The StarLight International Software-Defined Network Exchange (SDX): Emerging Innovations In Global Interdomain SDN Services, Architecture, Capabilities and Technologies

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, PI-iGENI, PI-OMNINet (www.startap.net/starlight)
Co-PI Chameleon (www.chameleoncloud.org)
Jim Chen, Associate Director, International Center for Advanced Internet Research (www.icair.org)
Northwestern University

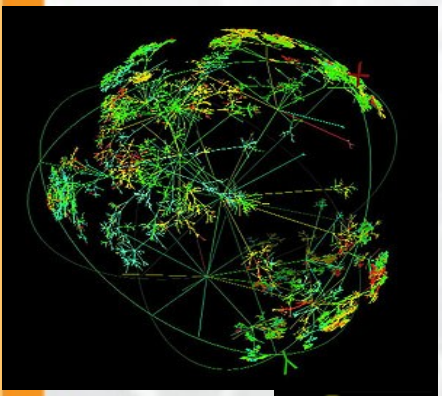

GLIF Technical Working Group Meeting

March 26-27, 2015

StarLight SDX Initiative

- GENI StarLight SDX Project Goal: To Provide Implementation Of Key Software and Hardware Components of Layer 2 SDN/OpenFlow Exchange Between GENI L2 Network Resources and Other Research Networks.
- Provide 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 AM API, GENI Stitching AMs, and the GENI Commercial Software Defined Exchange Point.
- Integrate Provided GENI tools With Experimenter Tools From Other Participating Networks.
- Demonstrate One Functioning Exchange With at Least Two Research Network Partners, Two Data-Intensive Science Campuses, and Multiple Experimenters On Multiple Participating Layer 2 Networks.

Software Defined Networking Exchanges (SDXs)

- With the Increasing Deployment of SDN In Production Networks, the Need for an SDN Exchange (SDX) Has Been Recognized.
- Current SDN Architecture Is Single Domain Centralized Controller Oriented – Many SDN Islands Being Created
- Many Motivations Exist for SDXs – e.g., Connecting SDN Islands, Enhanced Control Over Specific Traffic Flows, Resource Optimization, Network Function Virtualization -- ***New Services***
- Required Capabilities: Multi-Domain Distributed SDN Resource Discovery, Signaling Provisioning, Operations, and Fault Detection and Recovery.

SDN & The Transformation of Exchanges

- Transitions
  - From Static L3 BGP Exchanges
  - To Multi-Layer Multi-Services Exchanges
  - Foundations
    - Grid Networking Federation
    - OGF - GLIF NIS CS
    - Global Environment For Network Innovations (GENI)
  - Enabling Multiple Layers/Services and Hybrids
  - Deep Visibility Into All Flows
  - Direct Control Over All Flows
  - Significant New Capabilities for Customized Exchanges!
  - Including For Domain Sciences…

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
www.alma.nrao.edu

ANDRILL: Antarctic Geological Drilling
www.andrill.org

BIRN: Biomedical Informatics Research Network
www.nbirn.net

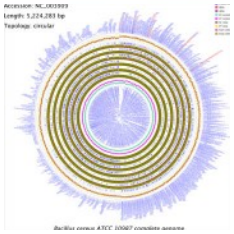

CAMERA metagenomics
camera.calit2.net

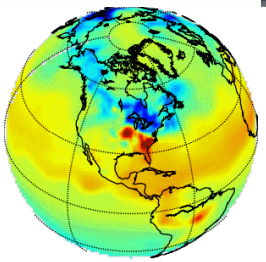

Carbon Tracker
www.esrl.noaa.gov/gmd/ccgg/carbontracker

CineGrid
www.cinegrid.org

LHCONE
www.lhcone.net

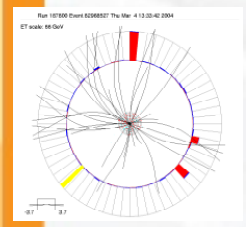

DØ (DZero)
www-d0.fnal.gov

GEON: Geosciences Network
www.geongrid.org

GLEON: Global Lake Ecological Observatory Network
www.gleon.org

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

OOI-CI
ci.oceanobservatories.org

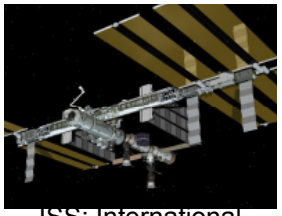

ISS: International Space Station
www.nasa.gov/station

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

LIGO
www.ligo.org

IVOA: International Virtual Observatory
www.ivoa.net

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

PRAGMA

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

TeraGrid
www.teragrid.org

OSG
www.opensciencegrid.org

Globus Alliance
www.globus.org

SKA
www.skatelescope.org

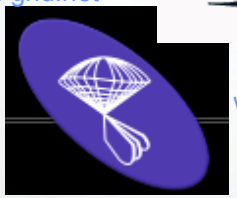

Sloan Digital Sky Survey
www.sdss.org

XSEDE
www.xsede.org

Compilation By Maxine Brown

Many Years Ago, the GLIF Community Began Developing The Network Service Interface (NSI) Architecture With the Open Grid Forum (OGF) To Enable Path Resource Sharing Among GOLEs Which Have Multiple Control Frameworks

