StarLight SDX A Software Defined Network Exchange for Global Science Research and Education

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)
Co-Director, StarLight (www.startap.net/starlight)
PI IRNC: RXP: StarLight SDX
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, (maxine@uic.edu)
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

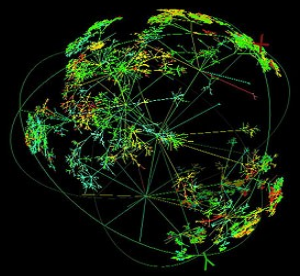

Global LambdaGrid Workshop
Prague, Czech Republic
Sept 29-30, 2015

iCAIR

STARLIGHT SDX

Program: NSF IRNC

- National Science Foundation Program
- Directorate for Computer & Information Science & Engineering (CISE)
- Division of Advanced Cyberinfrastructure
- NSF 14-554 International Research Network Connections (IRNC)
- Infrastructure and Innovation of U.S. R&E Open Exchange Points (IRNC: RXP)

IRNC: RXP: StarLight SDX Key Participants

- PI Joe Mambretti, Director, International Center for Advanced Internet Research
- Northwestern University, Director, Metropolitan Research and Education Network
- Co-Director, StarLight,
- Co-PI Tom DeFanti, Research Scientist, (tdefanti@soe.ucsd.edu)
- California Institute for Telecommunications and Information Technology (Calit2),
- University of California, San Diego
- Co-PI Maxine Brown, Director
- Electronic Visualization Laboratory, University of Illinois at Chicago
- Co-PI Jim Chen, Associate Director, International Center for Advanced Internet Research
- Northwestern University
- Senior Personnel
- 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
- John Graham, Senior Development Engineer Calit2 UCSD
- 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
- Also, Other Members of the StarLight Consortium, Multi National and International Partners

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 Over Optics – World's "Largest" 10G/100G Exchange

First of a Kind Enabling Interoperability At L1, L2, L3

Also, StarWave Multi-100 Gbps Exchange

View from StarLight

Abbott Hall, Northwestern University's Chicago Campus

iCAIR

STARLIGHT SDX

StarLight Software Defined Network Exchange (SDX)

- The StarLight SDX Will Provide The Services, Architecture, and Technologies Designed To Provide Scientists, Engineers, and Educators With Highly Advanced, Diverse, Reliable, Persistent, and Secure Networking Services, Enabling Them to Optimally Access Resources in North America, South America, Asia, South Asia (including India), Australia, New Zealand, Europe, the Middle East, North Africa, And Other Sites Around the World.
- The StarLight SDX Initiative Will Undertake Continued innovation and Development of Advanced Networking Services and Technologies.

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

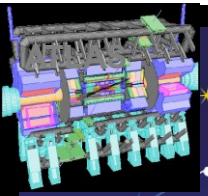

ALMA: Atacama Large Millimeter Array

ANDRILL: Antarctic Geological Drilling
www.andrill.org

BIRN

BIRN: Biomedical Informatics Research Network
www.nbirn.net

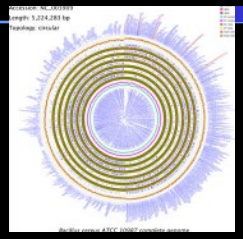

CAMERA metagenomics
camera.calit2.net

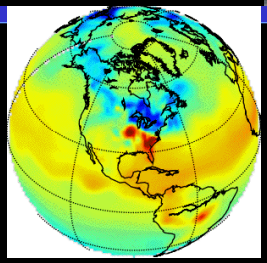

Carbon Tracker
www.esrl.noaa.gov/gmd/ccgg/carbontracker

CineGrid™
CineGrid
www.cinegrid.org

LHCONE
www.lhcone.net

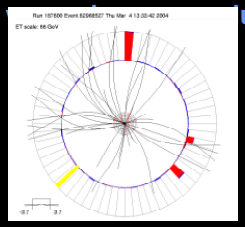

DØ (DZero)
www-d0.fnal.gov

GEON

GEON: Geosciences Network
www.geongrid.org

GLEON: Global Lake Ecological Observatory Network

OOI
OCEAN OBSERVATORIES INITIATIVE
CYBERINFRASTRUCTURE
Providing a link between ocean research and discovery

OOI-CI
ci.oceanobservatories.org

ISS: International Space Station
www.nasa.gov/station

CLASS

Comprehensive Large-Array Stewardship System
www.class.noaa.gov

LIGO
www.ligo.org

PRAGMA

Pacific Rim Applications and Grid Middleware Assembly
www.pragma-grid.net

WLCG
Worldwide LHC Computing Grid

WLCG
lcg.web.cern.ch/LCG/public/

TeraGrid™
TeraGrid
www.teragrid.org

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

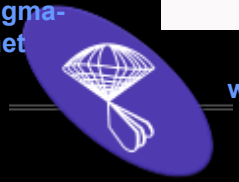

Sloan Digital Sky Survey
www.sdss.org

XSEDE

XSEDE
www.xsede.org

Compilation By Maxine Brown

STARLIGHT SDX

The Global Lambda Integrated Facility: a Global Programmable Resource

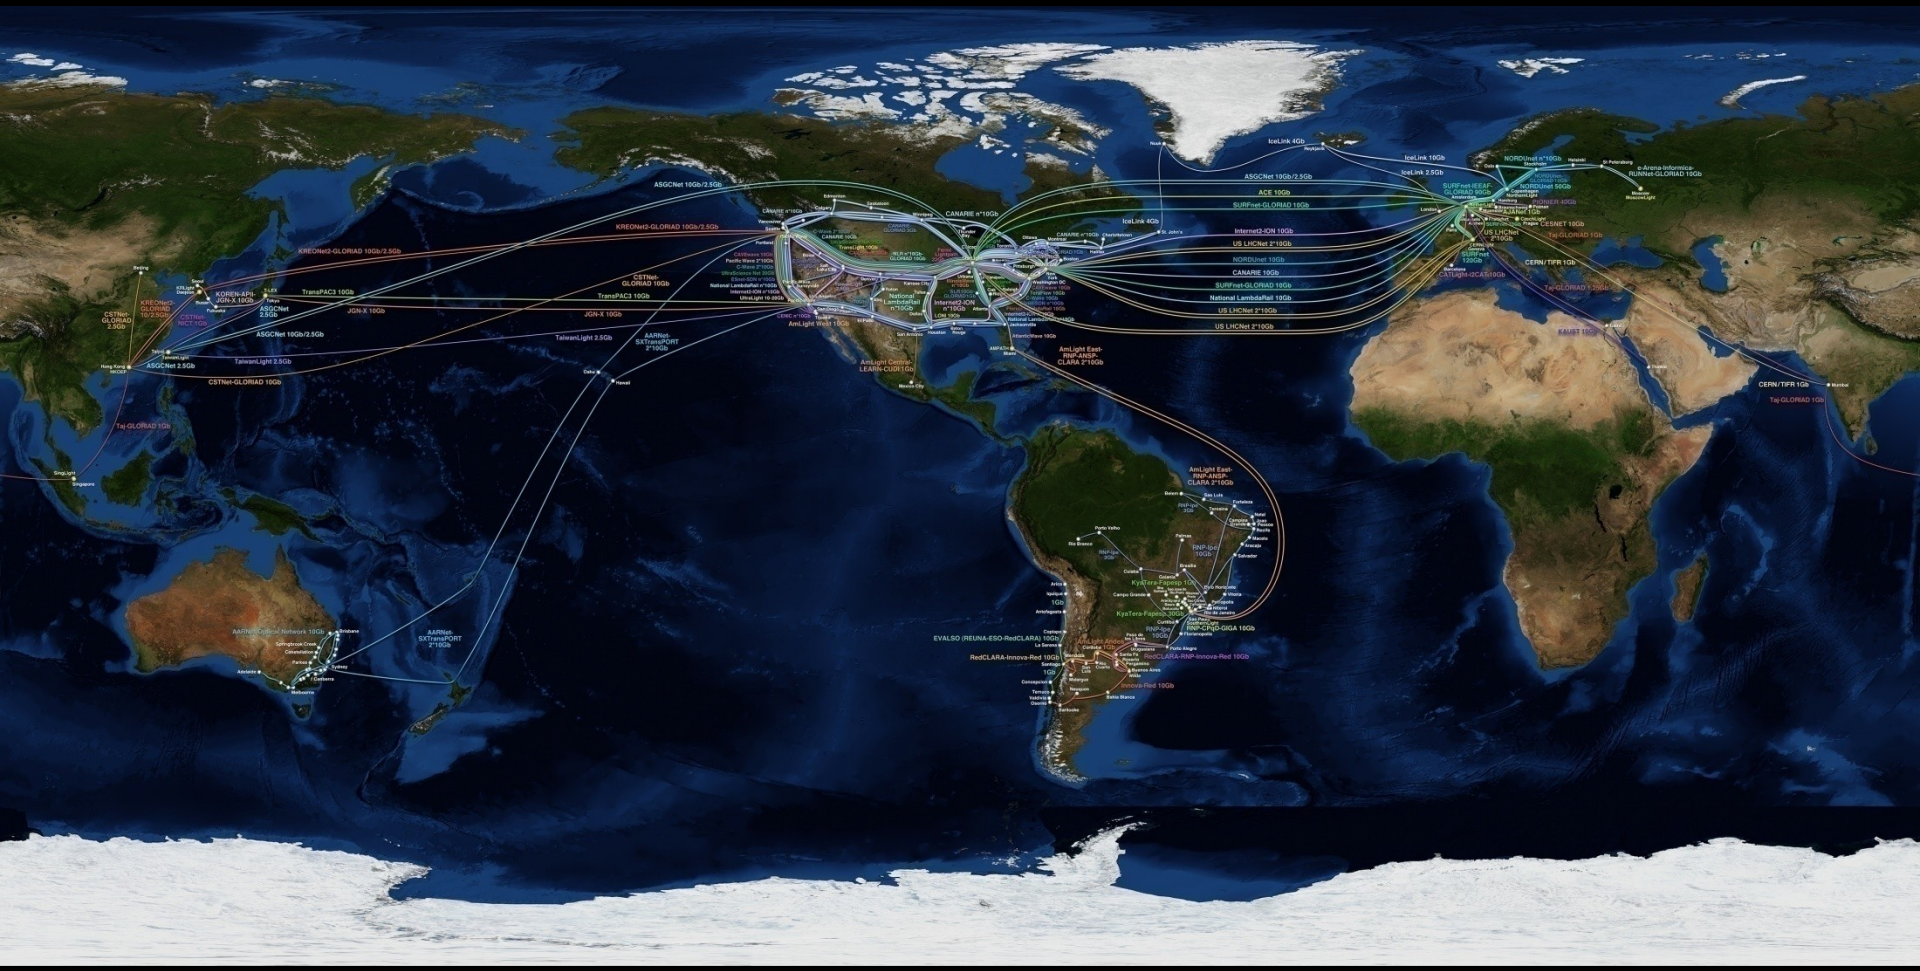

iCAIR

STARLIGHT SDX

Automated GOLE Fabric

Source: GLIF Auto GOLE Group

AutoGOLE Topology

Source: John MacAuley

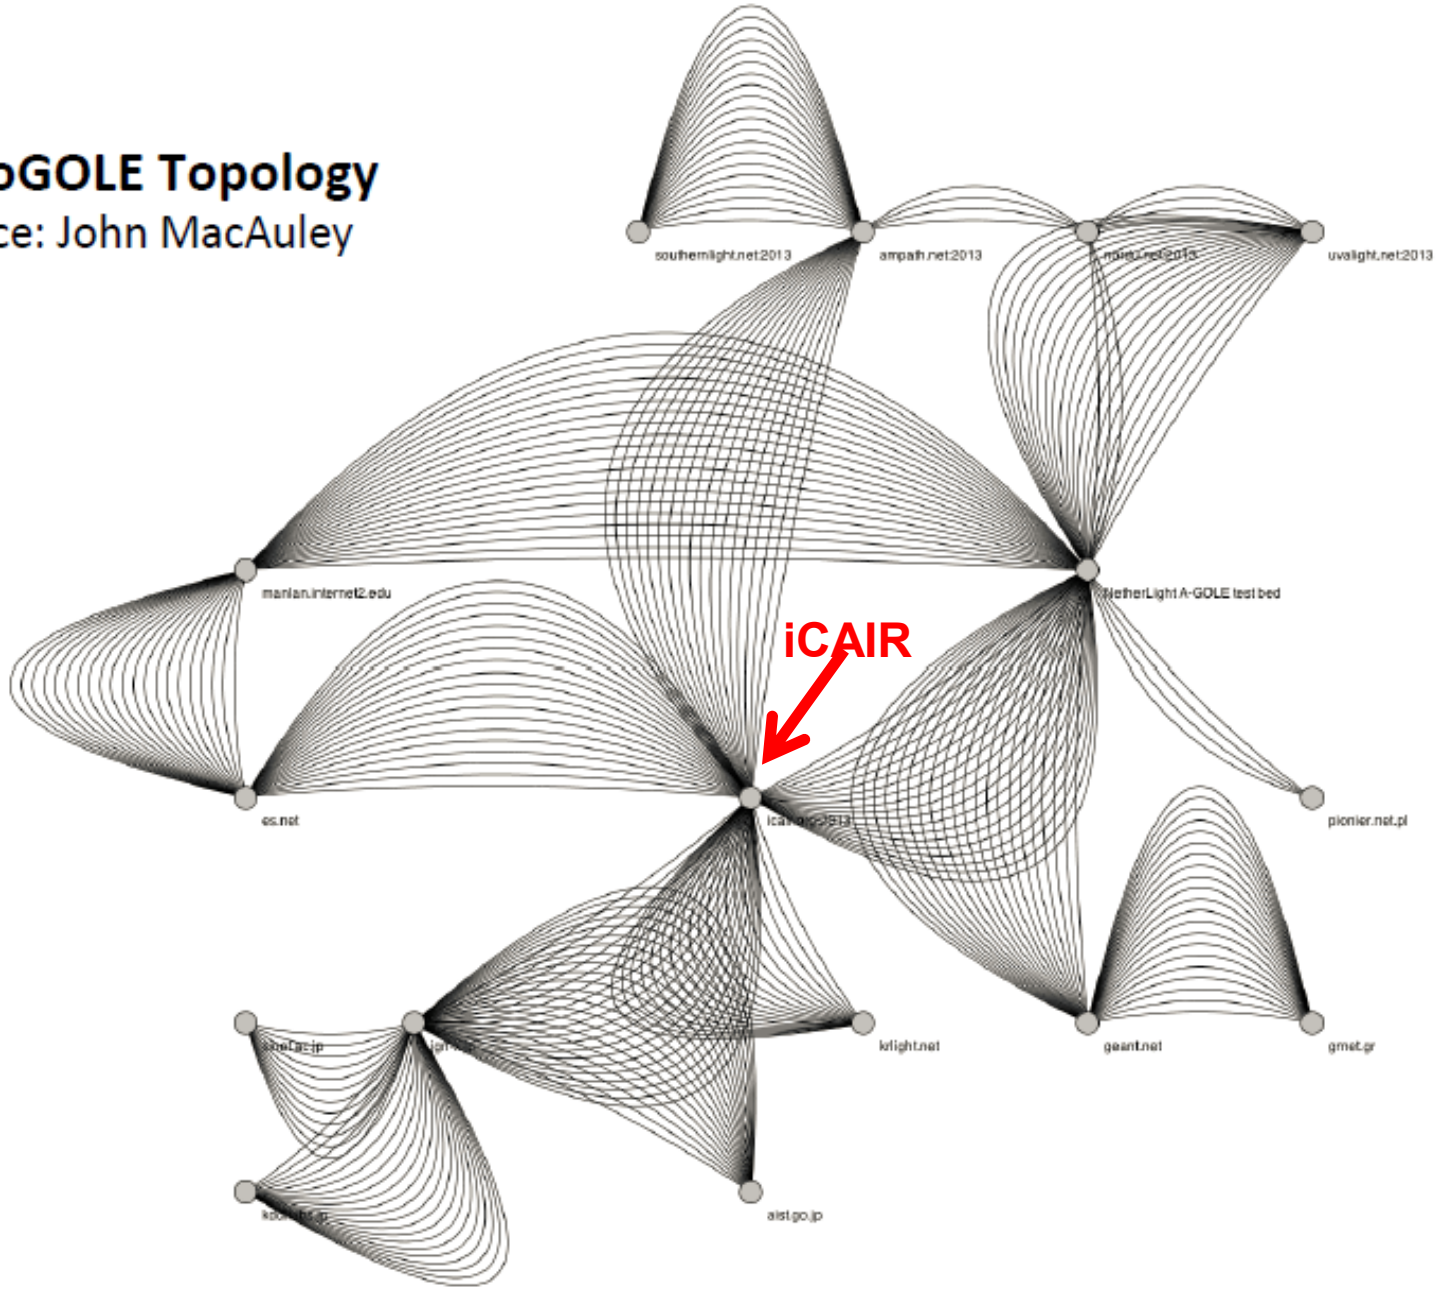

Source: John MacAuley

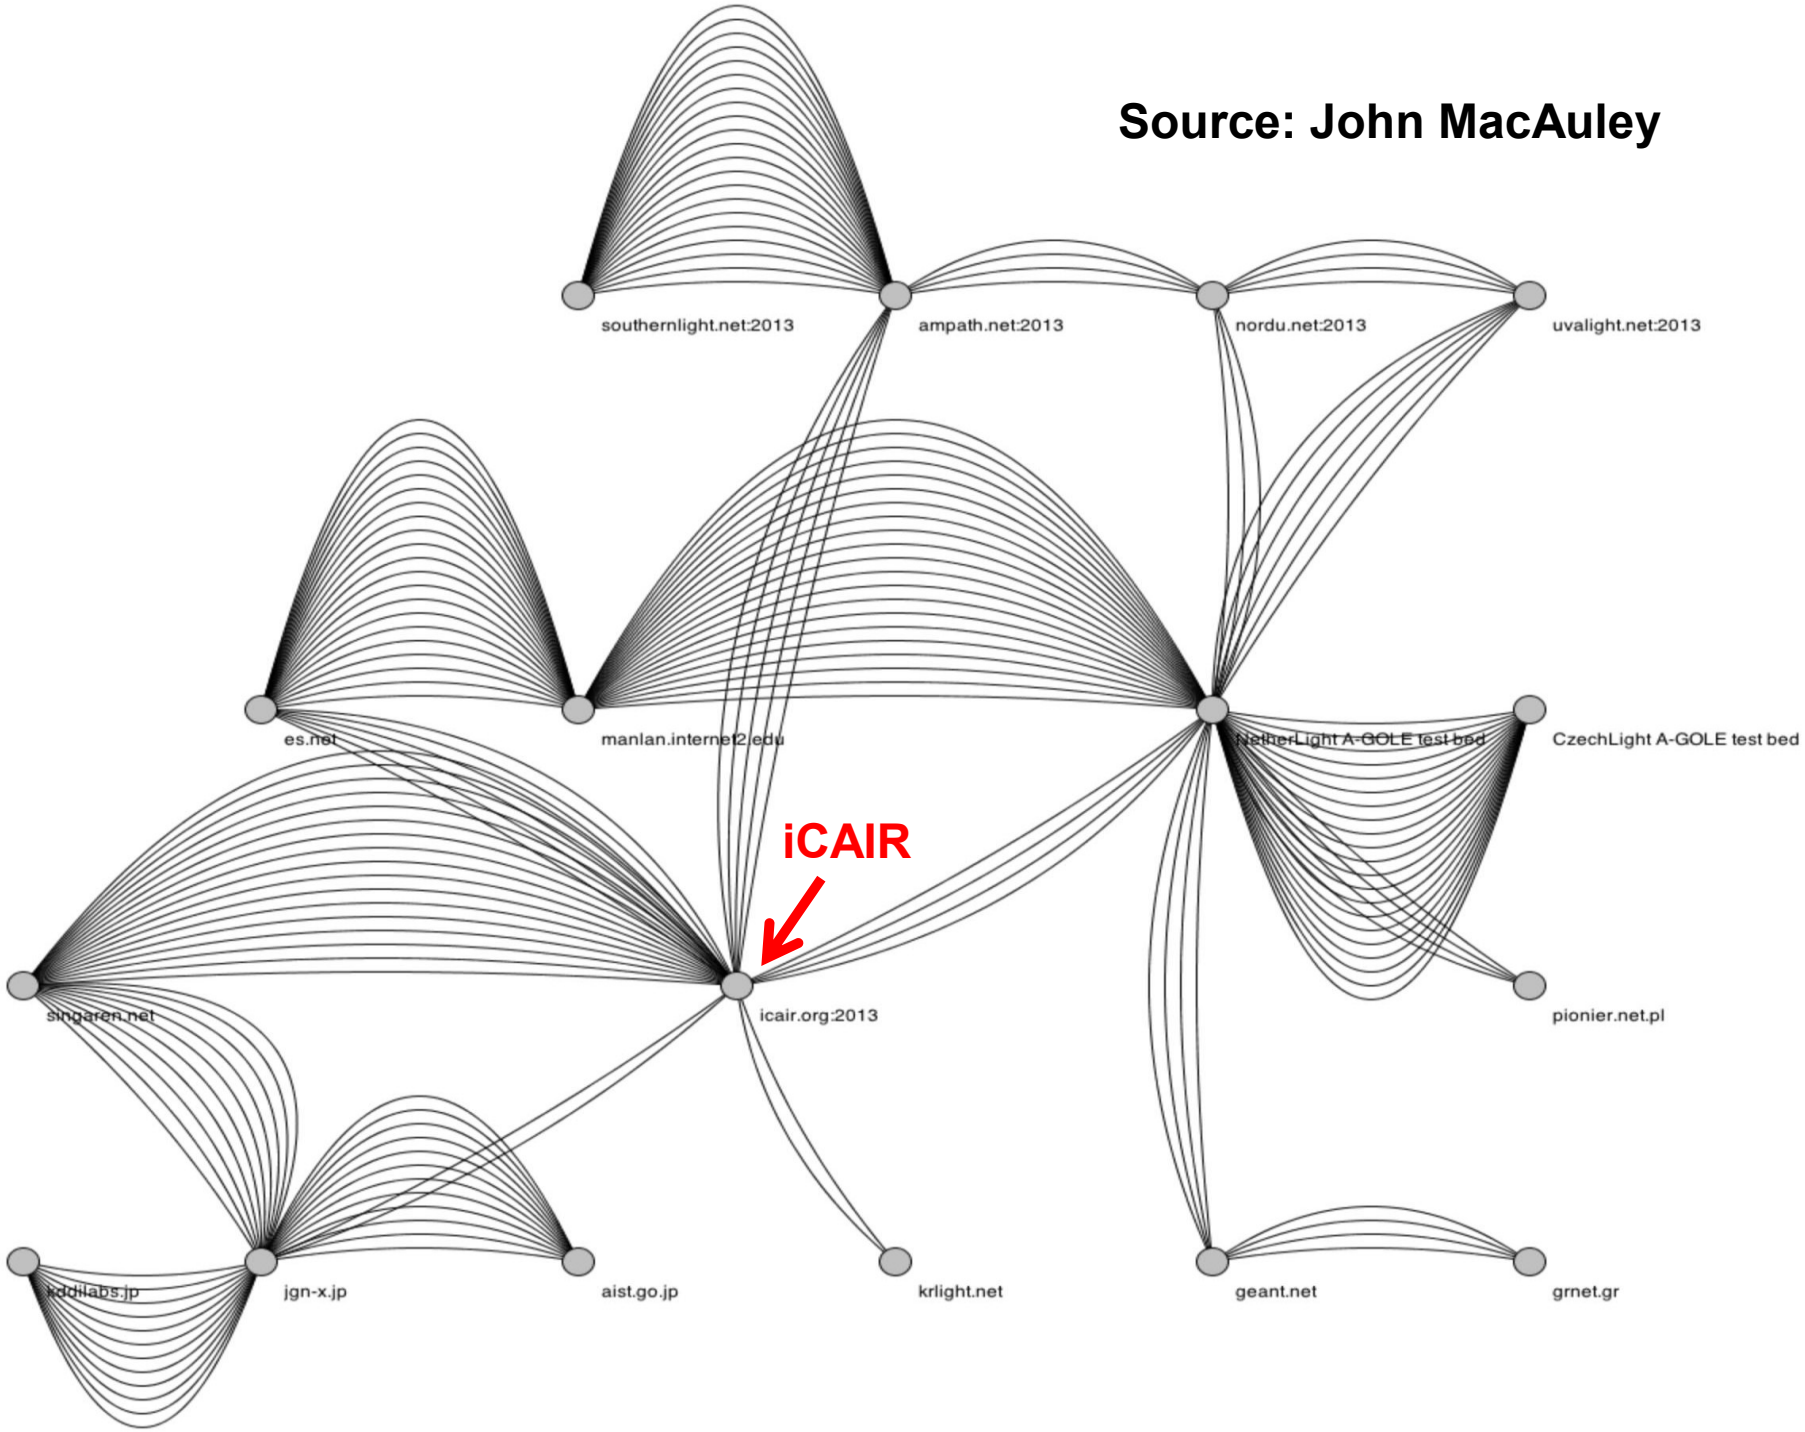

Tasks/Goals For 2014 Design and Implement ofNSI – 1st Showcased At SC14

Work items 2014

Item	Description	Due	Leading organization
Authentication / Authorization	Creating a AAI framework that allows secure setup of services	TNC2014	SURFnet (Hans Trompert)
Topology Exchange	Creating a mechanism that exchanges topology descriptions of GOLEs automatically	SC'14	ESnet, UvA (Chin Guok, Miroslav Zivkovic)
Retagging capabilities	Describing what's necessary to implement retagging capabilities inside the AutoGOLE fabric – also creating a plan for implementing	SC'14	Group effort
SDN/OpenFlow inside the AutoGOLE	It's foreseen that AutoGOLE NRMs could be talking OpenFlow to actual hardware. This item results in deployment of an OpenFLow controller speaking NSIv2 inside the AutoGOLE	Q4	iCAIR (Jim Chen, Joe Mambretti)
Operational items	Creating concepts on strengthening operations, implementing these	Q4	Tangui Coulouarn to look for someone to lead (uniform) perational issues

AutoGOLE Dashboard (In Development)

