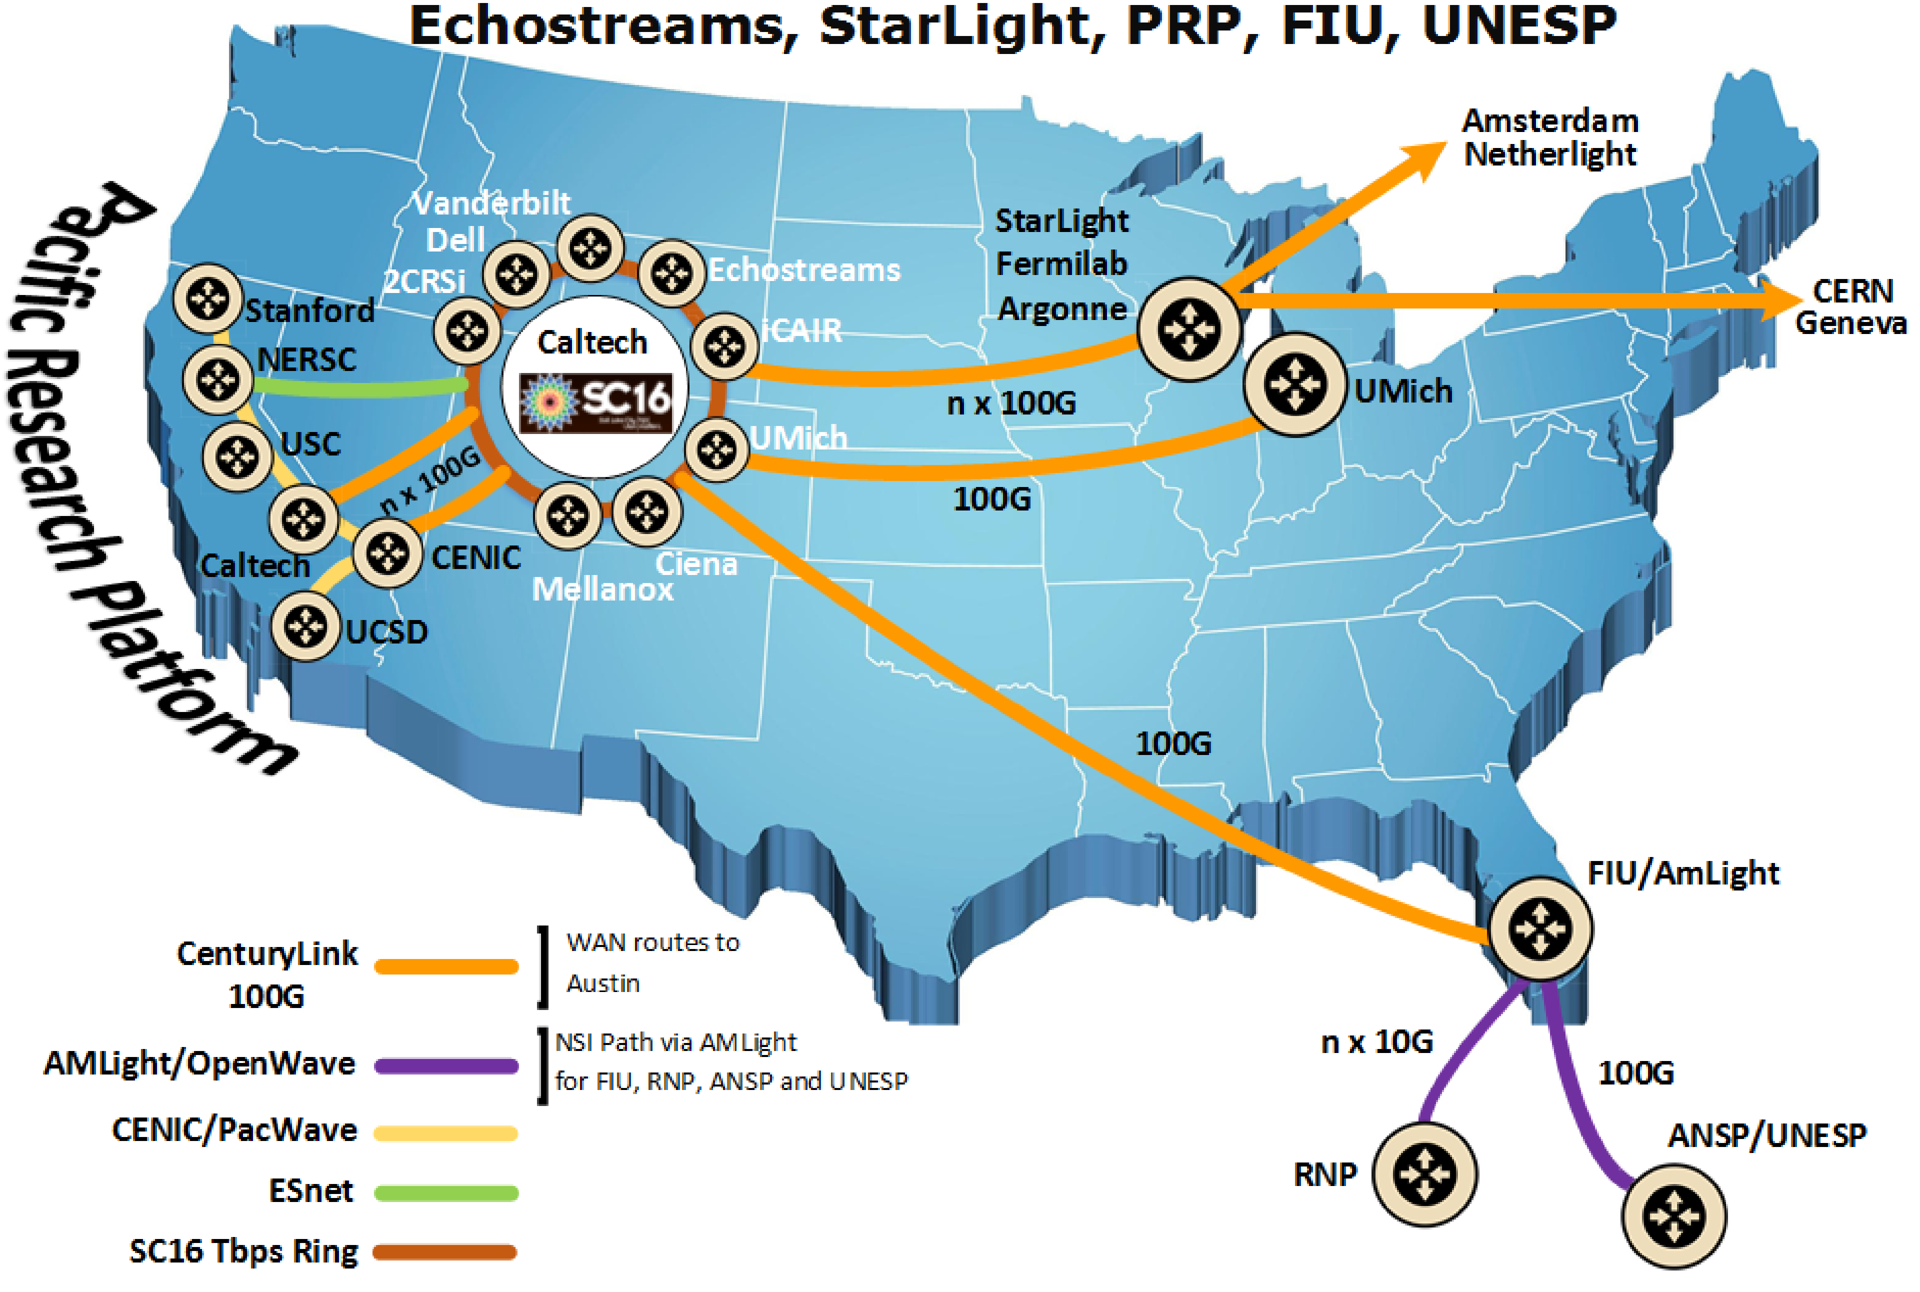

Demonstration Fermi National Accelerator Laboratory & iCAIR: Multicore-aware Data Transfer Middleware (MDTM) Project (DOE)

- **The Multicore-aware Data Transfer Middleware (MDTM) Project**
  - Collaborative effort by Fermilab and Brookhaven National Laboratory
  - Funded by DOE's Office of Advanced Scientific Computing Research (ASCR)

MDTM aims to accelerate data movement toolkits on multicore systems

Office of Science U.S. DEPARTMENT OF ENERGY Fermilab BROOKHAVEN NATIONAL LABORATORY

Multicore-aware Data Transfer Middleware (MDTM) Project (DOE)

- Issue: Existing Data Transfer Tools Do Not Efficiently Exploit Multicore Devices Using Common OS Implementations, Especially With Regard To NUMA Systems.
- The Large Gap Between OS Processes And Multicore Hardware Designs Results In Network I/O Inefficient – Only “Best Effort” Handling Of Processor Threads for Data Transfer.
- Basic Processes Are Generic: No Distinctions - Type of Transfer, Type of Service, Thread Locality, Anomalies, Requirements, Priorities, Dynamic Changes, etc
- OS Tuning Alone Cannot Resolve Performance Problems

MDTM Architecture

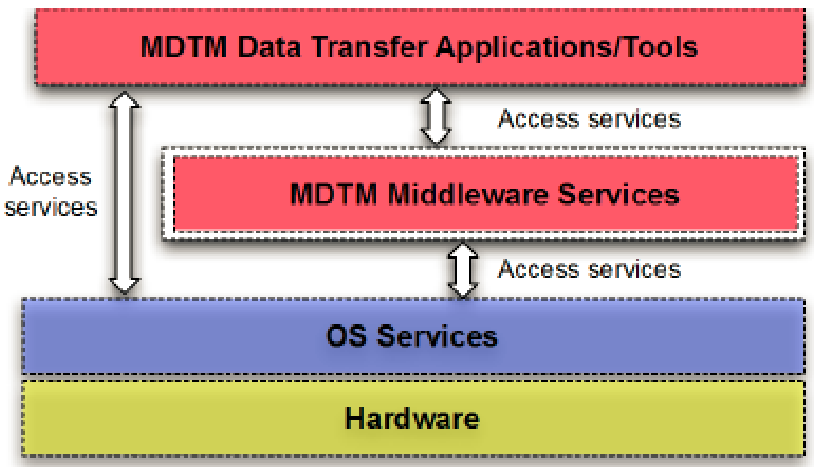

MDTM Architecture

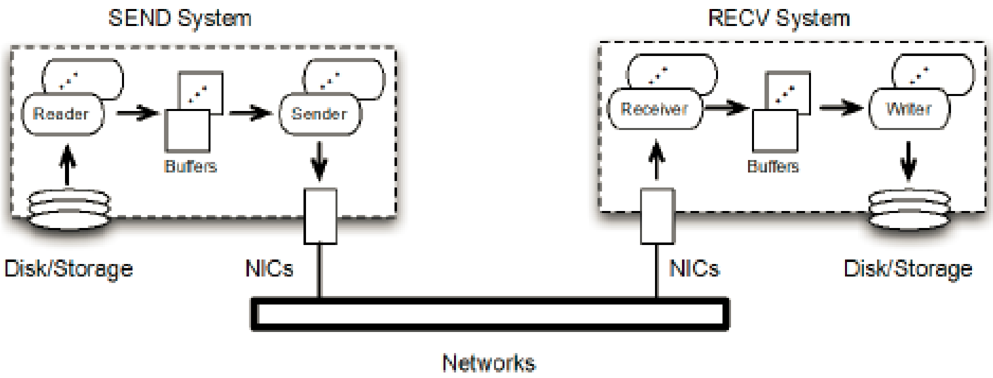

MDTM Data Transfer Model

MDTM consists of two components:

- **MDTM data transfer application (BNL)**
  - Adopts an I/O-centric architecture that uses dedicated threads to perform network and disk I/O operations
- **MDTM middleware services (FNAL)**
  - Harness multicore parallelism to scale data movement toolkits on host systems

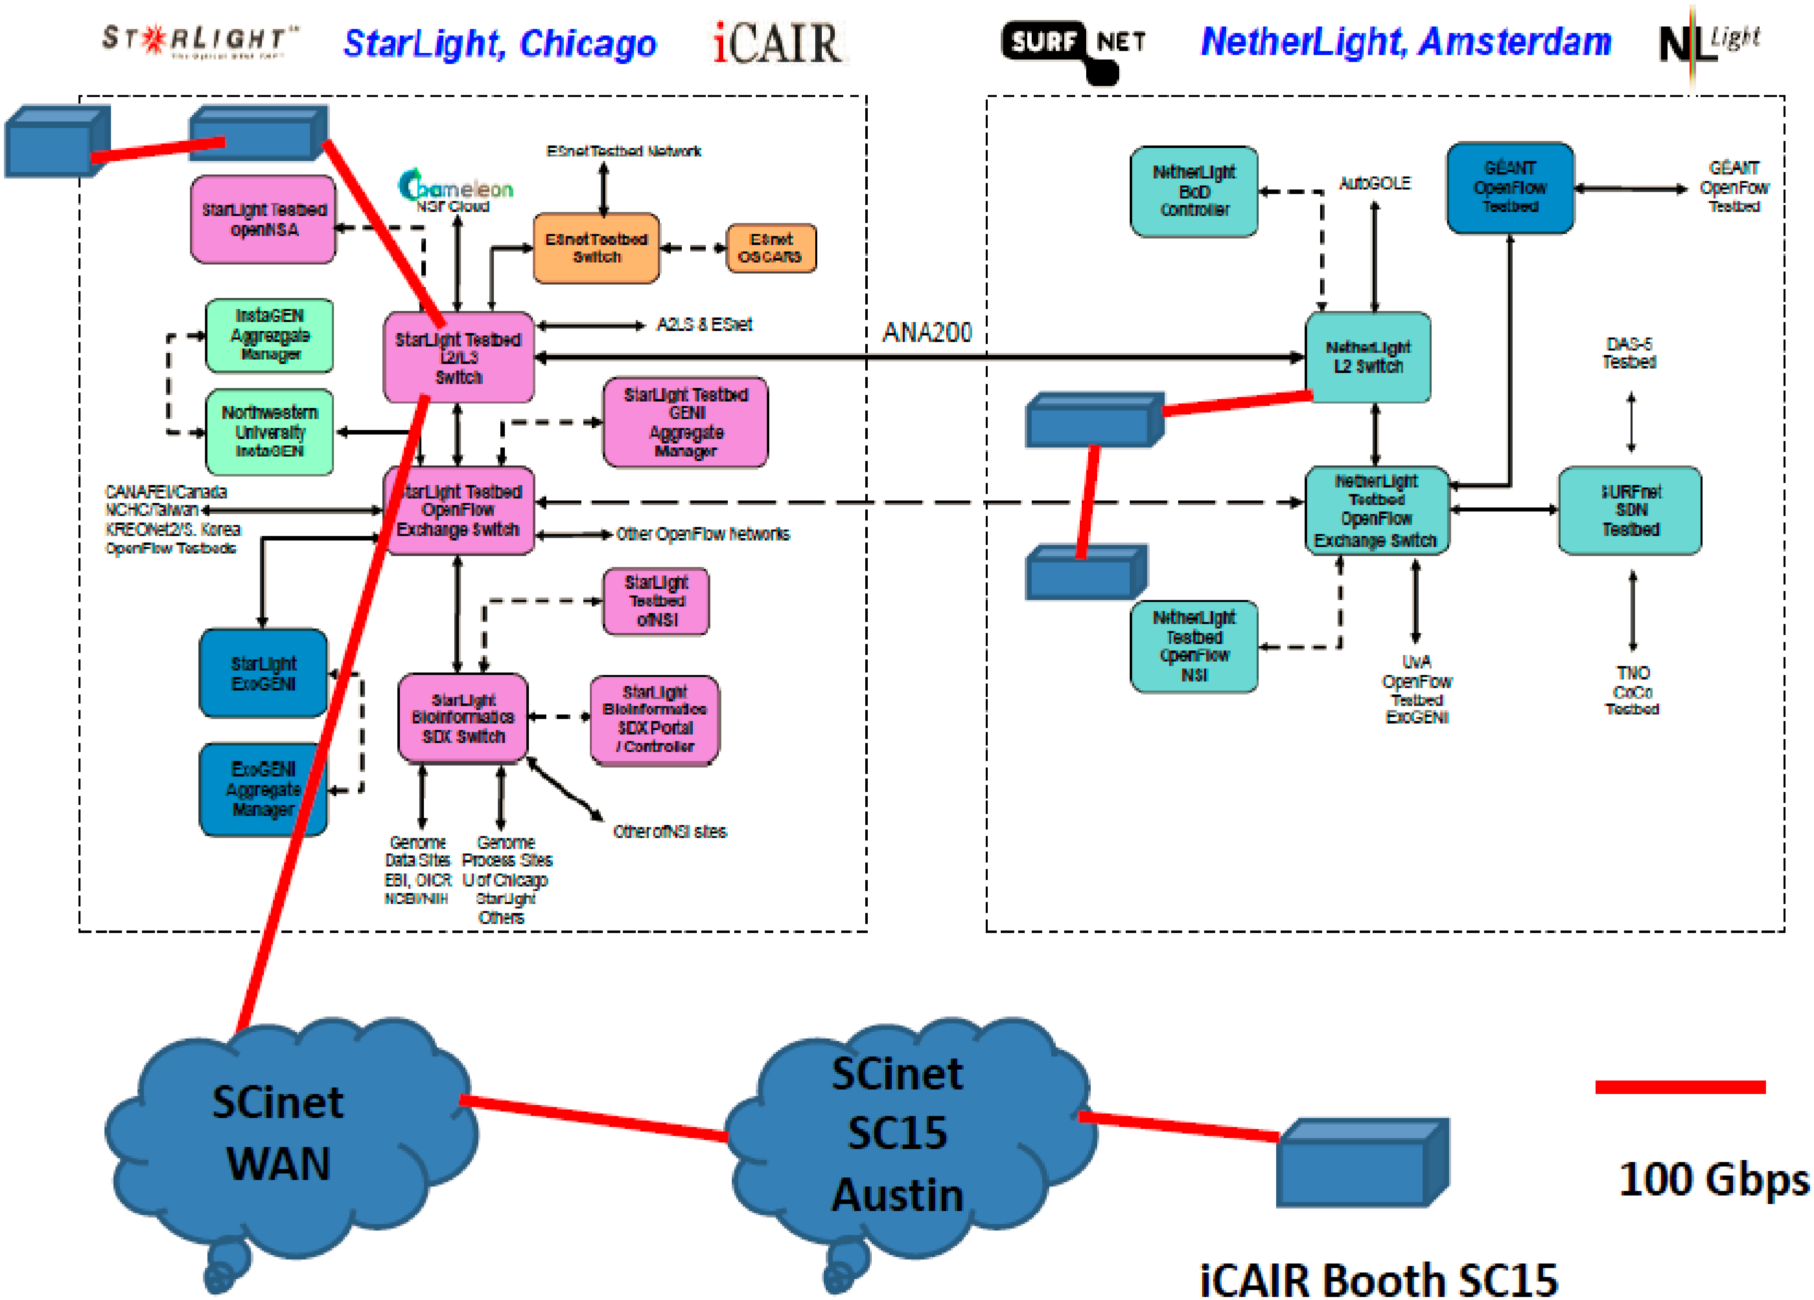

StarLight, Chicago
iCAIR
NetherLight, Amsterdam
StarLight Testbed openNSA
Chameleon NSF Cloud
ESnet Testbed Network
ESnet Testbed Switch
ESnet OSCARS
InstaGEN Aggregate Manager
StarLight Testbed L2/L3 Switch
A2LS & ESnet
Northwestern University InstaGEN
StarLight Testbed GENI Aggregate Manager
CANAREI/Canada
NCHC/Taiwan
KREONet2/S. Korea
OpenFlow Testbeds
StarLight Testbed OpenFlow Exchange Switch
Other OpenFlow networks
StarLight Testbed ofNSI
StarLight ExoGENI
StarLight Bioinformatics SDX Switch
StarLight Bioinformatics SDX Portal / Controller
ExoGENI Aggregate Manager
Other ofNSI sites
Genome Data Sites EBI, OICR NCBI/NIH
Genome Process Sites U of Chicago StarLight Others
ANA200
NetherLight BoD Controller
AutoGOLE
GÉANT OpenFlow Testbed
GÉANT OpenFow Testbed
NetherLight L2 Switch
DAS-5 Testbed
NetherLight Testbed OpenFlow Exchange Switch
SURFnet SDN Testbed
NetherLight Testbed OpenFlow NSI
UvA OpenFlow Testbed ExoGENI
TNO CoCo Testbed
SCinet WAN
SCinet SC15 Austin
100 Gbps
iCAIR Booth SC15