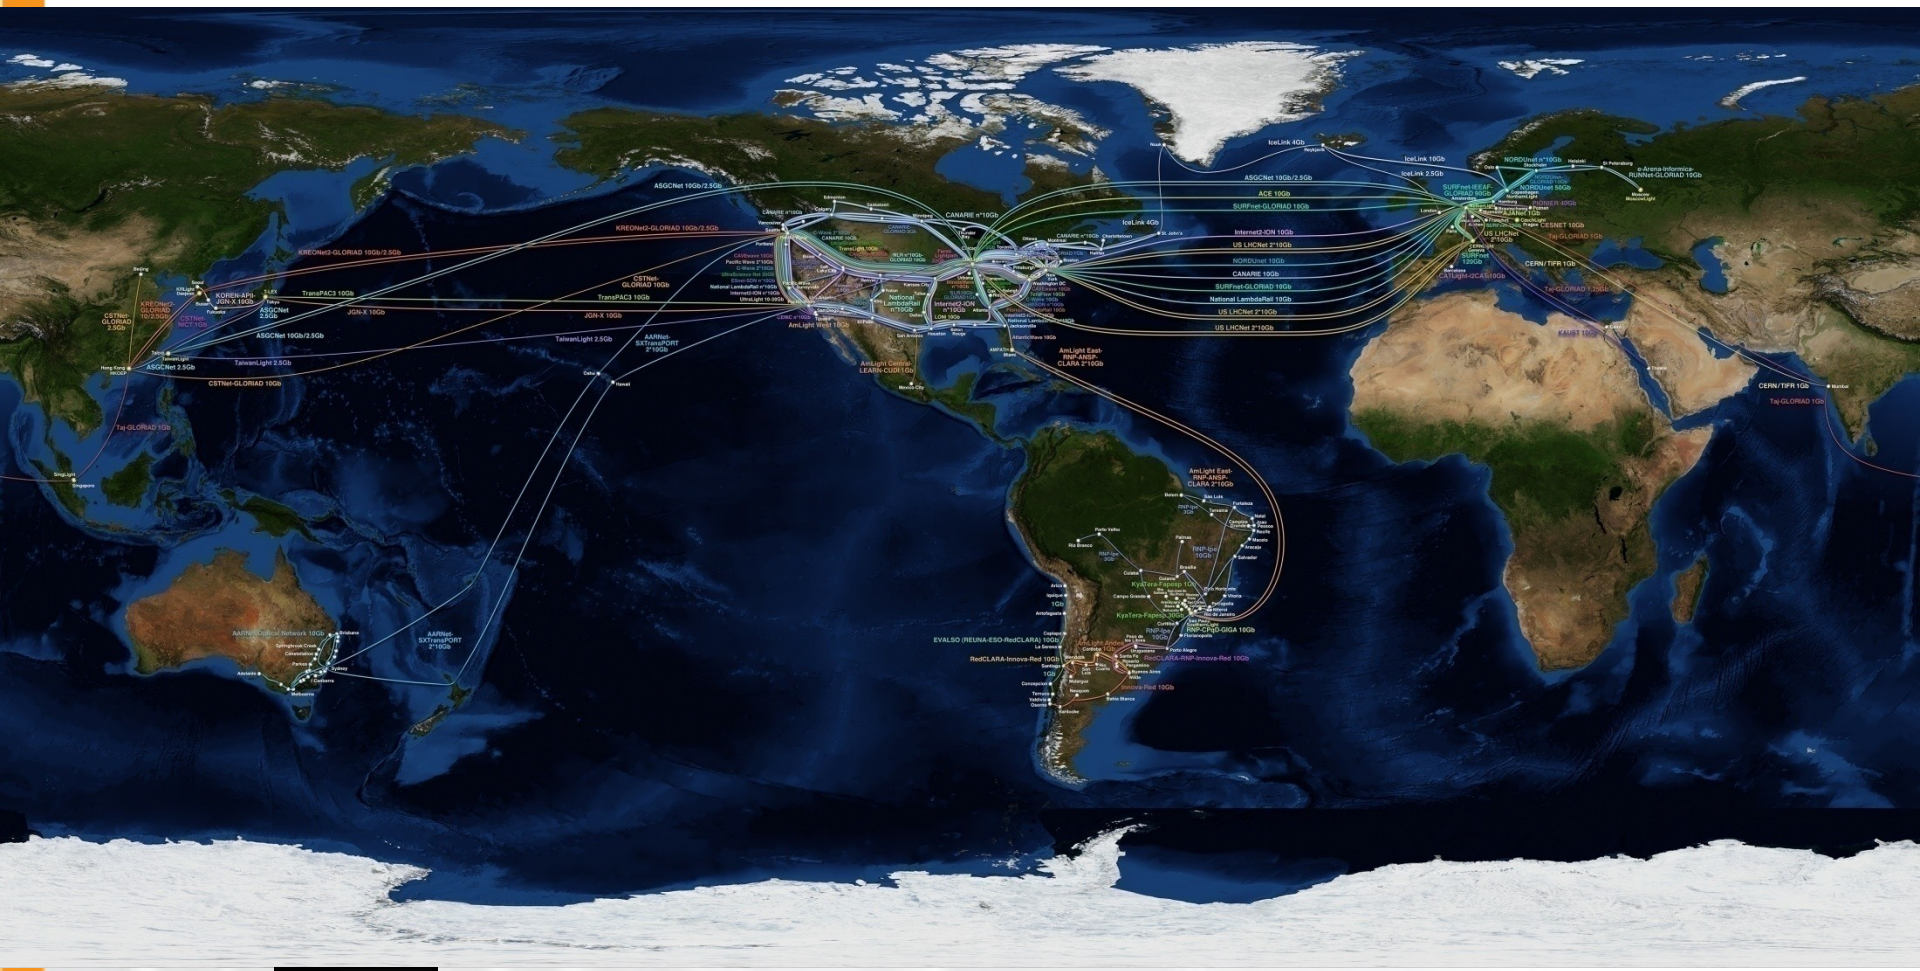

Automated GOLE Fabric

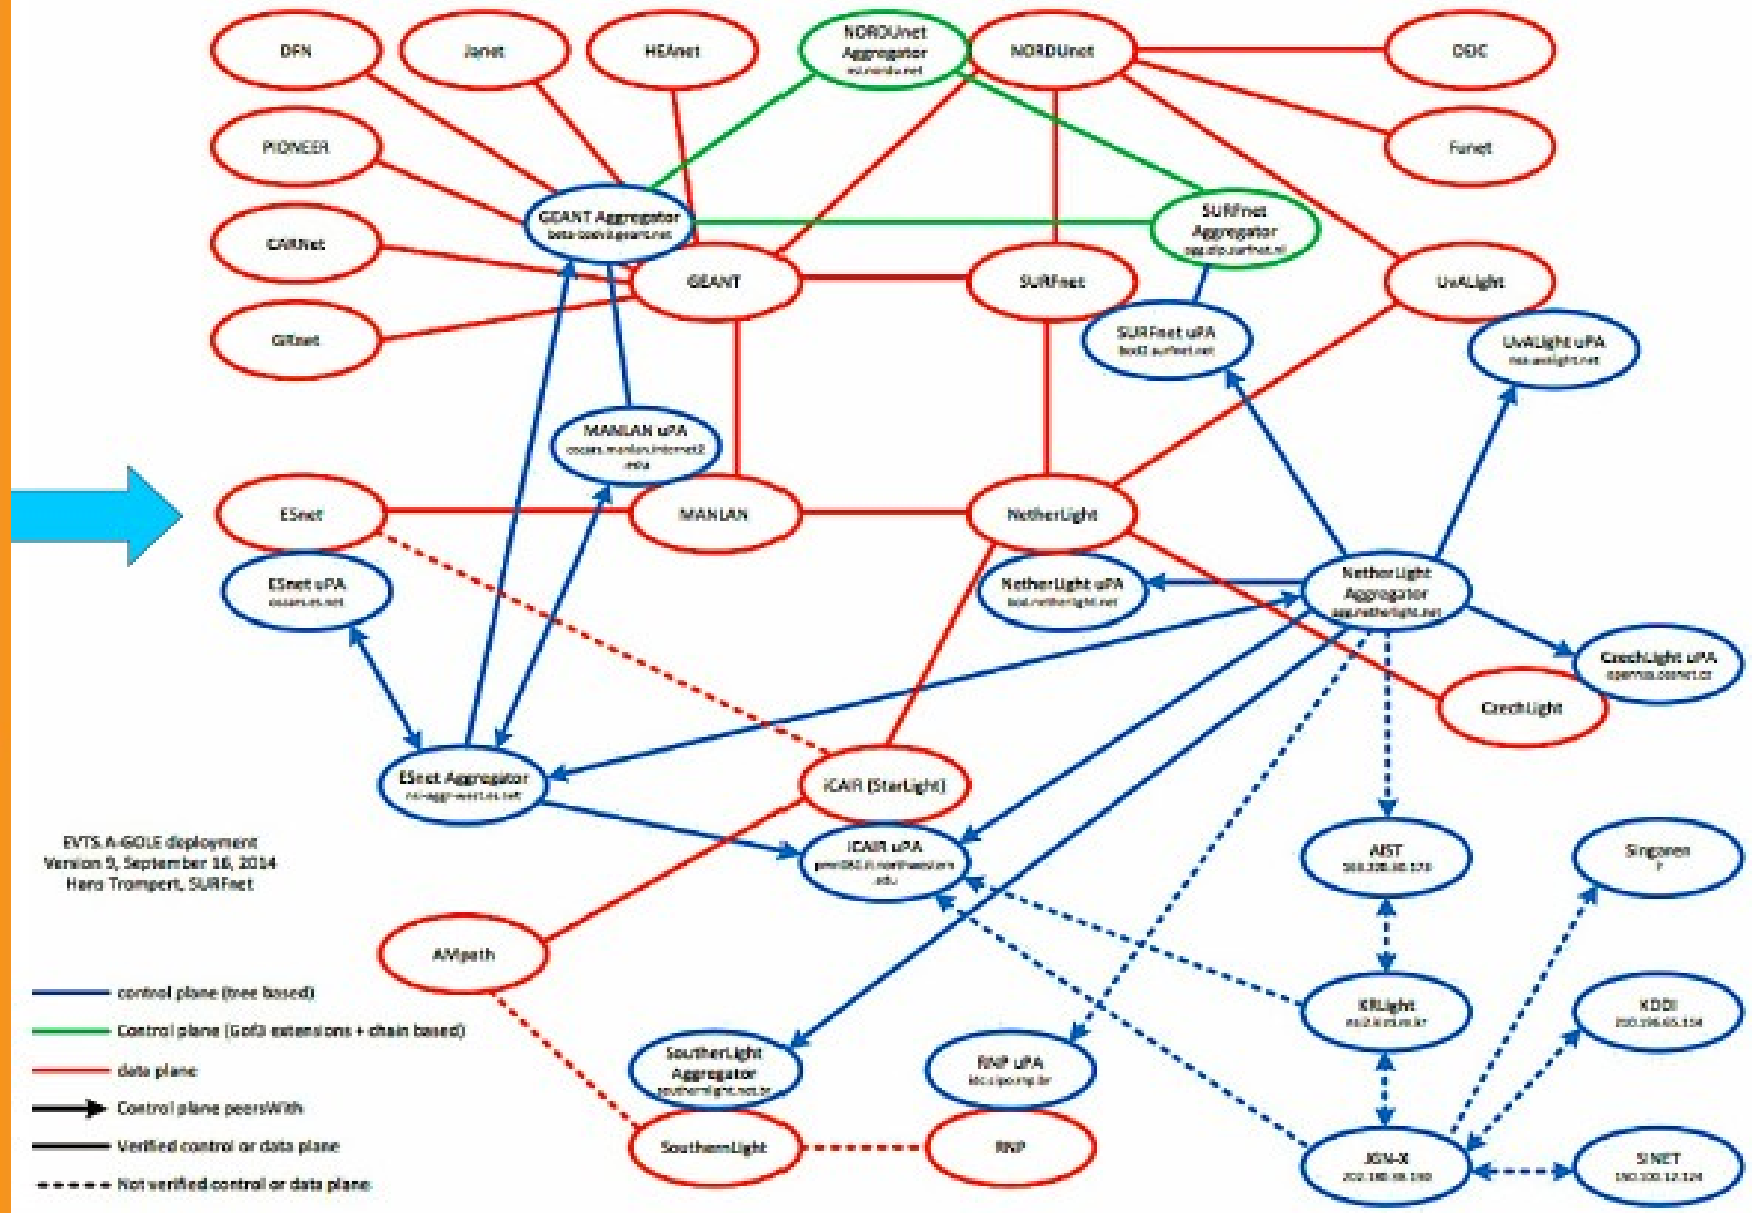
DFN
Janet
HEAnet
NORDUnet Aggregator
NORDUnet
DOE
PIONEER
Funet
GÉANT Aggregator
CARNet
GÉANT
SURFnet
SURFnet Aggregator
UvALight
GRNet
SURFnet uPA
UvALight uPA
MANLAN uPA
ESnet
MANLAN
NetherLight
ESnet uPA
NetherLight uPA
NetherLight Aggregator
CzechLight uPA
CzechLight
ESnet Aggregator
iCAIR (StarLight)
iCAIR uPA
AIST
Singaren
AMPath
KRLight
KDDI
SouthernLight Aggregator
RNP uPA
SouthernLight
RNP
JGN-X
SINET
EVTS.A-GOLE deployment
Version 9, September 16, 2014
Hans Trompert, SURFnet
control plane (tree based)
Control plane (GofS extensions + chain based)
data plane
Control plane peersWith
Verified control or data plane
Not verified control or data plane