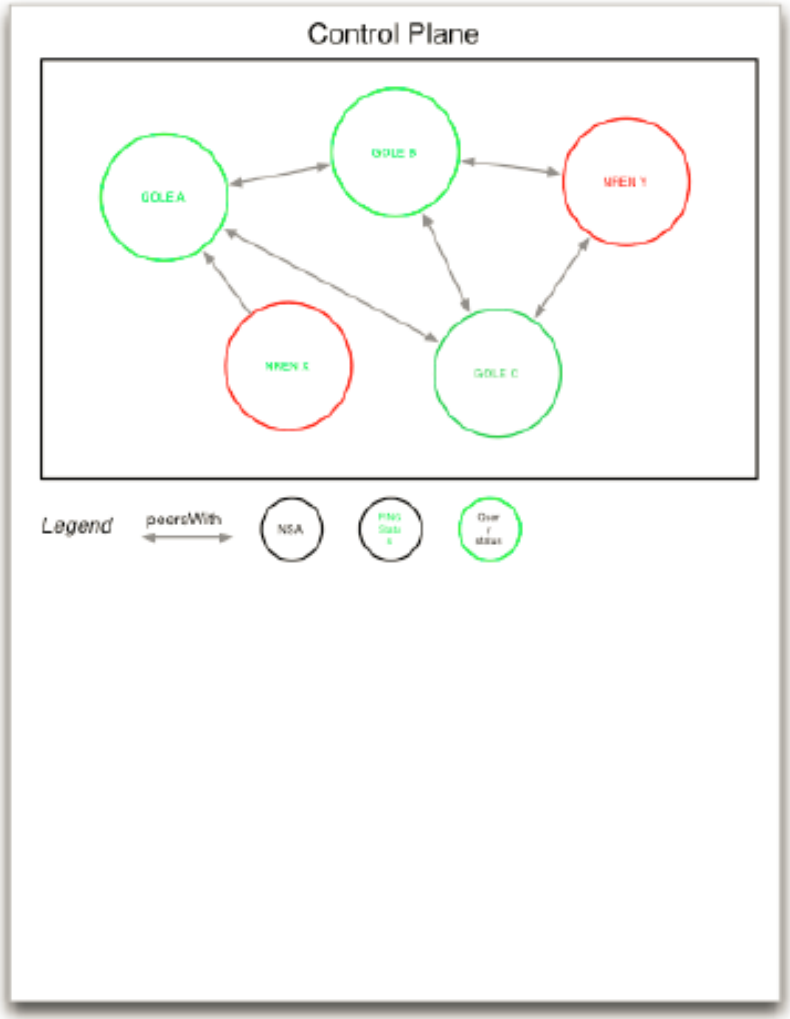

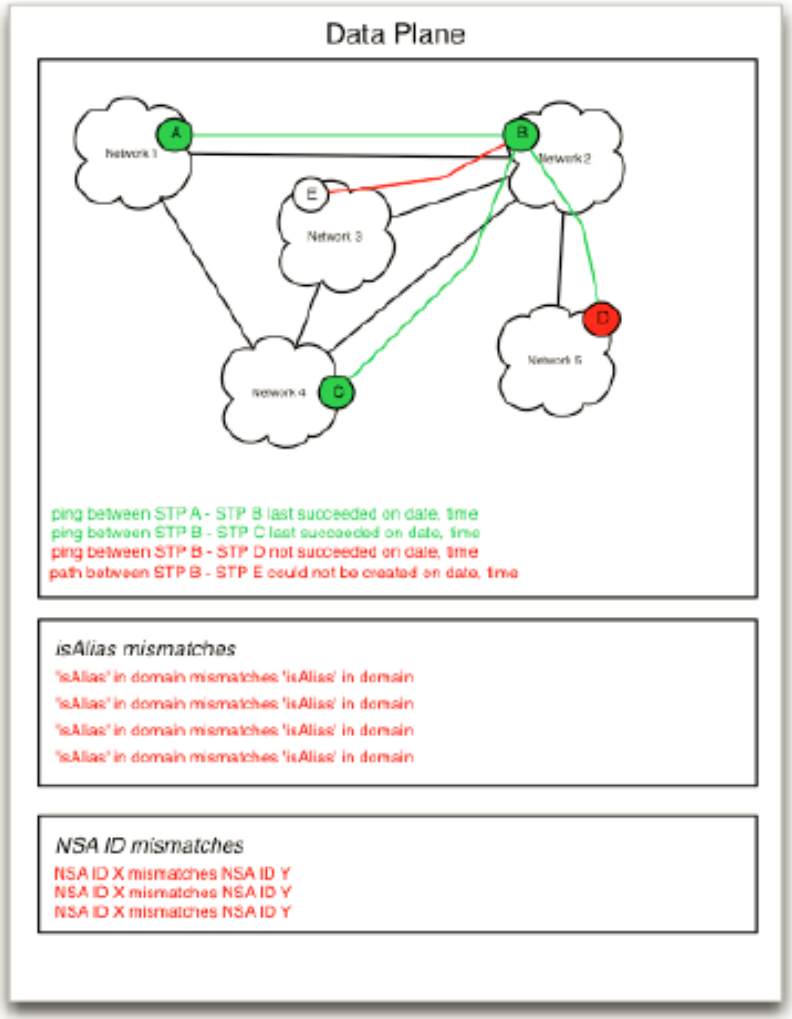

Source: Gerben van Malenstein, SURFnet

STARLIGHT SDX

AutoGOLE Dashboard Current Version

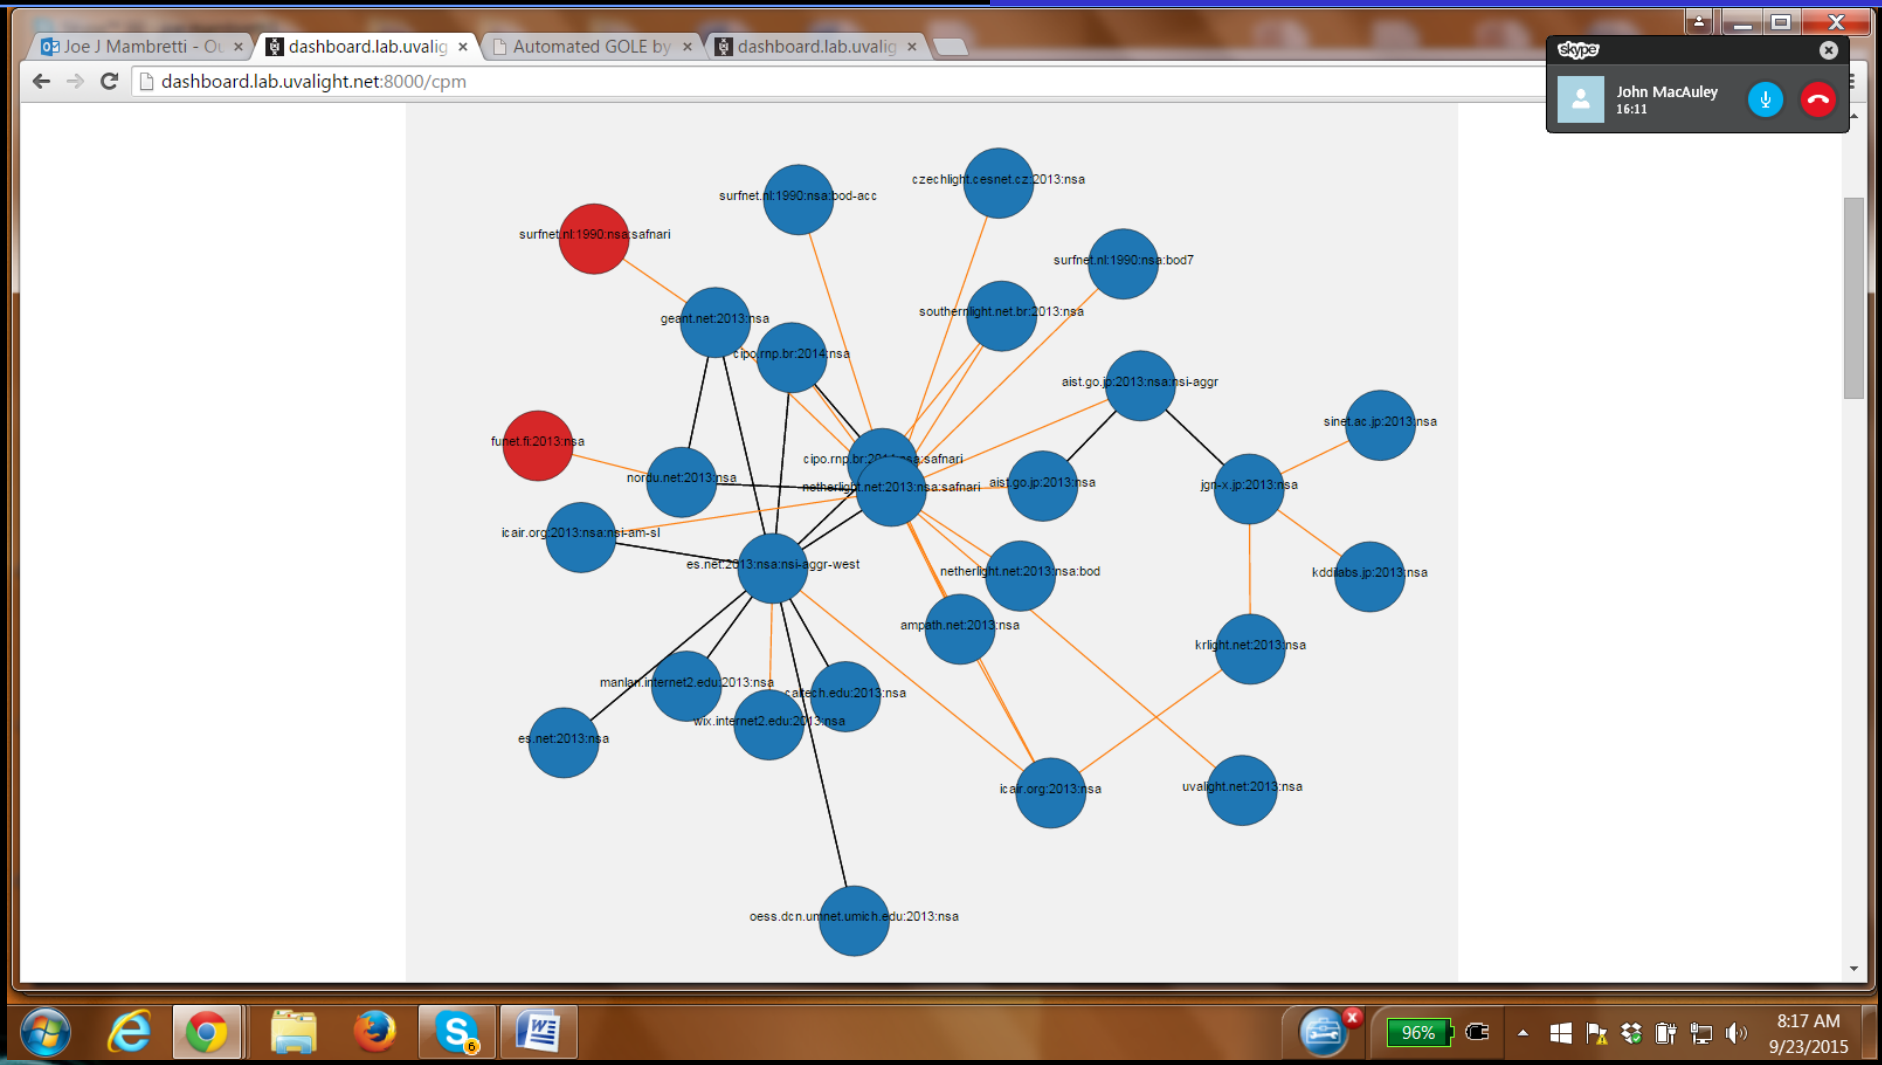

Source: Gerben van Malenstein, SURFnet

STARLIGHT SDX

Benefits of SDN

- **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) A More Comprehensive, Graulated View Into Network Capabilities and Resources**
  - **b) Many More Dynamic Provisioning and Adjustment Options, Including Those That Are Automatic and Implemented In Real Time**
  - **c) Faster Implementations of many New and Enhanced Services**
  - **d) Enabling Applications, Edge Processes and *Even Individuals* To Directly Control Core Resources;**
  - **e) Substantially Improved Options For Creating Customizable Networks**
  - **f) Enhanced Operational Efficiency and Effectiveness.**
  - **And Much, Much More!**

STARLIGHT SDX

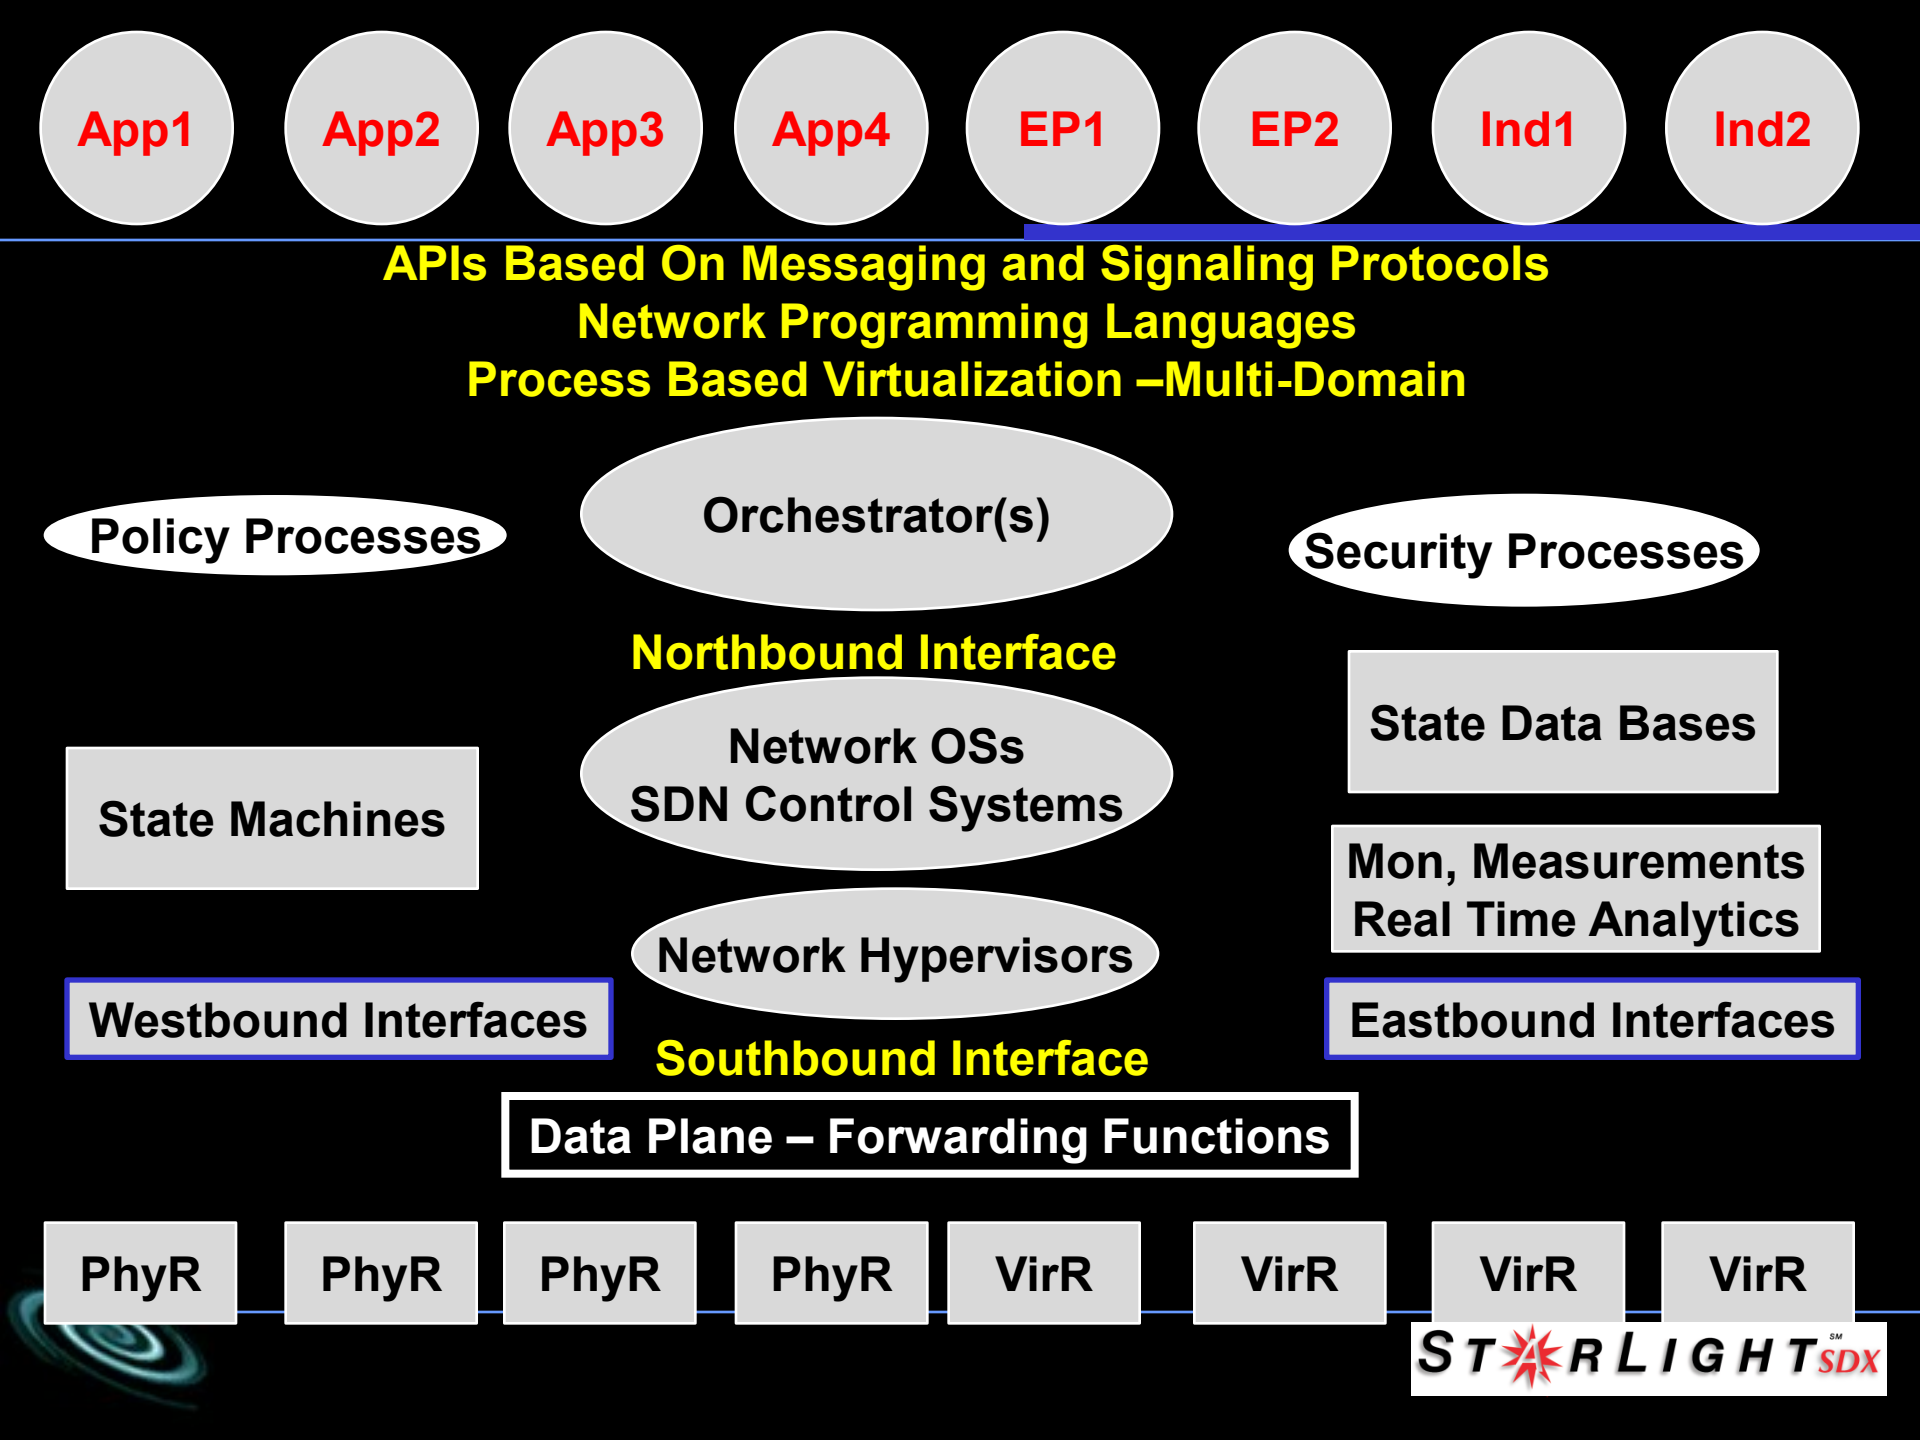
App1
App2
App3
App4
EP1
EP2
Ind1
Ind2
APIs Based On Messaging and Signaling Protocols
Network Programming Languages
Process Based Virtualization –Multi-Domain
Policy Processes
Orchestrator(s)
Security Processes
Northbound Interface
State Data Bases
State Machines
Network OSs
SDN Control Systems
Mon,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

Federation Should Be Able To Cascade To All Architectural Components

- Hybrid Networking Services (Multi-Service, Multi-Layer, Multi-Domain)
- Network Programming Languages (e.g., P4, Frenetic)
- Abstraction Definitions
- APIs
- AP/Service Signaling and Policy Bundling
- Policy Bundle Distribution
- Primitives
- BGP Extensions and Substitutes
- NDL Schema
- Orchestration Processes

STARLIGHT SDX

Other Architectural Components 2

- Northbound Interfaces
- Network OSs
- Network Hypervisors
- State Information Data Bases
- Data Modeling Languages (e.g., YANG)
- Controller Federation Processes
- Hybrid Services/Services Federation/Services Chaining
- Southbound Interfaces
- Eastbound Interfaces
- Westbound Interfaces

Other Architectural Components 3

- Data Plane Processes
- Network Function Virtualization (NFV)
- Measurements
- Real Time Analytics
- Distributed Virtual NOC Operations

Software Defined Networking Exchanges (SDXs)

- With the Increasing Deployment of SDN In Production Networks, the Need for an SDN Exchange (SDX) Has Been Recognized.
- Many Motivations Exist for SDXs
  - Bridging SDNs (Which Are Single Domain & Centralized Controller Oriented)
  - Granulated Engineering Over Flows
  - High Degrees Of Exchange Customization
- Required: Capabilities for Multi-Domain Distributed SDN Resource Discovery, Signaling, Provisioning, Federation, Operational Functions, Fault Detection and Recovery
- These Are Fairly Challenging Issues

STARLIGHT SDX

Selected SDX Architectural Attributes

- Control and Network Resource APIs
- Multi Domain Integrated Path Controllers (With Federation)
- Controller Signaling, Including Edge Signaling
- SDN/OF Multi Layer Traffic Exchange Services
- Multi Domain Resource Advertisement/Discovery
- Topology Exchange Services
- Multiple Highly Customized Services At All Layers
- Granulated Resource Access (Policy Based), Including Through Edge Processes, Including To individual Streams
- Foundation Resource Programmability
- Various Types of Gateways To Other Network Environments
- Integration of OF and Non-OF Paths, Including 3rd Party Integration
- Programmability for Large Scale Large Capacity Streams

SDX As Recursive Virtual Switch

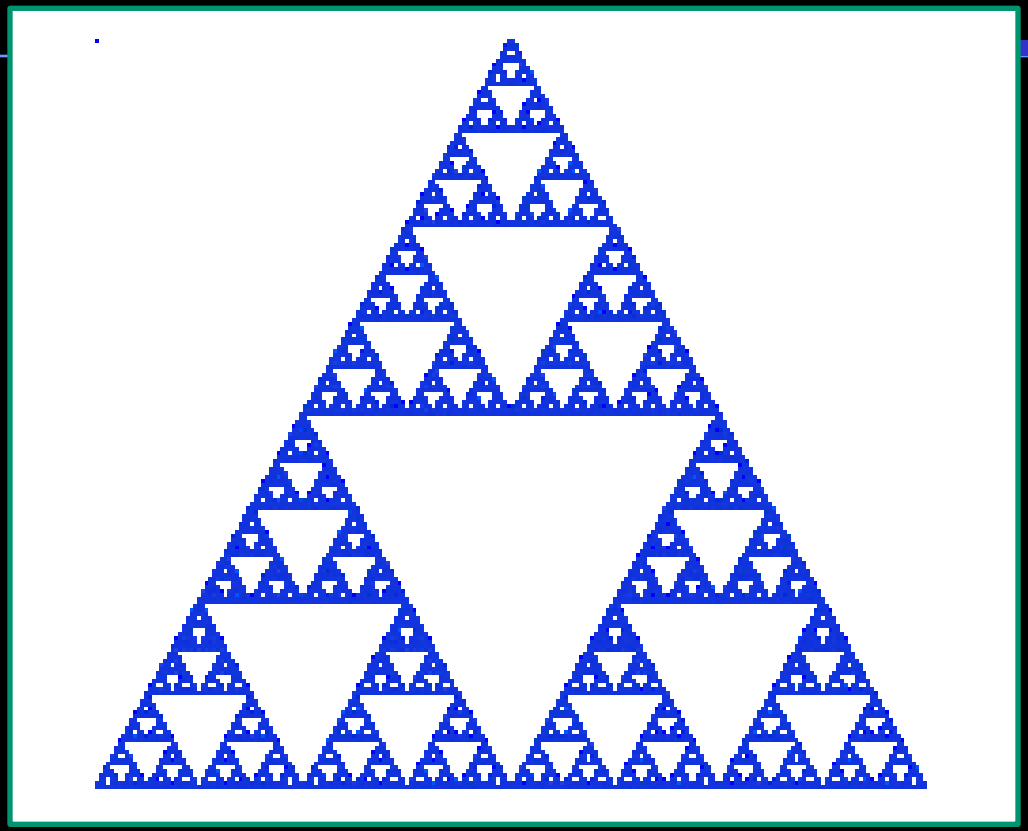

Sierpinski Triangle

Unlimited Number of
***Customized* Virtual Switches**
Within Macro Virtual Switch

STARLIGHT SDX

GLIF Based On SDXs Supporting Slice Exchanges

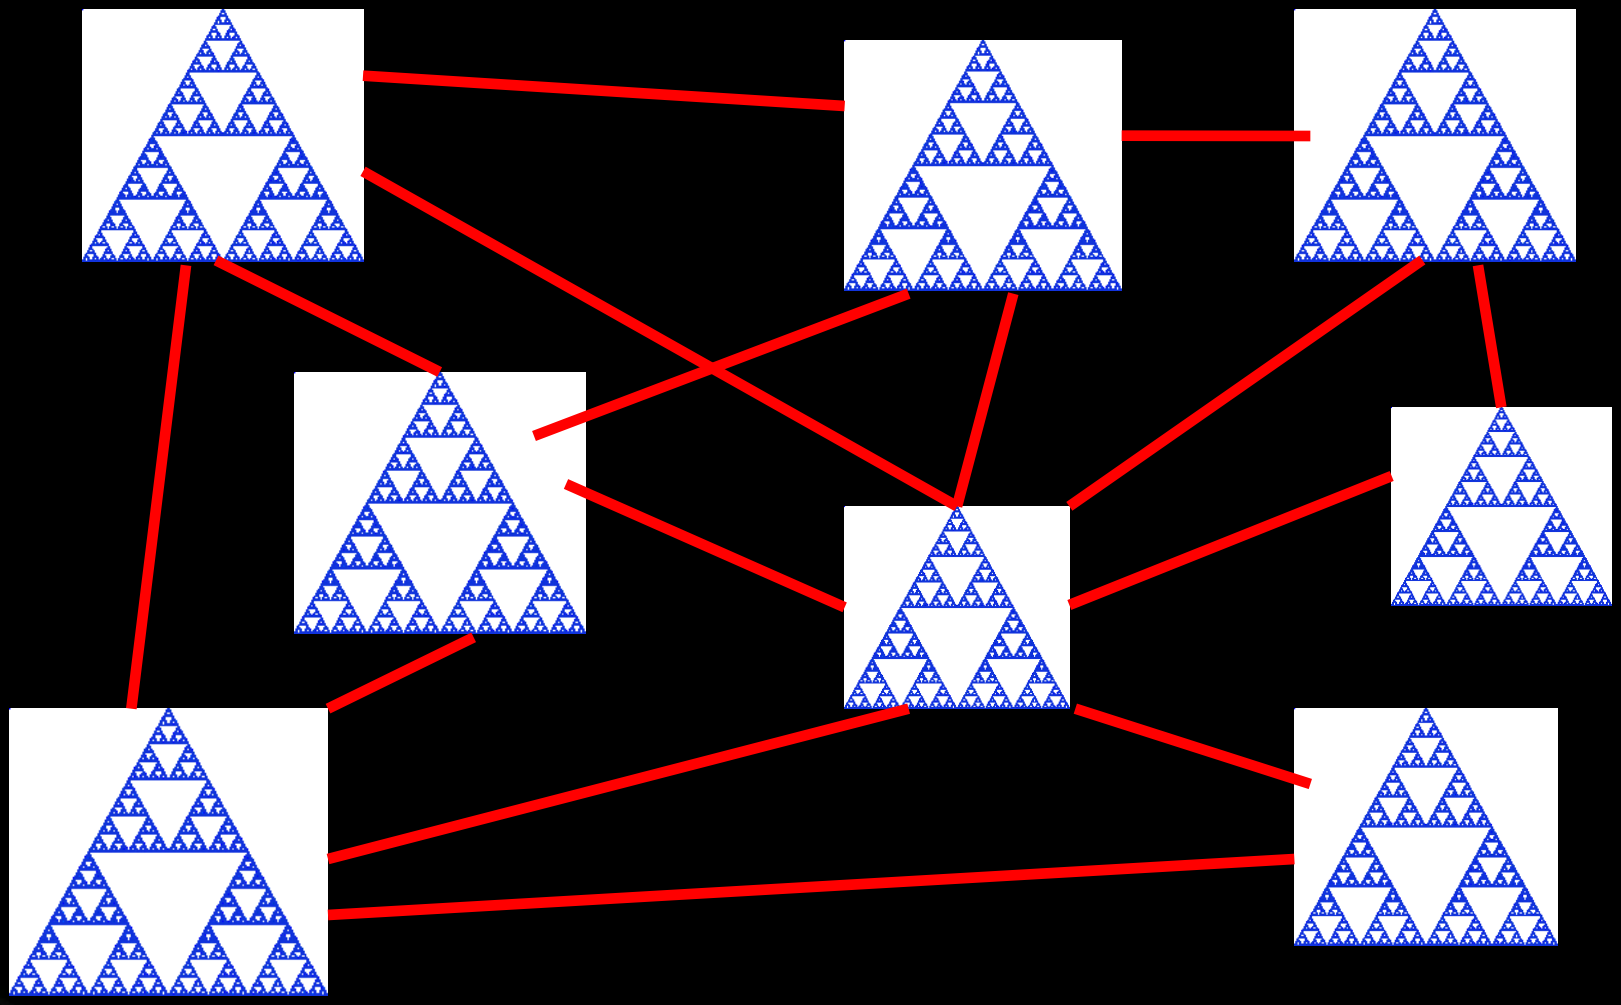

STARLIGHT SDX

Global Network Science: iGENI Consortium Uses The Global Lambda Integrated Facility As the Basis For a Distributed SDN/OpenFlow Testbed