StarLight SDX⇔ SURFnet SDX

100GE End to End Services based on Open Architecture (E2SOA) for Peta Scale Sciences: Prototype Set Up

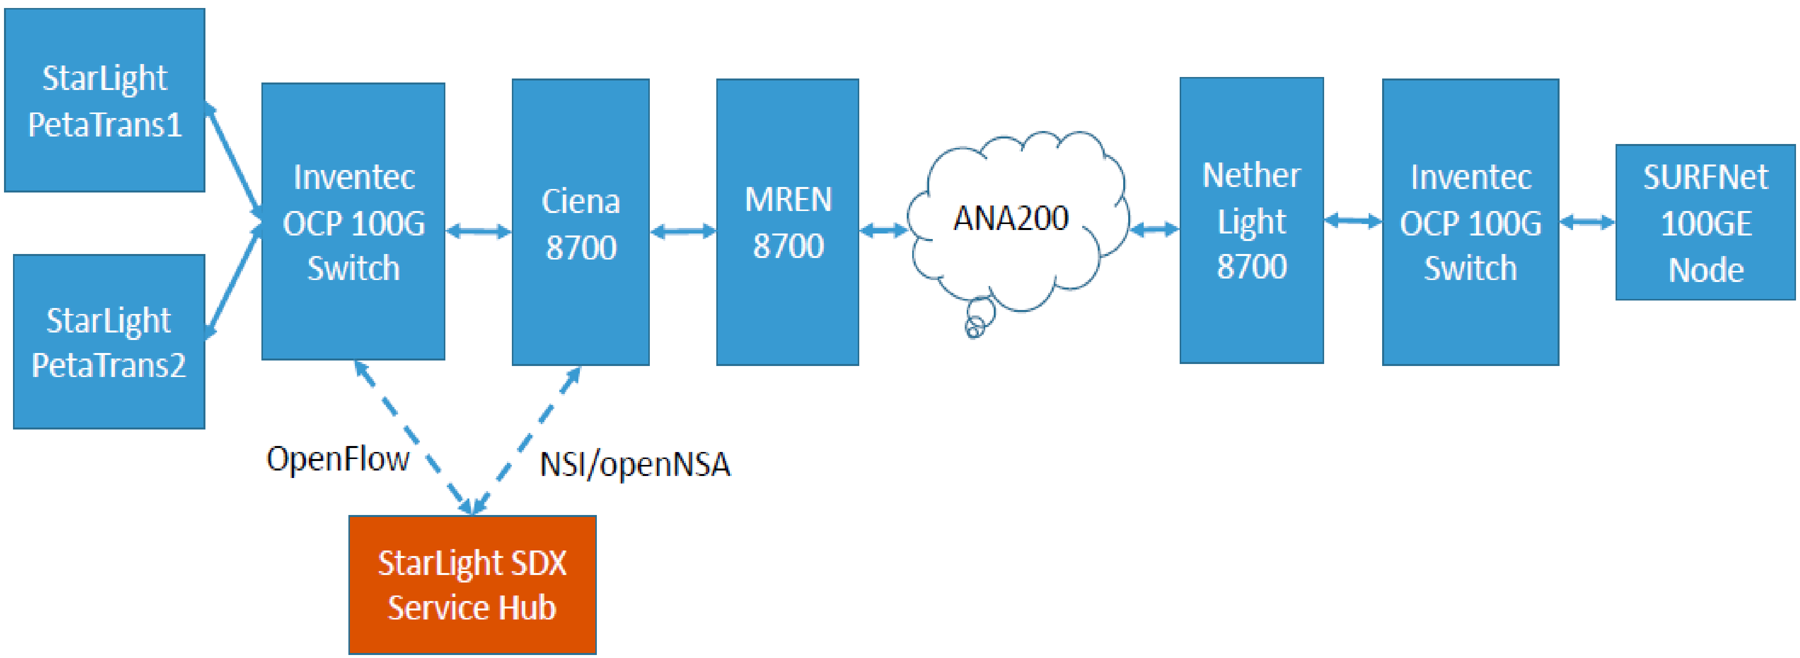

Multi-Tenant 100 GE SDX

Multi-Tenant 100GE Science Network Exchange
SC15 NRE Testing Phase

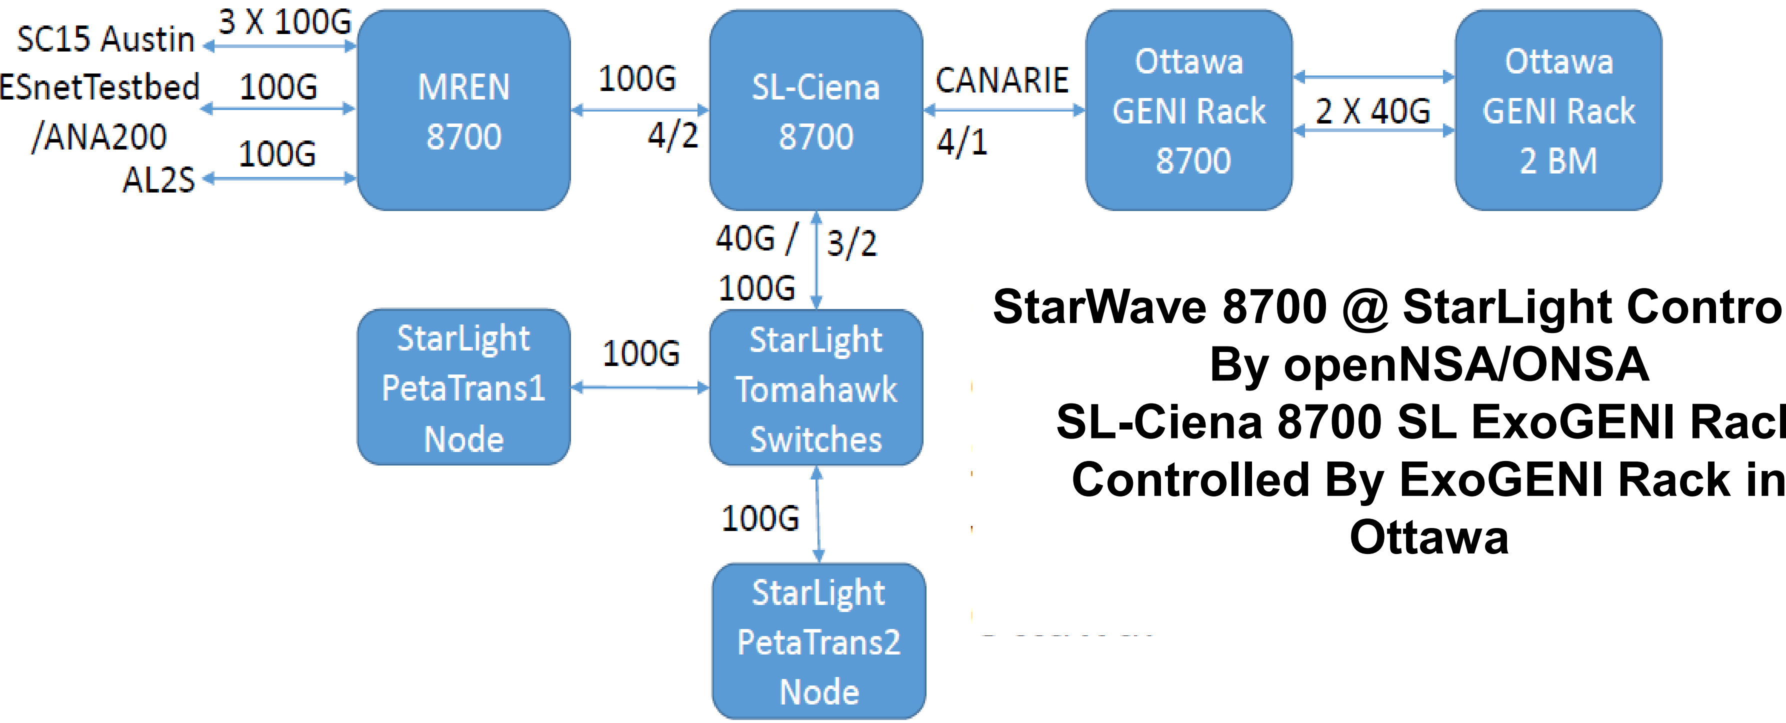

StarWave 8700 @ StarLight Controlled By openNSA/ONSA
SL-Ciena 8700 SL ExoGENI Rack Controlled By ExoGENI Rack in Ottawa