AutoGOLE Topology SC'13

Source: John MacAuley

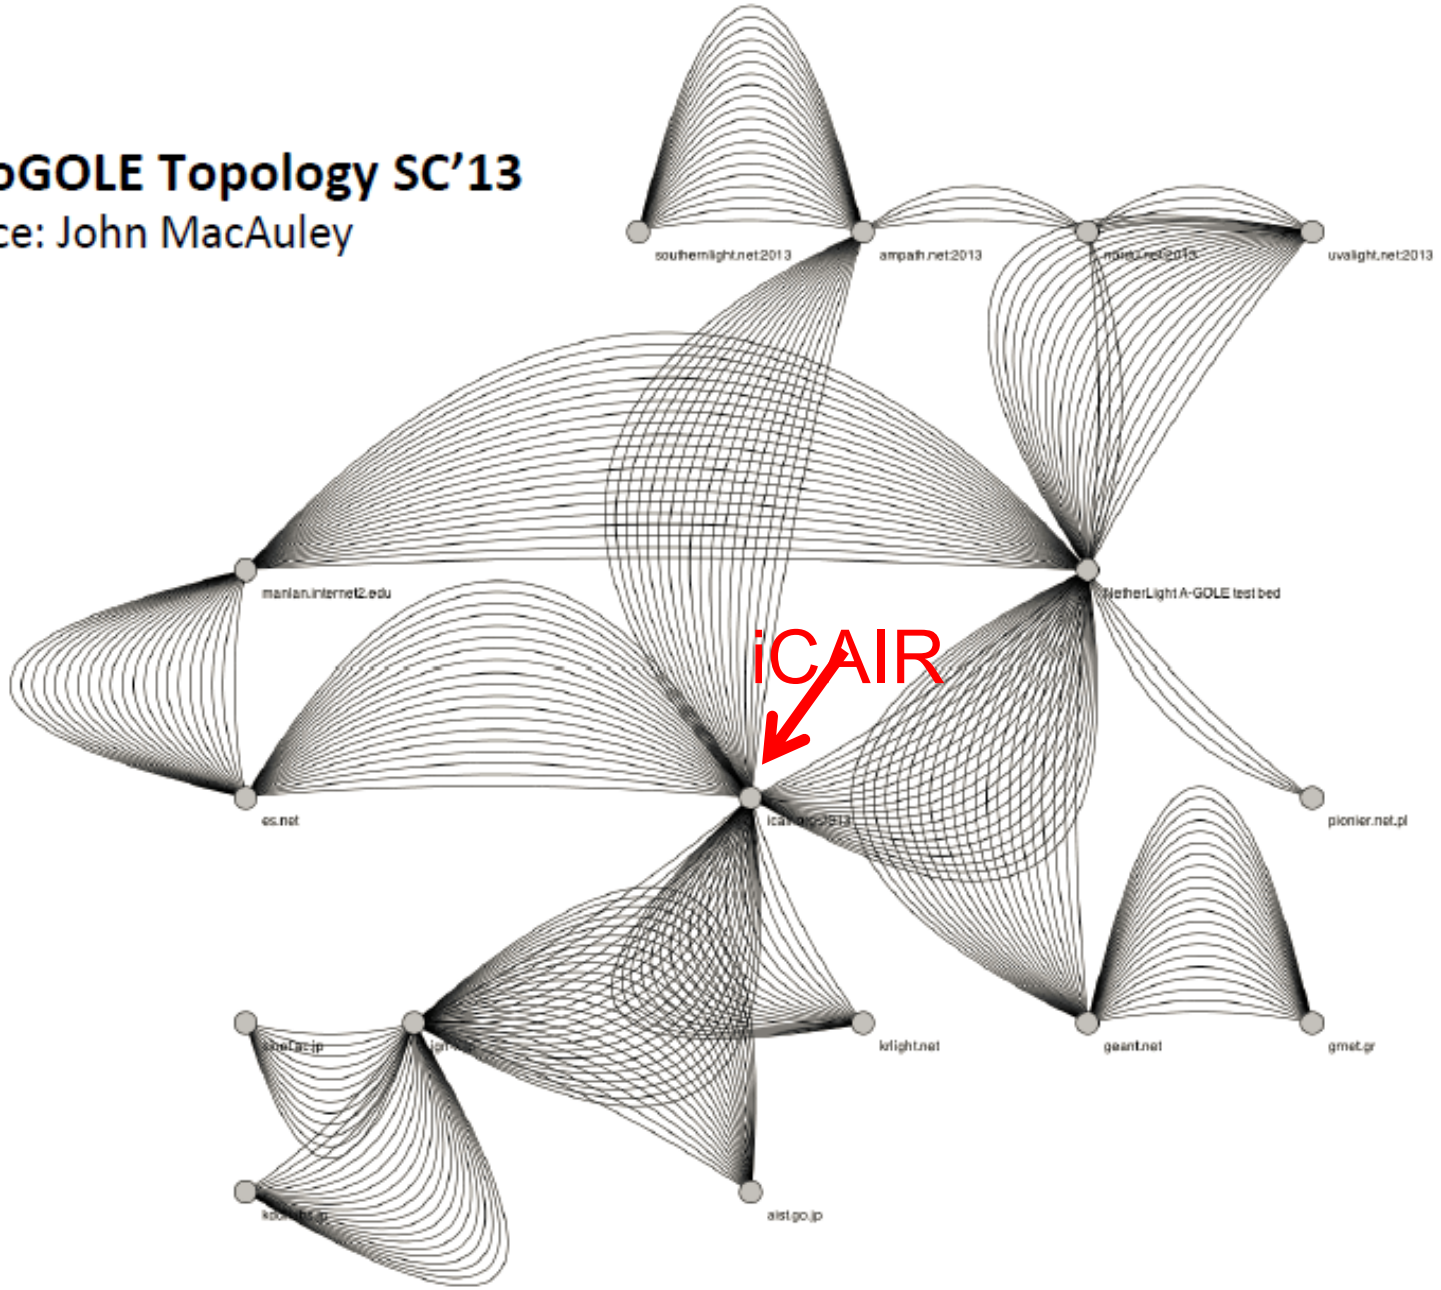

Tasks/Goals For 2014 Expansion To Incorporate SDN/OF

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


KUAS
NetFPGA
NOX
OpenFlow
NCKU
NetFPGA
NOX
OpenFlow
NOX
OpenFlow
NetFPGA
ETRI
NOX
OpenFlow
NetFPGA
KiSTi
TOMATO
OpenFlow
InstaGENI
University of Victoria
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
OpenFlow
AMPATH
NSI
OpenFlow
MAGPI
UEN
UTAH EDUCATION NETWORK
GENI AM
OpenFlow
CPqD
OpenFlow
RNP
NSI
NLLight
NSI
NOX
OpenFlow
NetFPGA
UNIVERSITY OF AMSTERDAM
NSI
ExoGENI
THE UNIVERSITY OF UTAH
VNode
OpenFlow
InstaGENI
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

StarLight – "By Researchers For Researchers"

StarLight is an experimental optically based infrastructure and proving ground for network services optimized for high-performance 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

Currently Undertaking Major Infrastruction Renovation

StarLight SDX (iSDX)

- At Highest Level, Appears As a Very Large Scale Virtual Switch
- Resources Can Be Segmented/Partitioned
- Architecture Is Informed By NSI, GENI, Related International Network Testbeds, Emerging Concepts
- Based On An Underlying Foundation of Programmable Resources
- Includes Specialized APIs for Provisioning
- This SDX is Being Showcased Through GLIF, GENI and SC14 Demonstrations