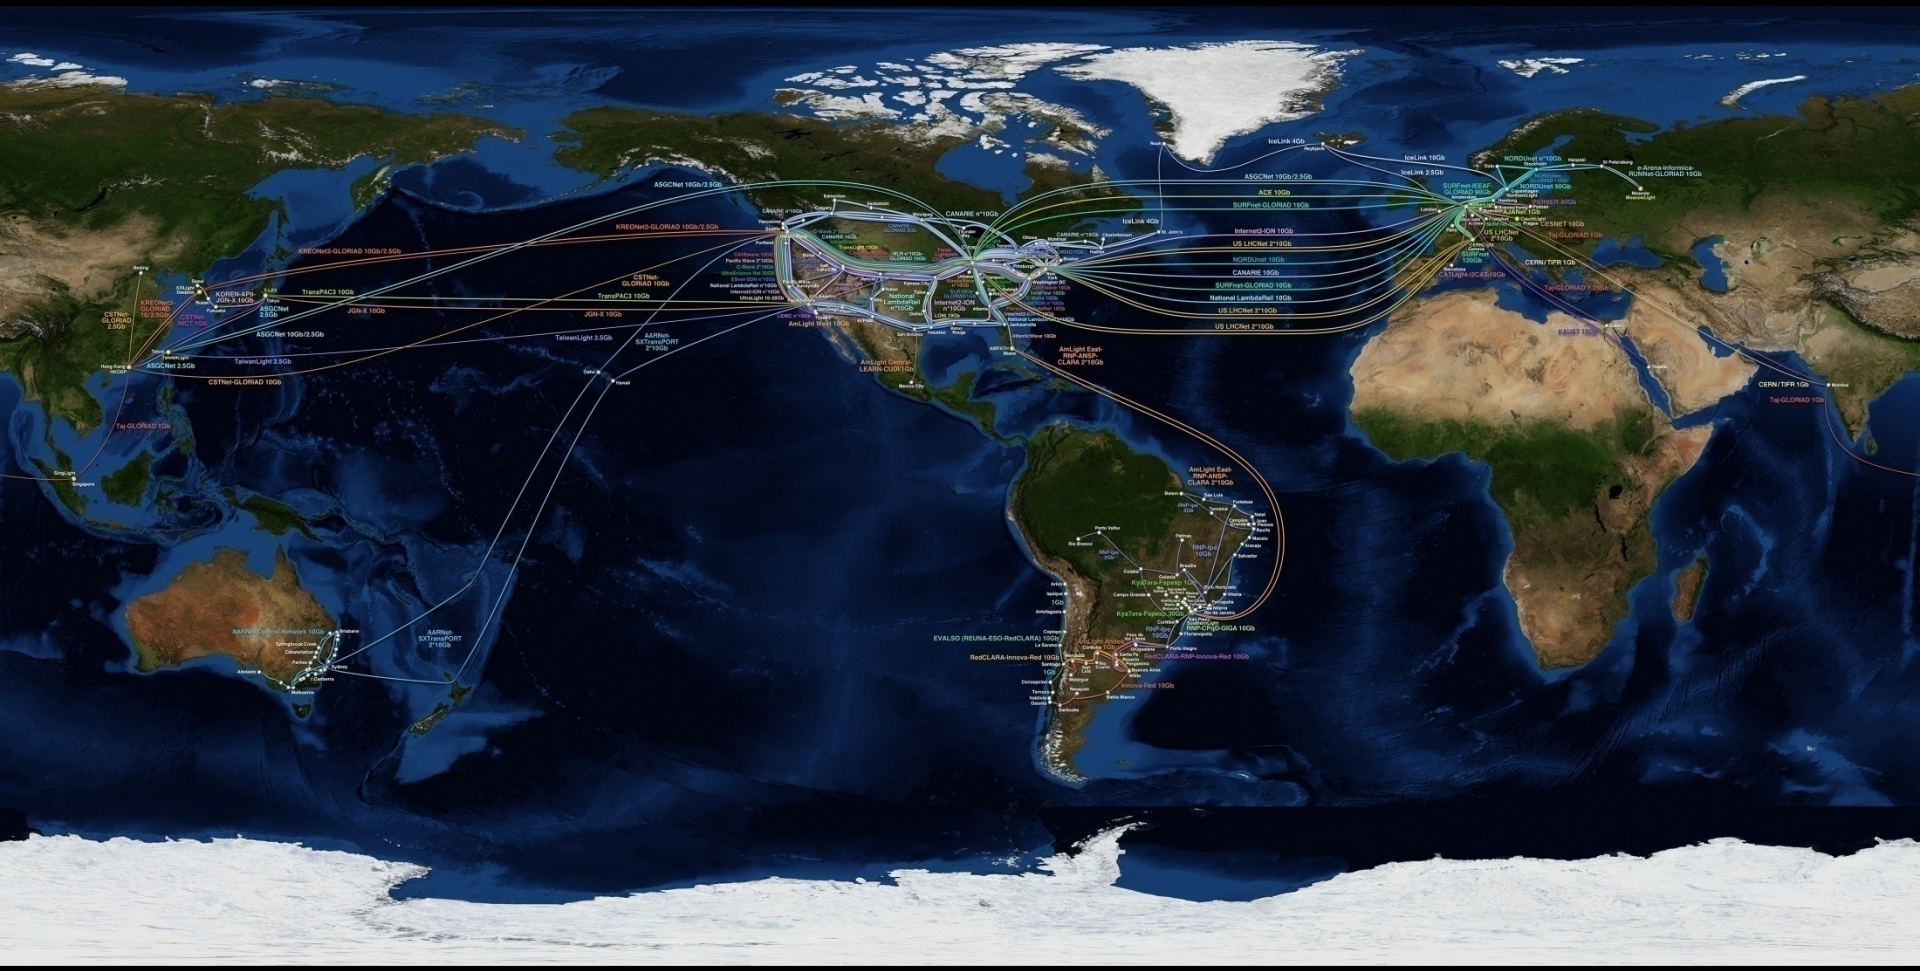

iCAIR

STARLIGHT SDX


KUAS
NetFPGA
NOX
OpenFlow
NCKU
NetFPGA
NOX
OpenFlow
ETRI
NOX
OpenFlow
NetFPGA
KiSTi
University of Victoria
TOMATO
OpenFlow
InstaGENI
CRC
NOX
OpenFlow
NetFPGA
cenic
InstaGENI
OpenFlow
JGN-X
NSI
東京大学
THE UNIVERSITY OF TOKYO
VNode
PACIFIC NORTHWEST GIGAPOP
aarnet
NICTA
ExoGENI
OpenFlow
NCHC
NetFPGA
NOX
OpenFlow
ofNSI
KRLight
NSI
canarie
APPLIED RESEARCH CENTER FOR COMPUTER NETWORKS
OpenFlow
GLORIAD
NSI
Northern Wave
SoX
Southern Crossroads
GENI AM
OpenFlow
InstaGENI
SingAREN
OpenFlow
NetFPGA
SingLight
MREN
STARLIGHT
The Optical STAR TAP
GENI AM
ExoGENI
OpenFlow
NSI
ofNSI
INTERNET2
GENI AM
OpenFlow
BEN
Breakable Experimental Network
ExoGENI
renci
ExoGENI
OpenFlow
CNC
OpenFlow
NORTHWESTERN UNIVERSITY
NetFPGA
NOX
OpenFlow
TOMATO
InstaGENI
Duke
UNIVERSITY
ExoGENI
OpenFlow
AMPATH
NSI
OpenFlow
MAGPI
uen
UTAH EDUCATION NETWORK
GENI AM
OpenFlow
THE UNIVERSITY OF UTAH
VNode
OpenFlow
InstaGENI
CPqD
OpenFlow
RNP
NSI
NLight
NSI
NOX
OpenFlow
NetFPGA
UNIVERSITY OF AMSTERDAM
NSI
ExoGENI
Princeton University
OpenFlow
InstaGENI
USP
InstaGENI
TOMATO
TECHNISCHE UNIVERSITÄT KAISERSLAUTERN
DFN
Deutsches Forschungsnetz
GEANT2
NSI
ja.net
University of BRISTOL
OpenFlow
University of Essex
OpenFlow
grnet
NSI
Belnet
dedicated connectivity
iMinds
CONNECT.INNOVATE.CREATE
InstaGENI

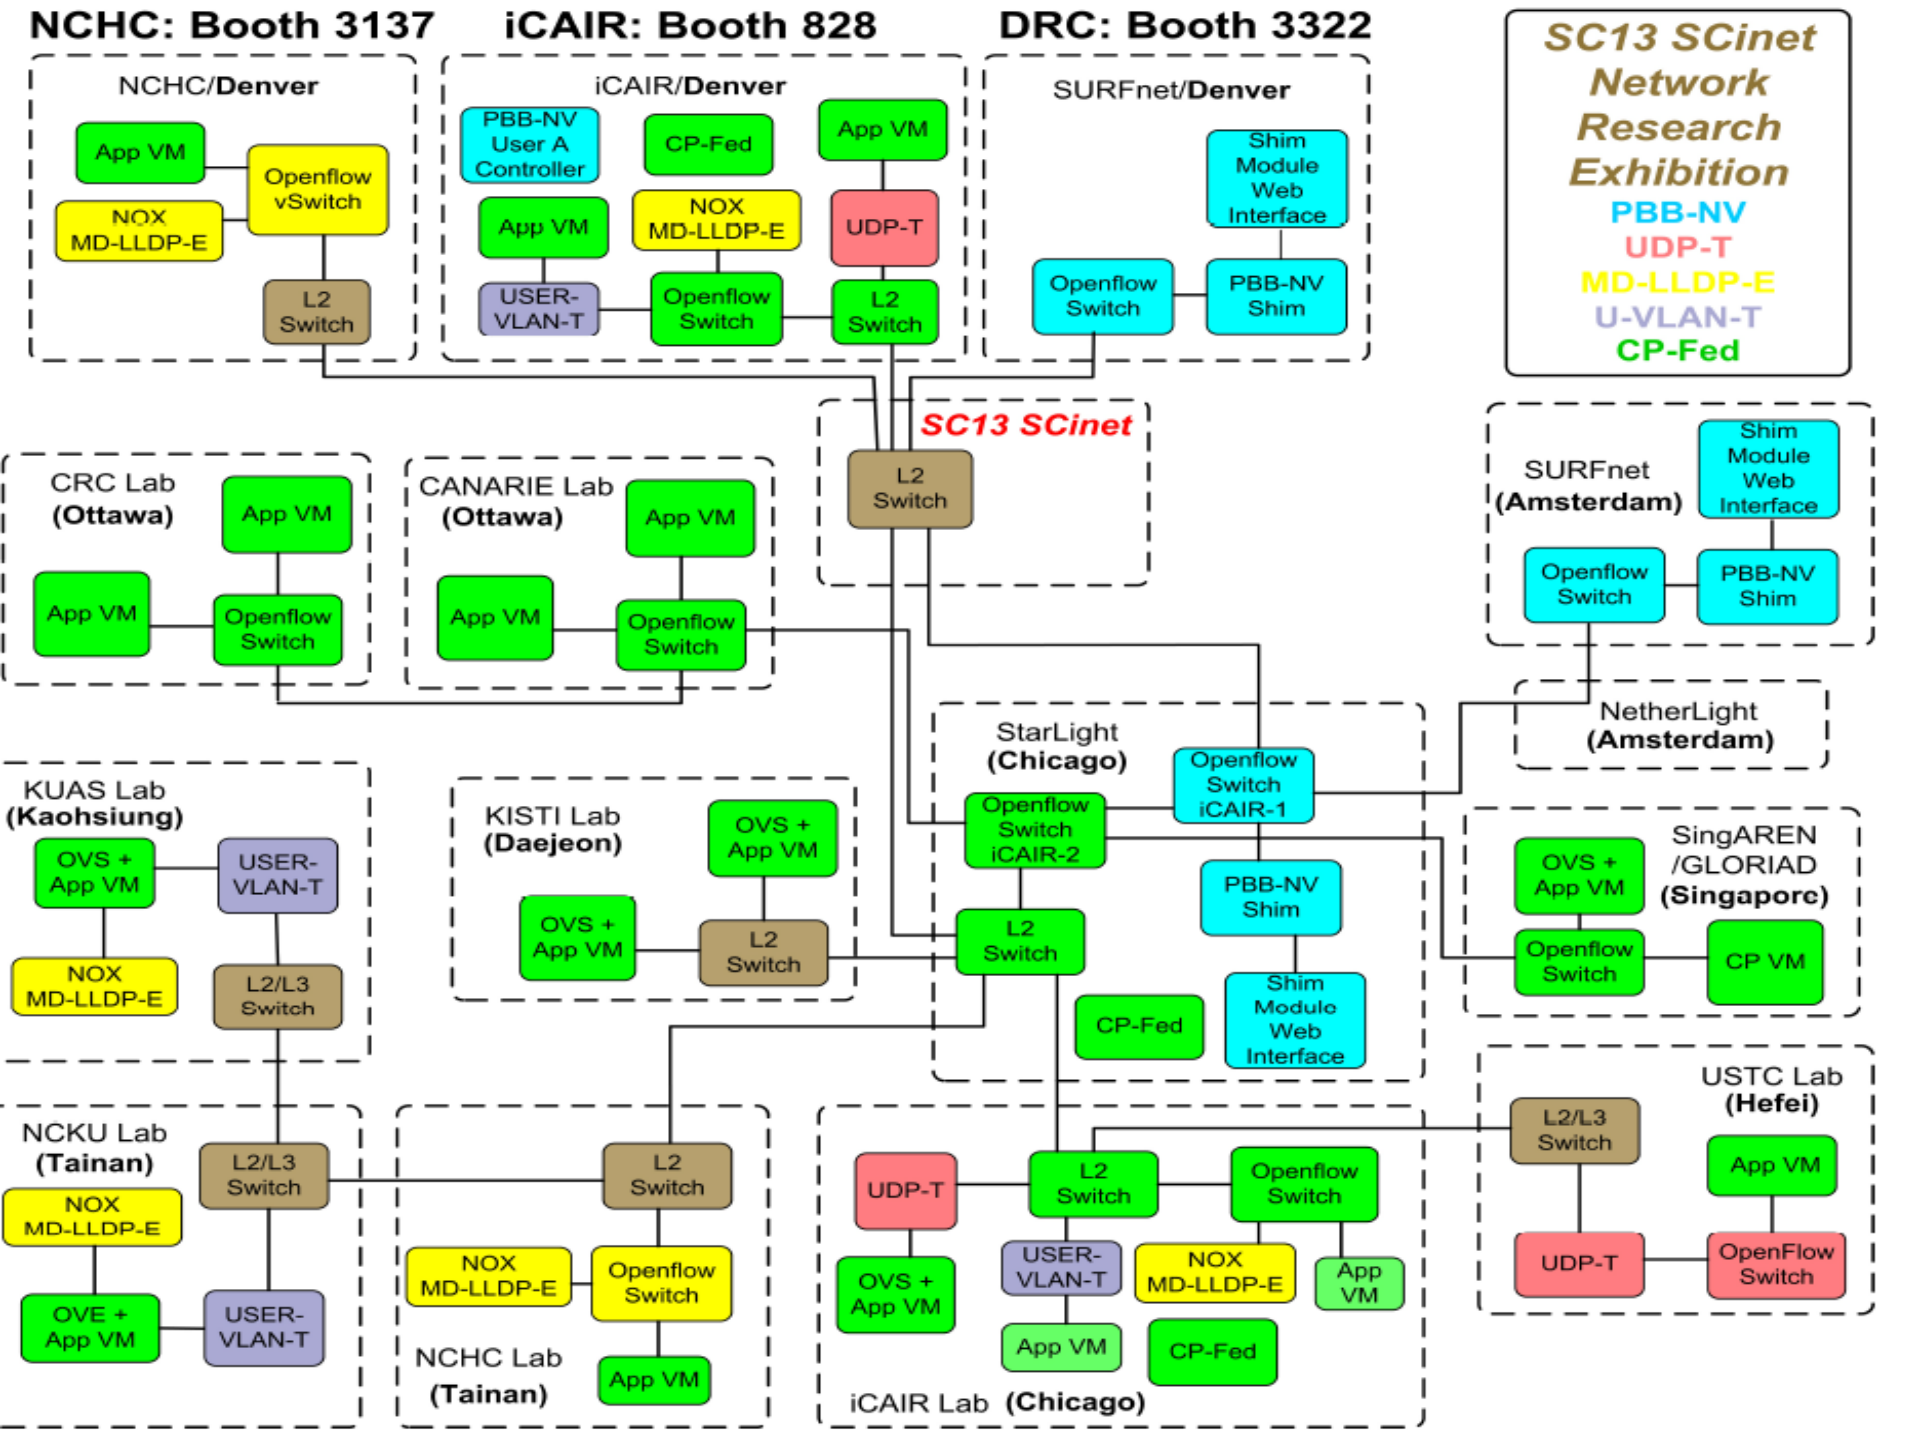
NCHC: Booth 3137
iCAIR: Booth 828
DRC: Booth 3322
SC13 SCinet Network Research Exhibition
PBB-NV
UDP-T
MD-LLDP-E
U-VLAN-T
CP-Fed
NCHC/Denver
App VM
Openflow vSwitch
NOX MD-LLDP-E
L2 Switch
iCAIR/Denver
PBB-NV User A Controller
CP-Fed
App VM
NOX MD-LLDP-E
UDP-T
App VM
USER-VLAN-T
Openflow Switch
L2 Switch
SURFnet/Denver
Shim Module Web Interface
Openflow Switch
PBB-NV Shim
SC13 SCinet
L2 Switch
CRC Lab (Ottawa)
App VM
App VM
Openflow Switch
CANARIE Lab (Ottawa)
App VM
App VM
Openflow Switch
SURFnet (Amsterdam)
Shim Module Web Interface
Openflow Switch
PBB-NV Shim
NetherLight (Amsterdam)
StarLight (Chicago)
Openflow Switch iCAIR-2
Openflow Switch iCAIR-1
PBB-NV Shim
L2 Switch
CP-Fed
Shim Module Web Interface
KUAS Lab (Kaohsiung)
OVS + App VM
USER-VLAN-T
NOX MD-LLDP-E
L2/L3 Switch
KISTI Lab (Daejeon)
OVS + App VM
OVS + App VM
L2 Switch
SingAREN /GLORIAD (Singapore)
OVS + App VM
Openflow Switch
CP VM
USTC Lab (Hefei)
L2/L3 Switch
App VM
UDP-T
OpenFlow Switch
NCKU Lab (Tainan)
L2/3 Switch
NOX MD-LLDP-E
OVE + App VM
USER-VLAN-T
L2 Switch
NOX MD-LLDP-E
Openflow Switch
NCHC Lab (Tainan)
App VM
UDP-T
L2 Switch
Openflow Switch
OVS + App VM
USER-VLAN-T
NOX MD-LLDP-E
App VM
App VM
CP-Fed
iCAIR Lab (Chicago)

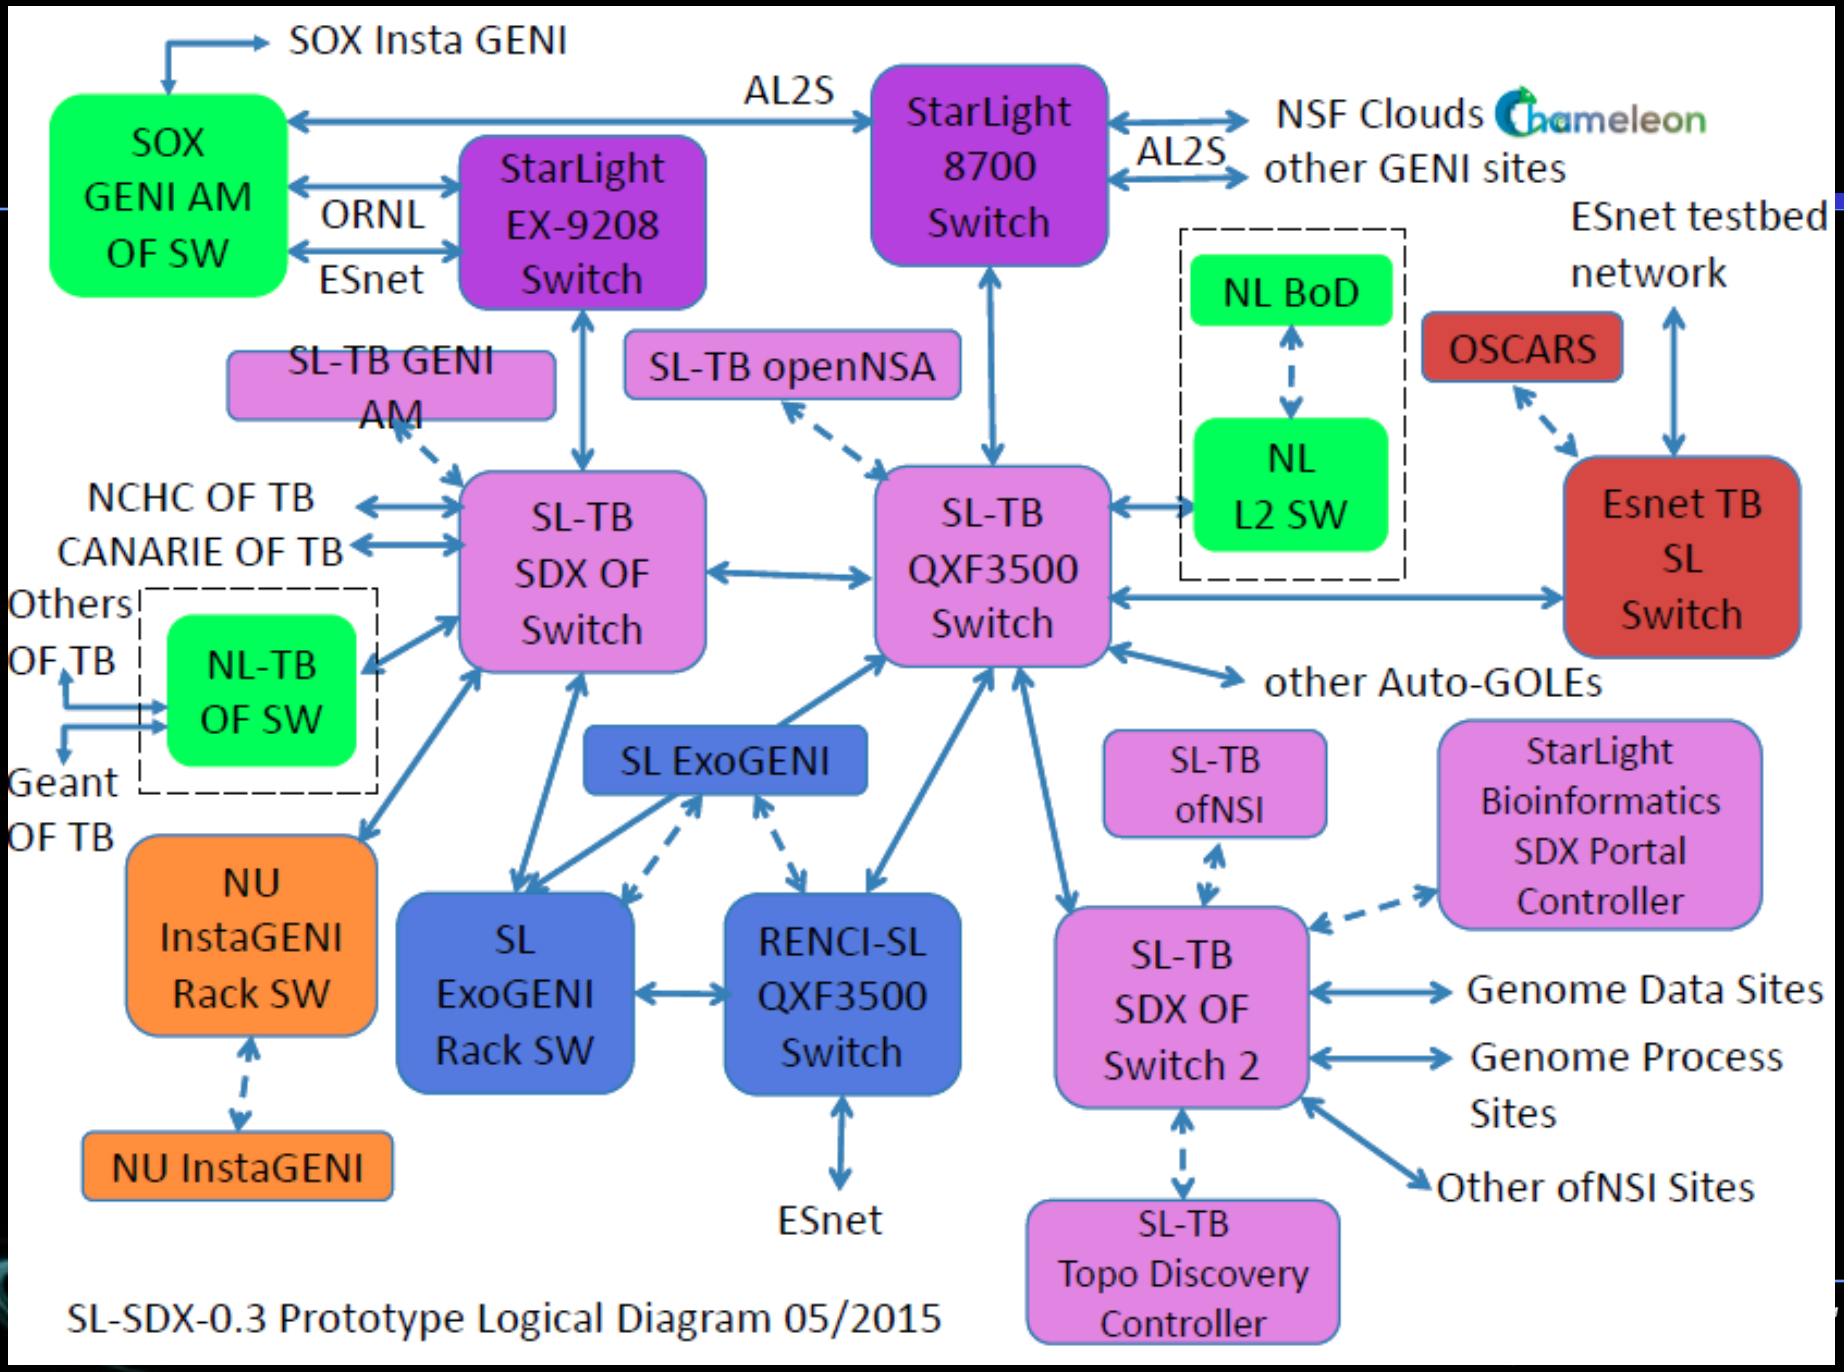
SOX Insta GENI
SOX GENI AM OF SW
AL2S
StarLight 8700 Switch
NSF Clouds
Chameleon
AL2S
other GENI sites
ORNL
ESnet
StarLight EX-9208 Switch
ESnet testbed network
NL BoD
OSCARS
SL-TB GENI AM
SL-TB openNSA
NL L2 SW
NCHC OF TB
CANARIE OF TB
SL-TB SDX OF Switch
SL-TB QXF3500 Switch
Esnet TB SL Switch
Others OF TB
NL-TB OF SW
other Auto-GOLEs
Geant OF TB
SL ExoGENI
SL-TB ofNSI
StarLight Bioinformatics SDX Portal Controller
NU InstaGENI Rack SW
SL ExoGENI Rack SW
RENCI-SL QXF3500 Switch
SL-TB SDX OF Switch 2
Genome Data Sites
Genome Process Sites
Other ofNSI Sites
NU InstaGENI
ESnet
SL-TB Topo Discovery Controller
SL-SDX-0.3 Prototype Logical Diagram 05/2015