NSI-OpenFlow Hybird Topology Exchange

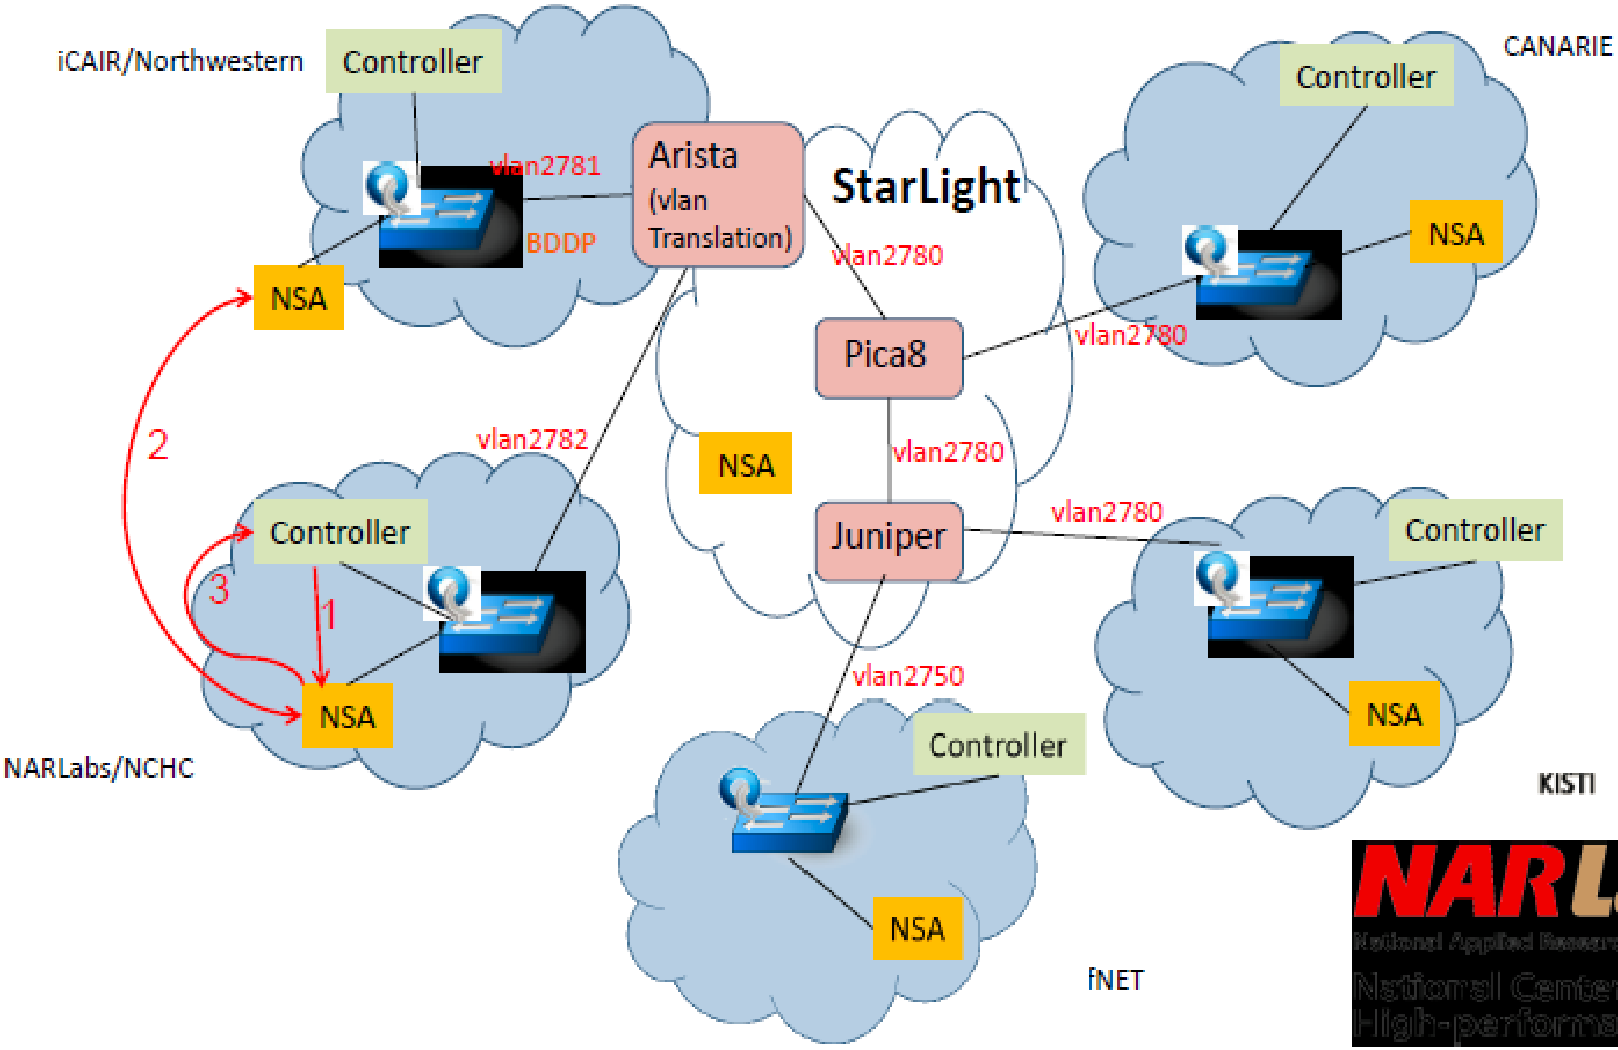

Coming Soon =>Taiwan SDX

Slice Exchange Showcase at GEC 21

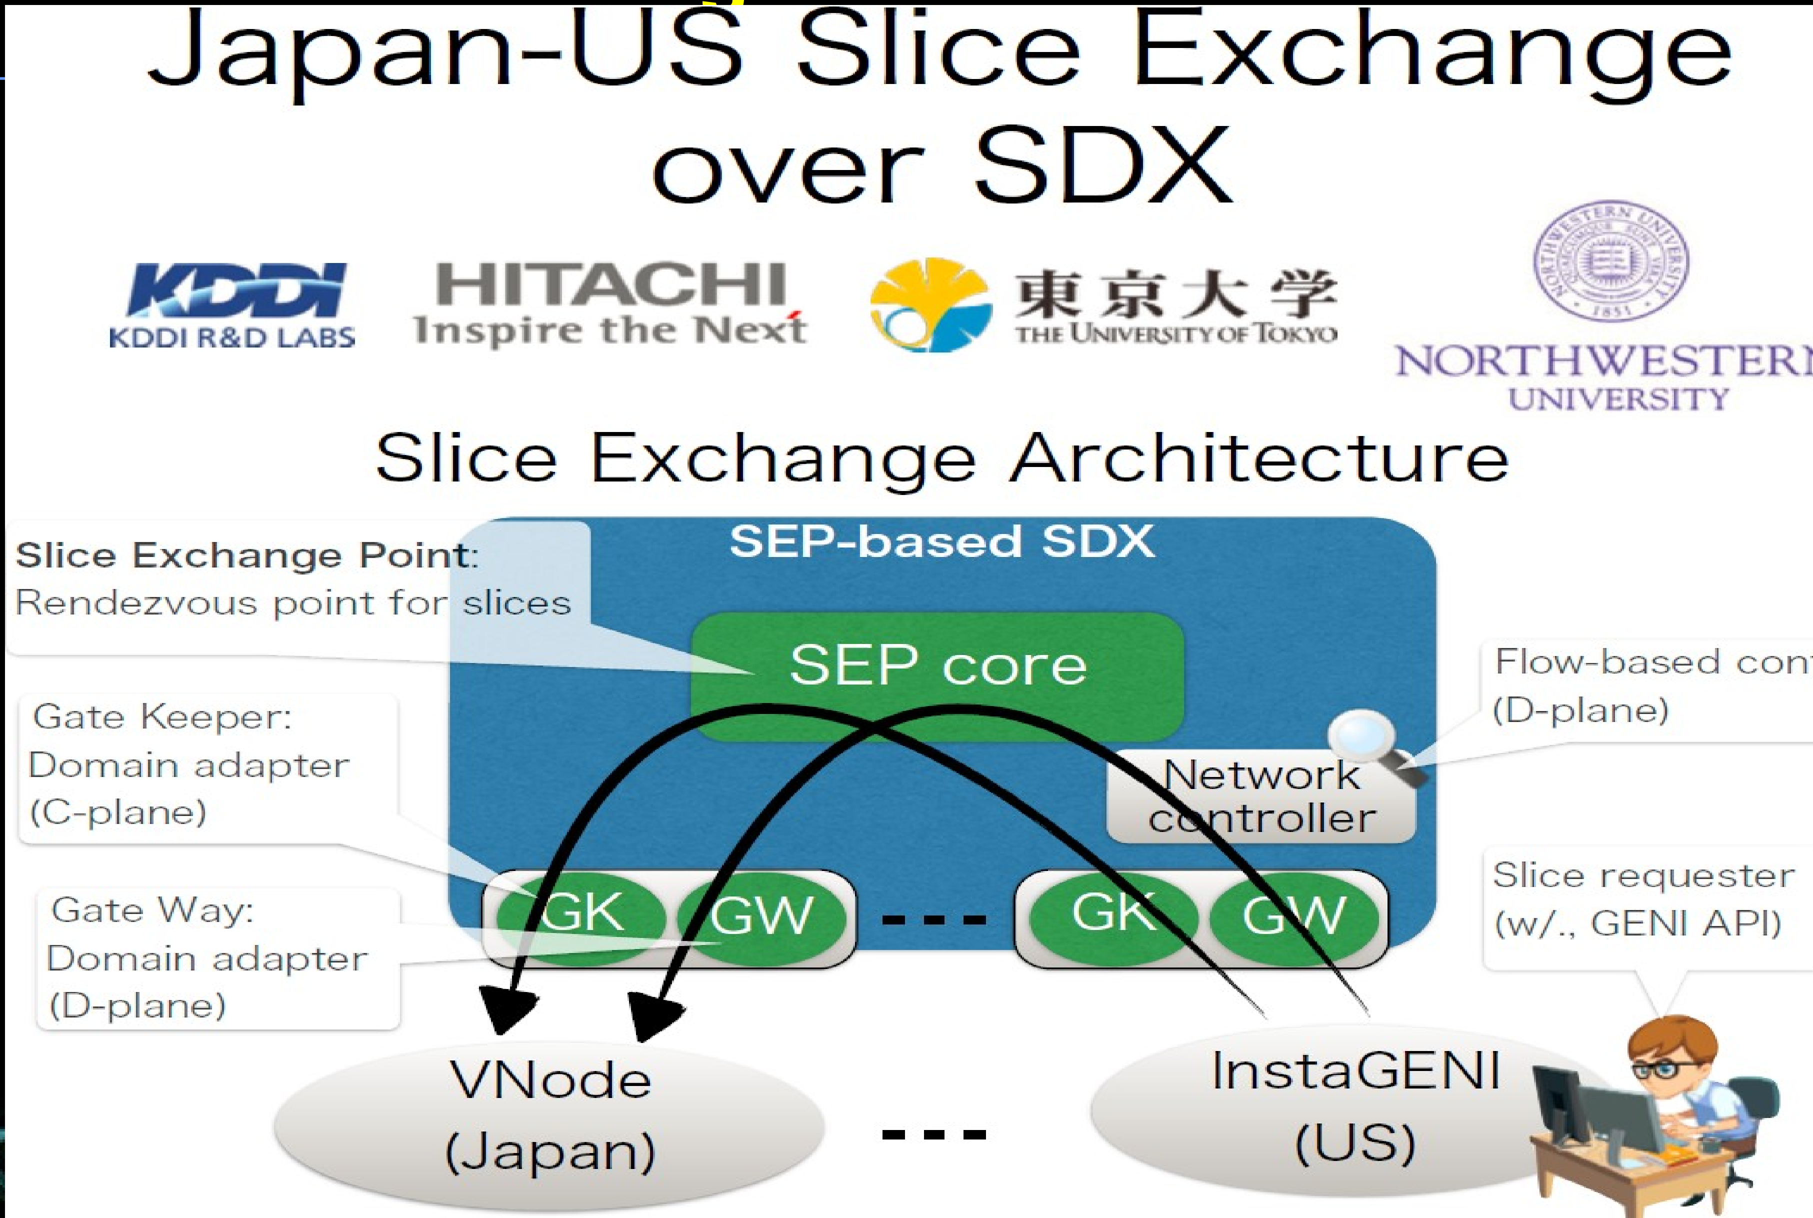

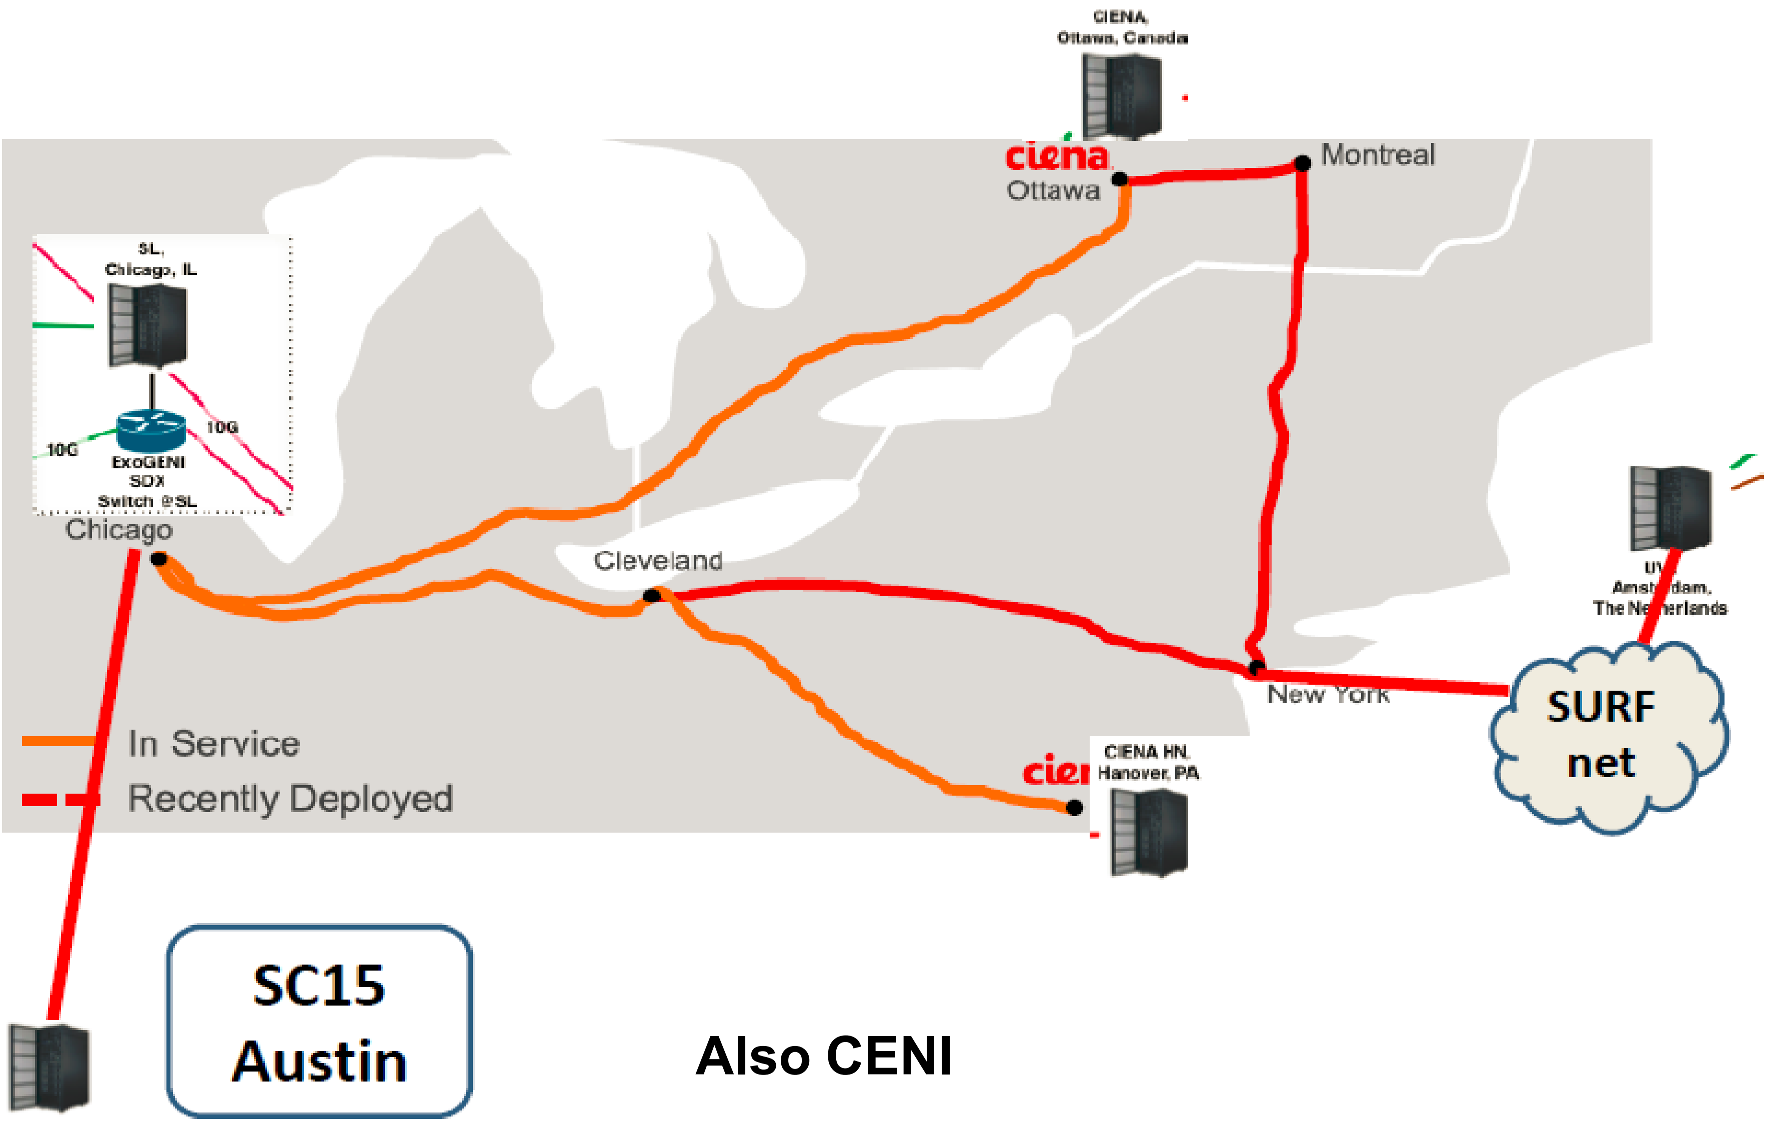
CIENA,
Ottawa, Canada
ciena
Ottawa
Montreal
SL,
Chicago, IL
10G
10G
ExoGENI
SDX
Switch @SL
Chicago
Cleveland
New York
SURF
net
CIENA HN,
Hanover, PA
In Service
Recently Deployed
SC15
Austin

Also CENI

An Industry First: June, 2016 – Ciena, Canada's Advanced Research and Innovation Network (CANARIE) and the StarLight International/National Communications Exchange Facility Consortium completed a 300 Gb/s 8QAM (Quadrature Amplitude Modulation) trial on a live network using Ciena's Waveserver.

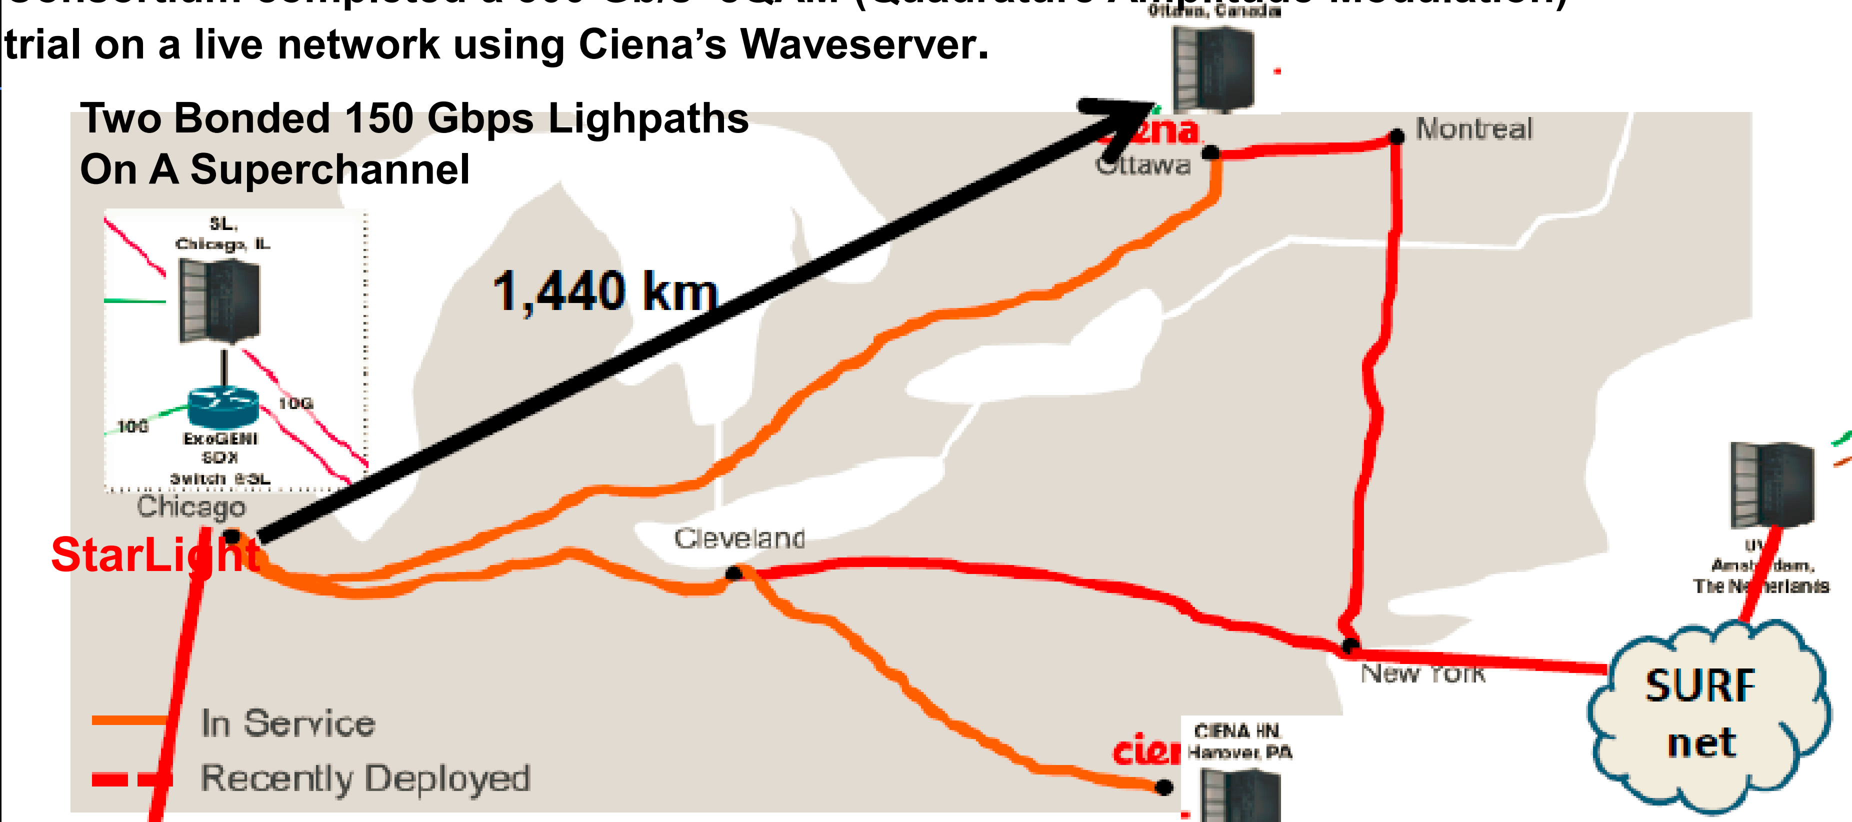

*__Demonstration of World's 1st 300 Gpbs Over 2 Lights Paths Over a Distance of 1,440 Kliometers, Supported By p__*artnership of Ciena, CANARIE and iCAIR

May 16 2016

New Open Exchange In Montreal

- New *R&E Open* Exchange Operated By CANARIE
- Montreal ⇔ 100 Gbps International Connectivity
- To Be Completed ~ October/November