SDX As Recursive Virtual Switch

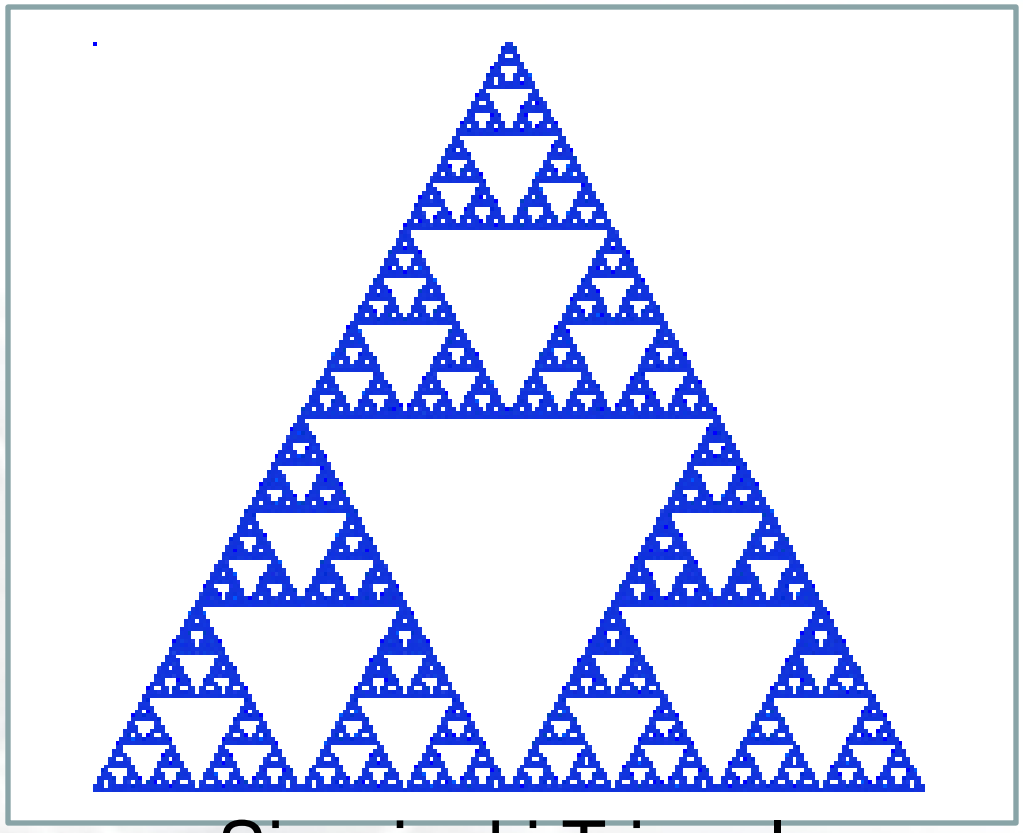

Sierpinski Triangle

Unlimited Number of Customized Virtual Switches Within Macro Virtual Switch

GLIF Based On SDXs Supporting Slice Exchanges

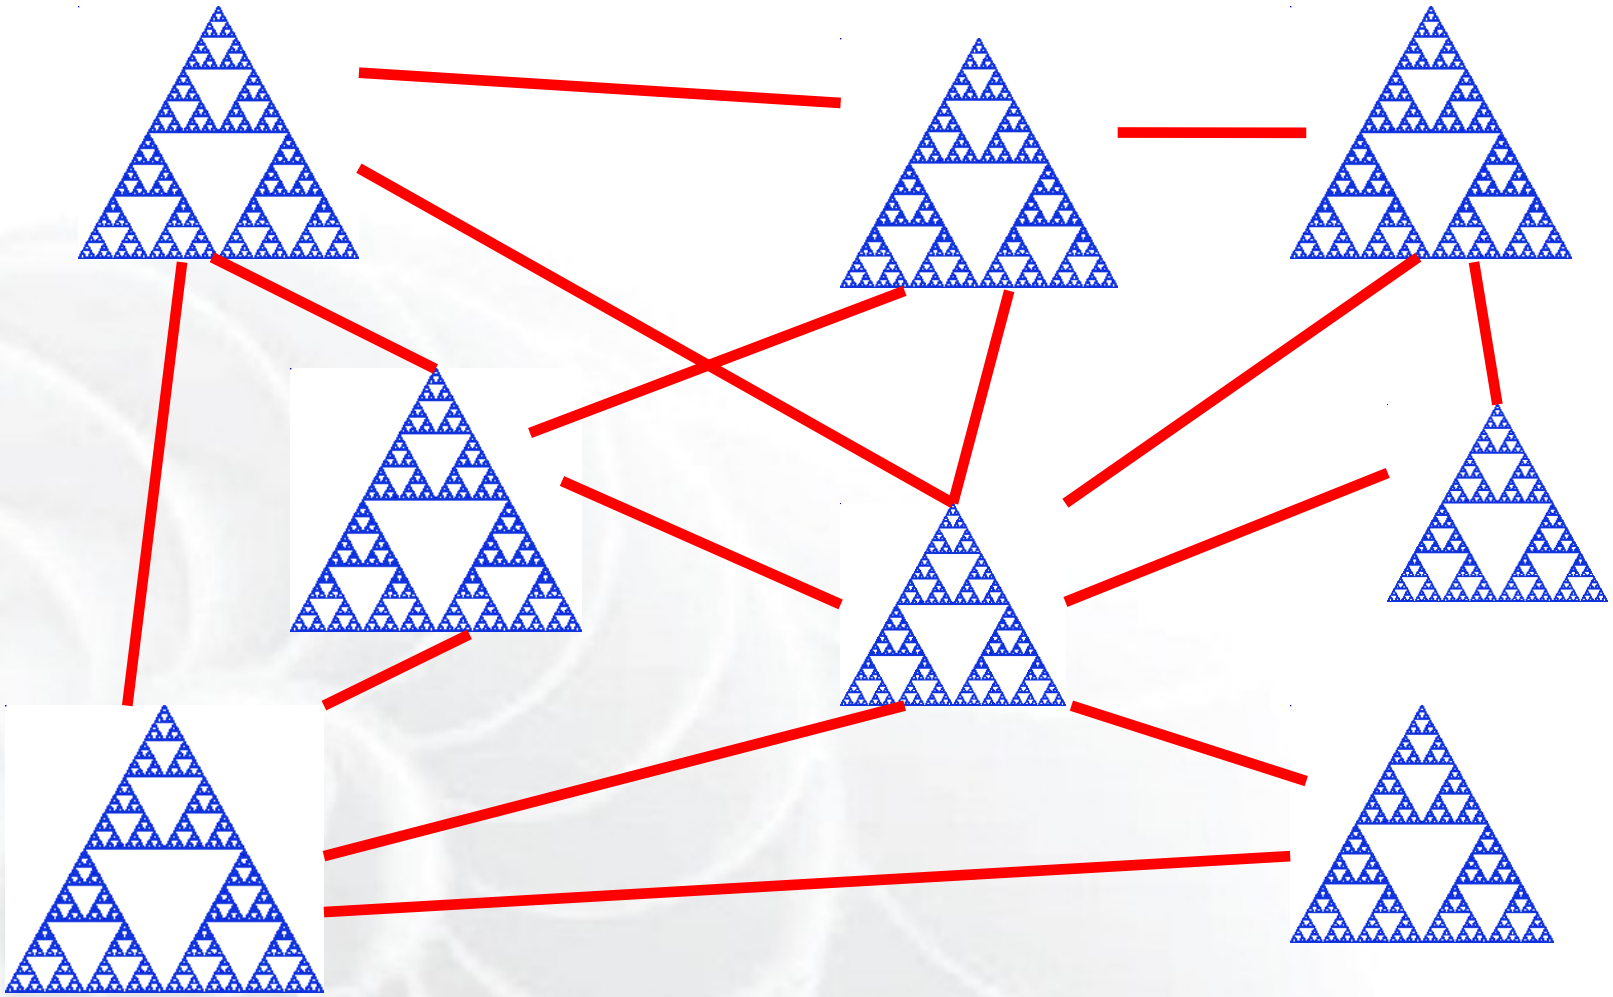

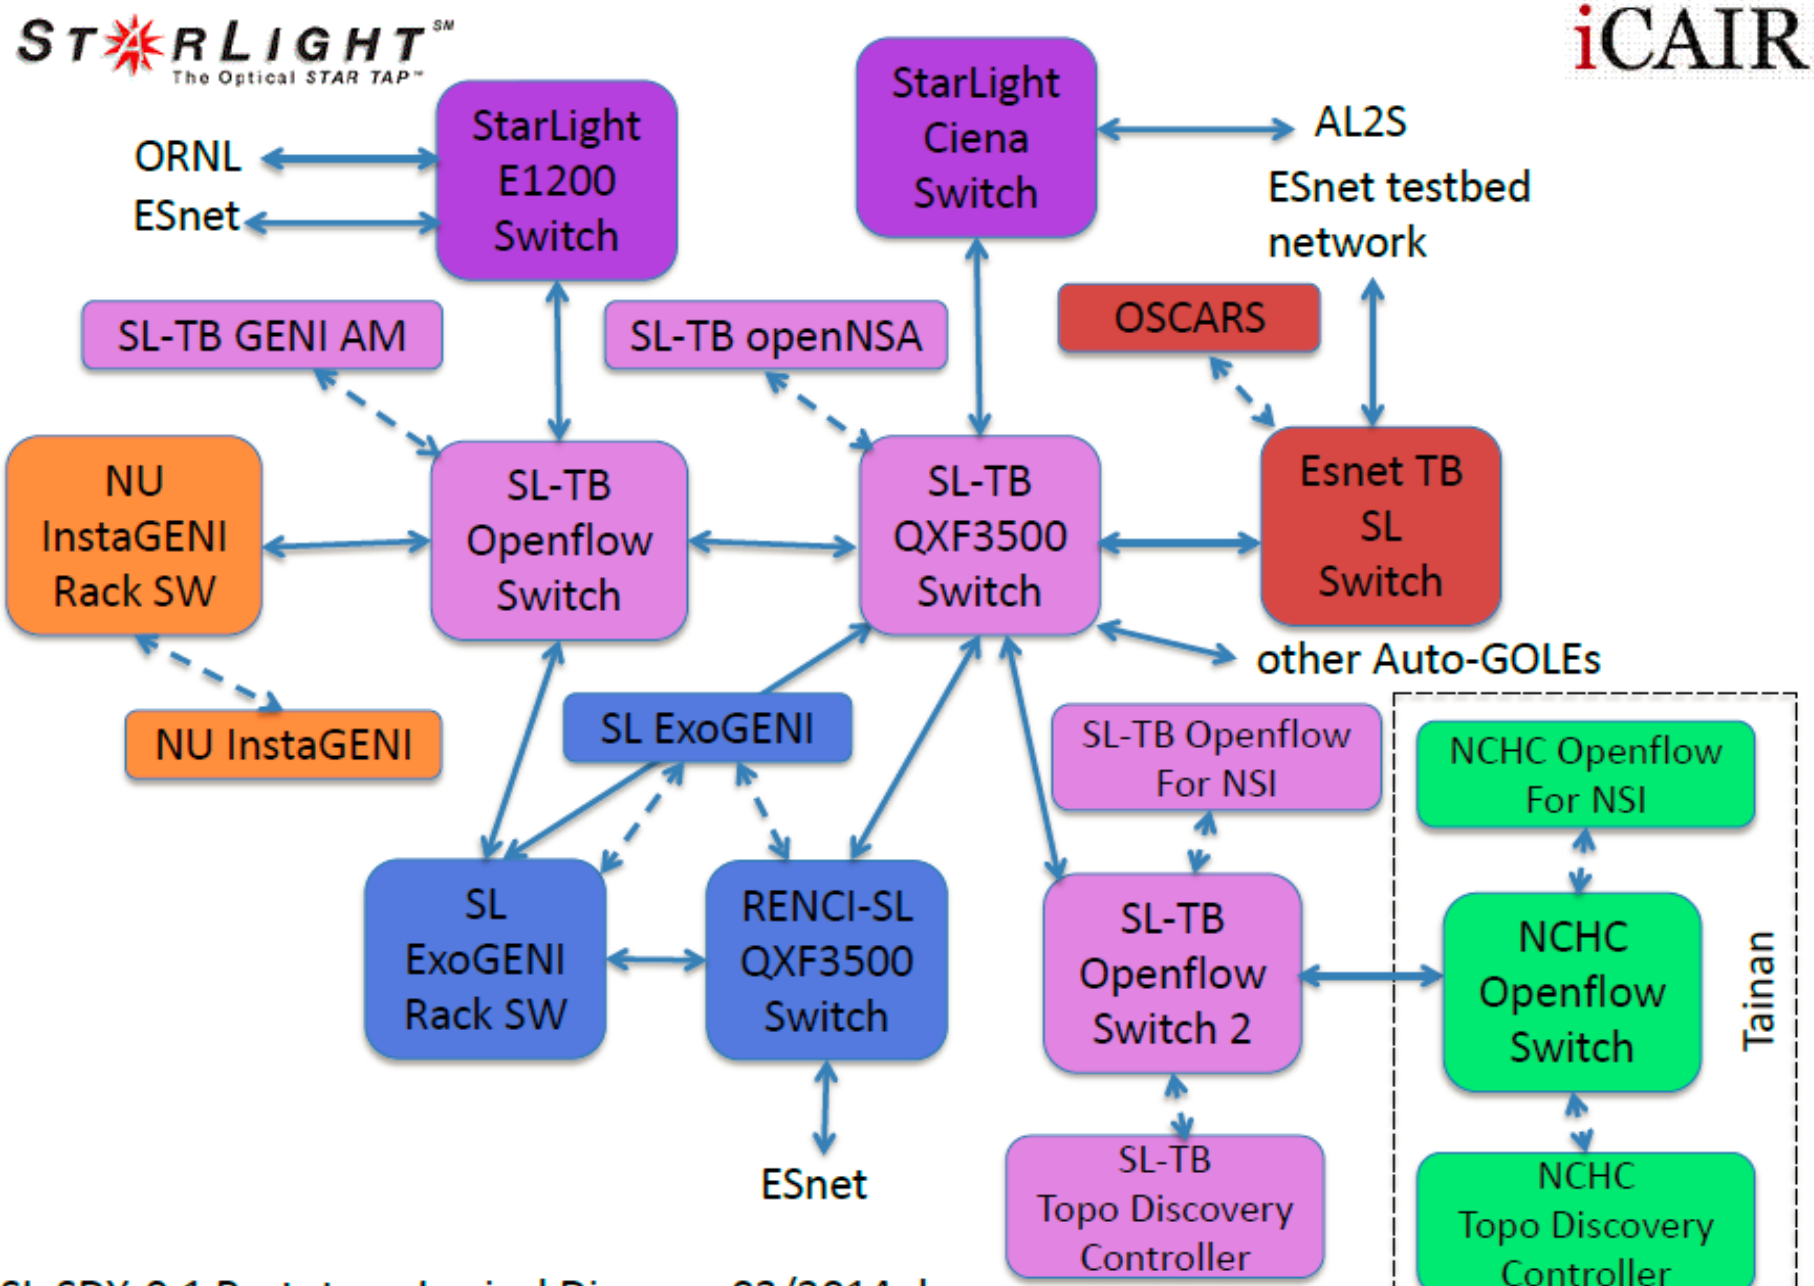