Science Use Case: Nowcasting With SDXs

Source: Mike Zink, UMass Amherst

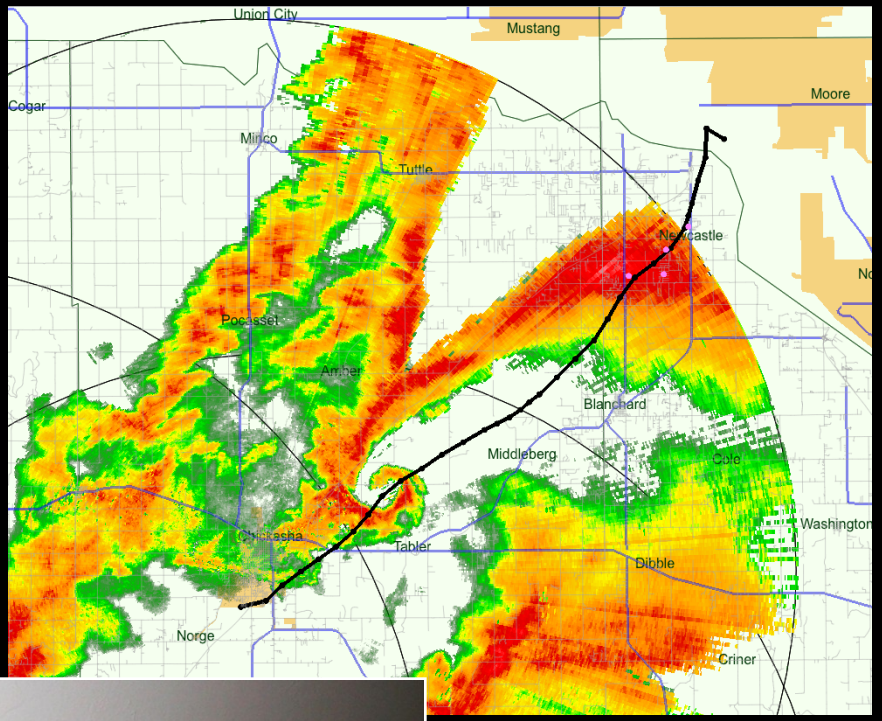

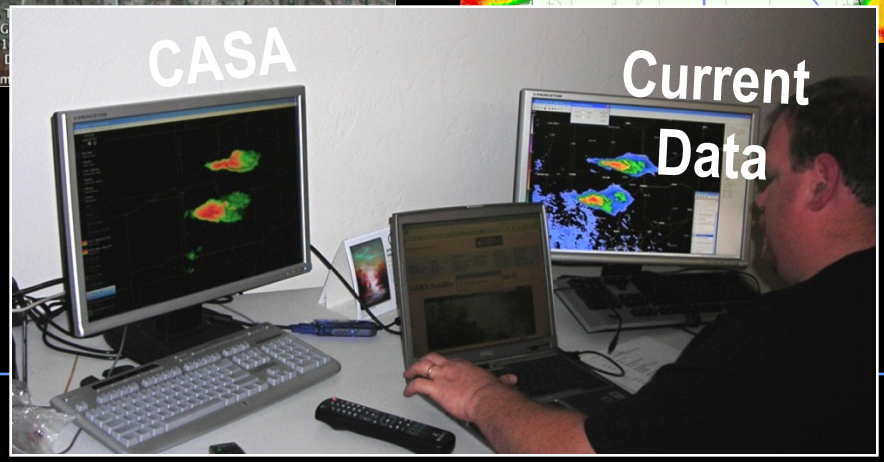

CASA Data, EM Decision-Making Protects First Responders and Public

STORMLIGHT SDX

Comparison With Existing System

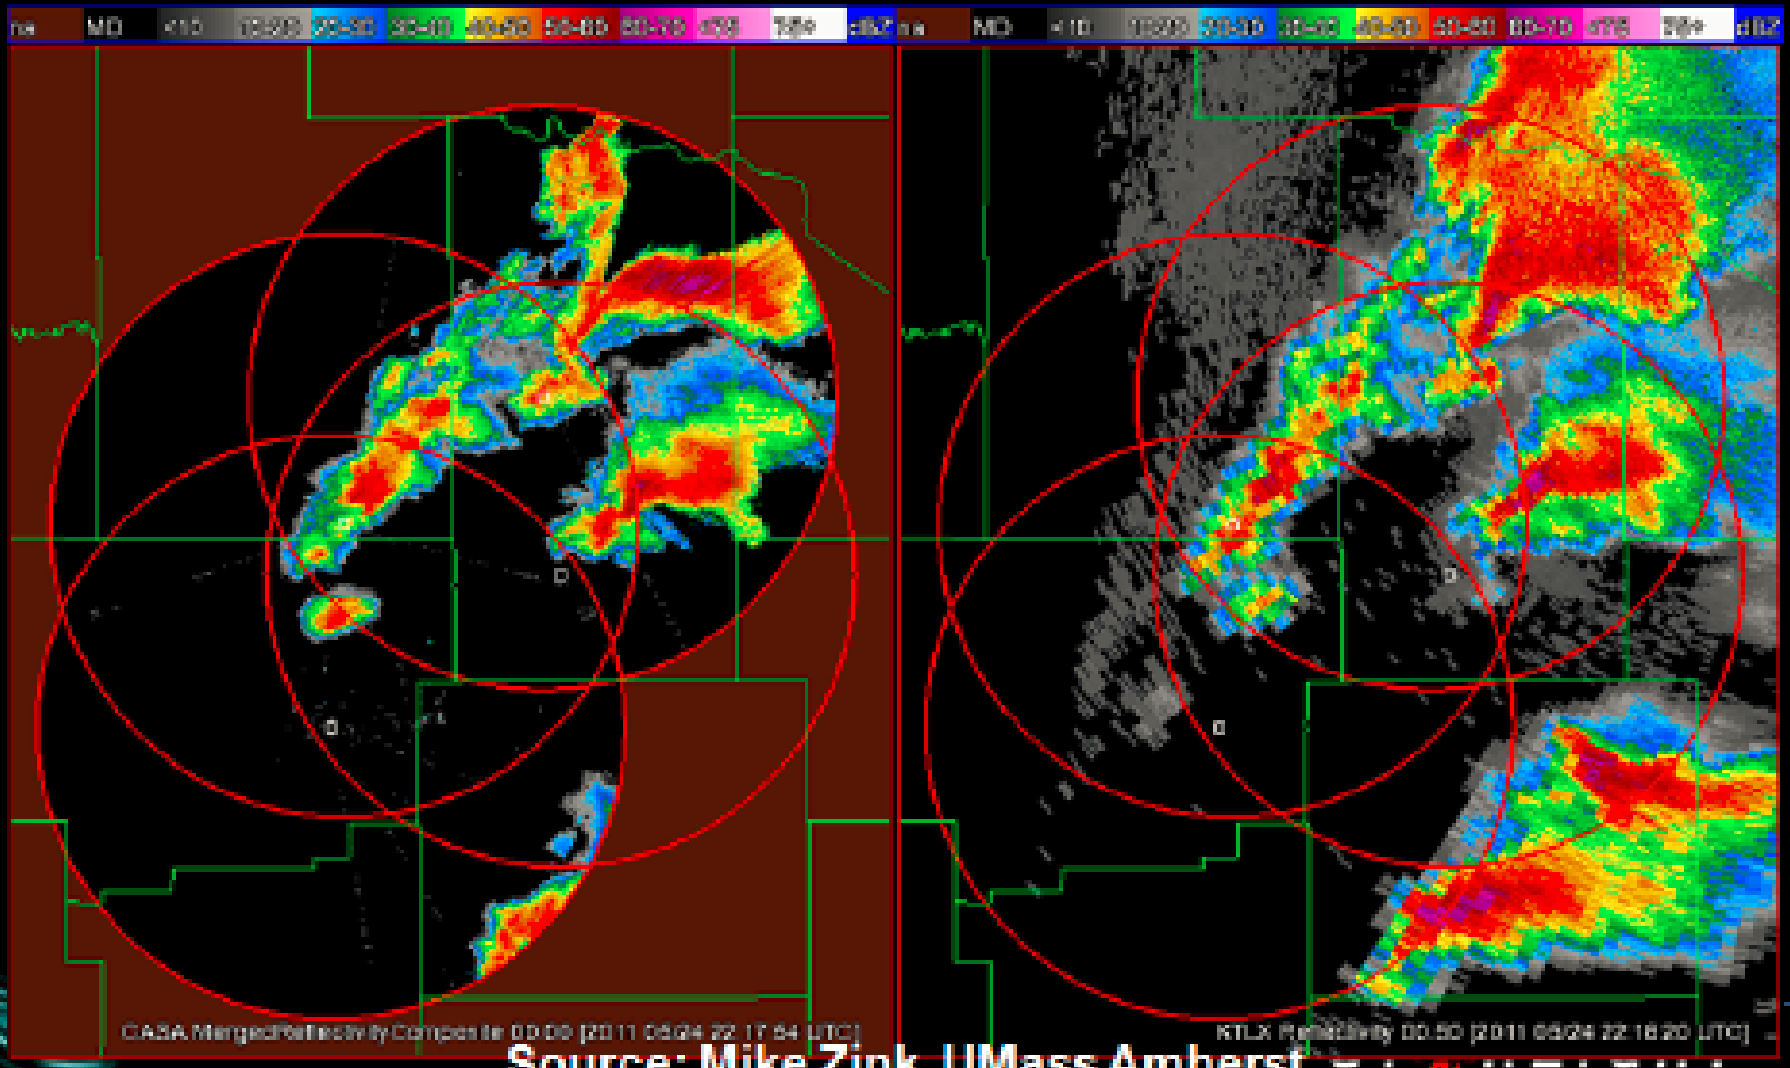

Source: Mike Zink, UMass Amherst

GENI SDX Demo Scenario

Simulated Radar (4)

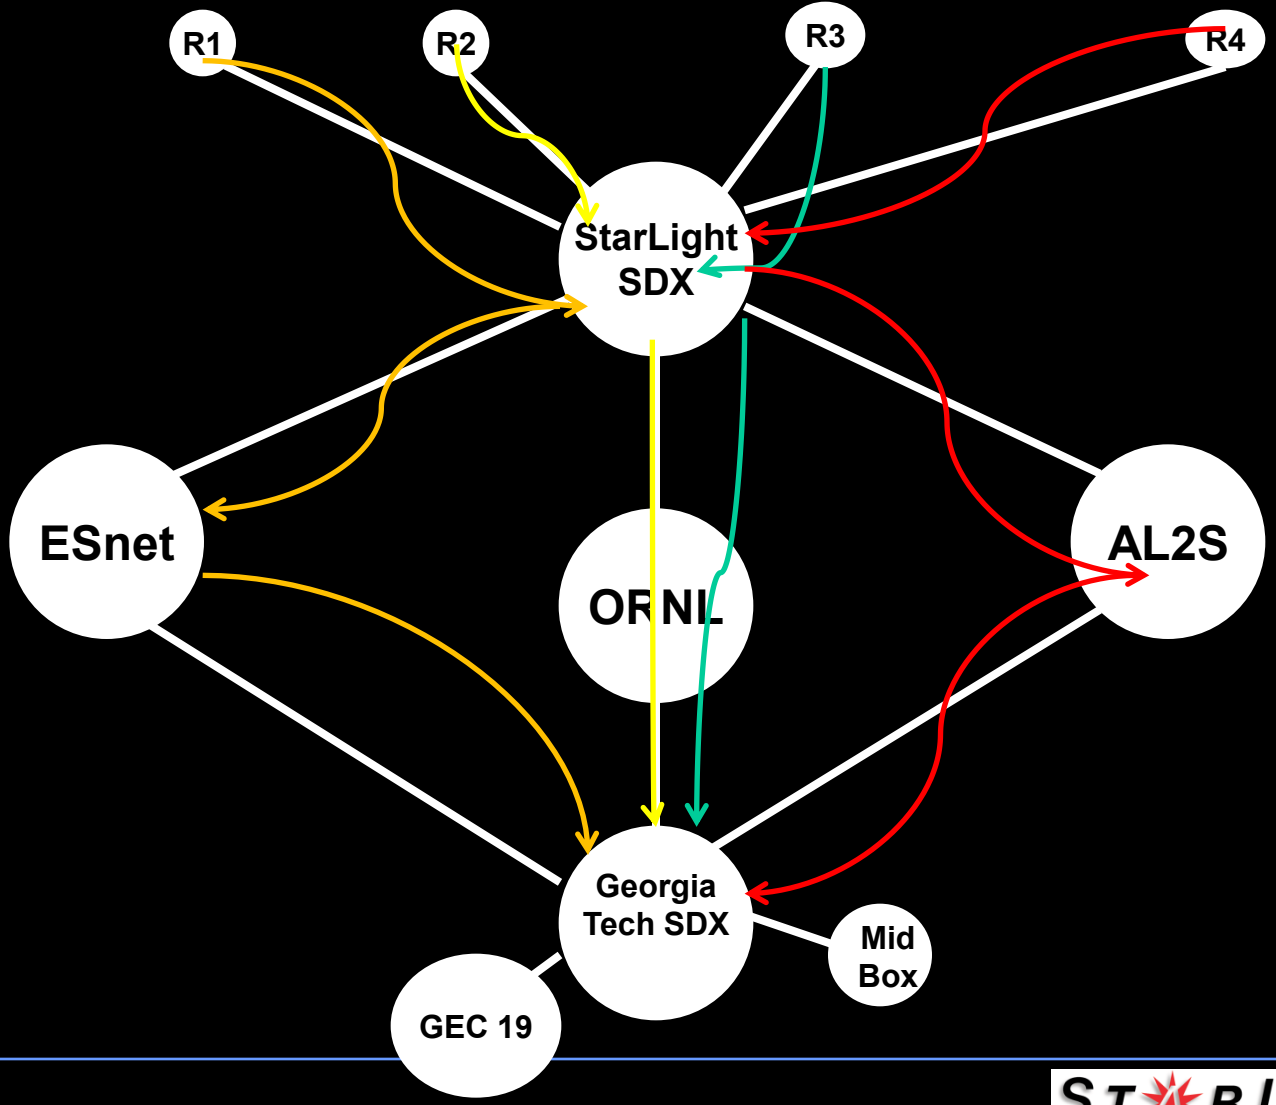

SDX StarLight⇔NetherLight

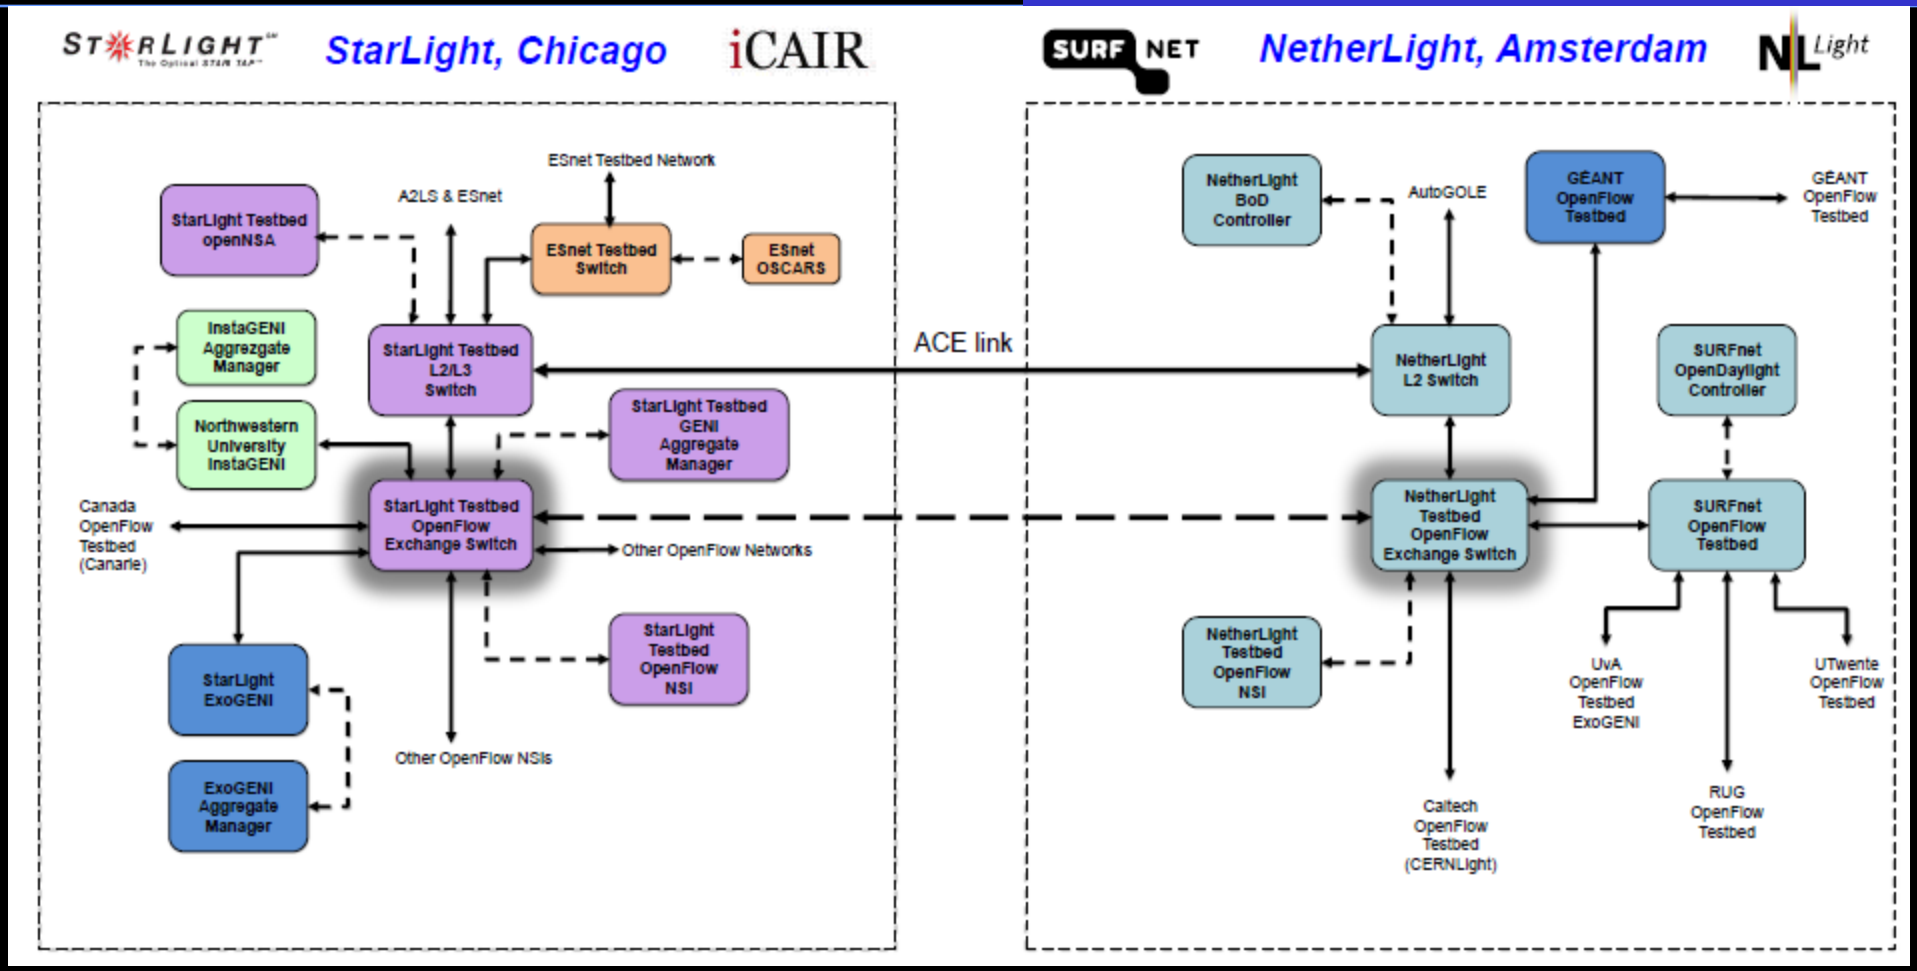

Ronald van der Pol, Joe Mambretti, Jim Chen, John Shillington

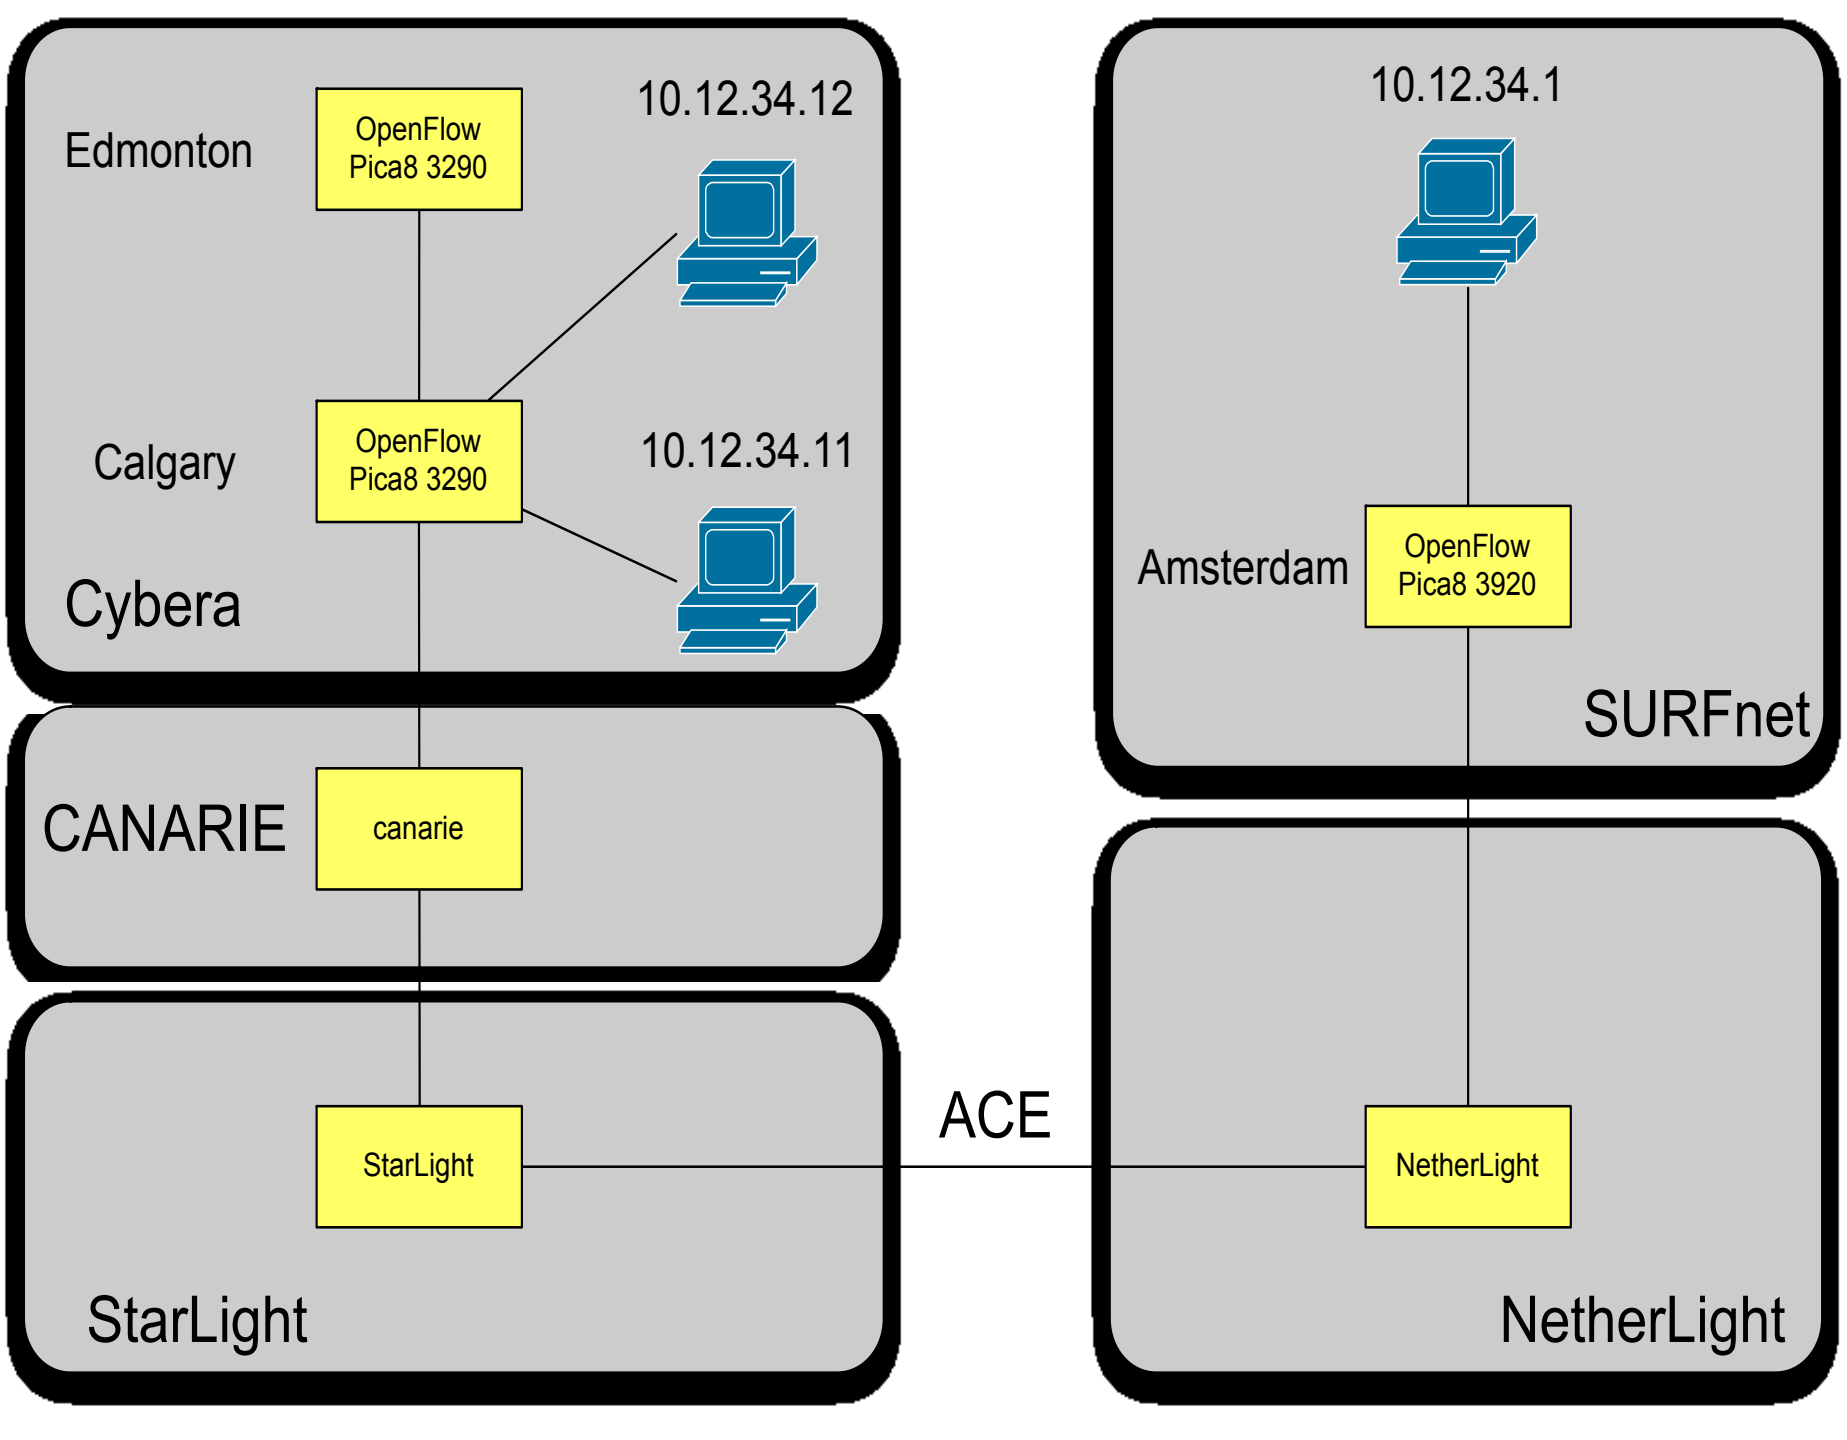
Edmonton
OpenFlow
Pica8 3290
10.12.34.12
Calgary
OpenFlow
Pica8 3290
10.12.34.11
Cybera
CANARIE
canarie
StarLight
StarLight
ACE
10.12.34.1
Amsterdam
OpenFlow
Pica8 3920
SURFnet
NetherLight
NetherLight

International Software-Defined Network Exchanges (iSDXs): A Demonstration of Global Capabilities