E2E Real Time Service Analytics Over 100 G Paths Using the Blue Planet F

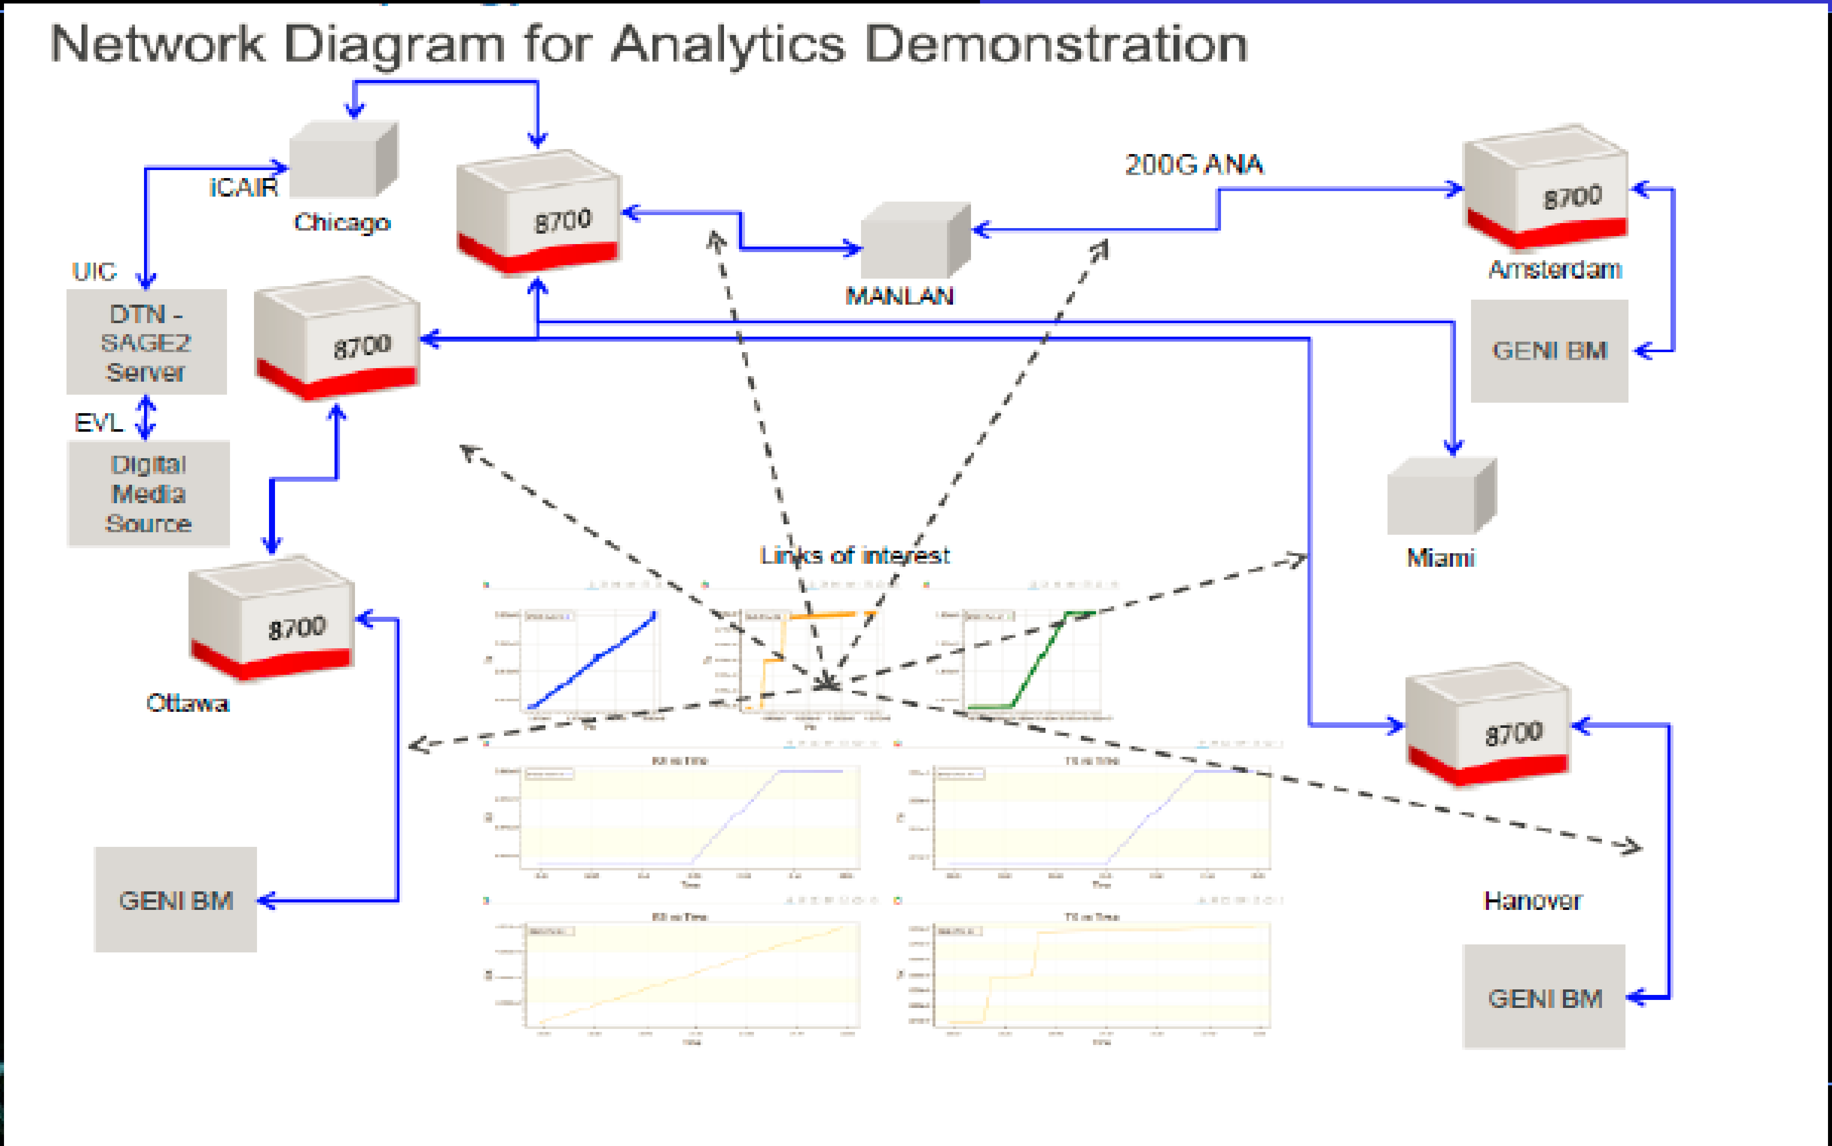

Project Matsu – NASA Collaboration (2009)

Open Science Data Cloud (2010)

Bionimbus Protected Data Cloud* (2013)

NCI Data Commons* (Nov 2015)

NOAA Data Commons (Nov 2015)

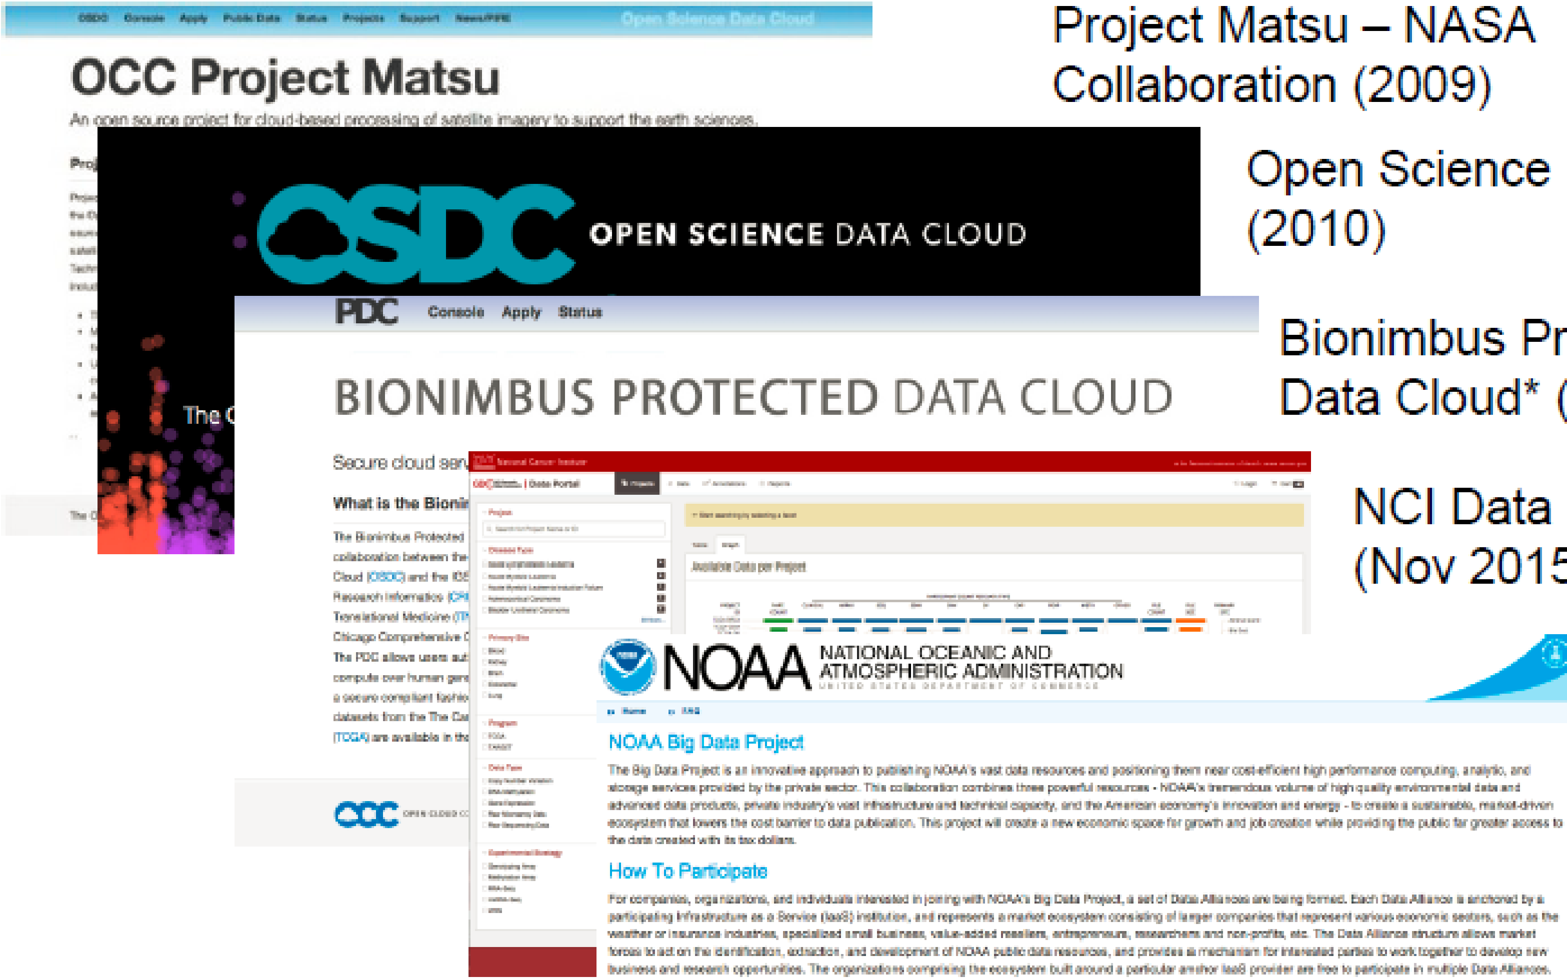

- OCC operated testbed from 2008 to 2012 to develop the technology.
- 760+ research projects supported since 2010.
- 250+ currently active researchers from 54 organizations from 14 countries.
- Over 18 million core hours used by allocation grantees in past year

*Operated under a subcontract from NCI / Leidos Biomedical to the University of Chicago with support from the OCC.

NOAA NATIONAL OCEANIC AND ATMOSPHERIC ADMINISTRATION
UNITED STATES DEPARTMENT OF COMMERCE

» Home » FAQ

NOAA Big Data Project

The Big Data Project is an innovative approach to publishing NOAA's vast data resources and positioning them near cost-efficient high performance computing, analytic, and storage services provided by the private sector. This collaboration combines three powerful resources - NOAA's tremendous volume of high quality environmental data and advanced data products, private industry's vast infrastructure and technical capacity, and the American economy's innovation and energy - to create a sustainable, market-driven ecosystem that lowers the cost barrier to data publication. This project will create a new economic space for growth and job creation while providing the public far greater access to the data created with its tax dollars.

How To Participate

For companies, organizations, and individuals interested in joining with NOAA's Big Data Project, a set of Data Alliances are being formed. Each Data Alliance is anchored by a participating Infrastructure as a Service (IaaS) institution, and represents a market ecosystem consisting of larger companies that represent various economic sectors, such as the weather or insurance industries, specialized small business, value-added resellers, entrepreneurs, researchers and non-profits, etc. The Data Alliance structure allows market forces to act on the identification, extraction, and development of NOAA public data resources, and provides a mechanism for interested parties to work together to develop new business and research opportunities. The organizations comprising the ecosystem built around a particular anchor IaaS provider are free to participate in multiple Data Alliances.

For more information, visit one of the NOAA Big Data Collaborators:

amazon webservices | Google Cloud Platform | IBM | Microsoft | OCC Open Cloud Consortium

- Public-private data collaborative announced April 21, 2015 by US Secretary of Commerce Pritzker.
- AWS, Google, IBM, Microsoft and Open Cloud Consortium will form five collaborations.
- We will develop an OCC/NOAA Data Commons.

STARLIGHT SDX

NCI Genomic Data Commons

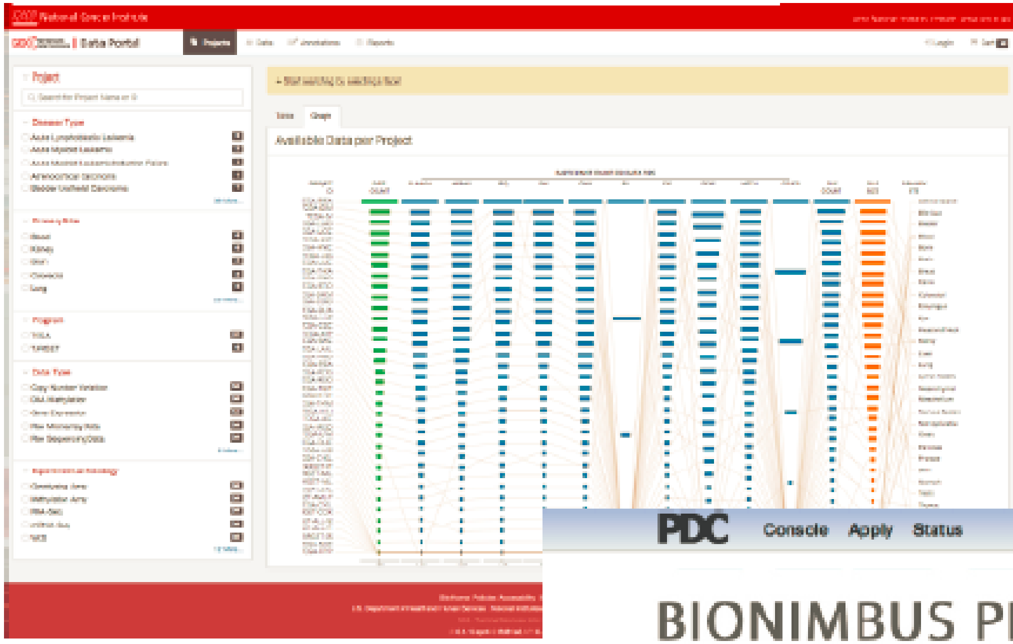

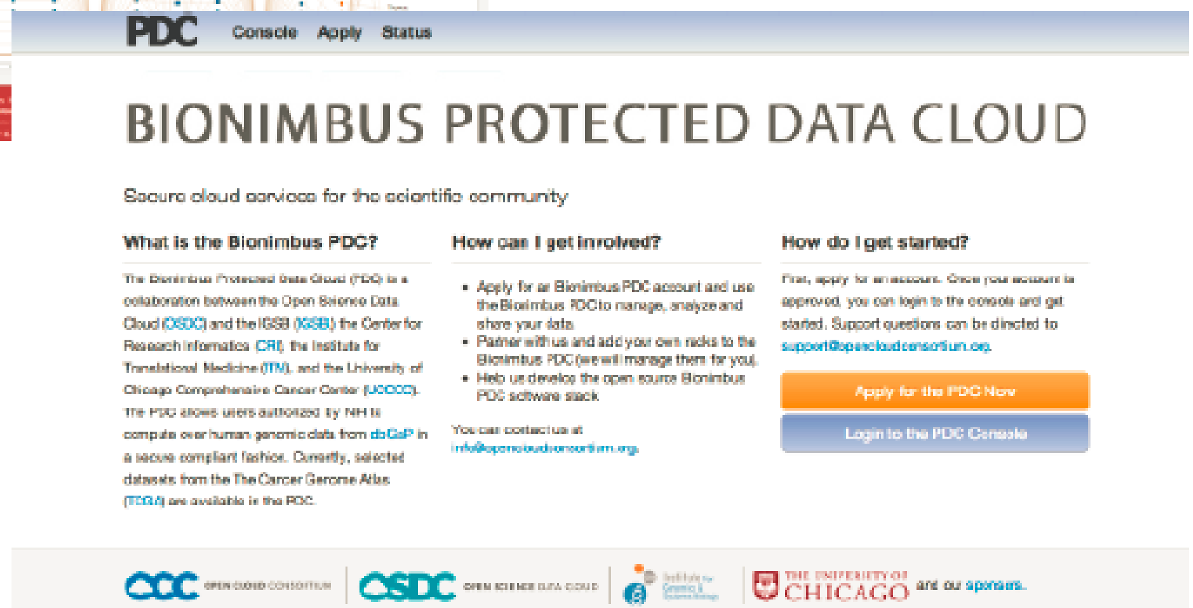

University of Chicago biomedical data commons developed in collaboration with the OCC.

STARLIGHT SDX

Genomic Data Commons (GDC) Announced

- The GDC, a Unified Data System for Sharing Genomic and Clinical Data Among Researchers, Was Launched June 6 With a Visit from Vice President Joe Biden To the Operations Center at the University of Chicago.
- An Initiative of the National Cancer Institute (NCI), the GDC Will Be a Core Component of the National Cancer Moonshot and the President's Precision Medicine Initiative (PMI)
- It Benefits from $70 million Allocated to NCI to Lead Efforts in Cancer Genomics as Part of PMI for Oncology.
- The GDC Will Centralize, Standardize and Make Accessible Data from Large-Scale NCI Programs
- The GDC Will be Part of a Global Ecosystem

Vice President Biden visiting the CDIS Data Commons Operation Center on June 6, 2016.