SL-SDX-0.1 Prototype Logical Diagram 03/2014 demo

Nowcast – Processing

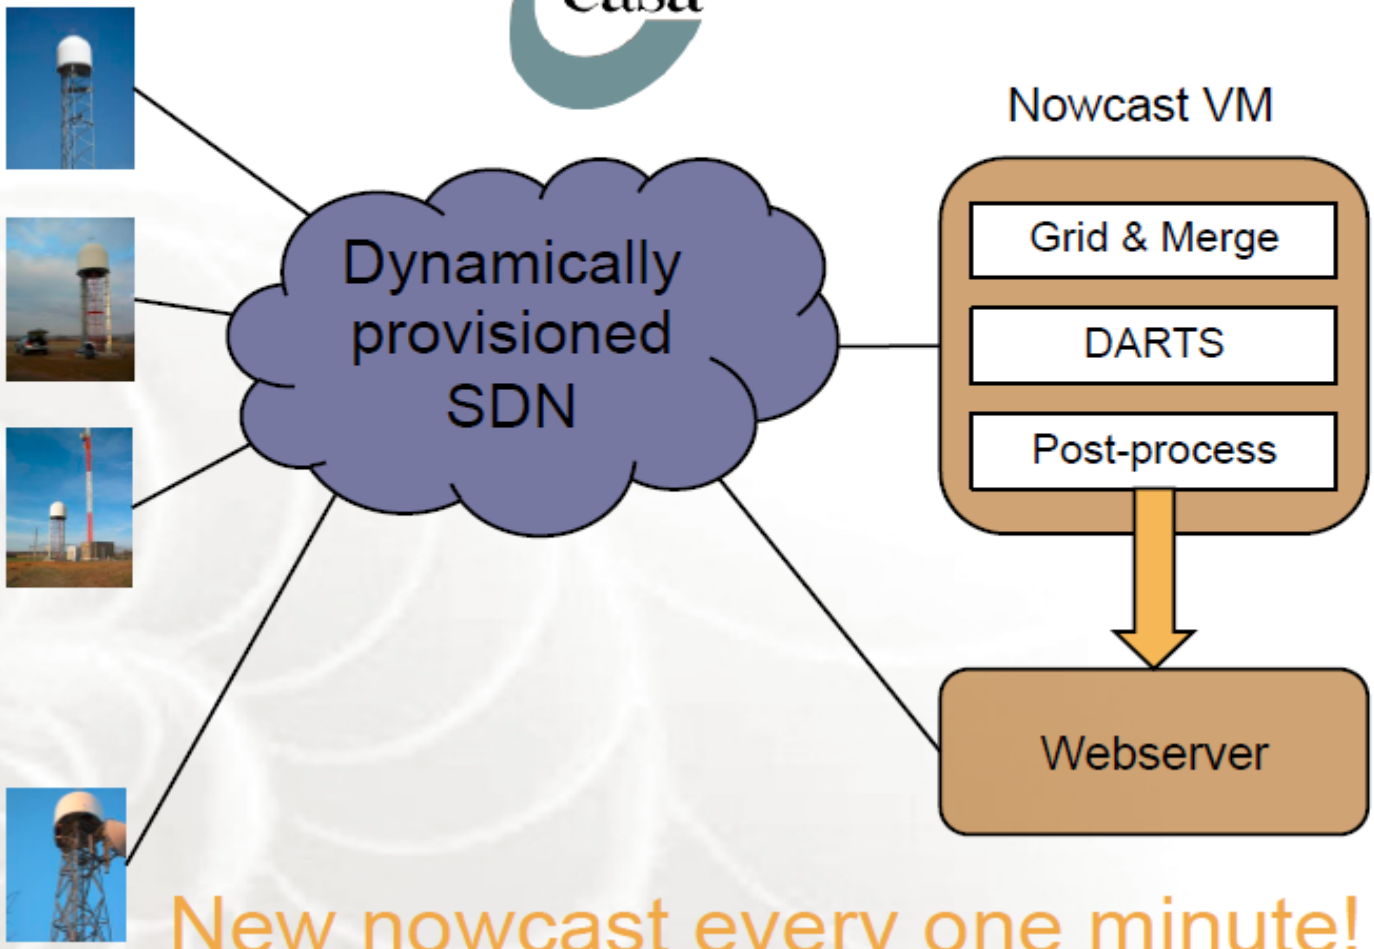

Source: Mike Zink

Comparison With Existing System

Slide by Mike Zink, UMass Amherst

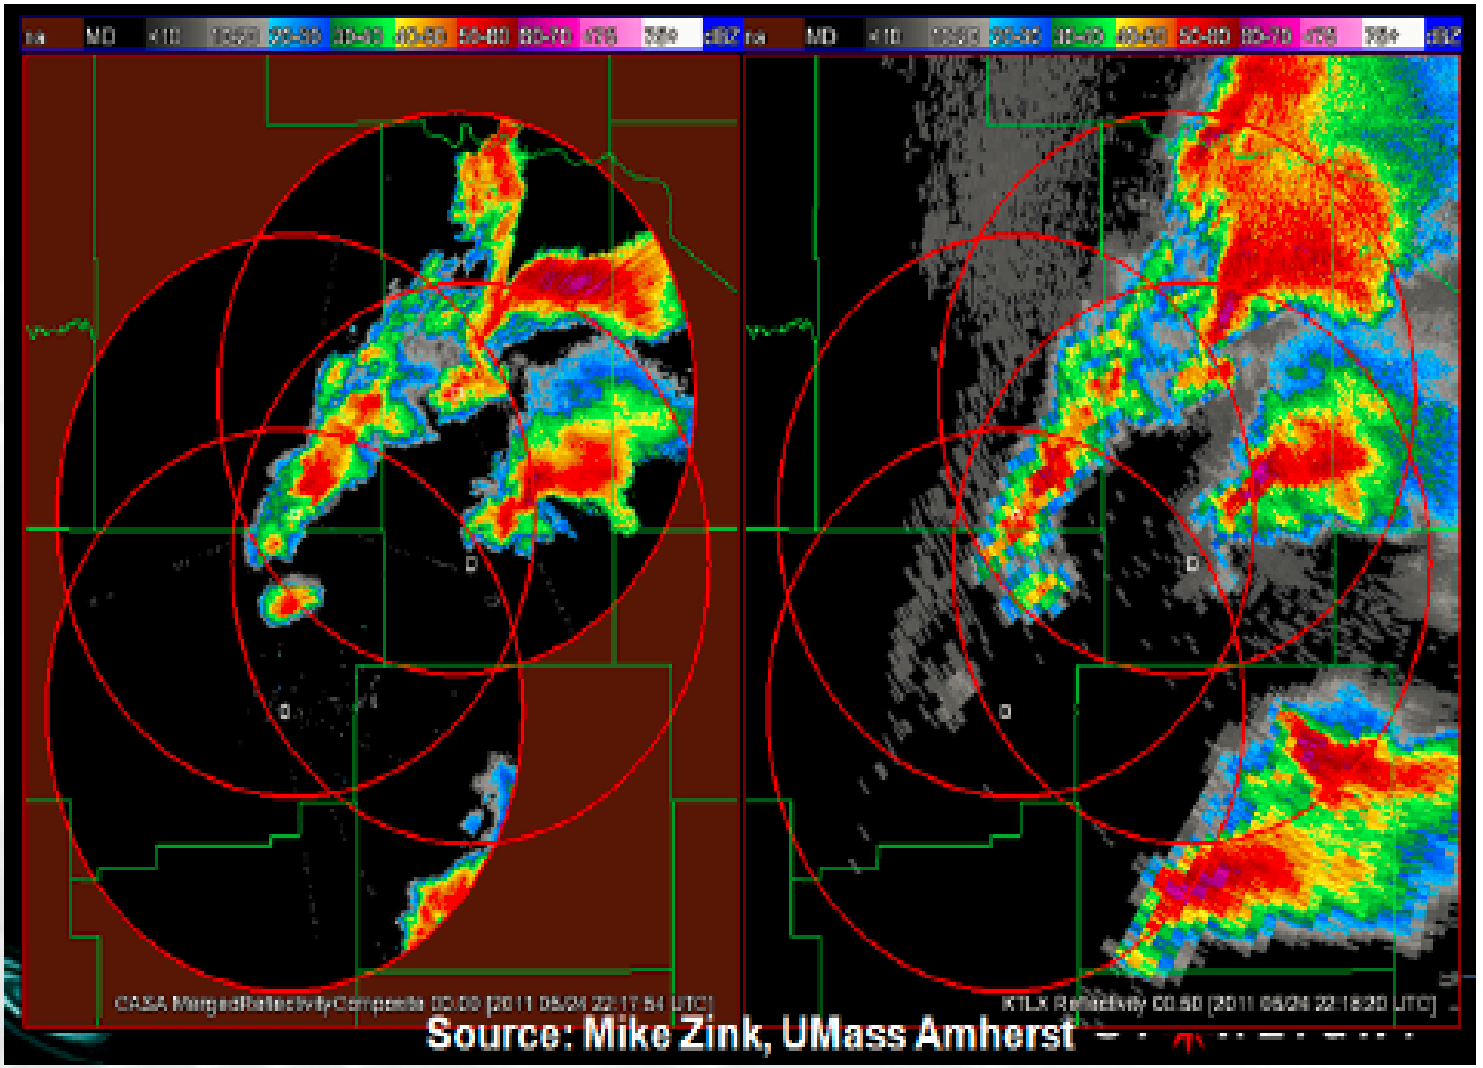

GENI SDX Demo Scenario 2

Simulated Radar (4)