Joe Mambretti, Jim Chen, Fei Yeh
International Center for Advanced Internet Research
Northwestern University, USA
Mike Zink, Divyashri Bhat
University of Massachusetts, Amherst, USA
Ronald Van der Pol
Surfnet, Netherlands
Grace Lee, WunYuan Huang, Te-Lung Liu
NARLabs, National Center for High Performance Computing,Taiwan
Thomas Tam, Herve Guy,
CANARIE, Canada
Alex Valiushko, John Shillington,
Cybera, Canada
Buseung Cho, KISTI
Republic of Korea
Michiaki Hayashi, KDDI Labs, Japan
Toshiaki Tarui, Hitachi, Japan
Aki Nakao, University of Tokyo, Japan
Steve Cotter, T. Charles Yun, Jamie Curtis, Andrej Ricnik
REANNZ, New Zealand
Josh Bailey, Google, New Zealand
Artur Binczewski Belter Bartosz Miłosz Przywecki Piotr Rydlichowski
Poznan Supercomputing and Networking Center, Poland
Russ Clark, Georgia Tech, USA

Global LambdaGrid Workshop
Queenstown, New Zealand
September 30-October 1, 2014

iCAIR

StarLight SDX

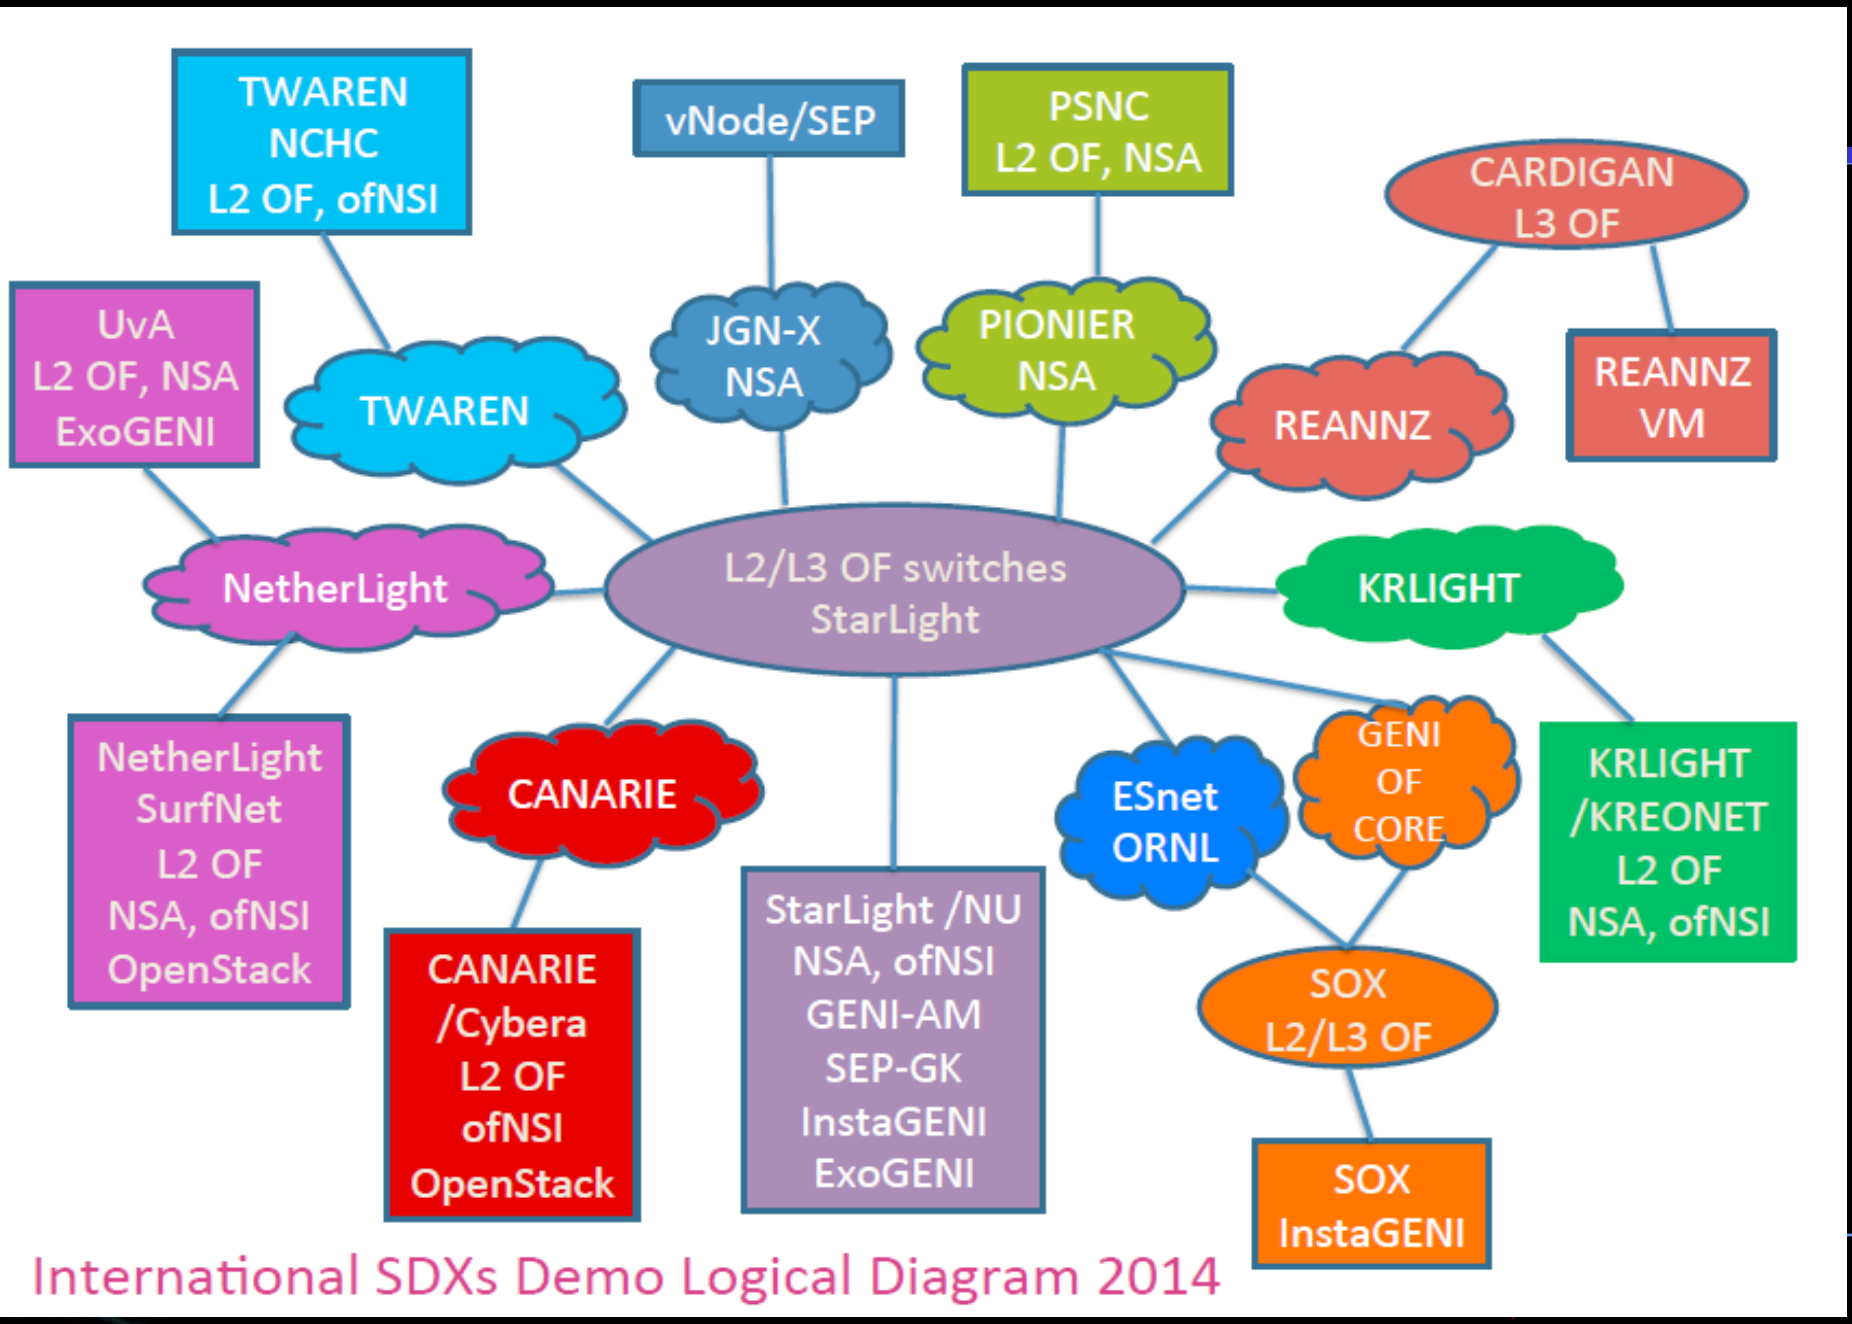

International SDXs Demo Logical Diagram 2014


Nowcast - SDX Demo
iCAIR
g1t

Slice Exchange Showcase at GEC 21

Japan-US Slice Exchange over SDX

KDDI R&D LABS

HITACHI Inspire the Next

東京大学 THE UNIVERSITY OF TOKYO

NORTHWESTERN UNIVERSITY

Slice Exchange Architecture

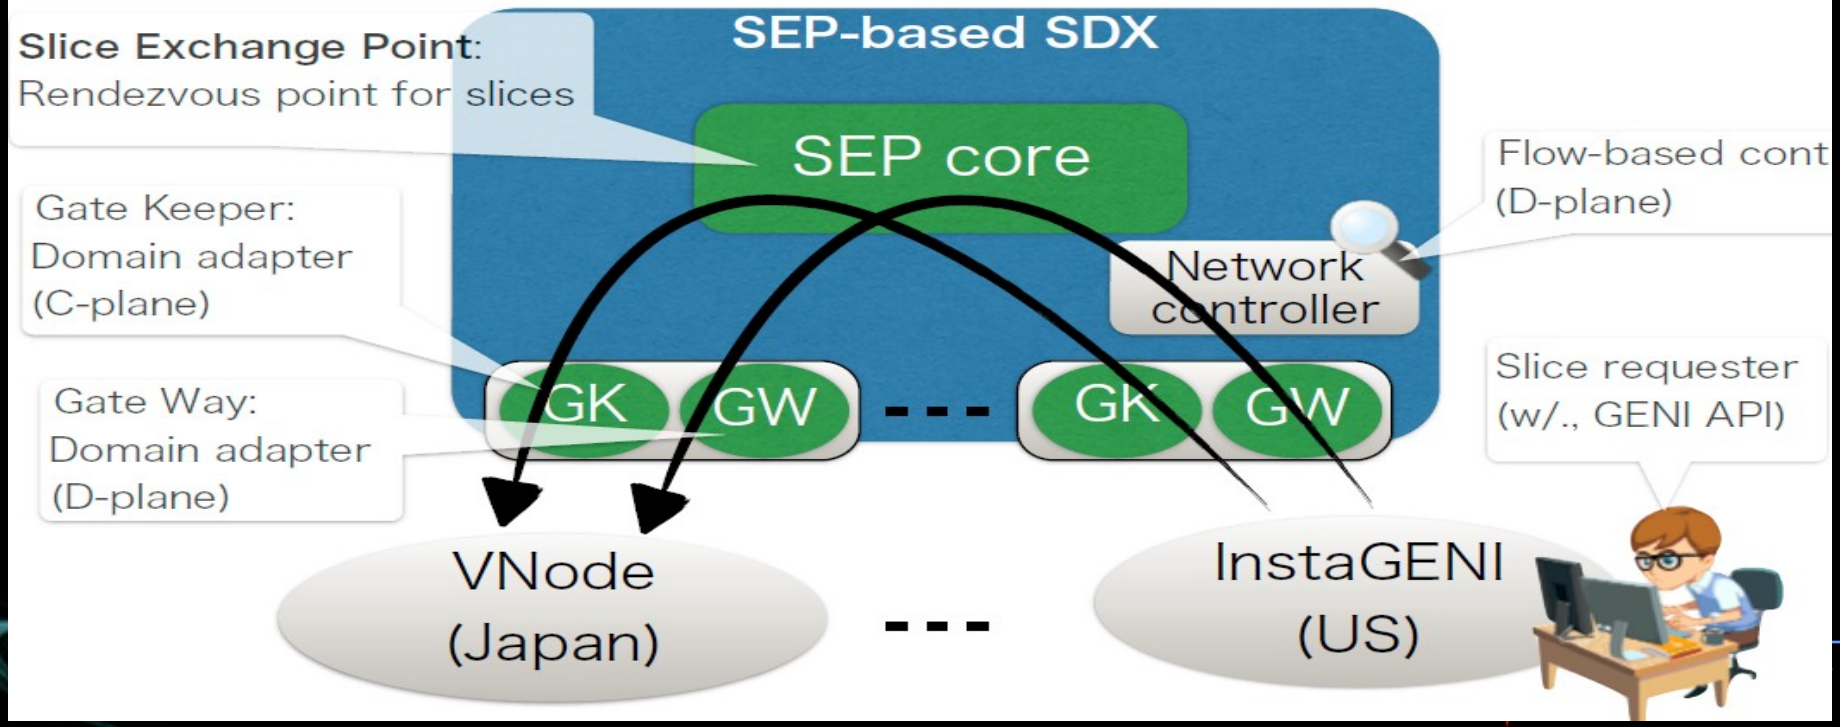

PetaTrans: Petascale Science Data Transfer

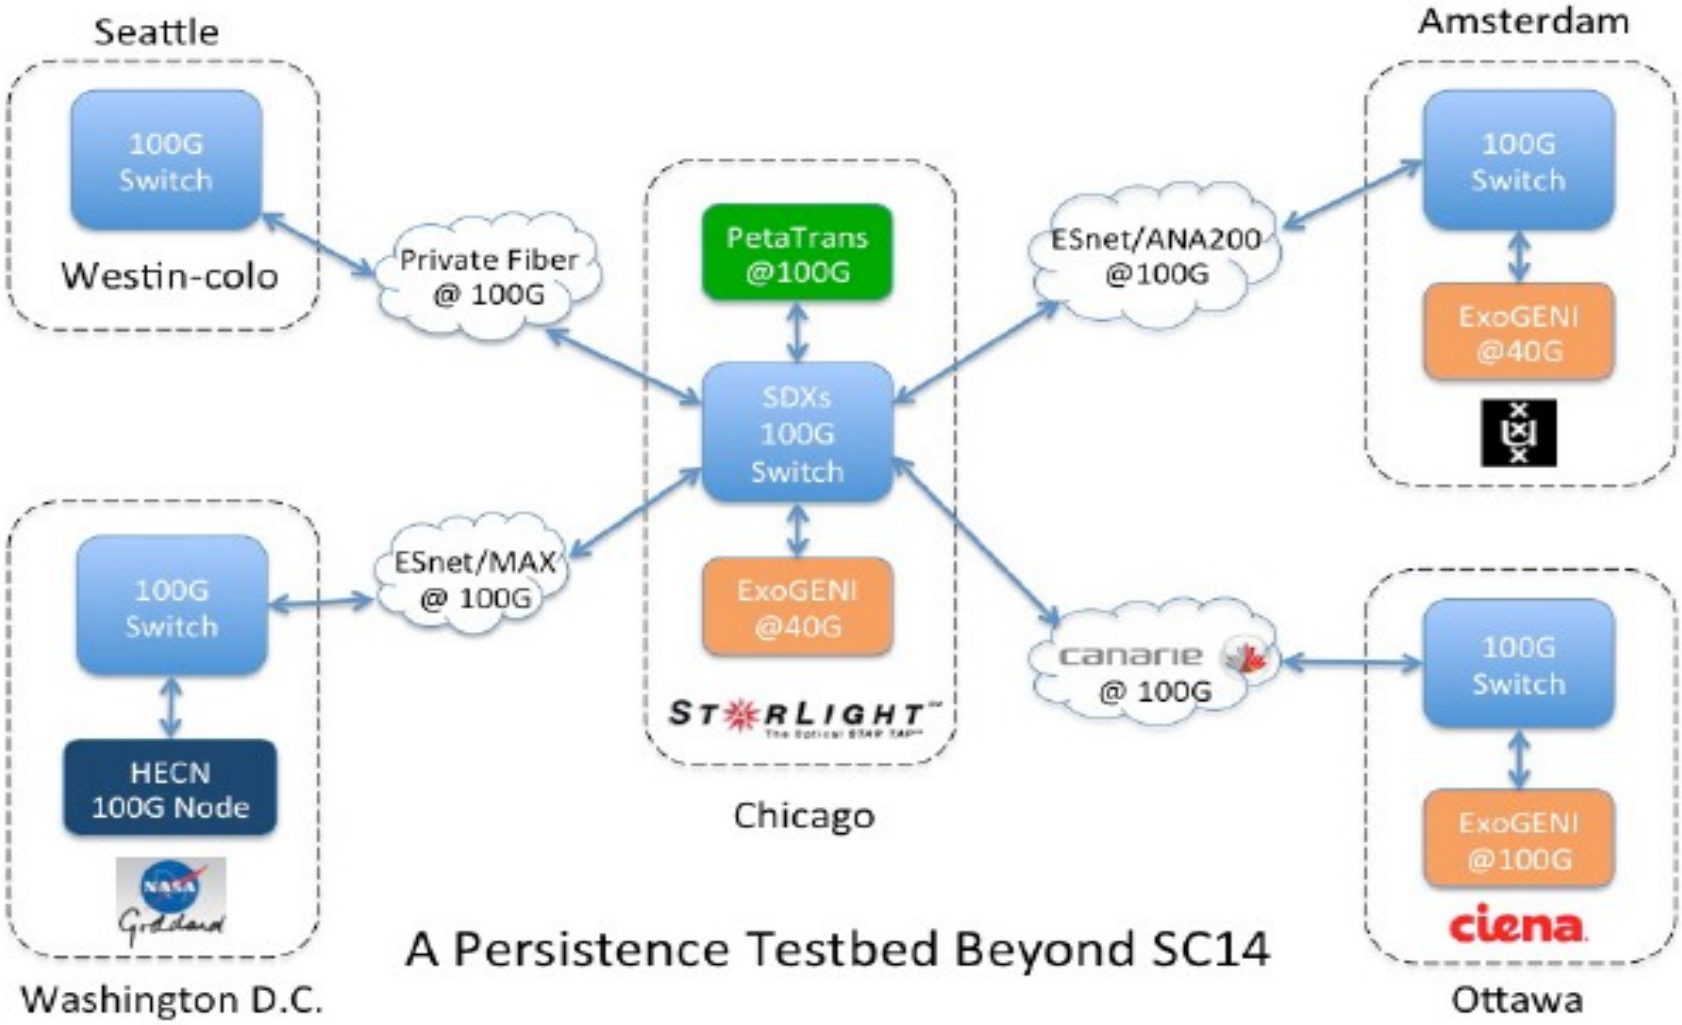

A Persistence Testbed Beyond SC14

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

NSI-OpenFlow Hybird Topology Exchange

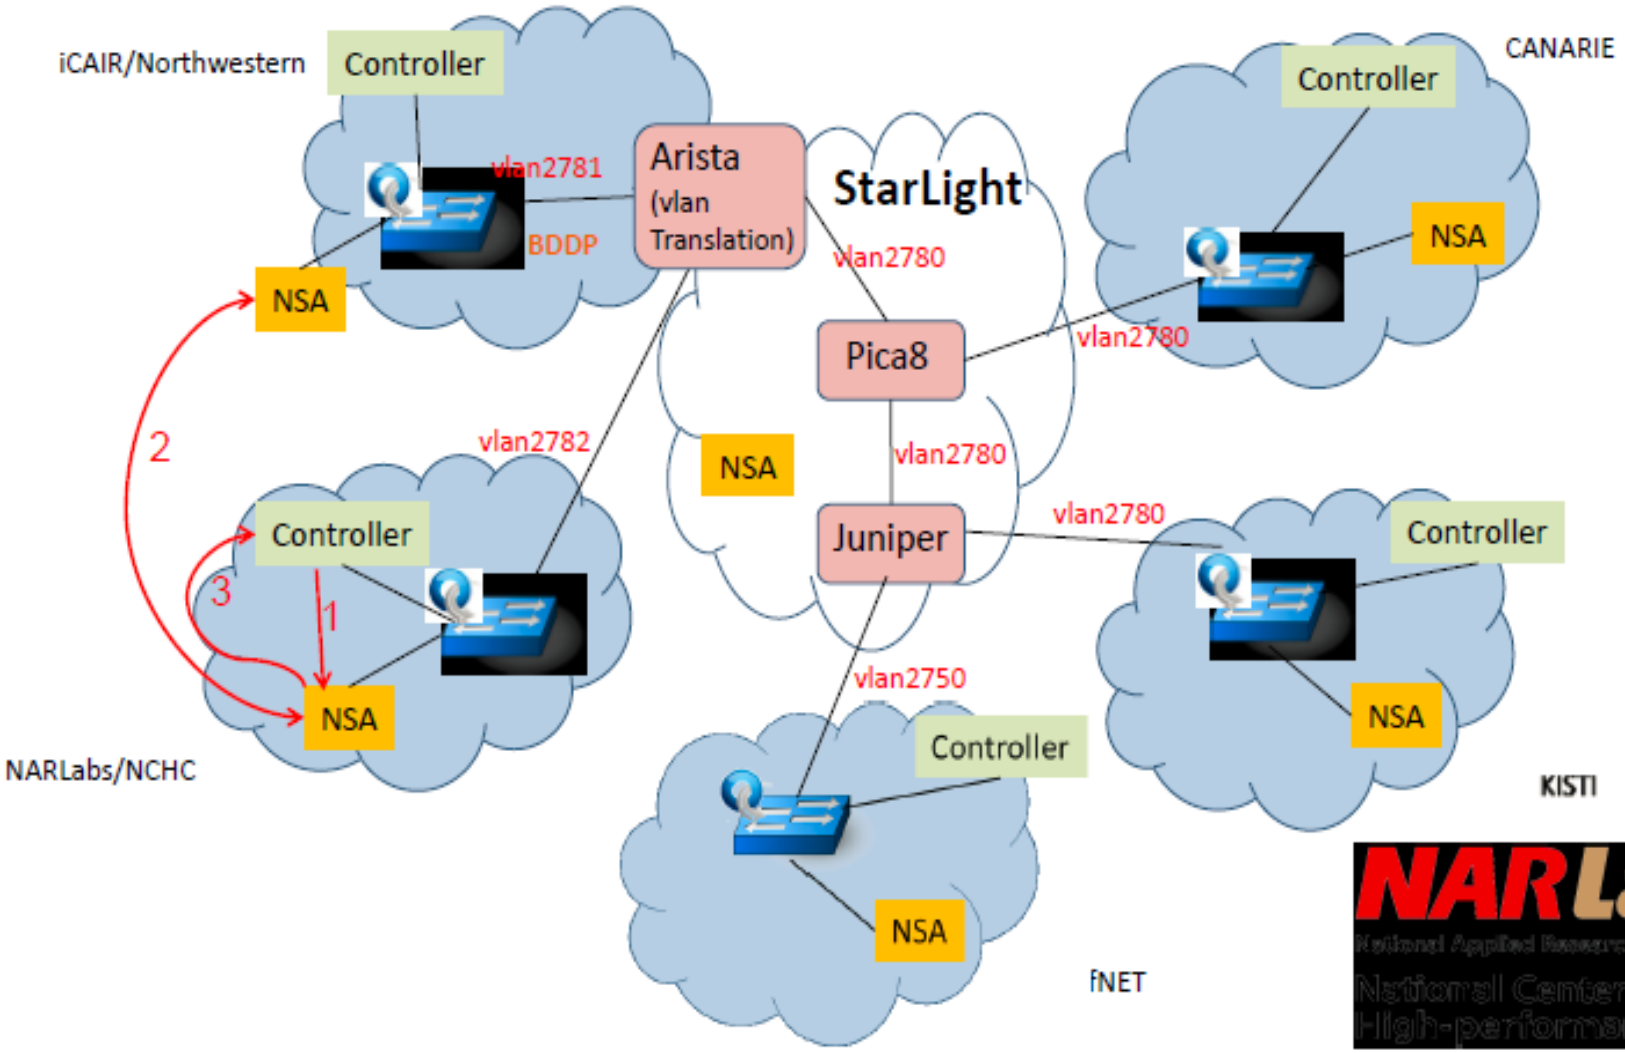

BI Data Flow Visualization (Inbound-Outbound) From SDSC To UoC

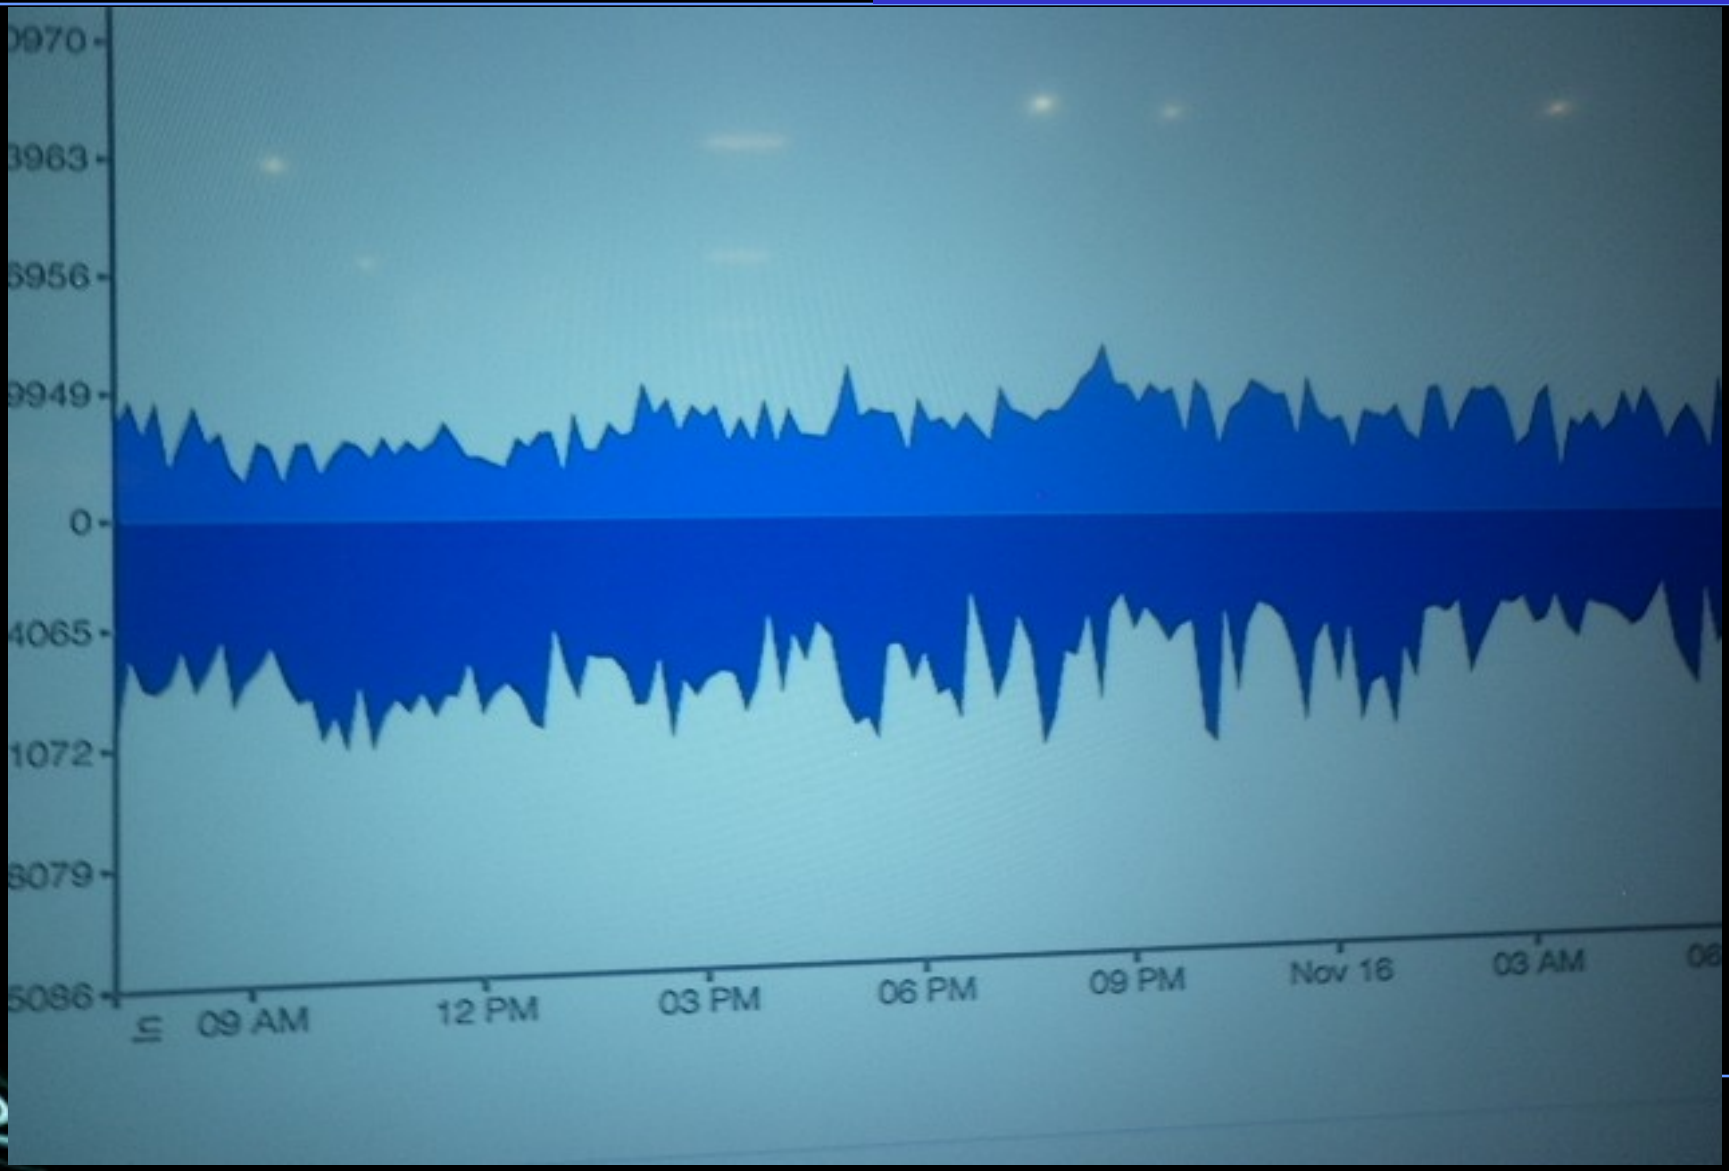

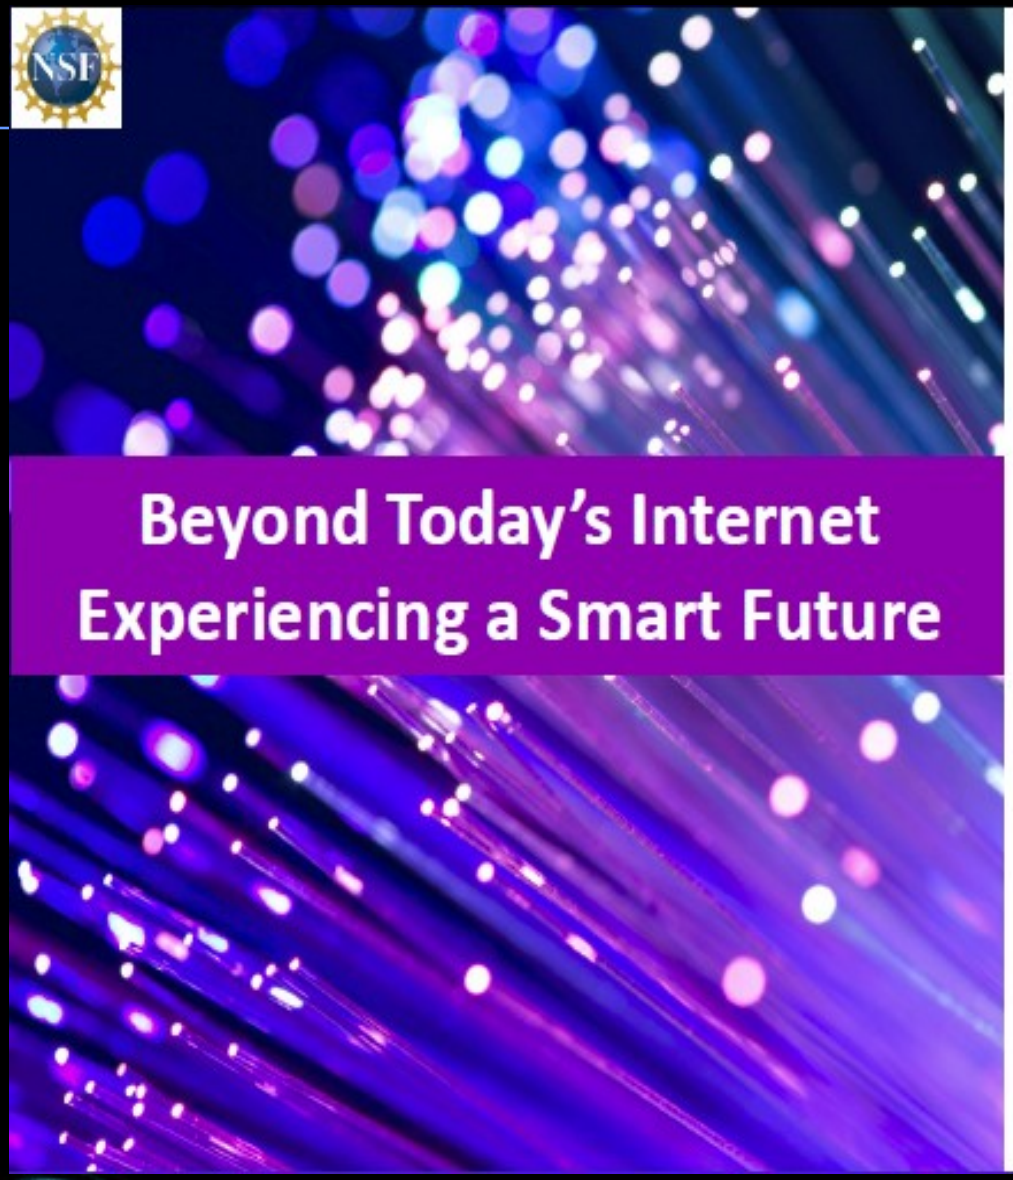

Prototype SDX Bioinformatics Exchange: Demonstrating an Essential Use-Case for Personalized Medicine