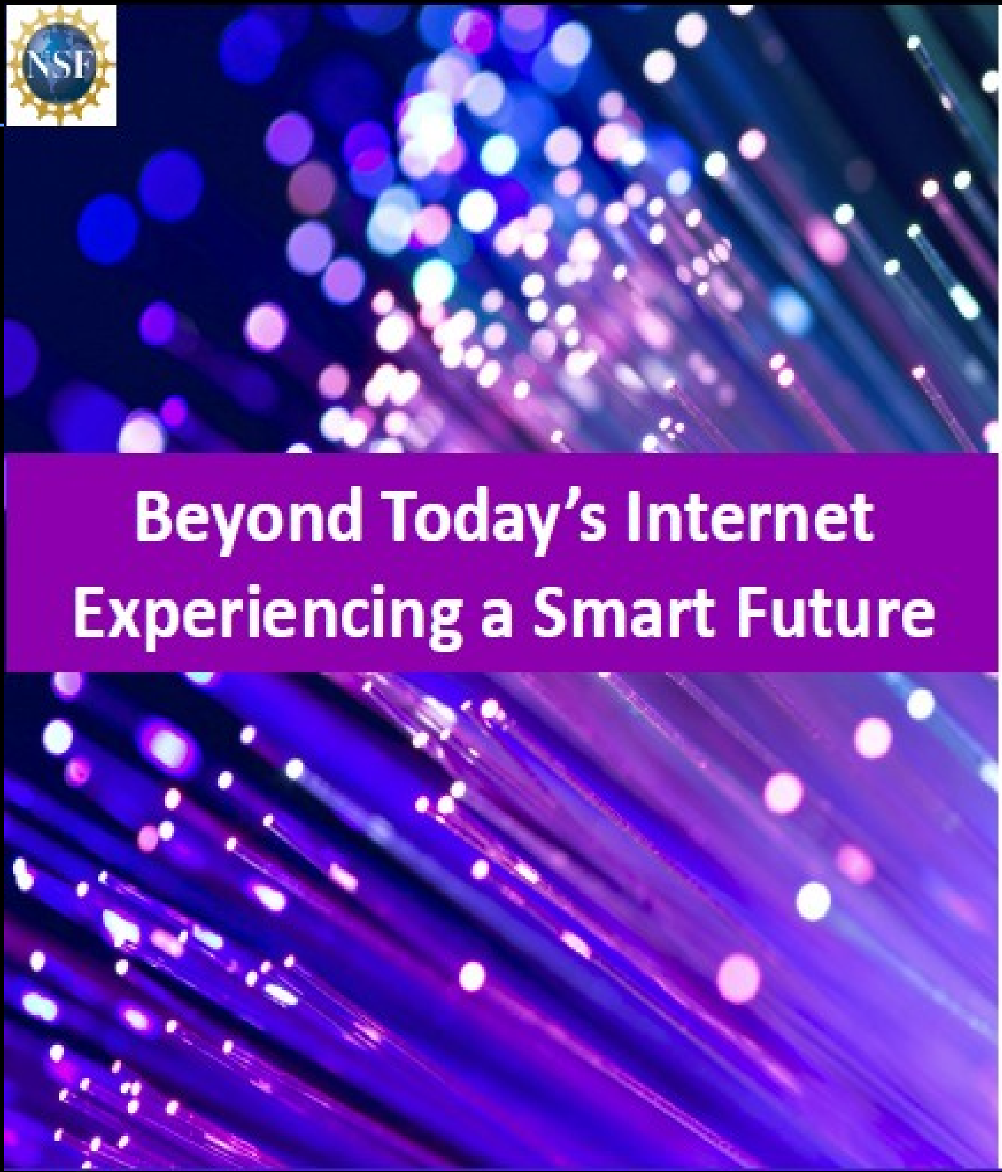

geni Exploring Networks of the Future

usignite

Prototype SDX Bioinformatics Exchange: Demonstrating an Essential Use-Case for Personalized Medicine