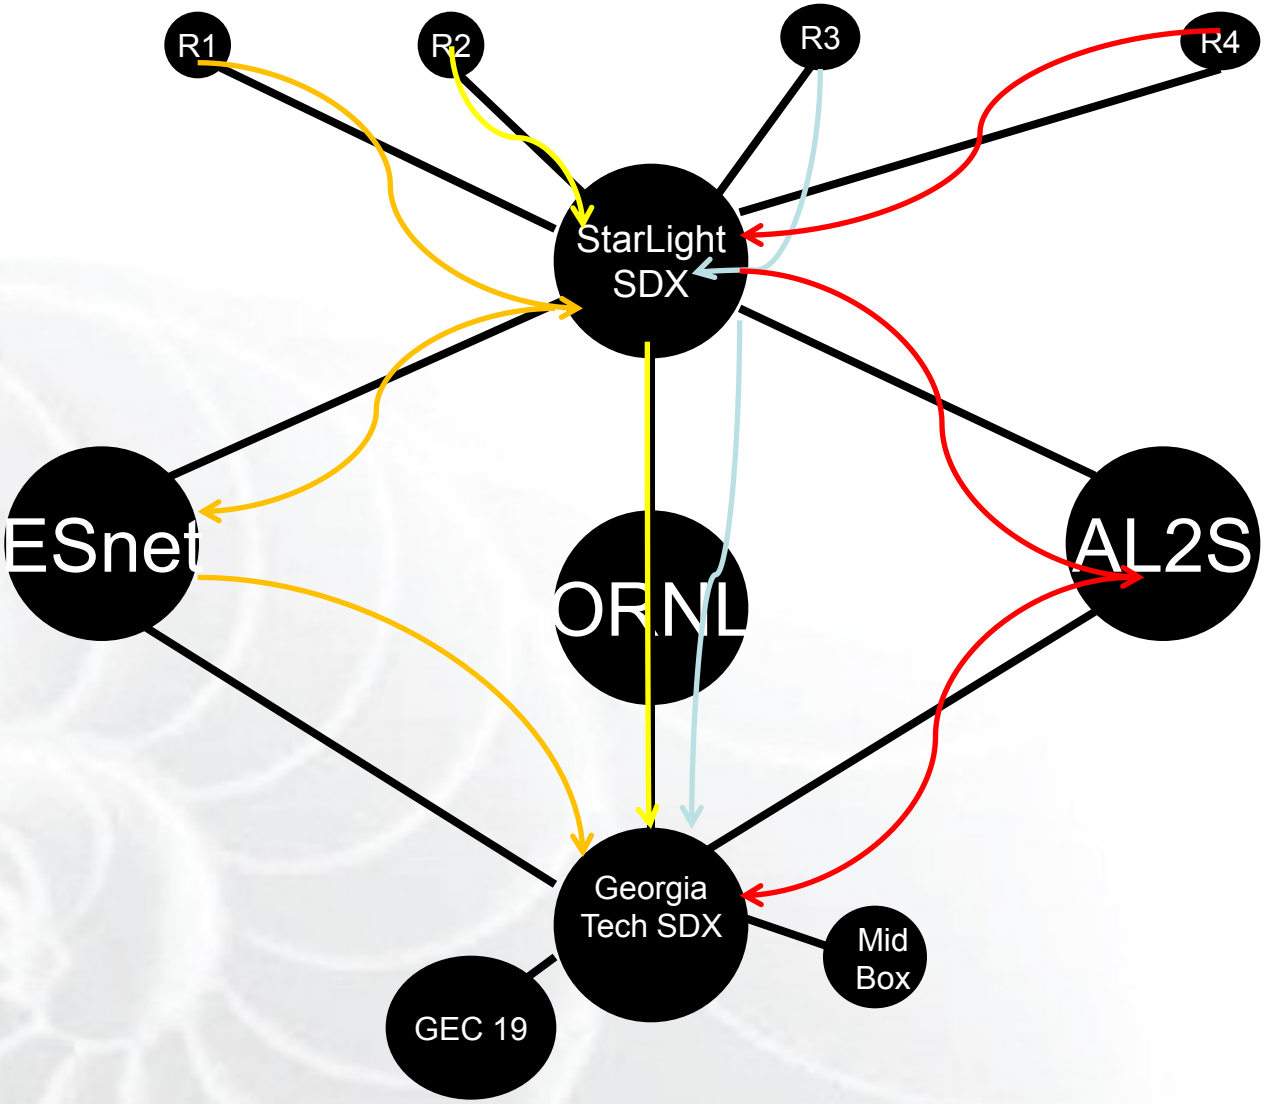

SDX StarLight⇔NetherLight

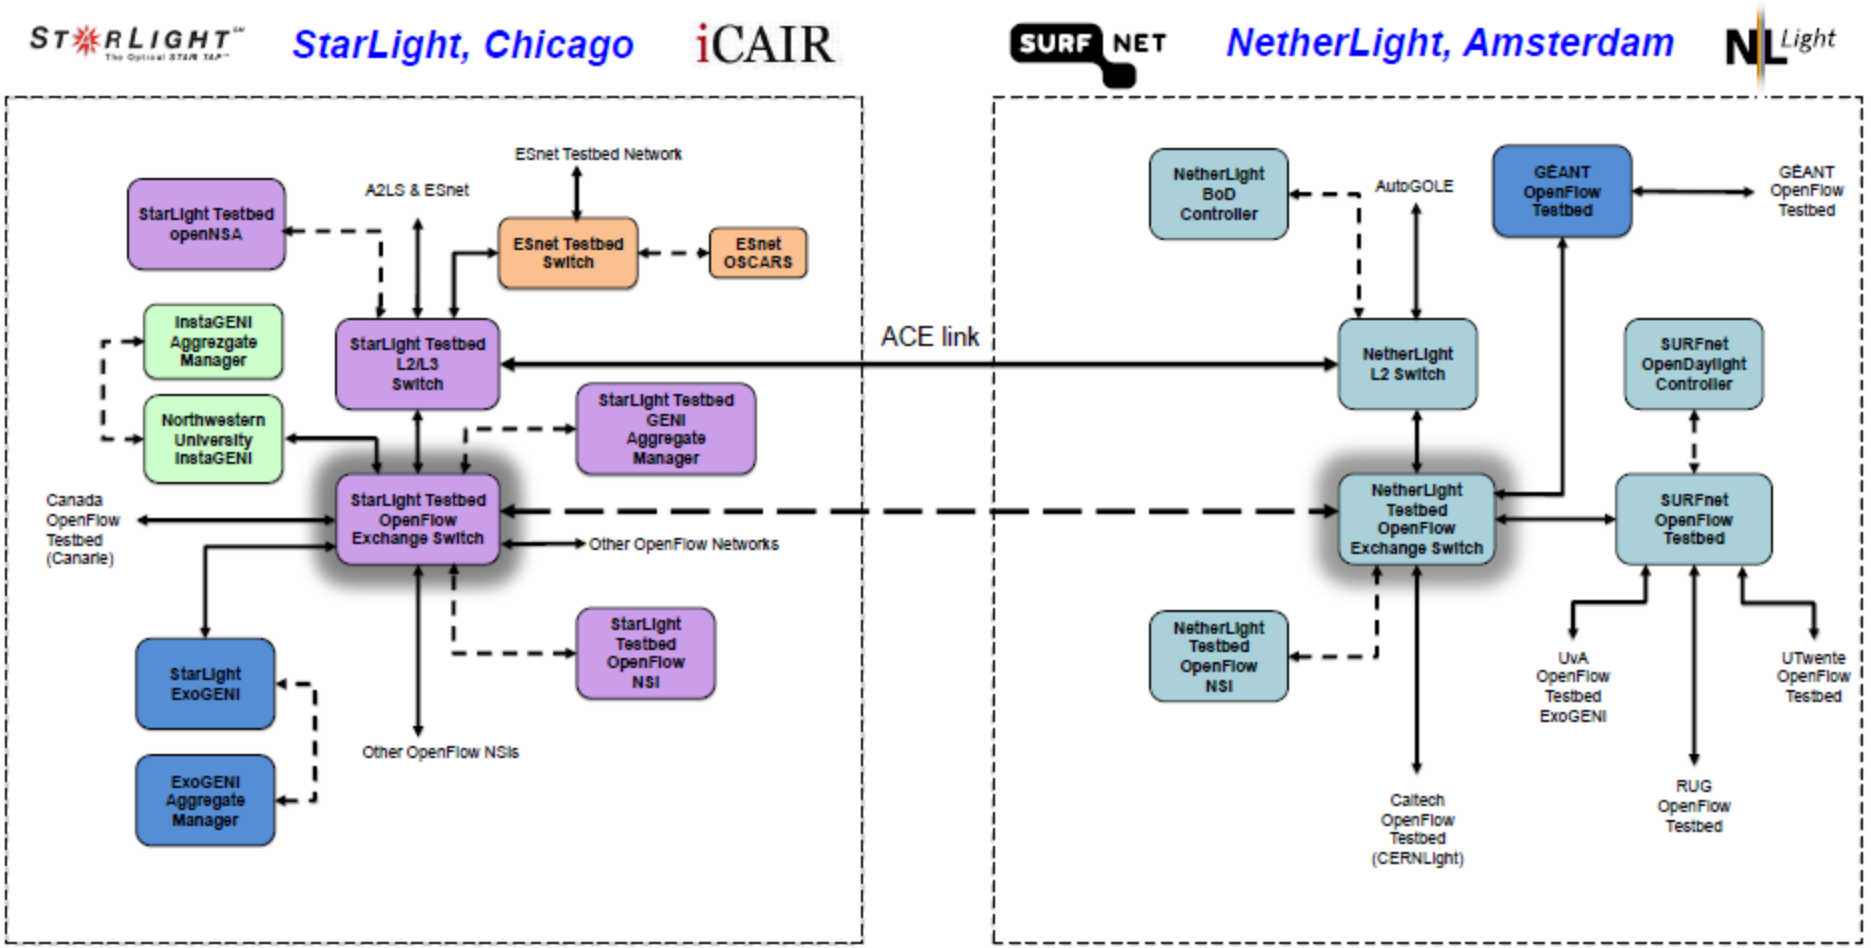

Ronald van der Pol, Joe Mambretti, Jim Chen, John Shillington

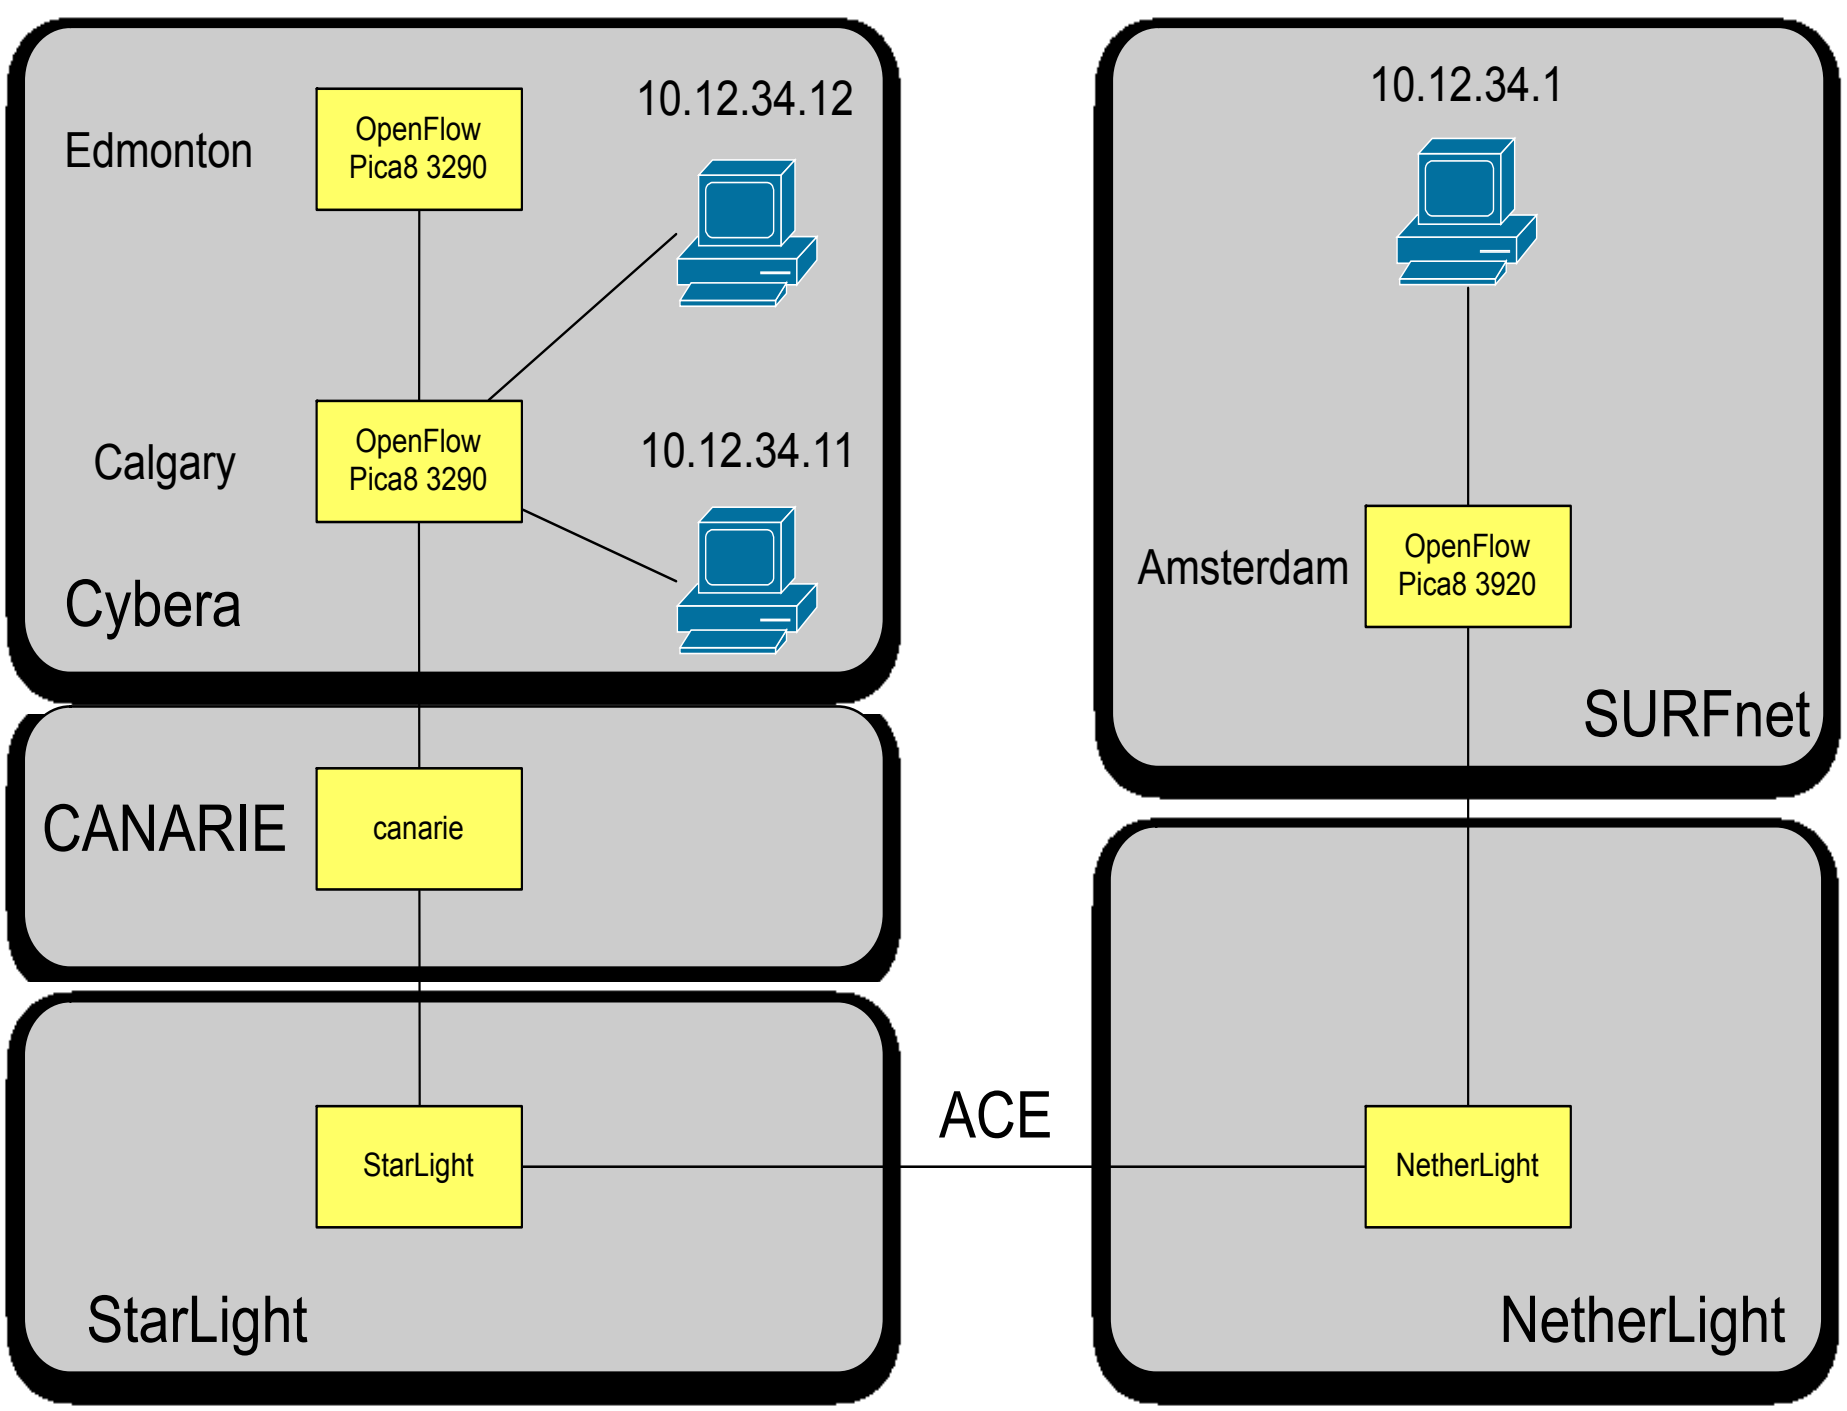
Edmonton
OpenFlow
Pica8 3290
10.12.34.12
Calgary
OpenFlow
Pica8 3290
10.12.34.11
Cybera
CANARIE
canarie
StarLight
StarLight
ACE
10.12.34.1
Amsterdam
OpenFlow
Pica8 3920
SURFnet
NetherLight
NetherLight