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

Future Vision: A Nationwide Virtual Comprehensive Cancer Center

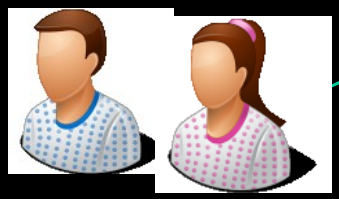

Patients

Hospitals, Doctors

Cloud Computation
Genomic Data Commons

Output: Data-Aware, Analytics-Informed Diagnosis, Prognosis, Optimal Treatment

STARLIGHT SDX

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

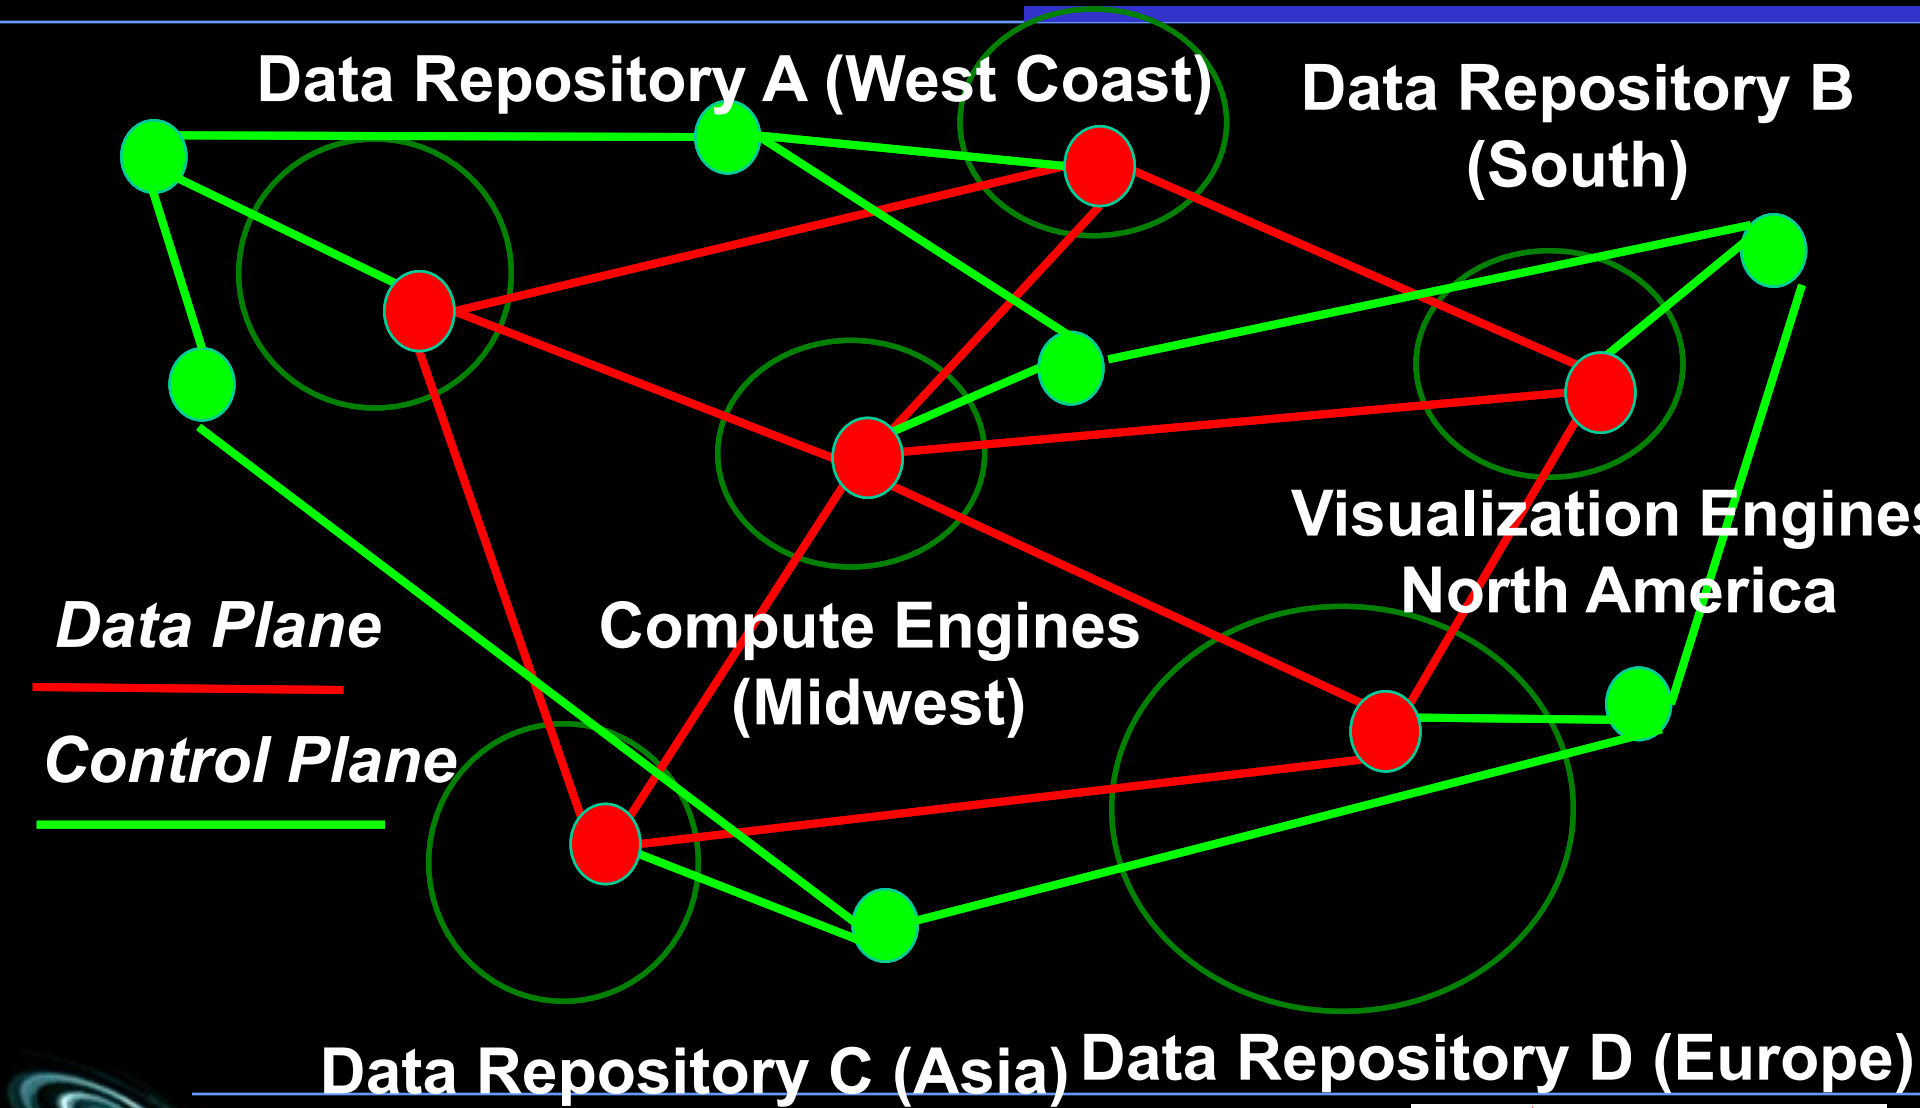

STARLIGHT SDX

GEC22 Bioinformatics SDXs Demo Network

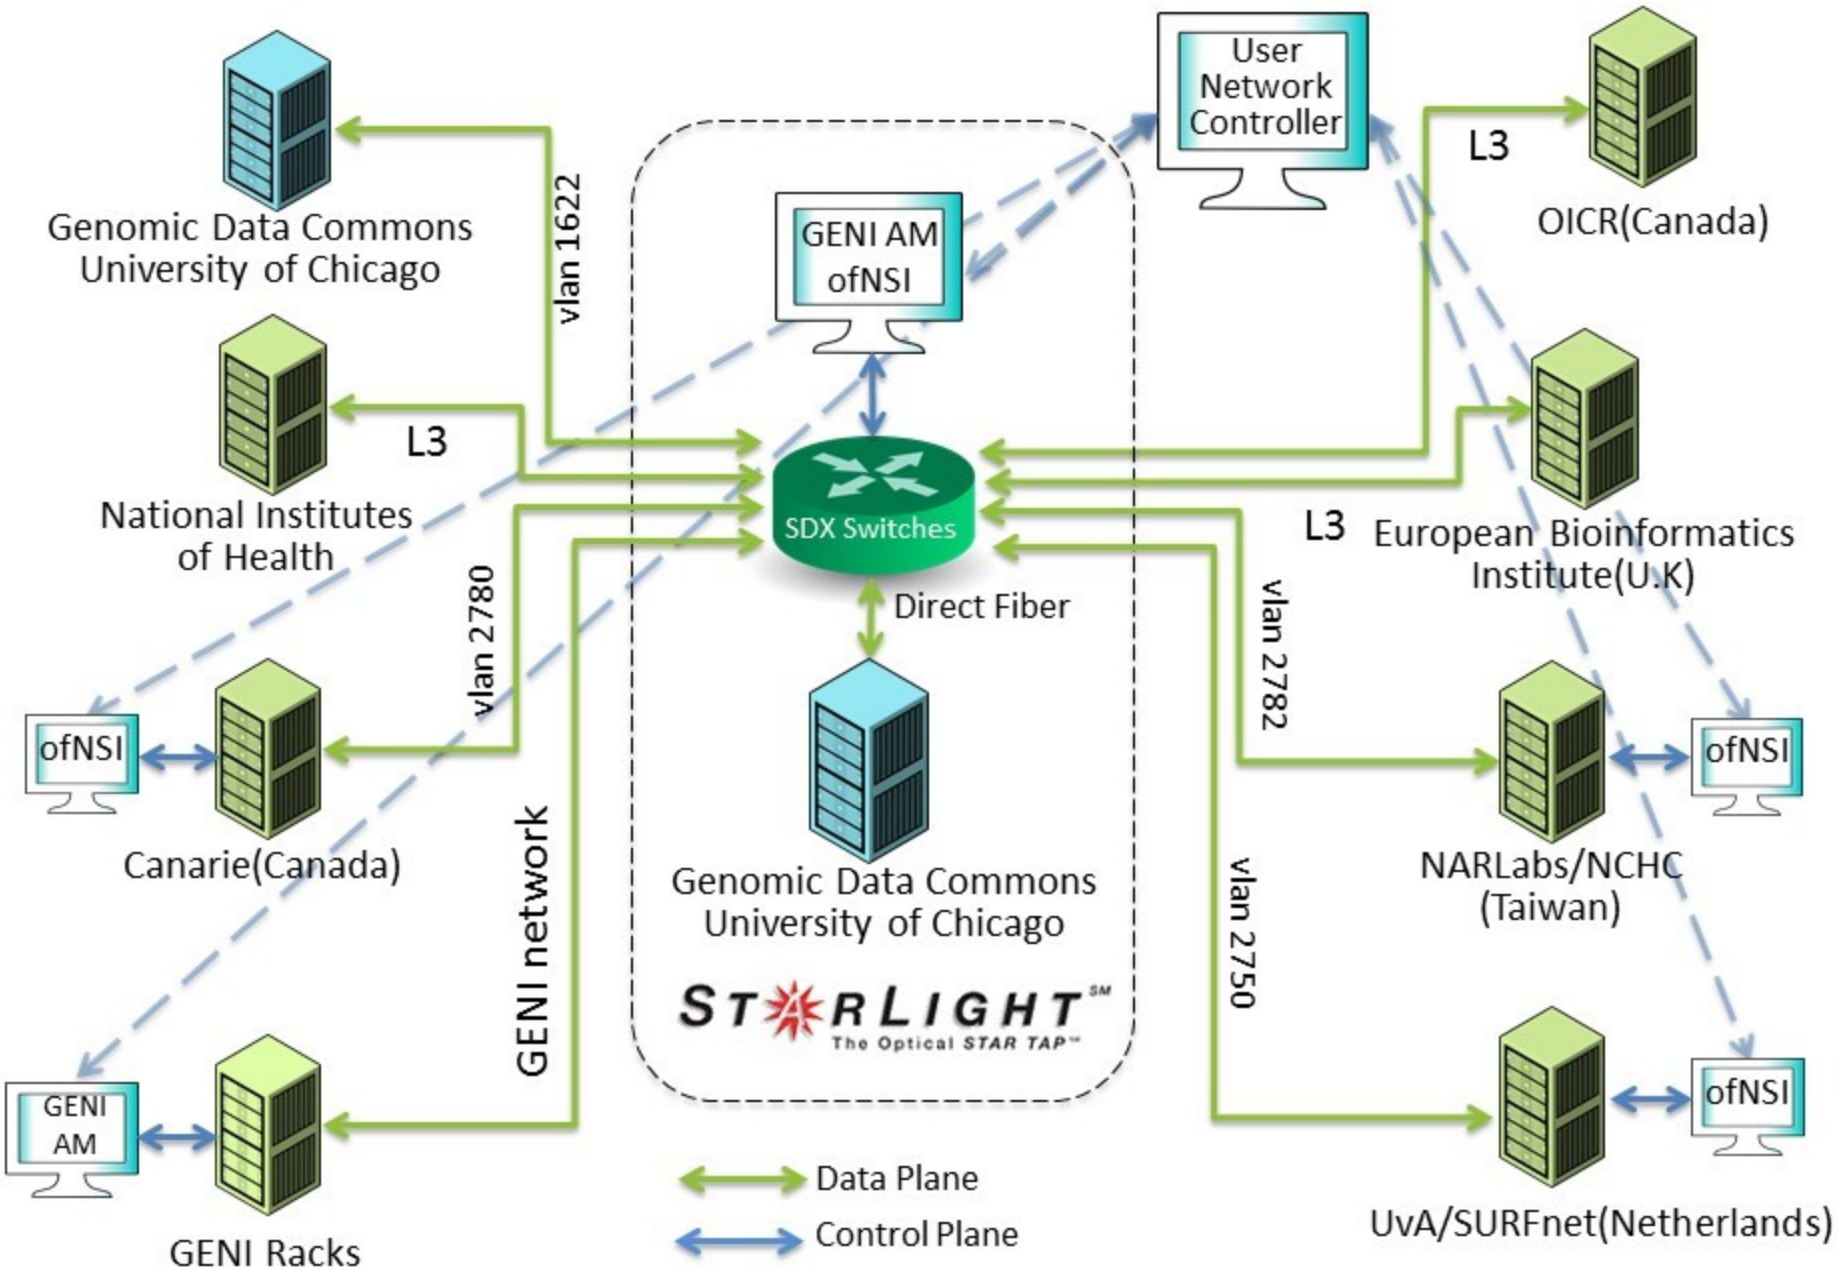

Genomic Data Commons Data Transfer

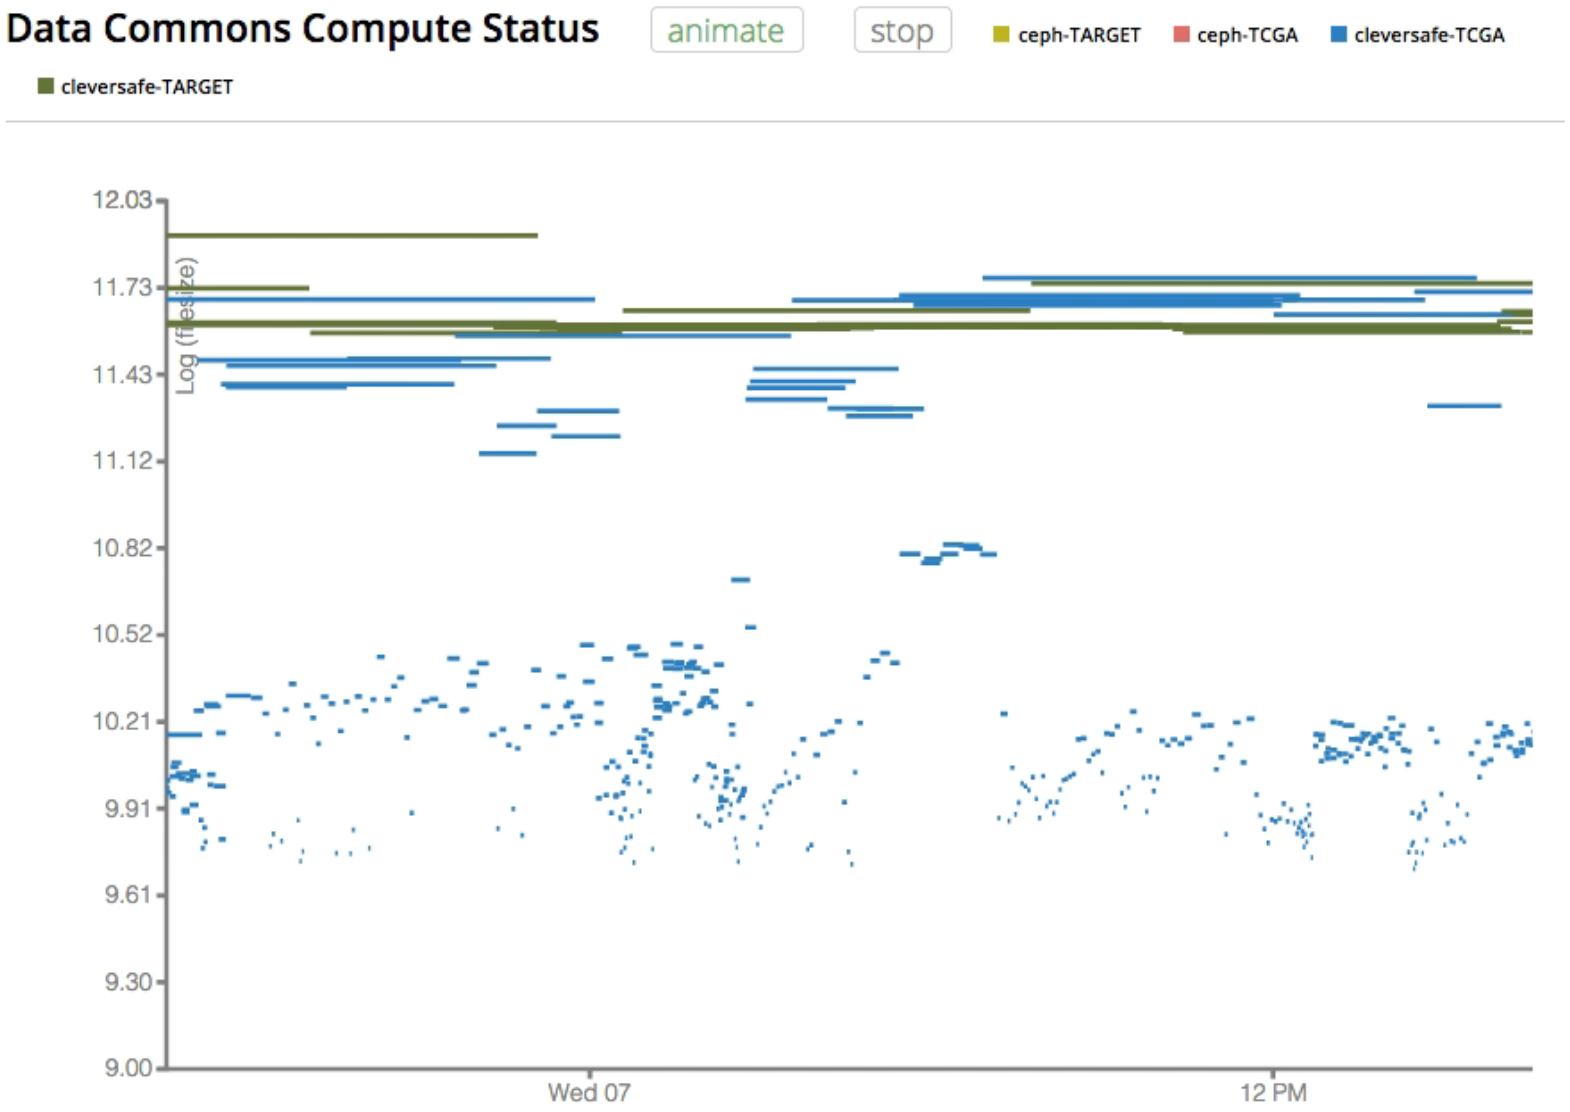

InstaGENI Network

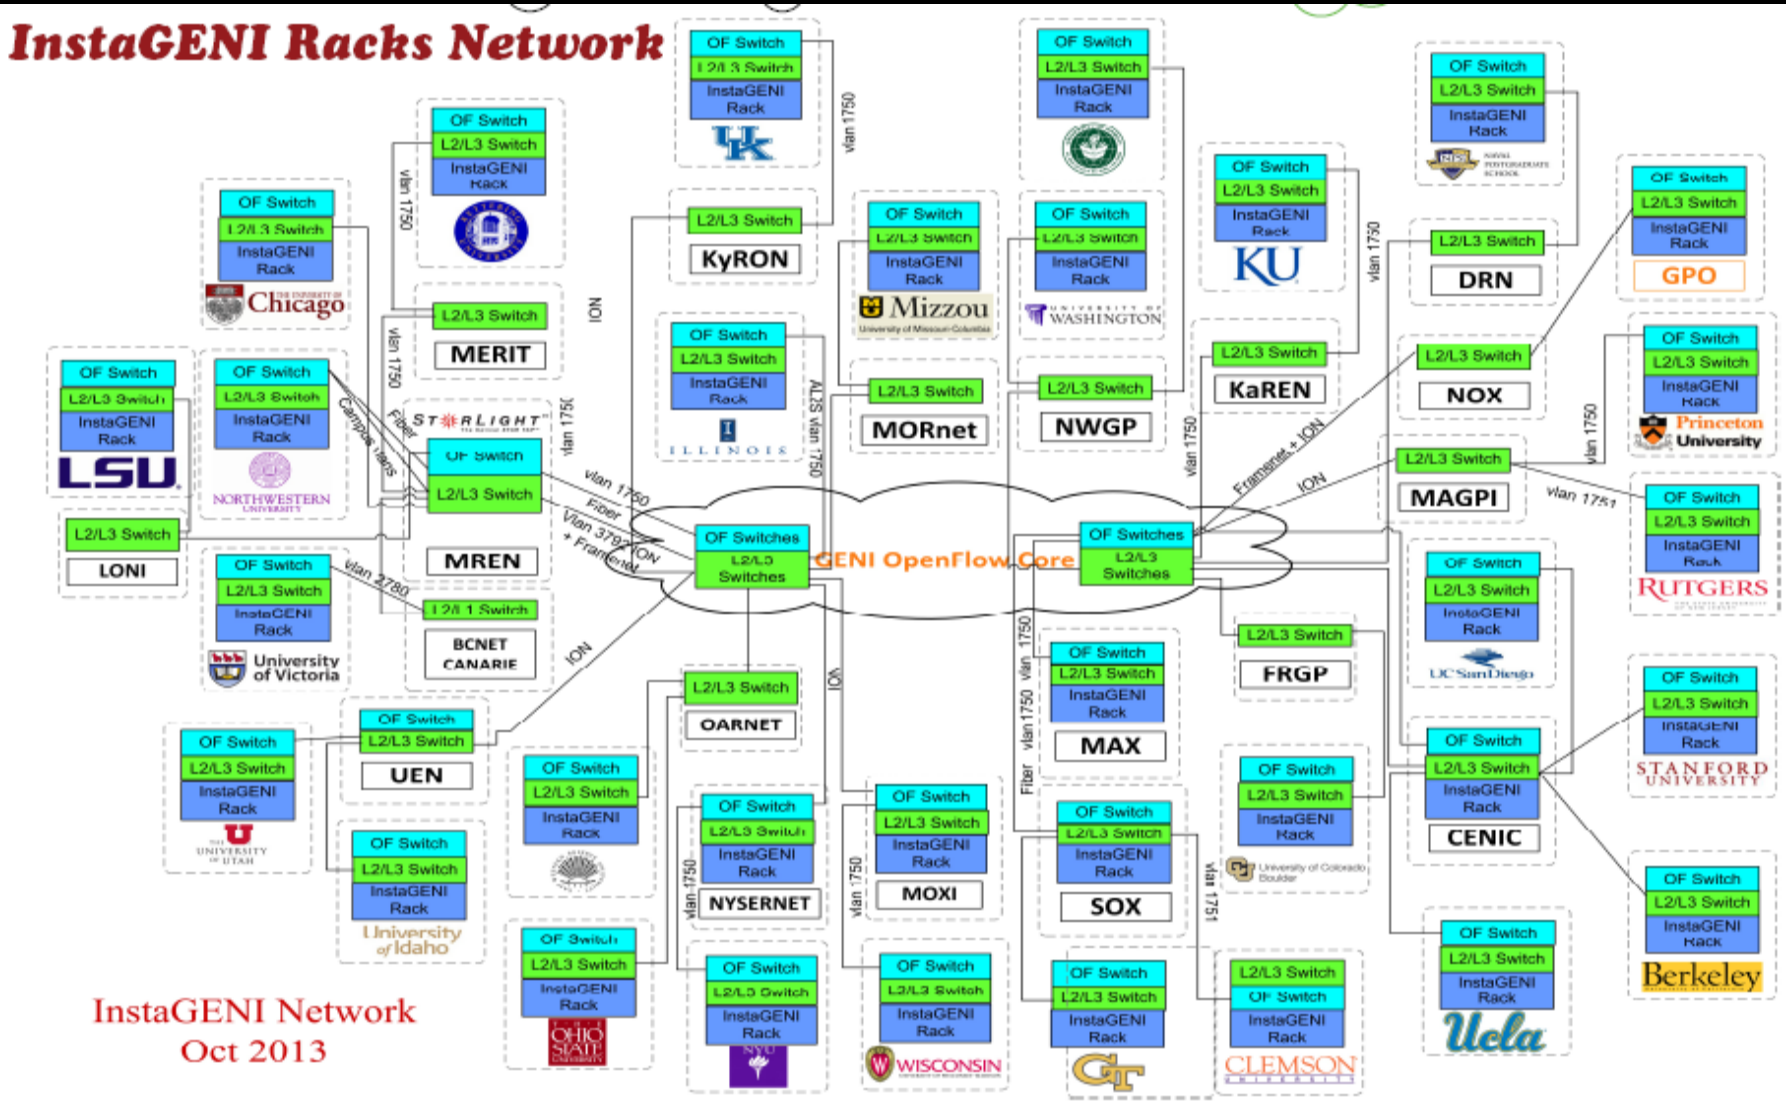

International ExoGENI Testbed

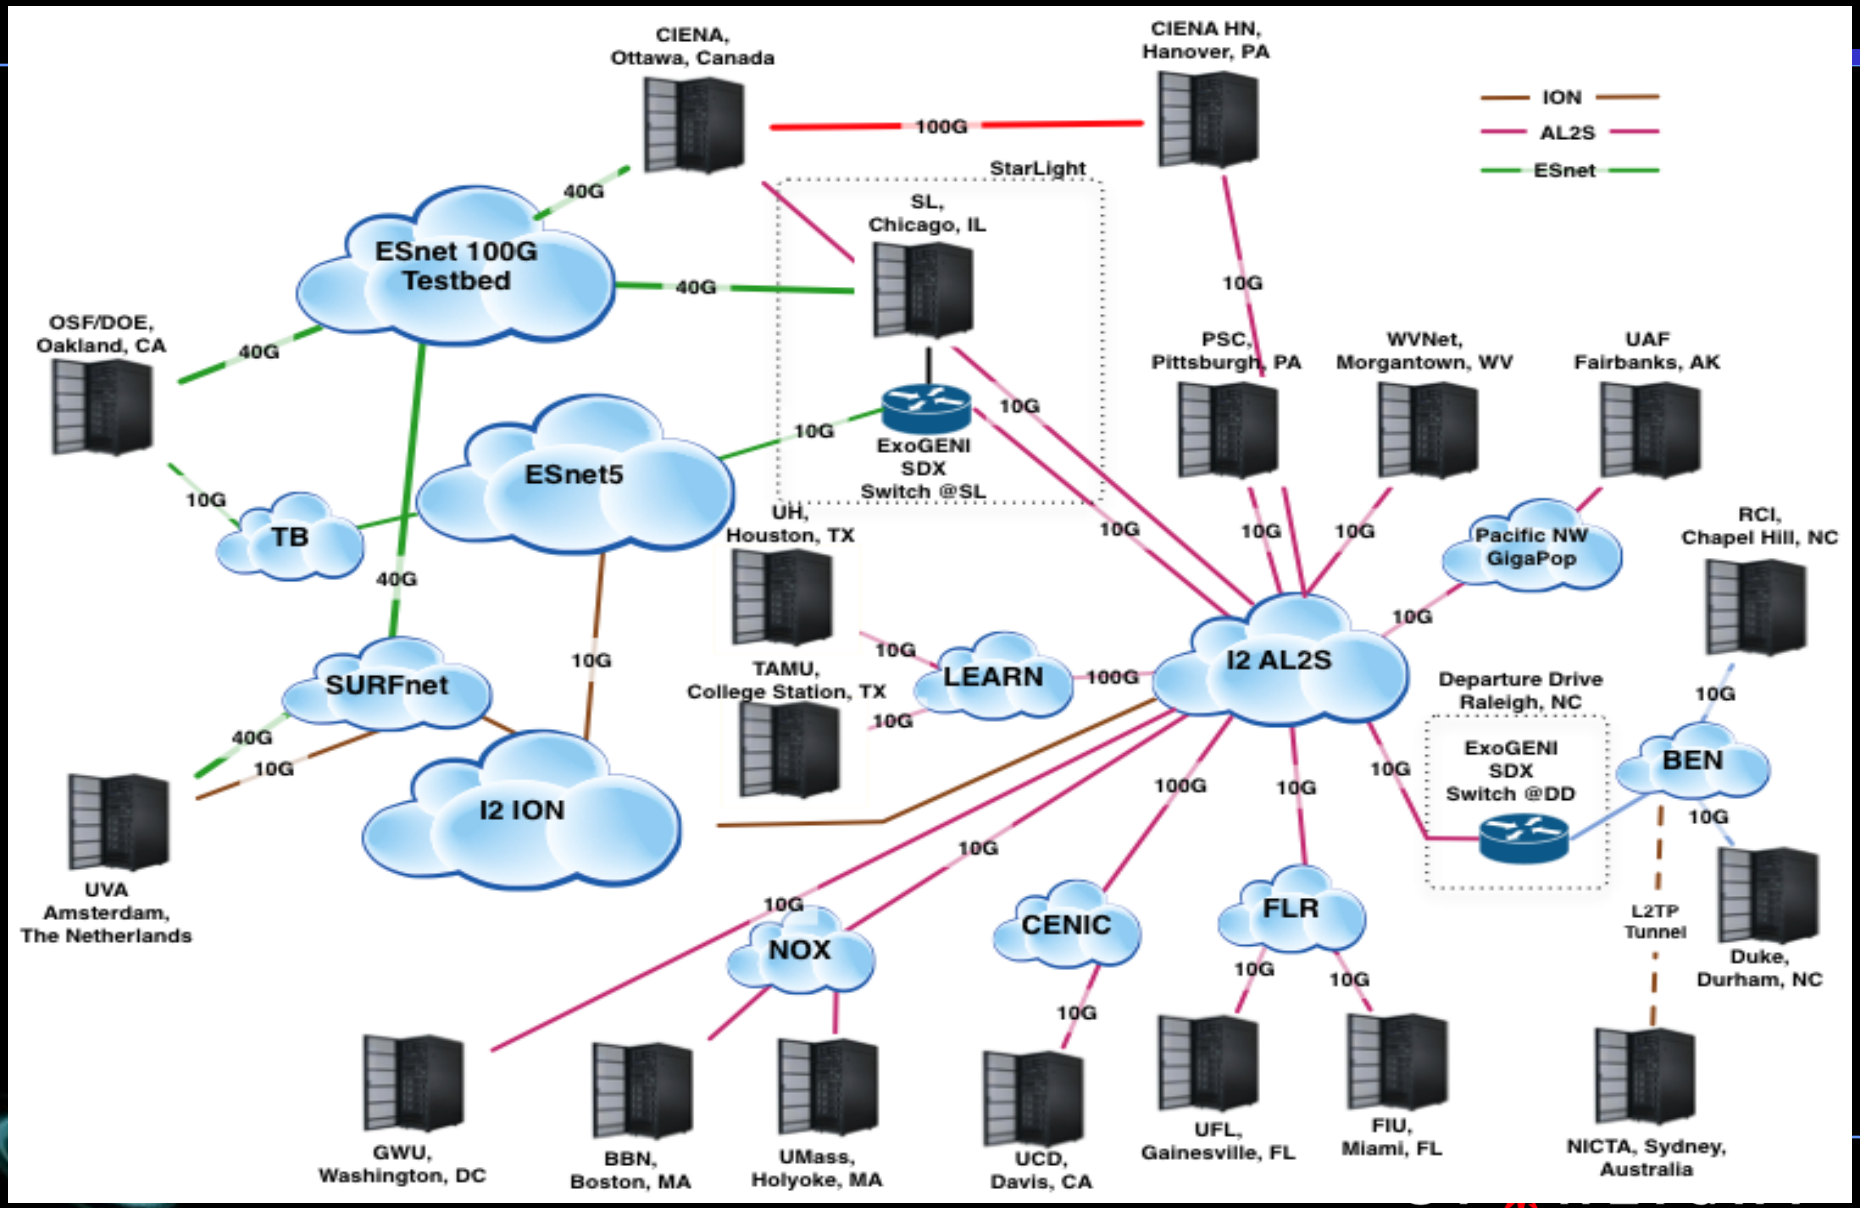

Another SDX Opportunity! An Experimental Testbed For Computer Science Research

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

AUGUST 29, 2014

Testbed to Support the Community's Research Challenges

The community builds the testbed, and afterwards the testbed will shape the community