Robert Grossman, Piers Nash, Allison Heath, Renuka Arya
University of Chicago

Joe Mambretti, Jim Chen
Northwestern University

THE UNIVERSITY OF CHICAGO MEDICINE

NORTHWESTERN UNIVERSITY

STARLIGHT SDX

Biomedical Data Commons: Flow Orchestration: Control Plane + Data Plane

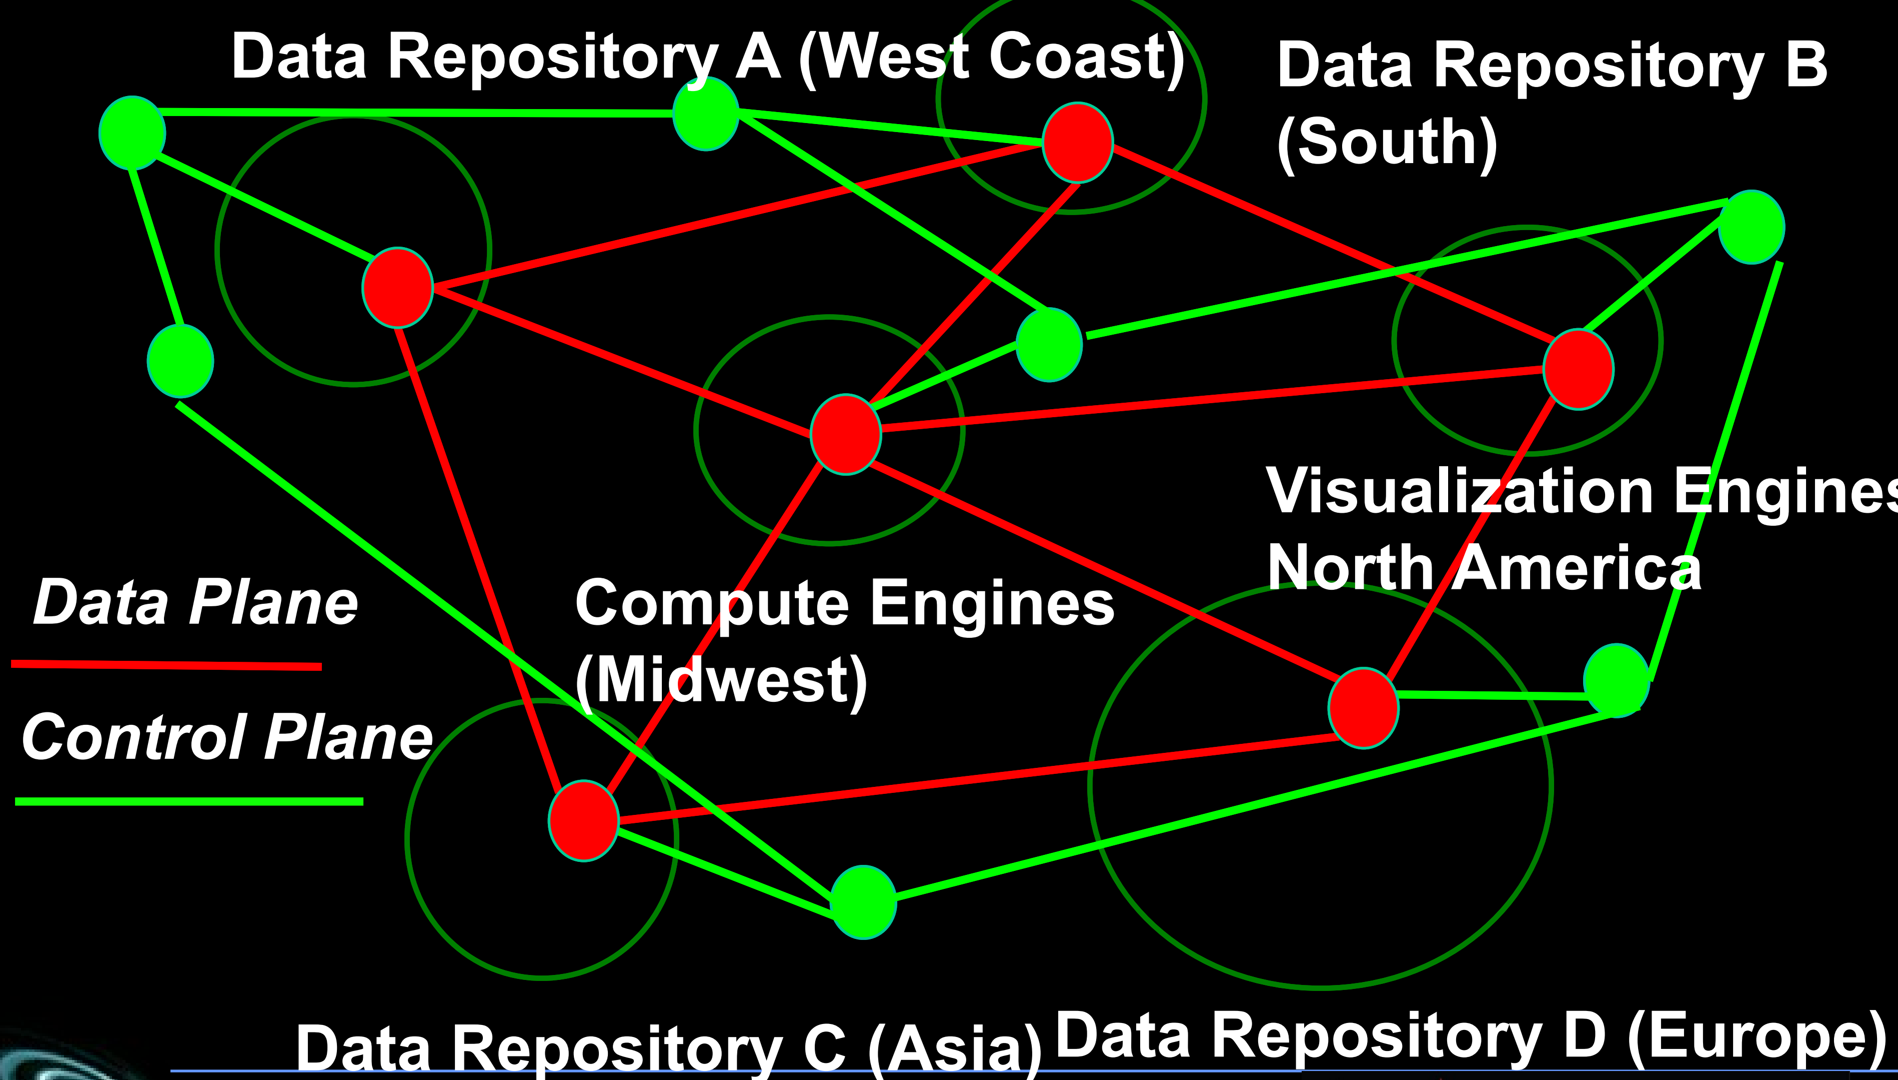

STARLIGHT SDX

Genomic Data Commons Data Transfer

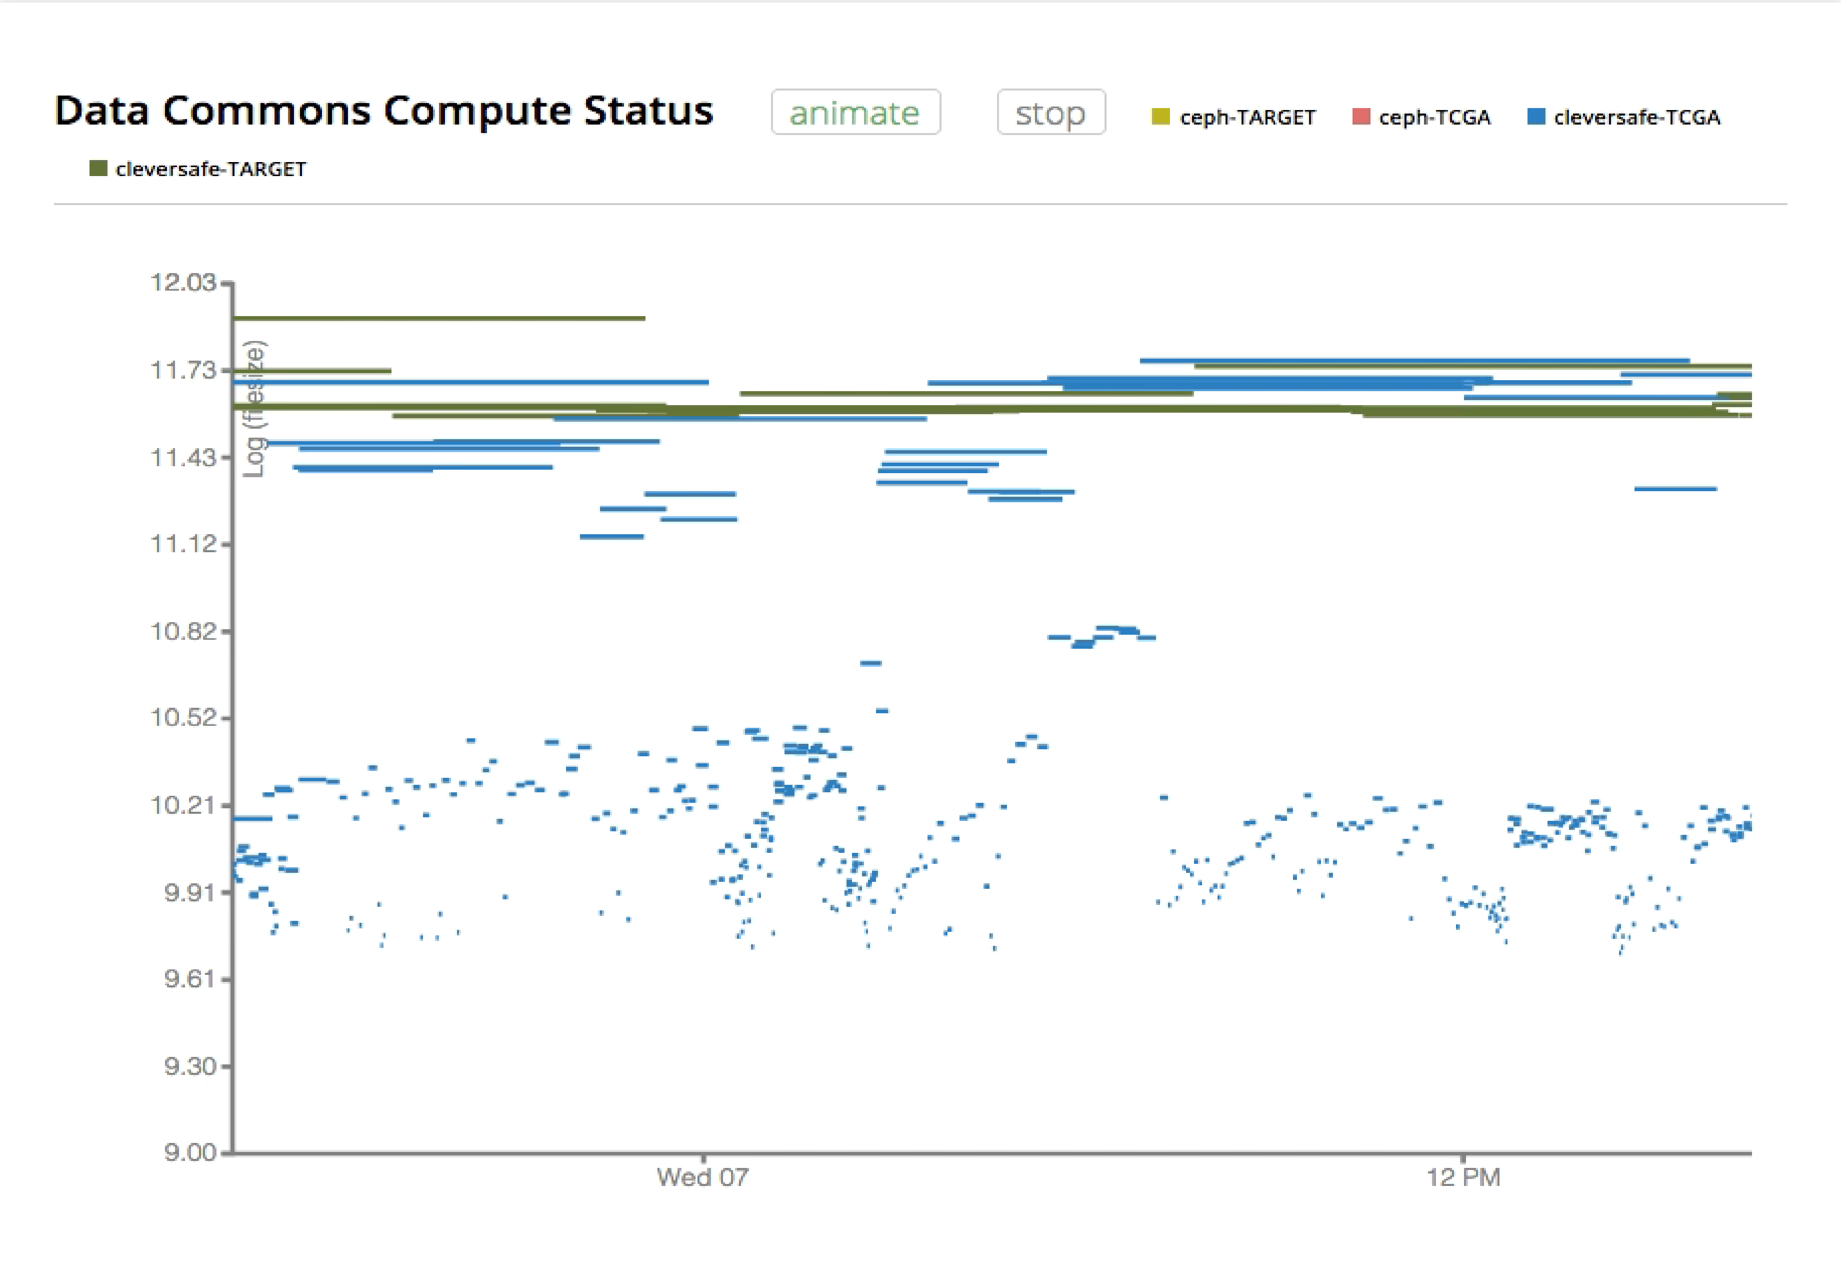

GEC22 Bioinformatics SDXs Demo Network

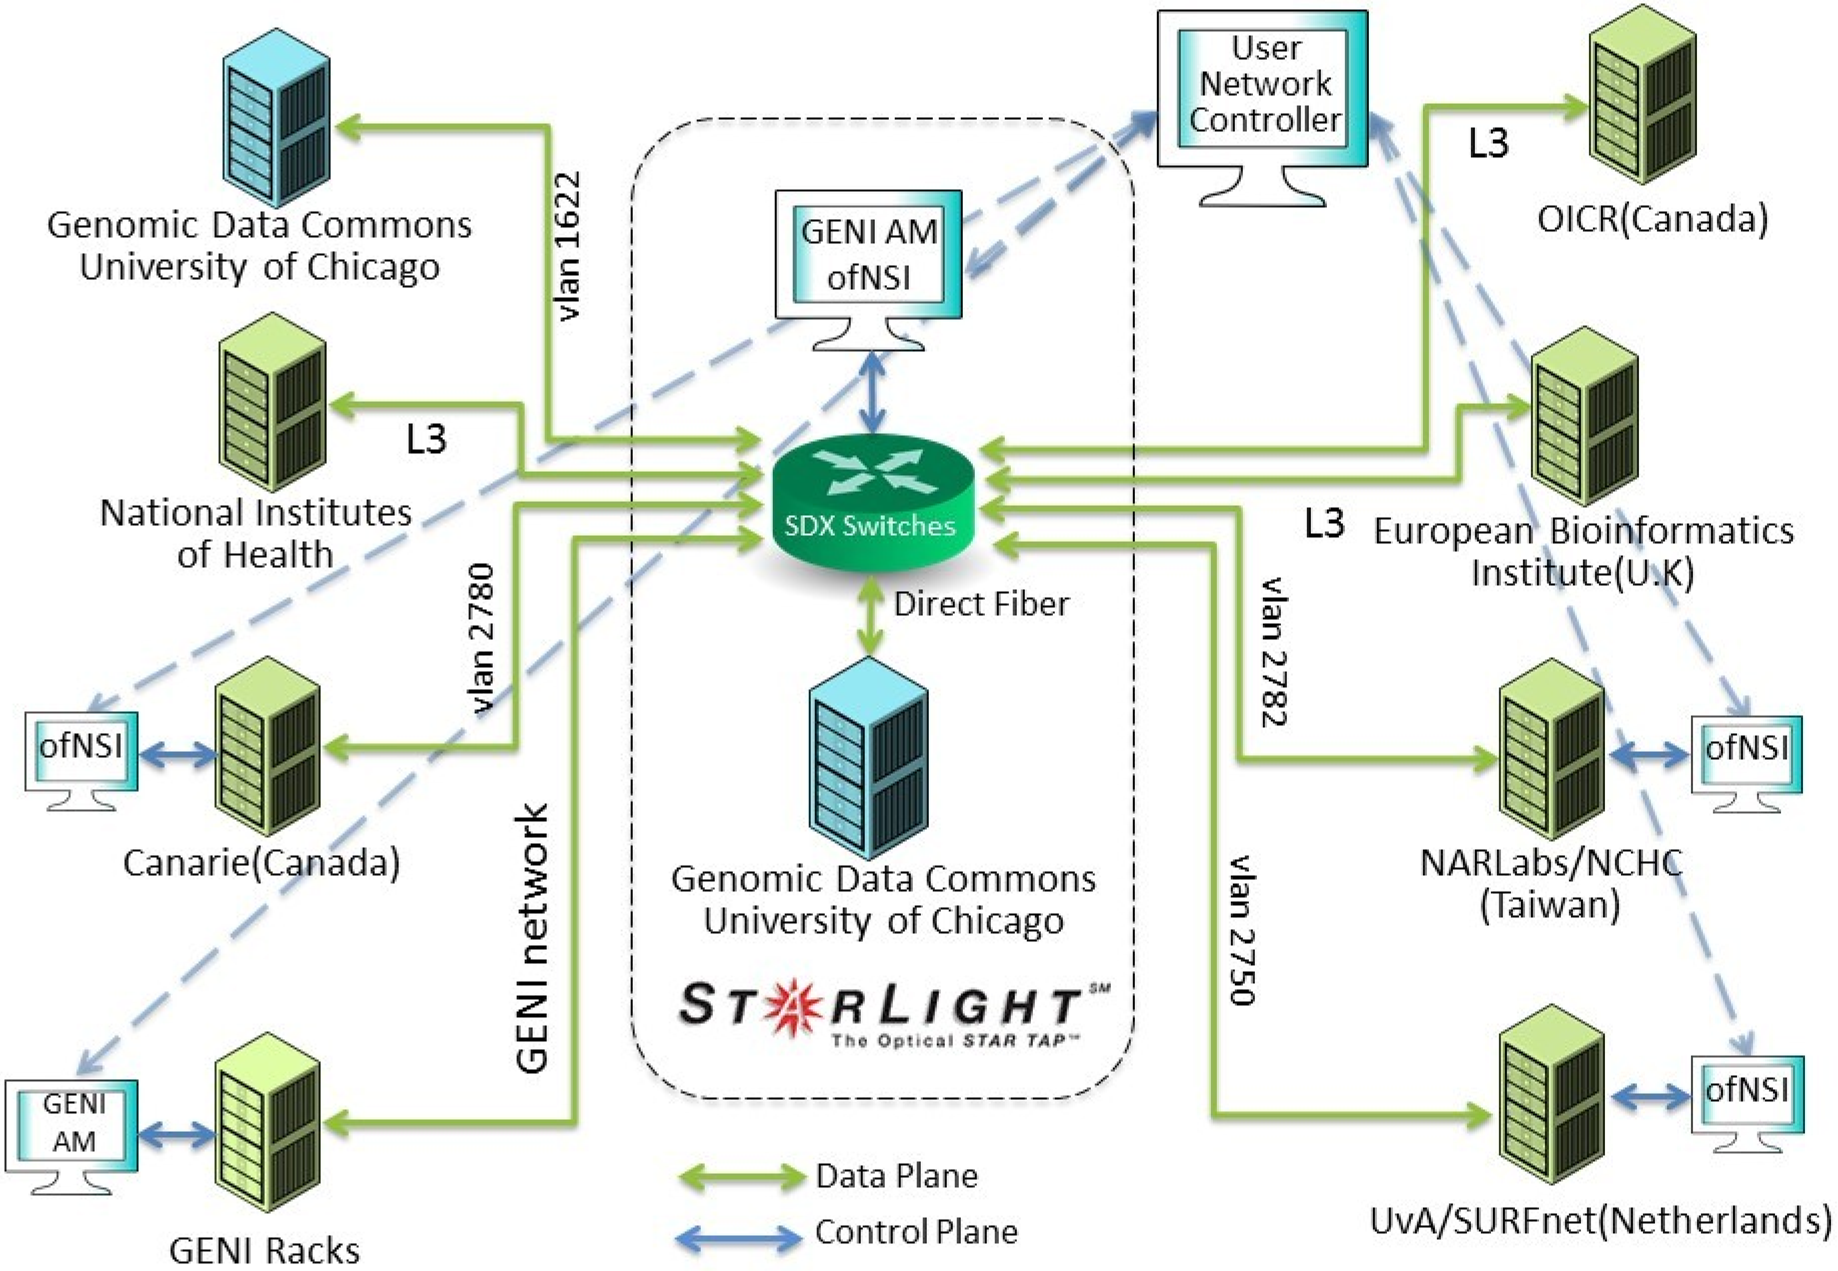

Parcel Based Collaboration

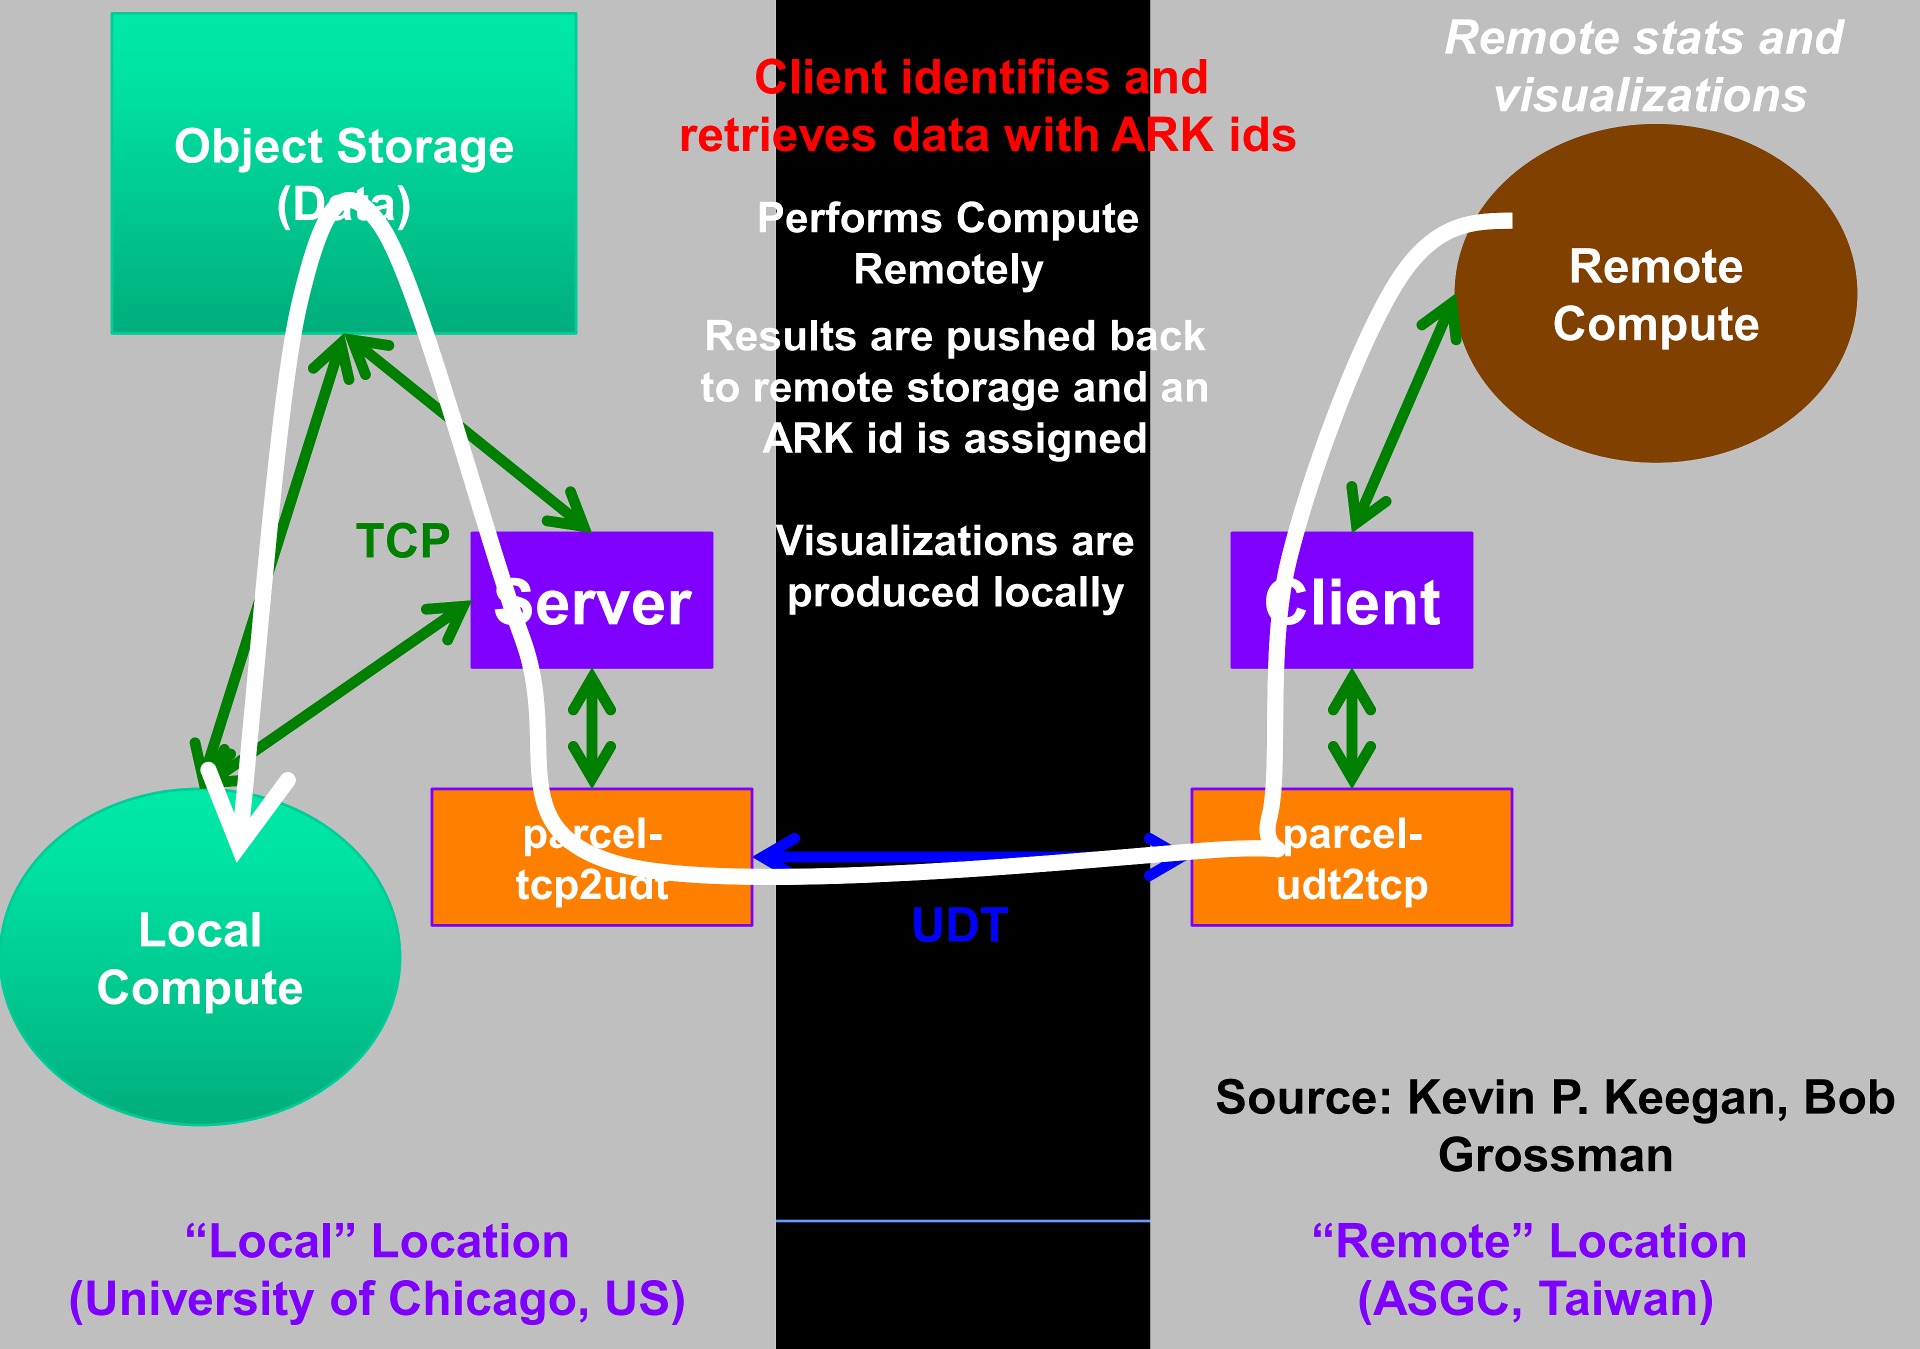

Planned US SDX Interoperable Fabric

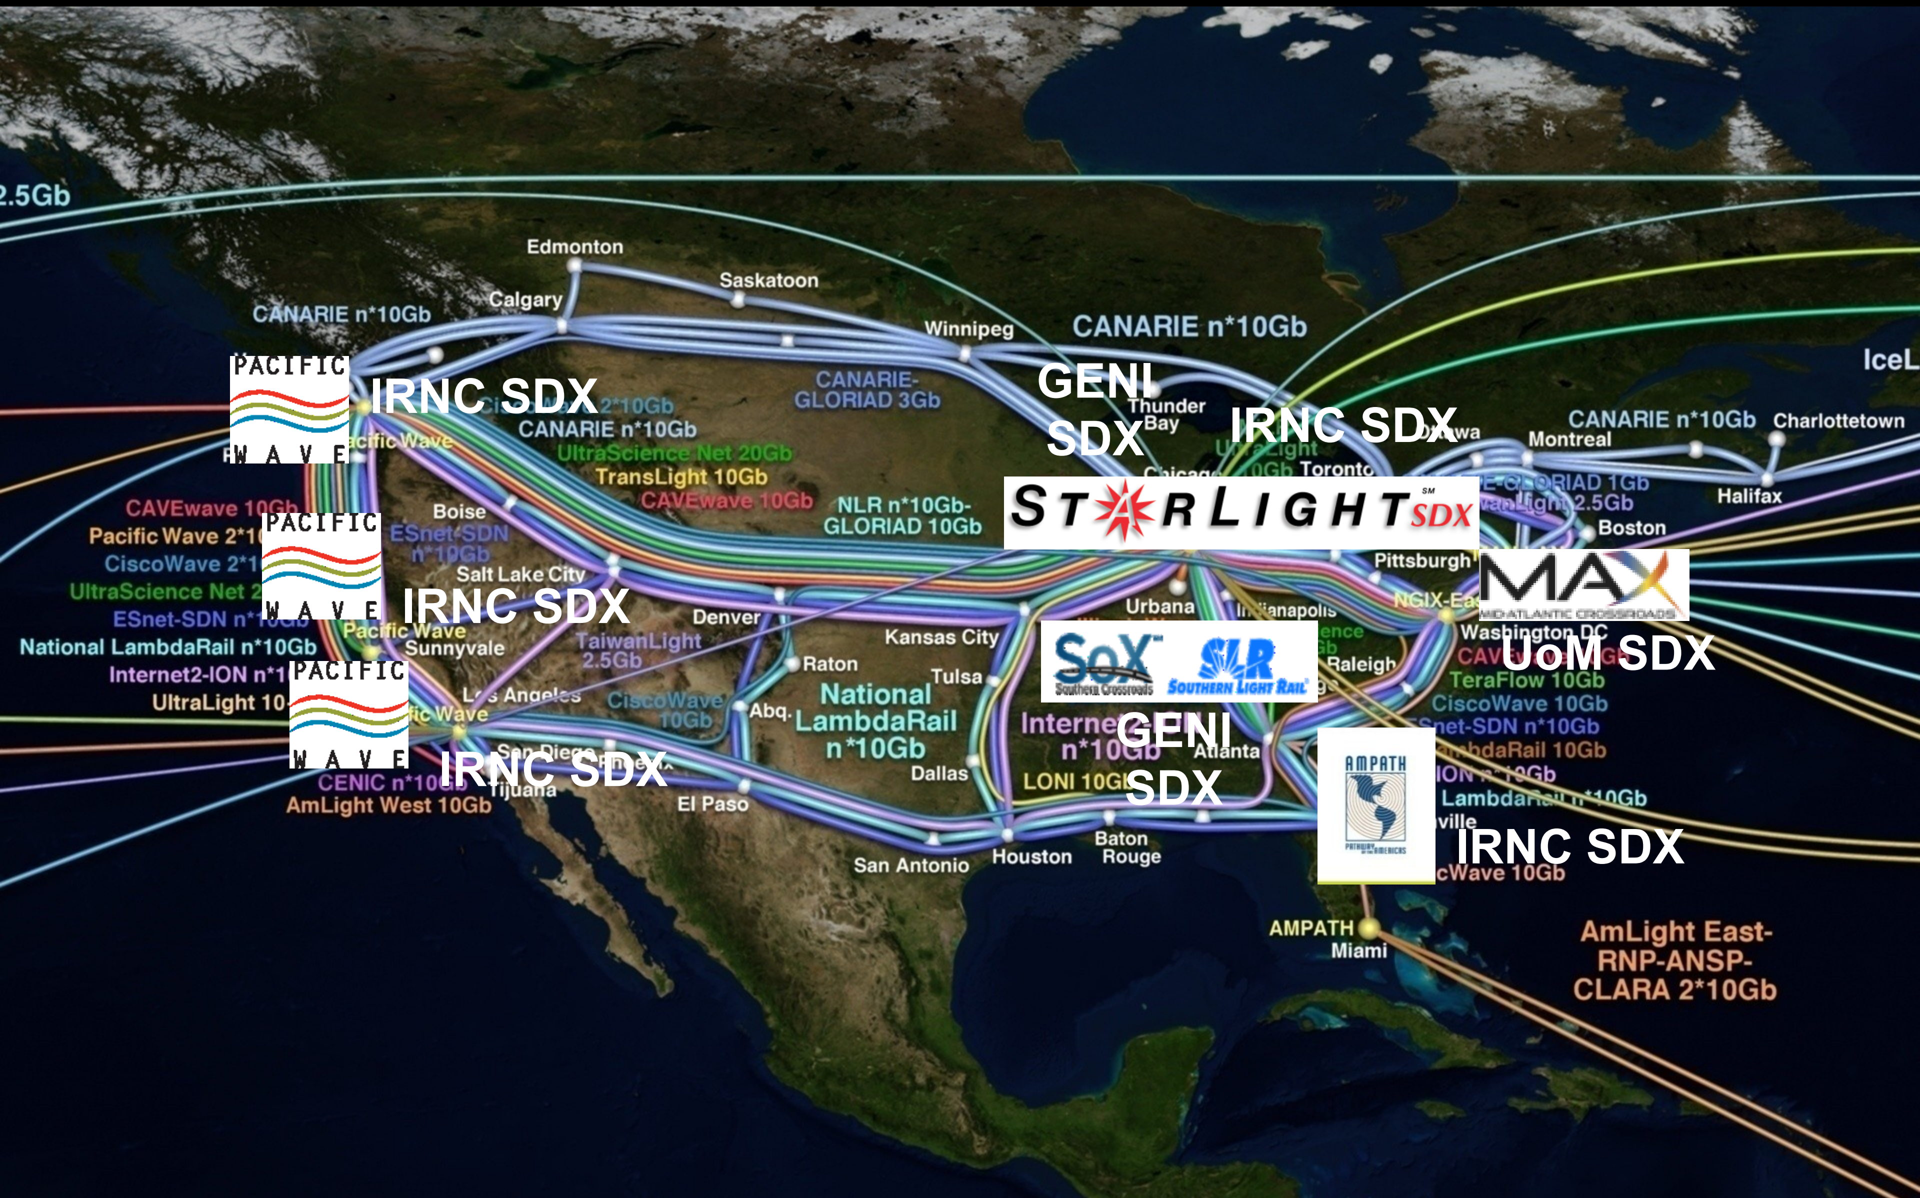

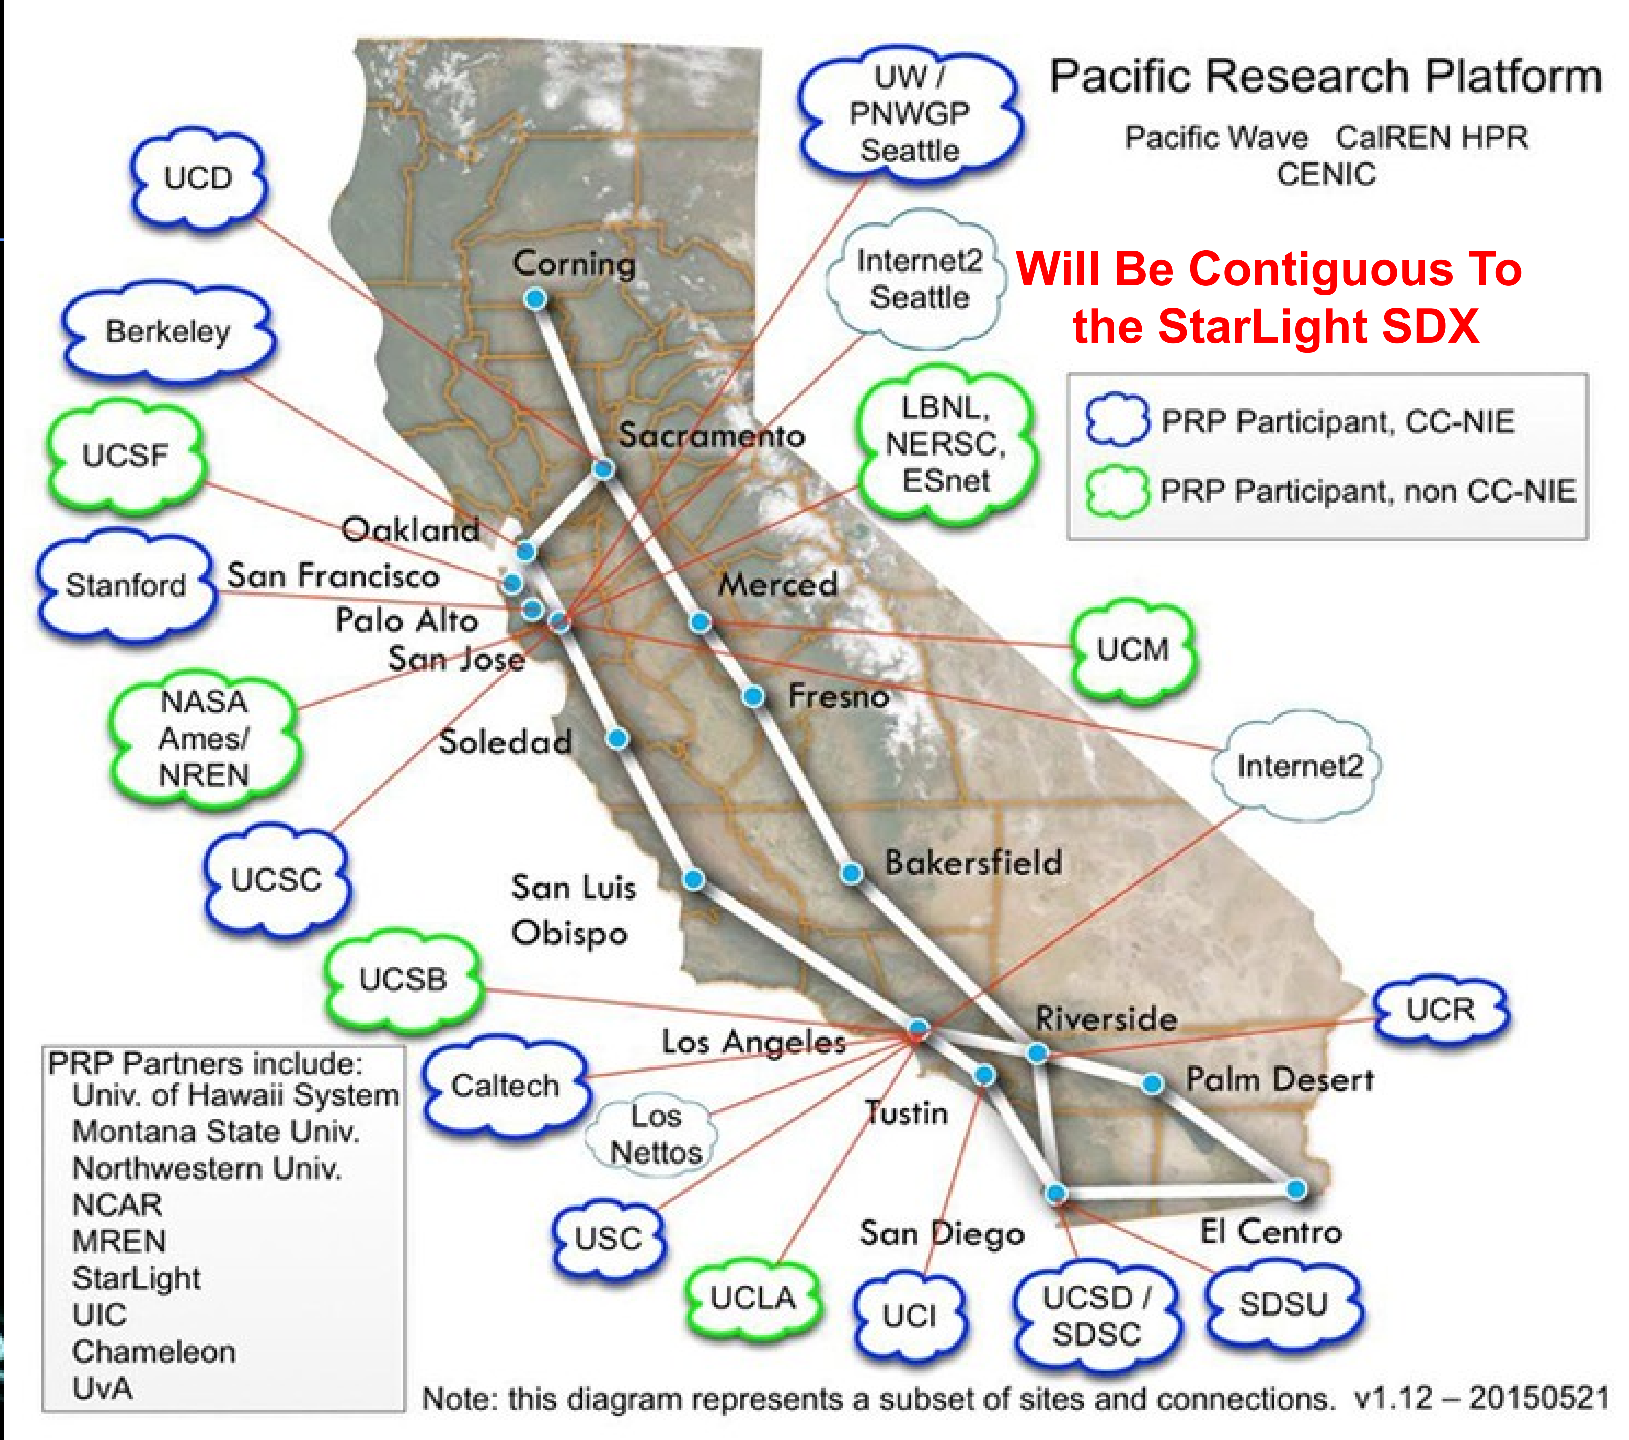
Pacific Research Platform
Pacific Wave CalREN HPR
CENIC
Will Be Contiguous To
the StarLight SDX
PRP Participant, CC-NIE
PRP Participant, non CC-NIE
UW / PNWGP Seattle
Internet2 Seattle
LBNL, NERSC, ESnet
UCD
Berkeley
UCSF
Stanford
NASA Ames/ NREN
UCSC
UCSB
Caltech
Los Nettos
USC
UCLA
UCI
UCSD / SDSC
SDSU
UCR
UCM
Internet2
Corning
Sacramento
Oakland
San Francisco
Palo Alto
San Jose
Merced
Fresno
Soledad
Bakersfield
San Luis Obispo
Los Angeles
Riverside
Palm Desert
Tustin
San Diego
El Centro
PRP Partners include:
Univ. of Hawaii System
Montana State Univ.
Northwestern Univ.
NCAR
MREN
StarLight
UIC
Chameleon
UvA
Note: this diagram represents a subset of sites and connections. v1.12 – 20150521

Next Step: Global Research Platform
Building on CENIC/Pacific Wave and GLIF

Current International GRP Partners

Global Research Platform

- A Emerging International Fabric
- A Specialized Globally Distributed Platform For Science Discovery and Innovation
- Based On State-Of-the-Art-Clouds
- Interconnected With Computational Grids, Supercomputing Centers, Specialized Instruments, et al
- Also, Based On World-Wide 100 Gbps Networks
- Leveraging Advanced Architectural Concepts, e.g., SDN/SDX/SDI – Science DMZs
- Ref: Demonstrations @ SC15, Austin Texas November 2015
- *New=> Global Research Platform 100 Gbps Network (GRPnet) On Private Optical Fiber Between PacificWave and StarLight via the PNWGP*

STARLIGHT SDX

SDX Services For Superchannels & DTNs

- **Many Architectural and Technology Options Exist For Both Superchannels and DTNs**
- **Different Types Of DTNs Are Being Developed By Many Research Organizations:**
  - ESnet
  - NASA
  - Naval Research Lab
  - iCAIR (PetaTrans DTN)
  - NCSA
  - SDSC
  - Fermi National Accelerator Laboratory
  - Brookhaven National Laboratory
  - Ciena Research
  - Dell Research
  - etc

STARLIGHT SDX

StarLight PetaTrans DT Nodes

PetaTrans Gen 1