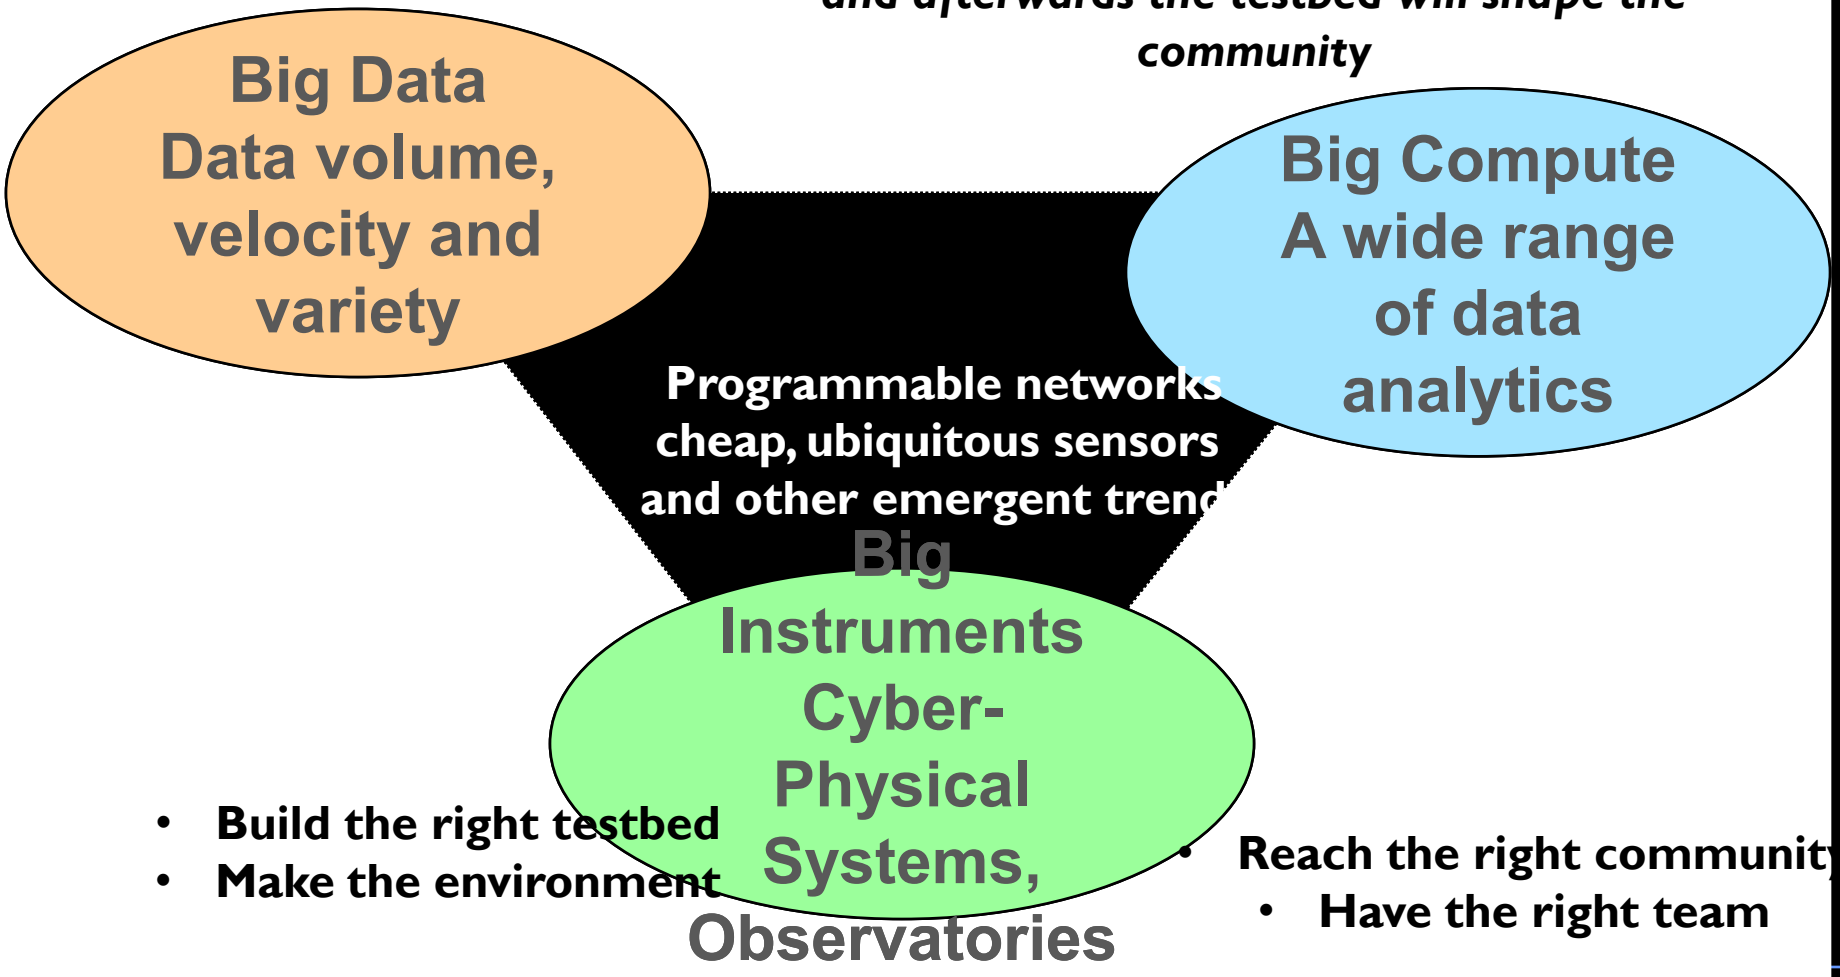

- **Build the right testbed**
- **Make the environment**

- **Reach the right community**
- **Have the right team**

Future Internet Research and Experimentation

What is FIRE?

FIRE Research

FIRE Experimental Facility

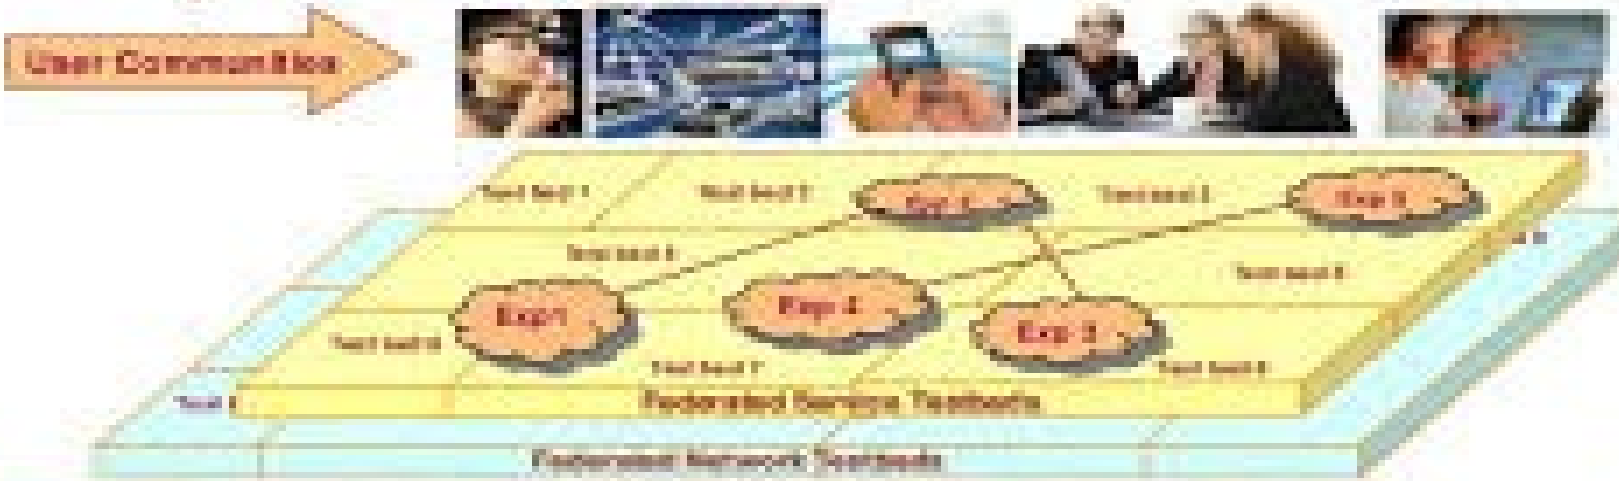

FIRE - Future Internet Research and Experimentation

FIRE⇔FIRE Federation Project⇔ GENI⇔Chameleon

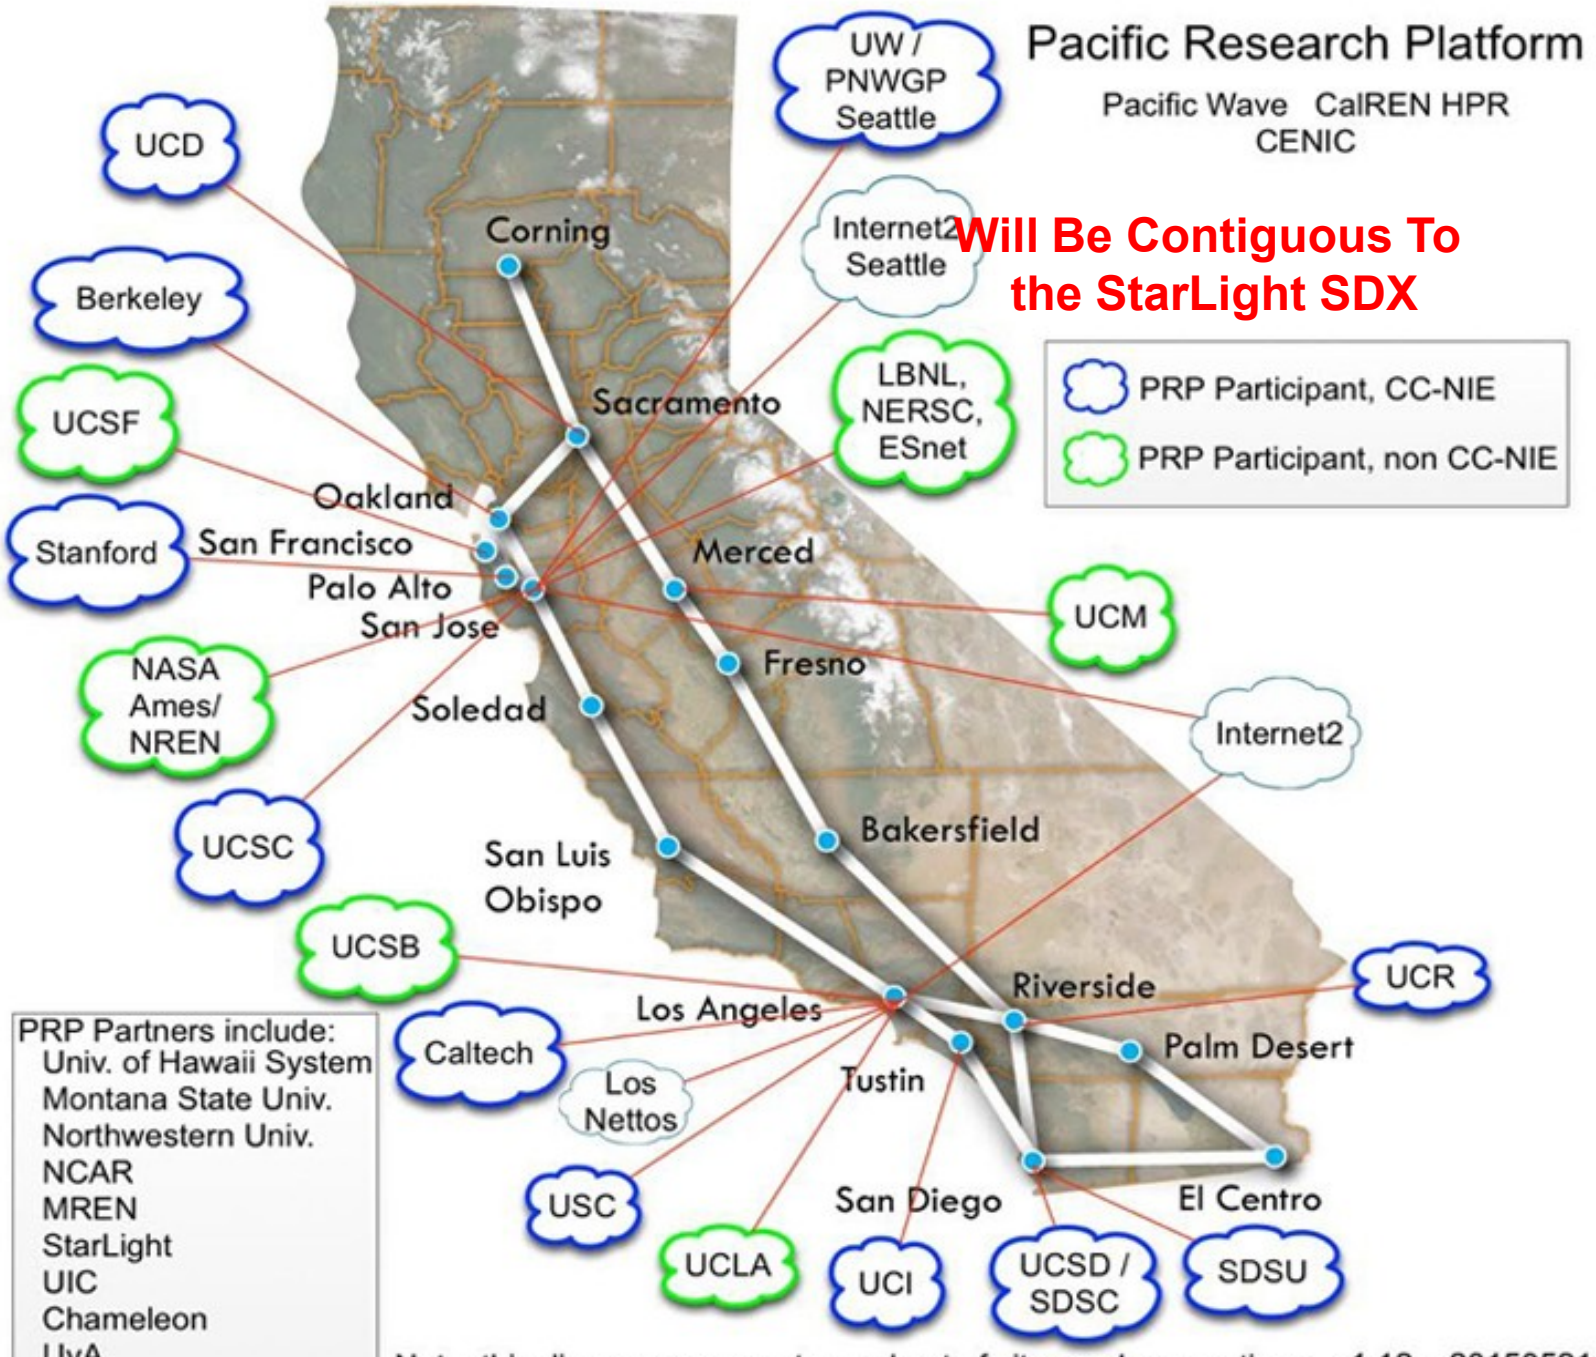
Pacific Research Platform
Pacific Wave CalREN HPR
CENIC
Will Be Contiguous To
the StarLight SDX
UW / PNWGP Seattle
Internet2 Seattle
LBNL, NERSC, ESnet
PRP Participant, CC-NIE
PRP Participant, non CC-NIE
UCD
Berkeley
UCSF
Stanford
NASA Ames/ NREN
UCSC
UCSB
Caltech
Los Nettos
USC
UCLA
UCI
UCSD / SDSC
SDSU
UCM
Internet2
UCR
Palm Desert
El Centro
Corning
Sacramento
Oakland
San Francisco
Palo Alto
San Jose
Soledad
Merced
Fresno
Bakersfield
San Luis Obispo
Los Angeles
Tustin
San Diego
Riverside
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