- 1 Supermicro X9DR3-F Dual SNB EP E5-2600
- 16 DIMM slots, 512GB Reg. ECC DDR3
- 1 CS8457 - Supermicro SC846A-R1200B 4U
- 2 Intel Xeon E5-2667 v2 3.30Ghz 8C 25M
- 16 Crucial 16GB DDR3 1866MHz ECC
- 3 LSI MegaRAID - LSI00330 - SAS 9271-8i
- 1 LSI MegaRAID SAS 9286-8e (SGL)
- 32 840 PRO 2.5" 250GB SATA III SSD
- 1 BF3205 - Any 2U JBOD (no motherboard)
- 1 CS8654 - Supermicro CSE-826BA-R920LPB,
- 1 MCX456A-ECAT ConnectX-4 VPI EDR IB (100Gb/s) and 100GbE, dual-port QSFP

PetaTrans Gen 2

1 Supermicro X10DRi
1 CS8457 - Supermicro 3U Chassis
2 Intel Xeon E5-2667 v2 3.30Ghz 8C 25M
8 16GB DDR4 2133MHz ECC
2 2.5" 256GB Internal SATA III SSD
1 MCX456A-ECAT ConnectX-4 VPI EDR IB (100Gb/s) and 100GbE, dual-port QSFP

100 Gbps Prototype Switching Components

Inventec 100G Switch and Pica8 Switch OS

- Inventec DCS7032Q28 Bare Metal Switch
- Switching chip TD2 Broadcom BCM56854
- X86 CPU Intel C2538-2.4GHz FH8065501516762S R1S9
- P2041 CPU Freescale P2041NSN7PNC 1.5GHz 1.0V FCPBGA780 FREESCALE
- 32 X QSFP28 ports, Support 10GE/25GE/40GE/100GE
- ONIE Support
- OSs: iCOS, BCM, ONL, PicOS & others

- Network operation system using user space standard Debian Linux environment
- OpenFlow 1.4 support though OVS 2.0
- CrossFlow mode, support Layer-2 / Layer-3 and OpenFlow simultaneously
- Support OpenFlow to control MPLS, GRE, NVGRE or VXLAN tunnels
- Support for all major OpenFlow controllers (OpenStack Neutron ML2, OpenDaylight, Ryu)
- Support VXLAN network virtualization and NAT
- ONIE support

Emerging WAN Architectural Models

- **Current Generic WAN Model**
  - ES⇔DC SW⇔DC Routers⇔WAN⇔DC Routers⇔DC SW⇔ES
- **Current Experimental Model Enabled By SDXs**
  - CES⇔DTN⇔DC SC SWs⇔WAN⇔DC SC SWs⇔DTN⇔CES
- **Potential Future WAN Model Enabled By SDXs**
  - ES⇔DTN⇔DC NGD⇔SDX⇔WAN⇔SDX<=>DC NGD⇔DTN⇔ES
  - ***NB: Attempted Only In Prototype Not In Production***

Key Emerging Enabling NGD Technology

- Support for L2/L3 Protocols and MAC Interfaces at 1 Gb/sec, 10 Gb/sec, 40 Gb/sec, and 100 Gb/sec
- a) Deep Buffers -- Gigabytes (V. Important For Large Scale Flows)
- vs Current General 16 MB
- b) Large Routing Tables (Scales To Millions of Routes & MAC Addresses)
- c) Ultra Fast Packet Processing
- d) Virtual SWs (Segmentable Switches)
- e) Support for Modular Switches (Can Link Line Cards Within Chassis To Create Scalable SW w/ Large # Of Ports)
- f) Quality of Service Features
- g) Options For Allocating Specific Capacity To Individual Subscribers Encapsulated Within Chassis
- h) Core Support for NFV
- i) Addresses Networks W/ High Variance/Mixes In Speeds (Can Generate Congestion)

KREONet2 and GLORIAD-KR
KISTI Daejeon ⇔ 100 G ⇔ StarLight

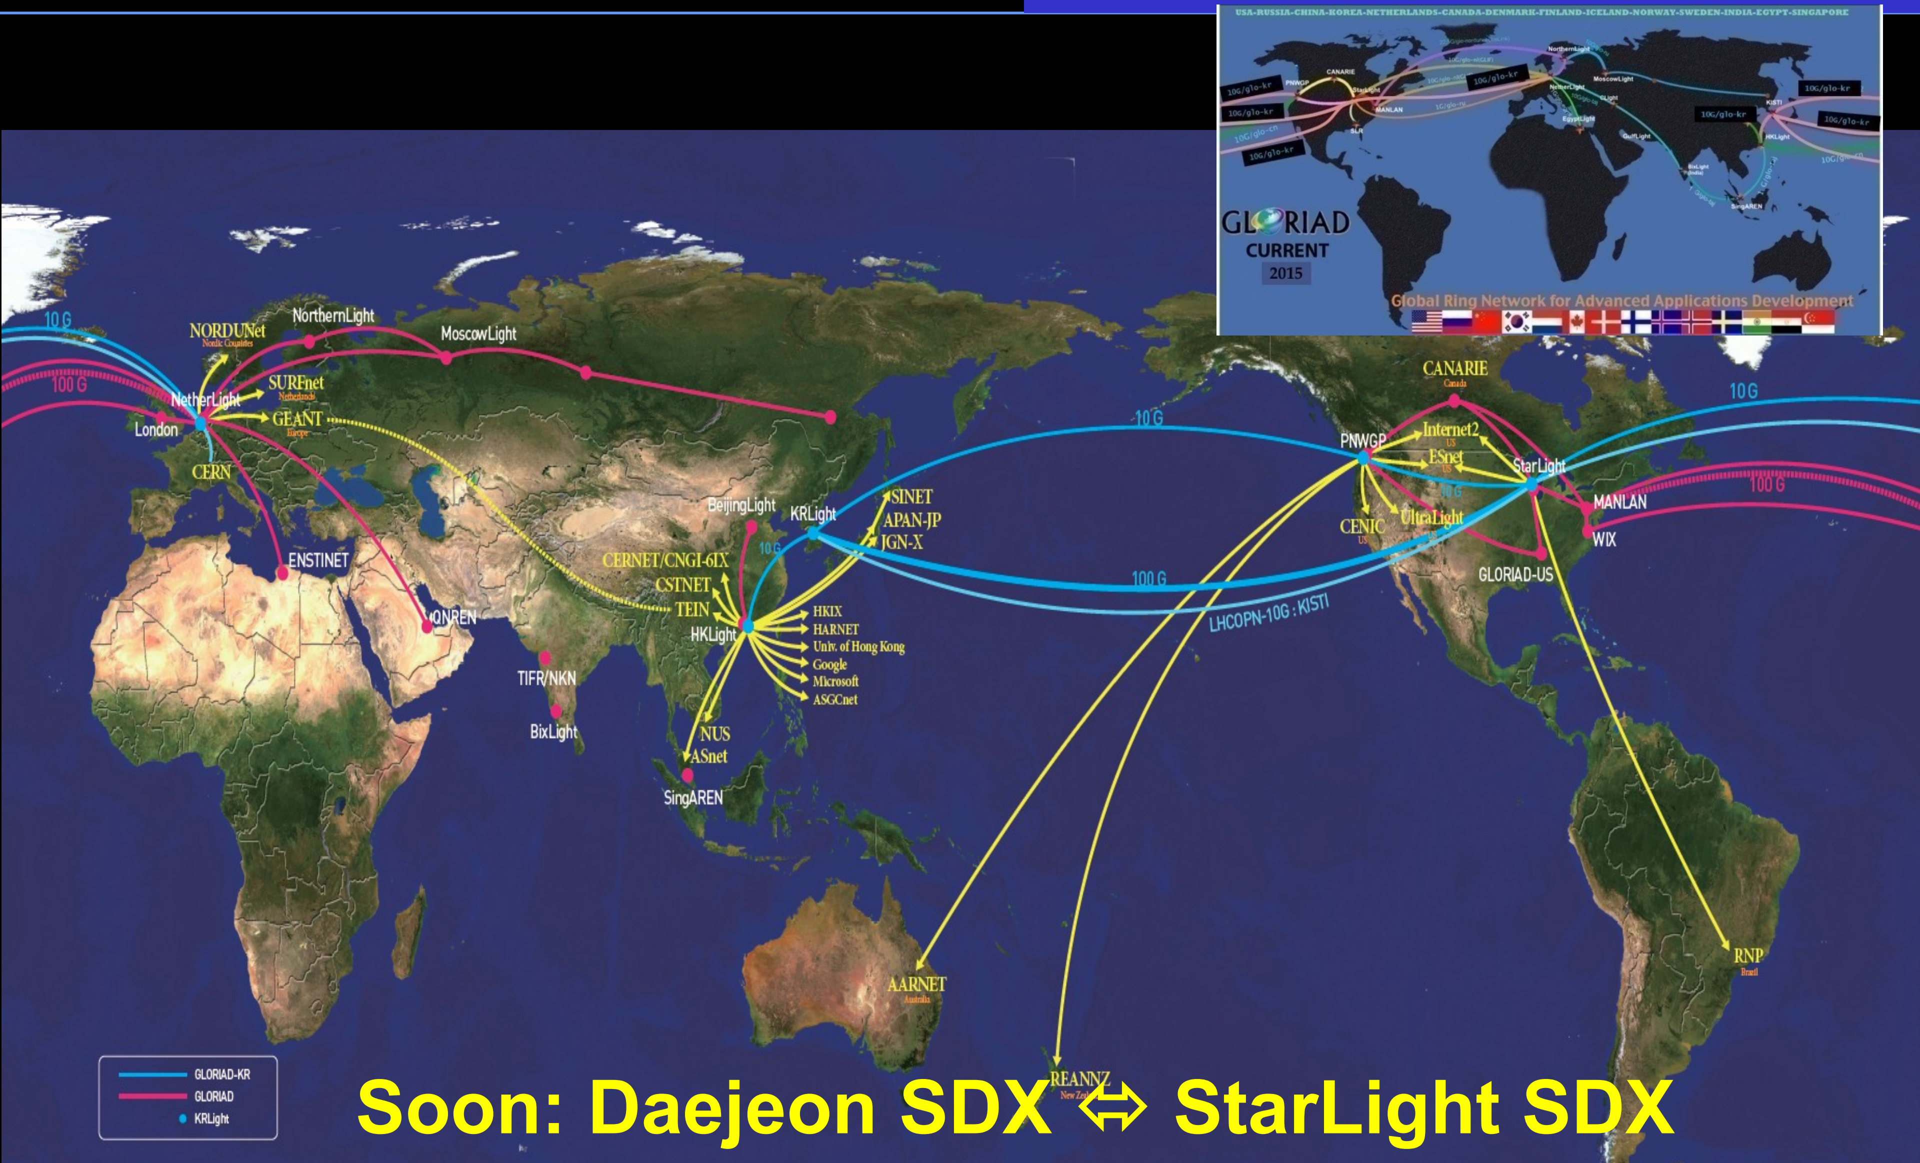

Soon: Daejeon SDX ⇔ StarLight SDX

KREONET vs. KREONET-S

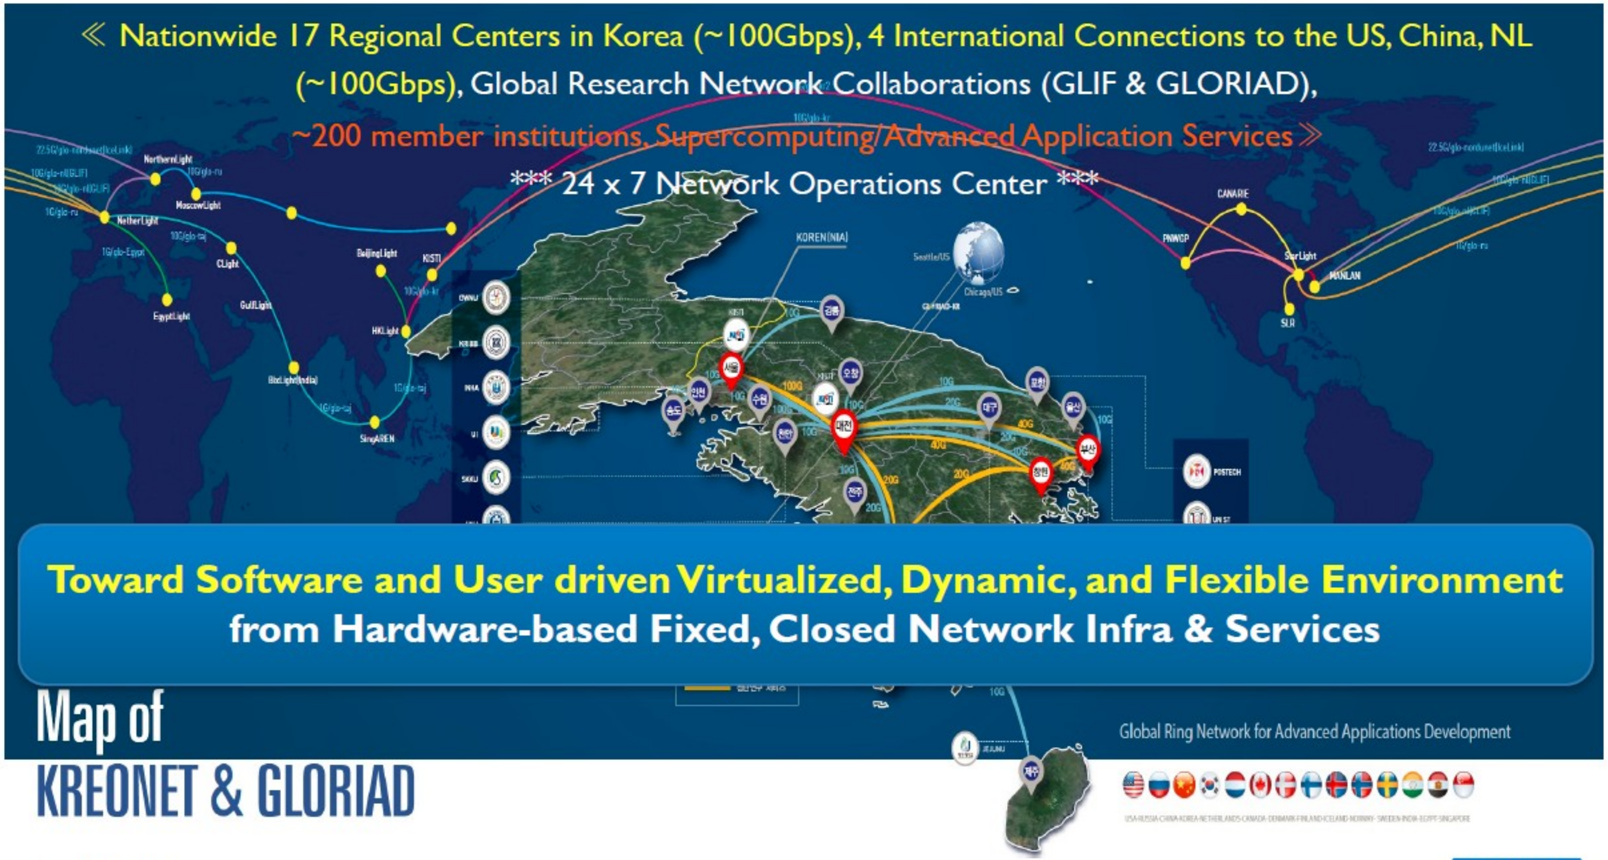

Toward Software and User driven Virtualized, Dynamic, and Flexible Environment from Hardware-based Fixed, Closed Network Infra & Services

Science Information Network (SINET)

SINET is a Japanese academic backbone network for more than 800 universities and research institutions, and current version, SINET4, started its operation in 2011.

- SINET4 covers all 47 prefectures.
- SINET4 covers 100% of national, 71% of public, and 53% of private universities.

(As of March 2014)

	National Universities	Public Universities	Private Universities	Junior Colleges	Colleges of Technology	Inter-Univ. Research Institutes	Others	Total
Number of Organizations	86	68	333	60	55	16	184	802

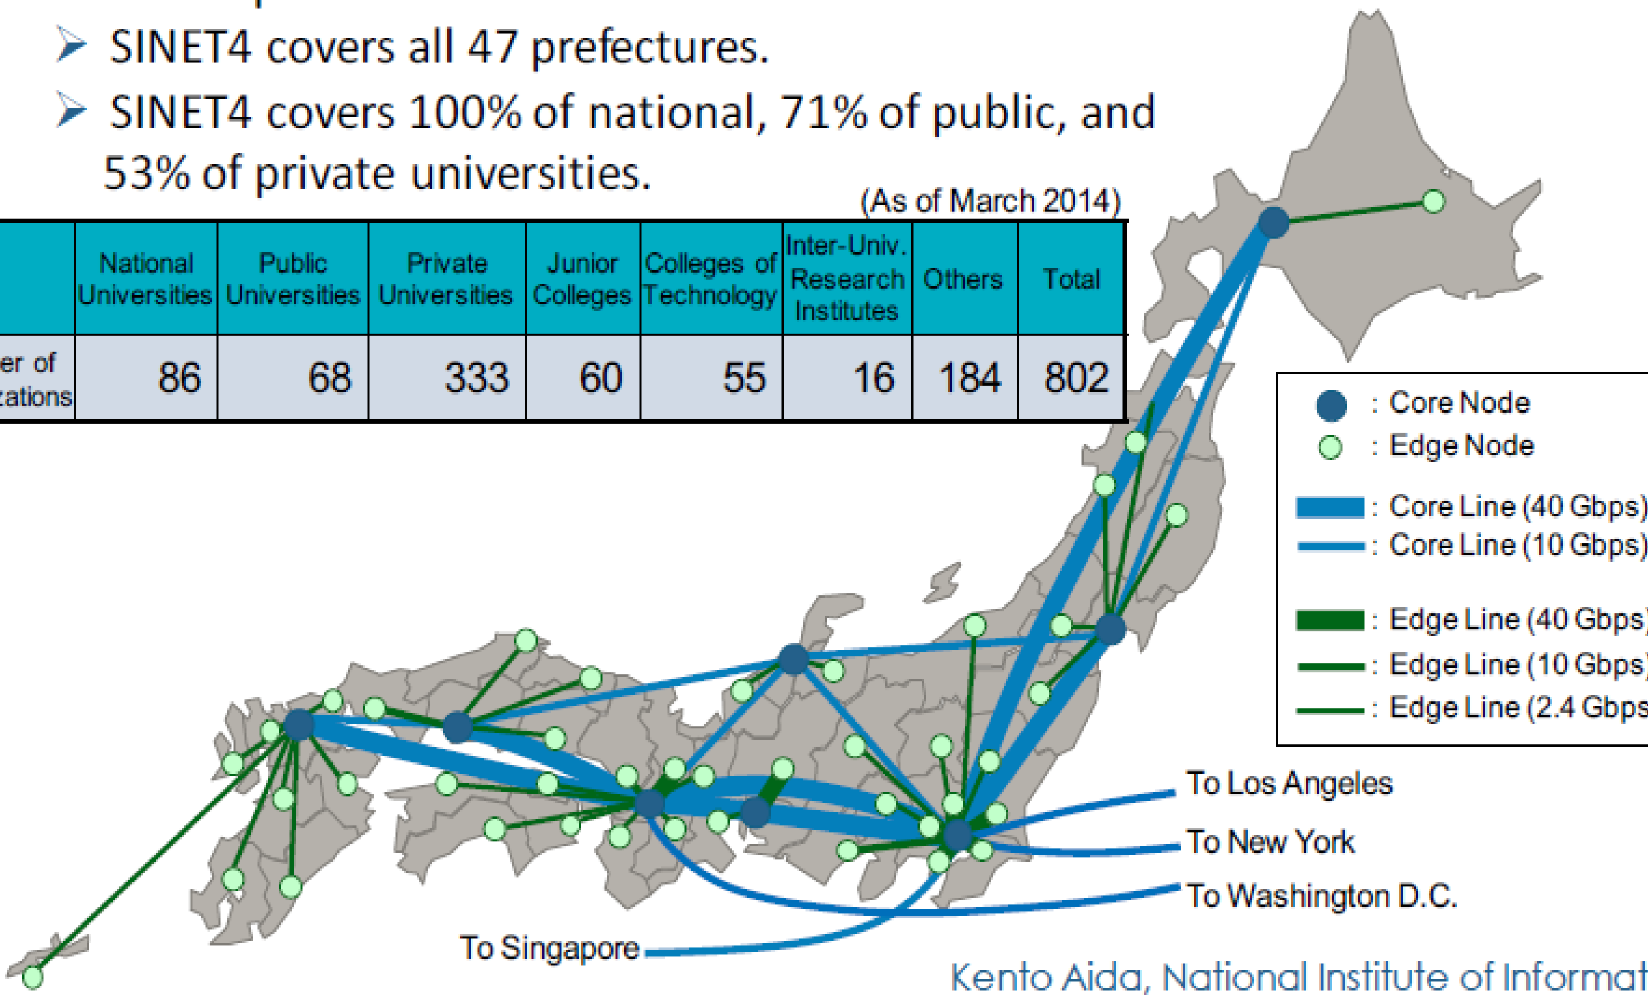

An Experimental Testbed For Computer Science Research

Chameleon

www.chameleoncloud.org

CHAMELEON:
A LARGE-SCALE, RECONFIGURABLE EXPERIMENTAL ENVIRONMENT FOR CLOUD RESEARCH

Principal Investigator: Kate Keahey

Co-PIs: J. Mambretti, D.K. Panda, P. Rad, W. Smith, D. Stanzione

Another SDX Opportunity – Especially For Federation!
SDX=> "Federation As A Service" Federated With GENI and Canadian SAVI

THE UNIVERSITY OF CHICAGO TACC NORTHWESTERN UNIVERSITY The Ohio State University UTSA NSF

STARLIGHT SDX

Testbed Federations

- GENI⬄FIRE
- GENI⬄Chameleon
- Chameleon⬄FIRE
- GENI⬄SAVI (Smart Applications On Virtual Infrastructure)
- Chameleon⬄SAVI
- GENI⬄Vnode⬄iCAIR Testbed
- Chameleon⬄ExoGENI⬄iCAIR Testbed
- GENI⬄CloudLab
- iCAIR Testbed⬄G-Lab Tomato
- In Process:
  - Chameleon⬄Grid'5000⬄GENI
  - GENI⬄Chameleon⬄ Open Science Data Cloud
  - Chameleon⬄XSEDE
  - GENI⬄MAKI ??
  - Etc.
  - "Federation-as-a-Service" ? Enabled By SDXs?

www.startap.net/starlight

**Thanks to the NSF, DOE, NIH, USGS, DARPA
Universities, National Labs,
International Partners,
and Other Supporters**

iCAIR

STARLIGHT SDX