FELIX

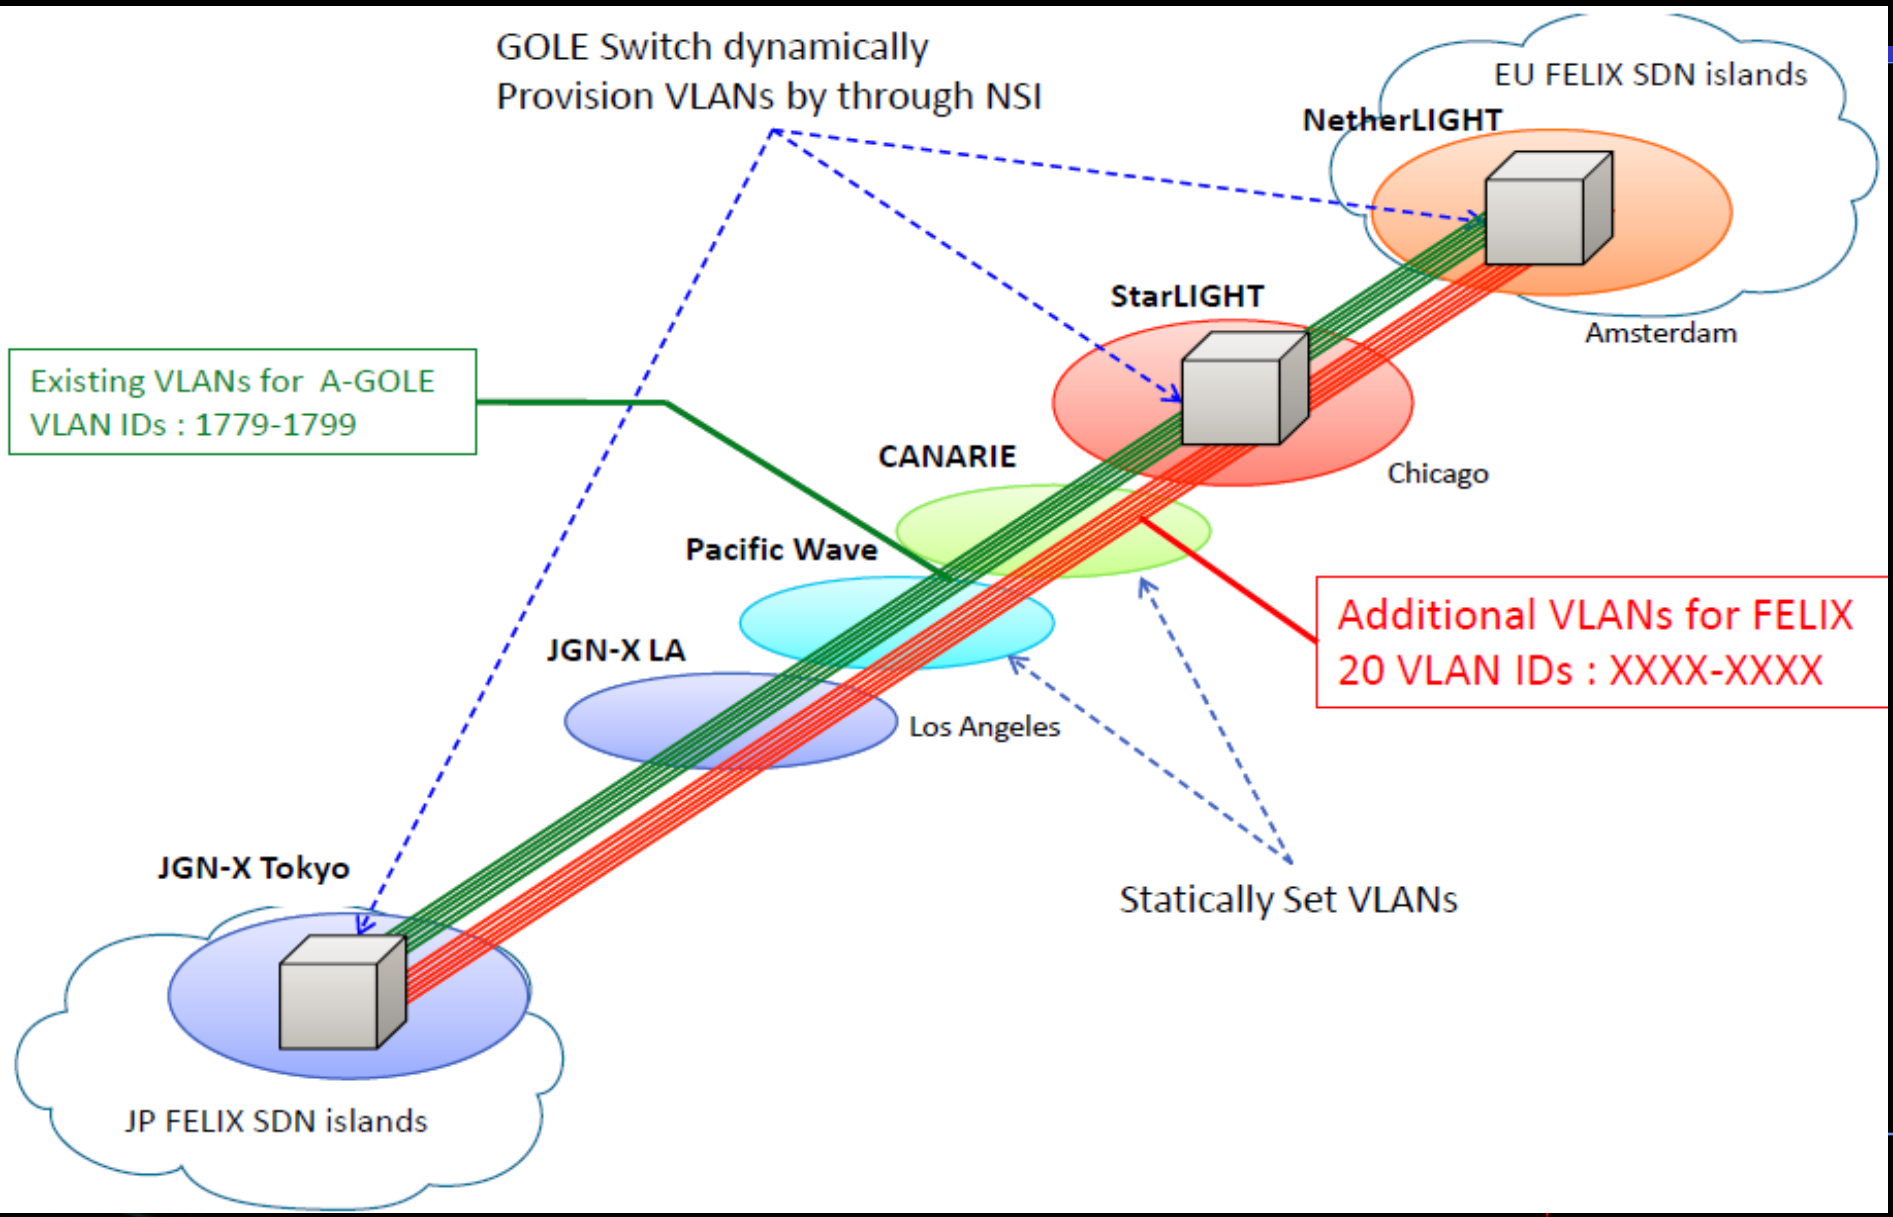
GOLE Switch dynamically
Provision VLANs by through NSI
EU FELIX SDN islands
NetherLIGHT
Amsterdam
StarLIGHT
Chicago
Existing VLANs for A-GOLE
VLAN IDs : 1779-1799
CANARIE
Pacific Wave
JGN-X LA
Los Angeles
Additional VLANs for FELIX
20 VLAN IDs : XXXX-XXXX
JGN-X Tokyo
Statically Set VLANs
JP FELIX SDN islands

GTS connectivity to GENI Facilities
(step 1: Static transit connection)

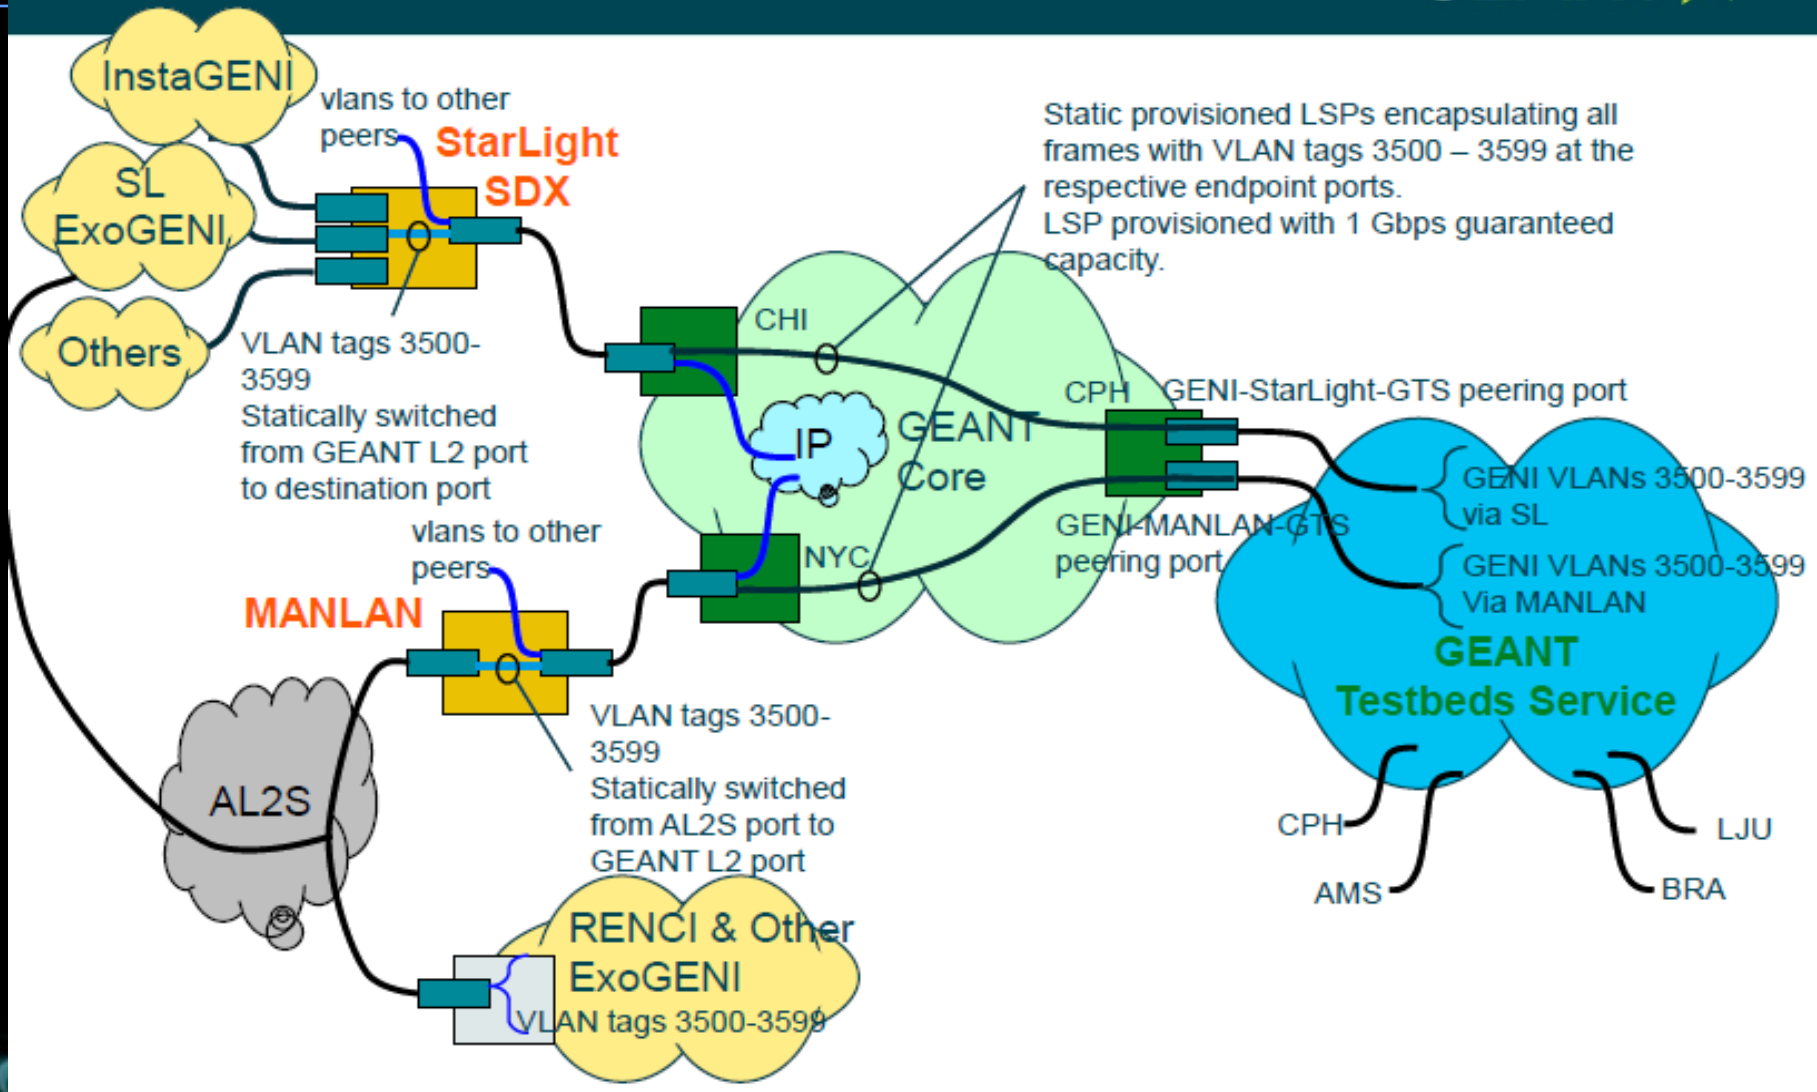

GTS connectivity to GENI Facilities
(step 2: Dynamic transit provisioning)

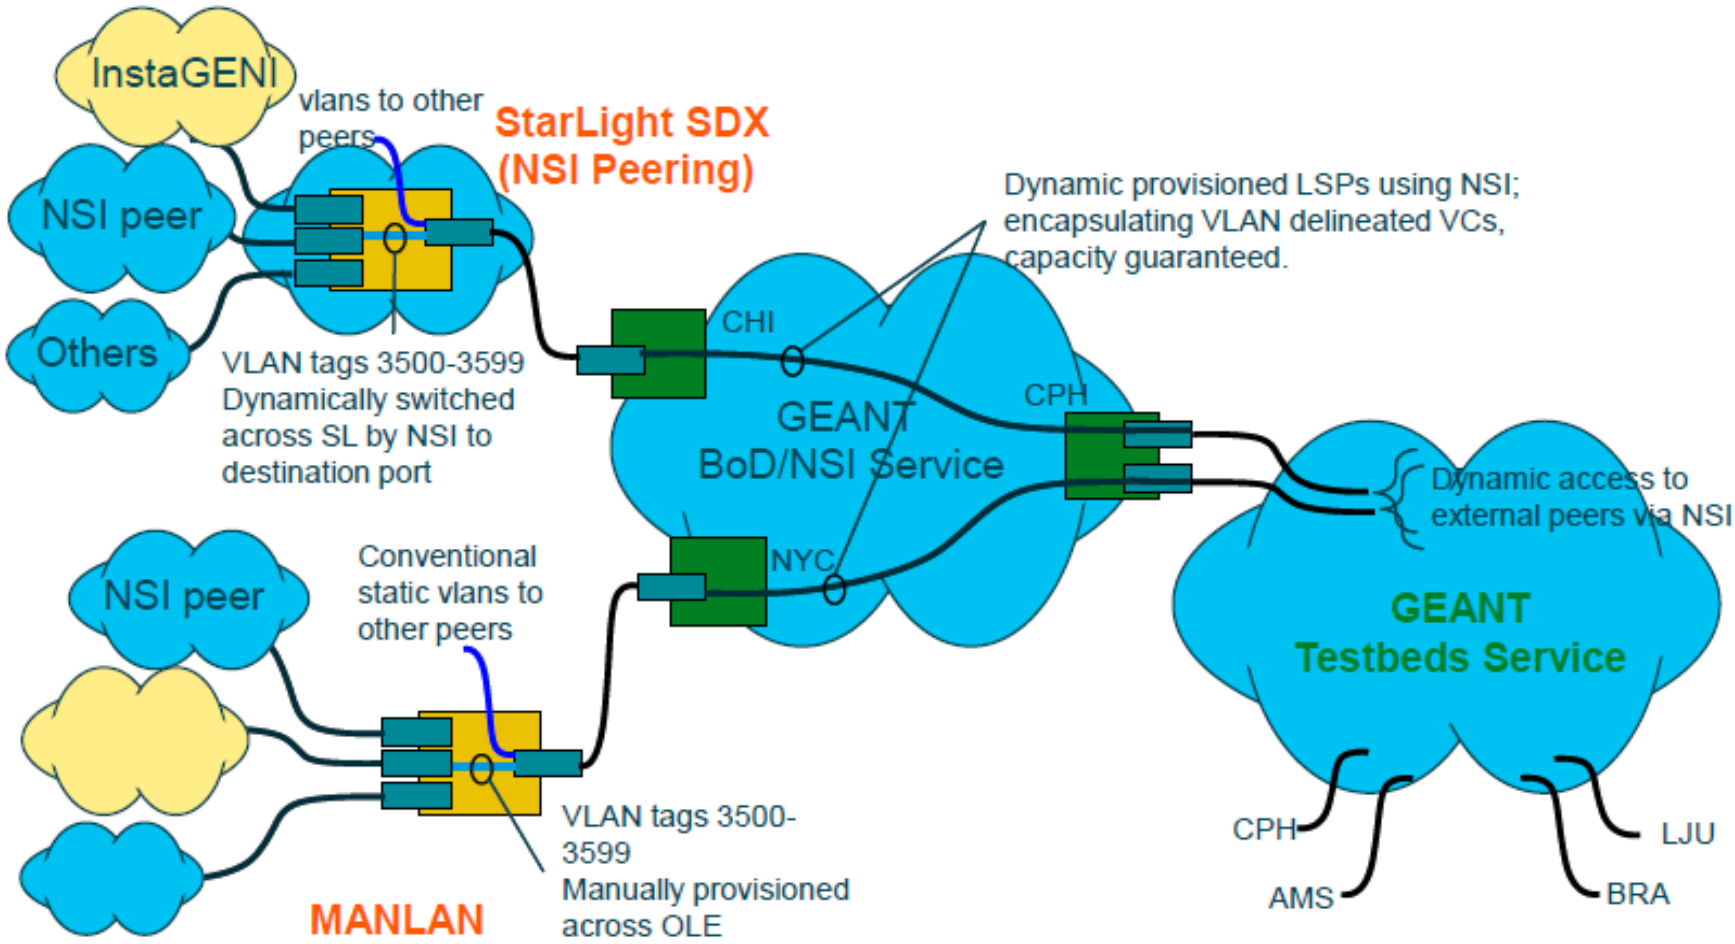

Large Scale File Transfers (~26 TB) for Global LambdaGrid Workshop 2015

- Much Thanks to
  - Michal Krsek
  - John Graham
  - Jan Ruzicka
  - Lukas Hejtmanek
  - Gerben van Malenstein
  - Tom Hutton
  - Et al
  - Phase A = UCSD⇔UvA
  - Phase B = UCSD ⇔ GLGW Prague

STARLIGHT SDX

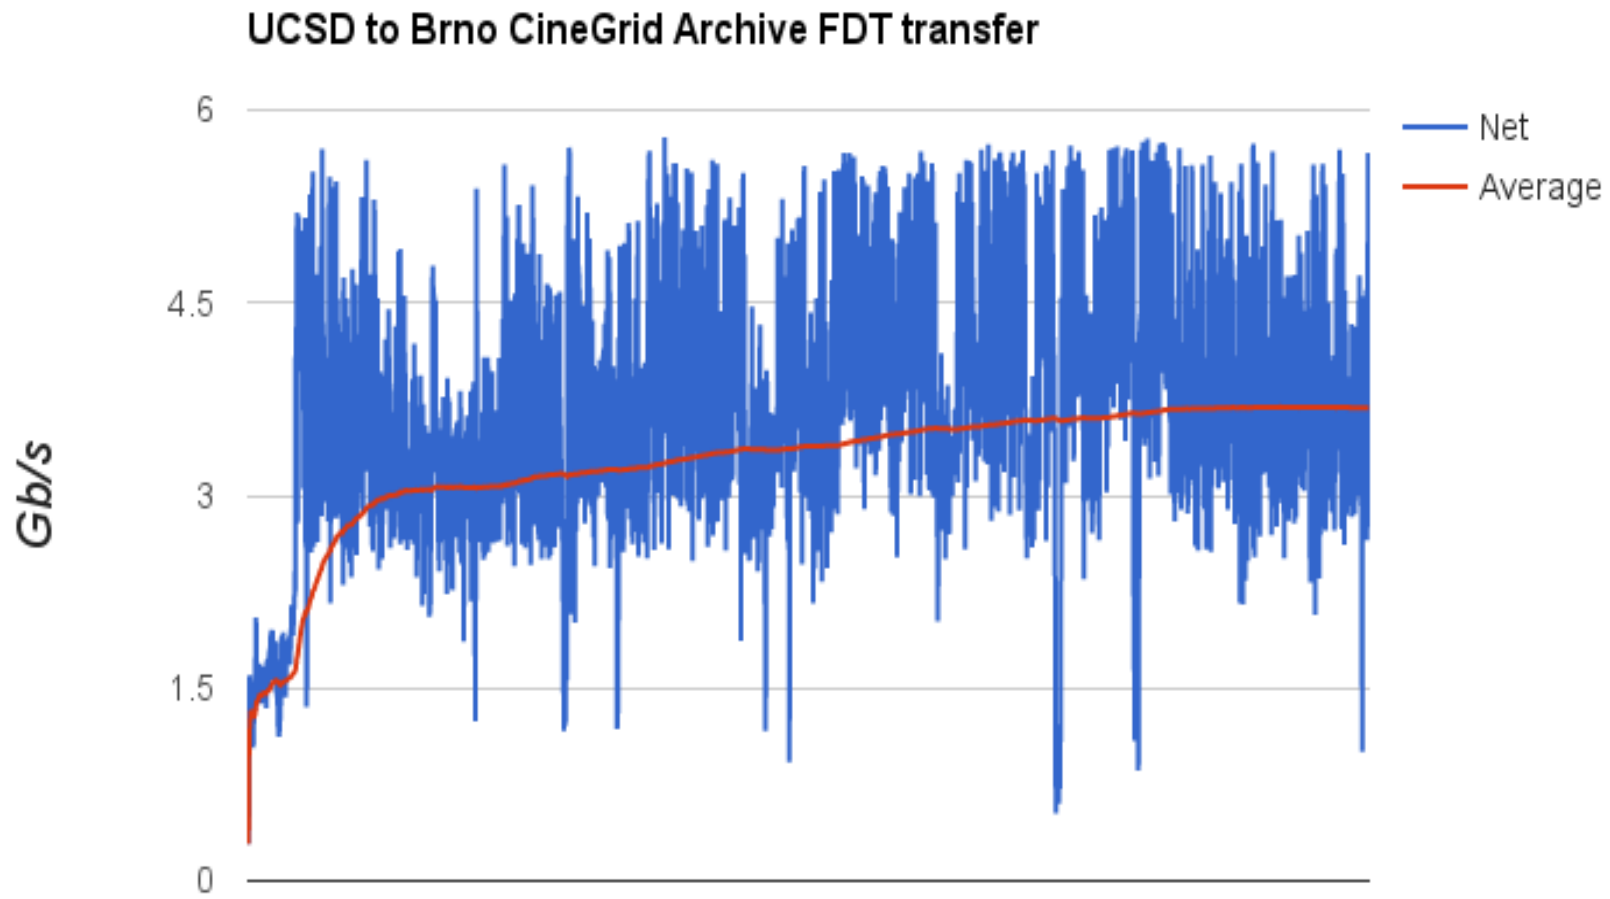
UCSD to Brno CineGrid Archive FDT transfer
Gb/s
6
4.5
3
1.5
0
Net
Average

router, R148, R148.cesnet.cz, 195.113.156.9

[Interfaces]

HundredGigE0/2/0/0

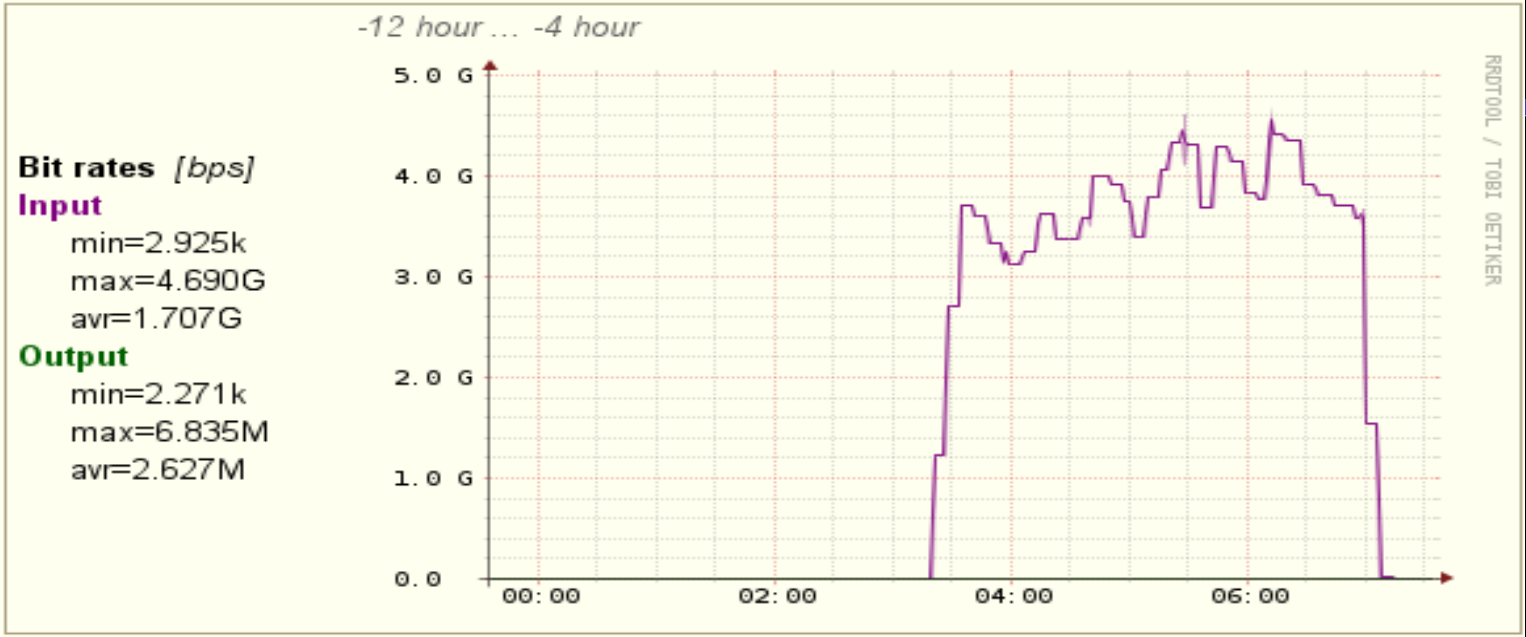

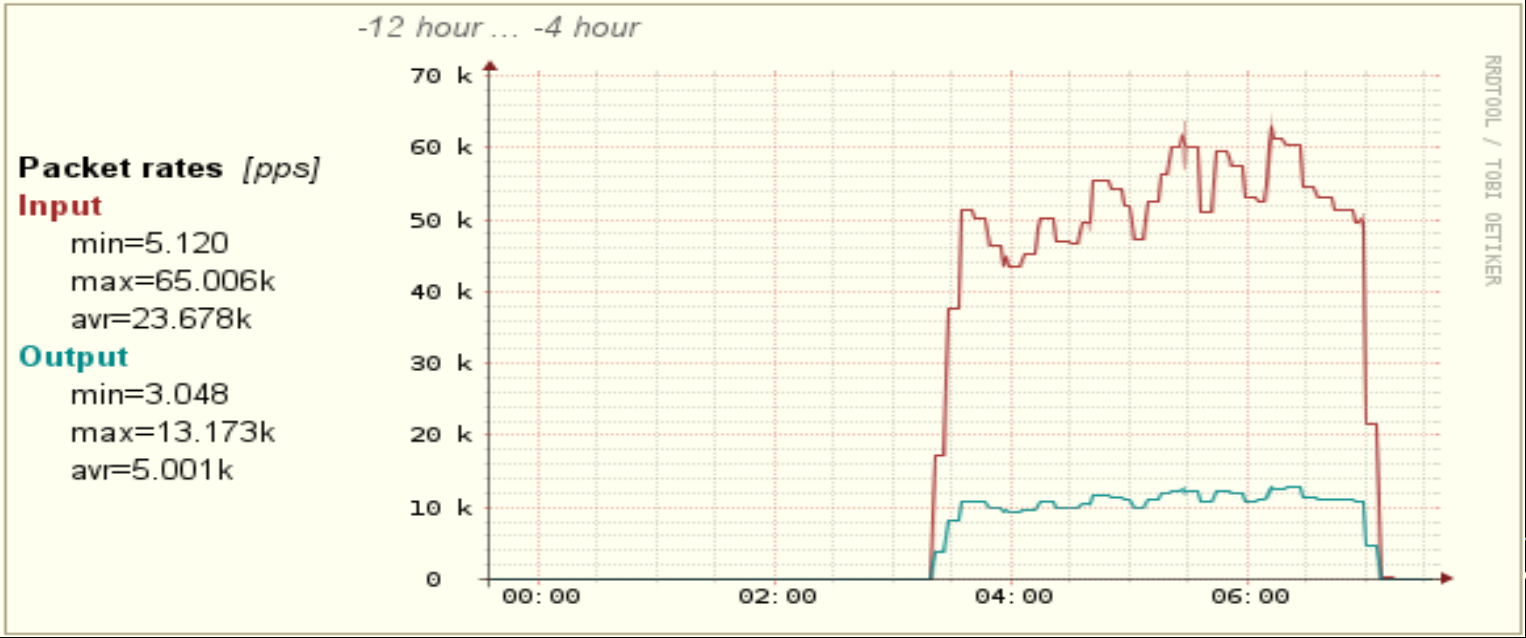

www.startap.net/starlight

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

iCAIR

STARLIGHT SDX