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

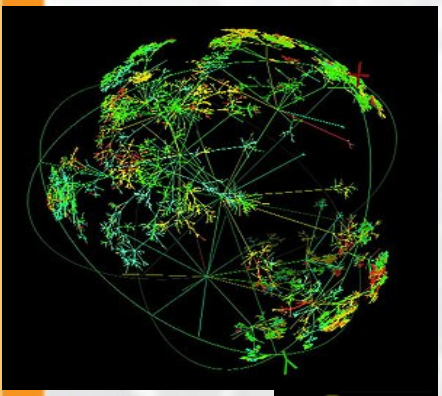

Global LambdaGrid Workshop
Queenstown, New Zealand
September 30-October 1, 2014

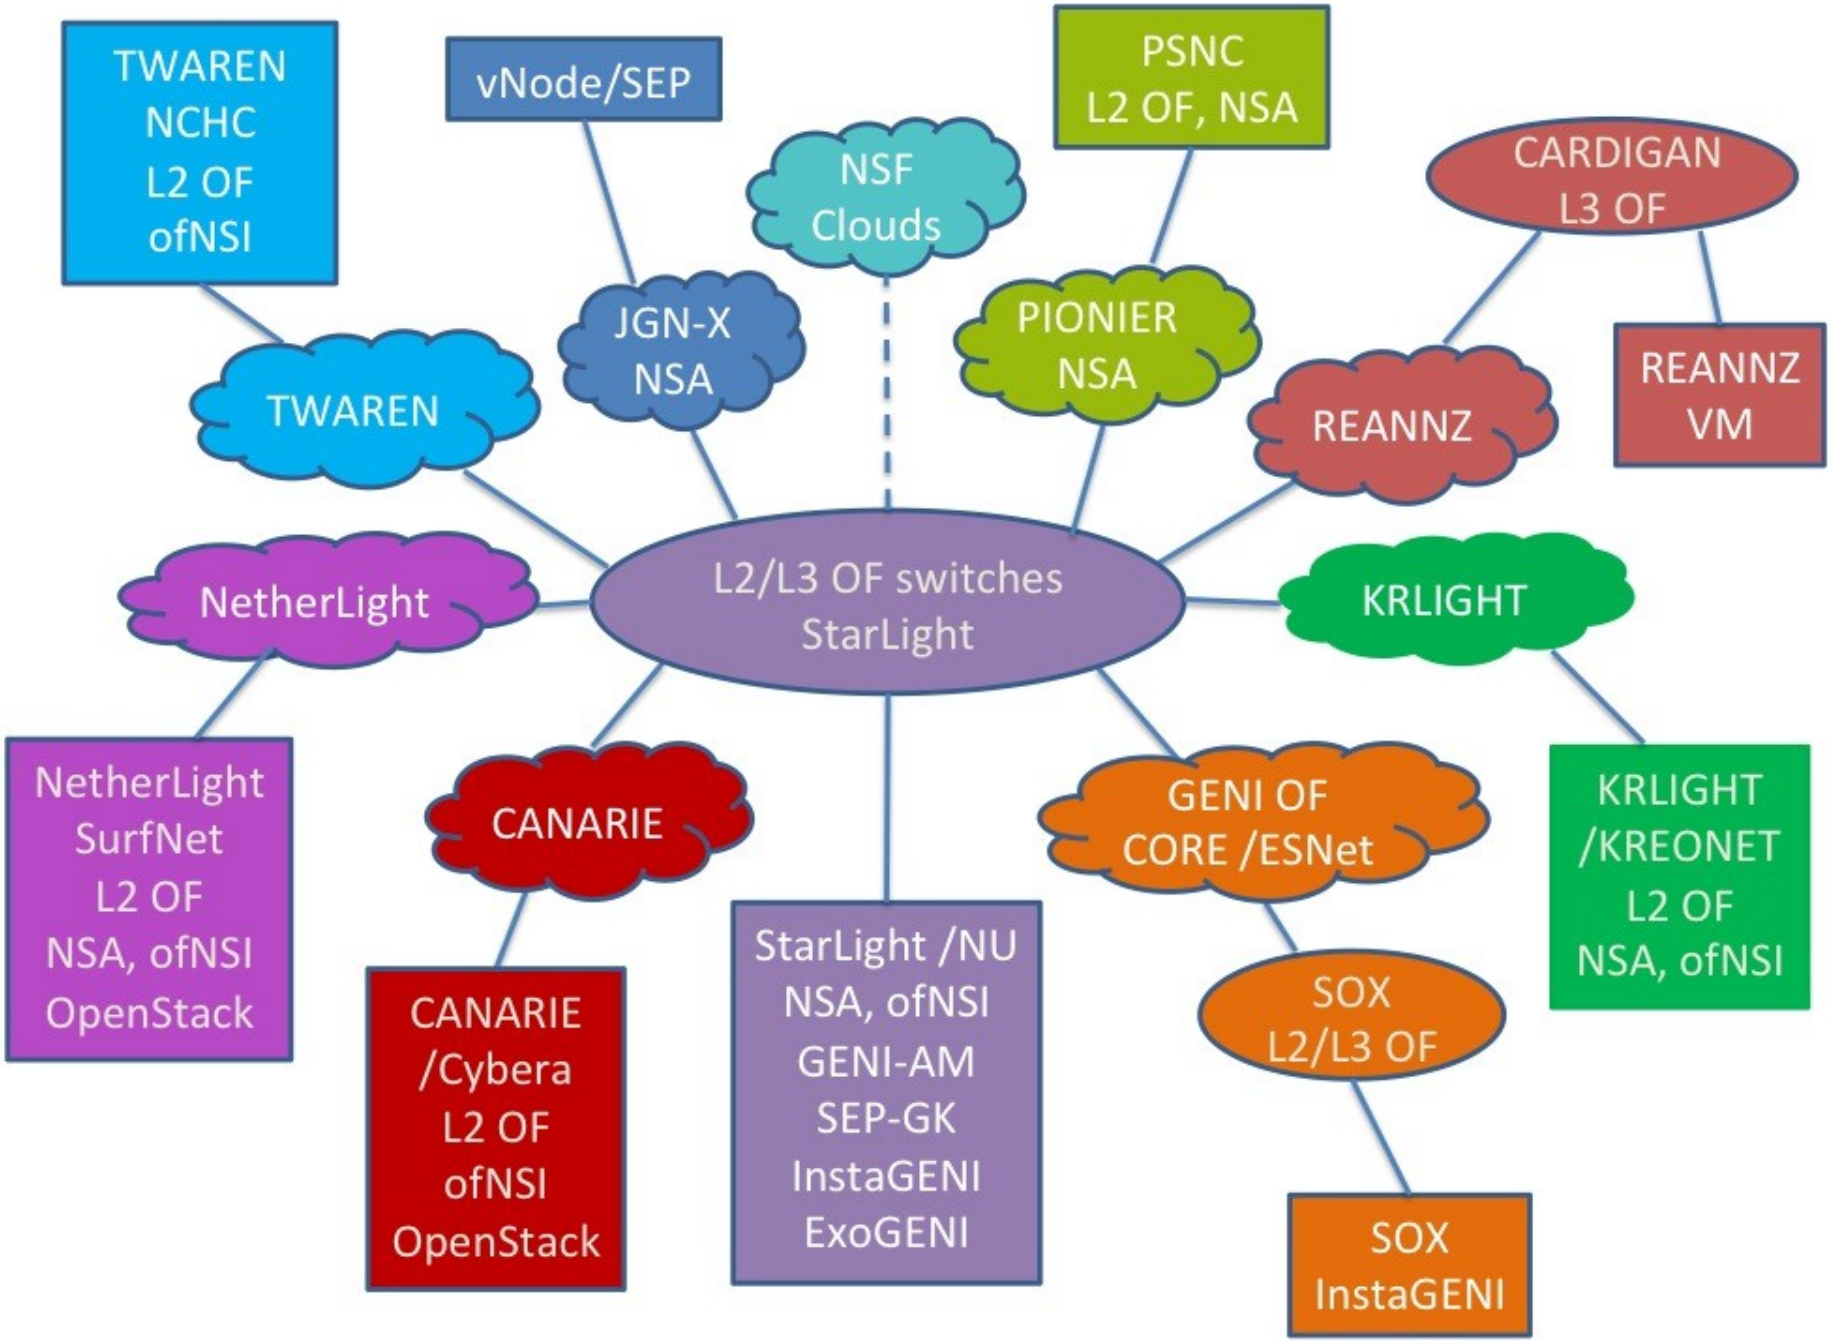
TWAREN
NCHC
L2 OF
ofNSI
vNode/SEP
NSF
Clouds
PSNC
L2 OF, NSA
CARDIGAN
L3 OF
JGN-X
NSA
PIONIER
NSA
TWAREN
REANNZ
REANNZ
VM
NetherLight
L2/L3 OF switches
StarLight
KRLIGHT
NetherLight
SurfNet
L2 OF
NSA, ofNSI
OpenStack
CANARIE
GENI OF
CORE /ESNet
KRLIGHT
/KREONET
L2 OF
NSA, ofNSI
CANARIE
/Cybera
L2 OF
ofNSI
OpenStack
StarLight /NU
NSA, ofNSI
GENI-AM
SEP-GK
InstaGENI
ExoGENI
SOX
L2/L3 OF
SOX
InstaGENI


Nowcast - SDX Demo
iCNR
g1t

Slice Exchange Showcase at GEC 21: Special Thanks To Aki Nakao

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

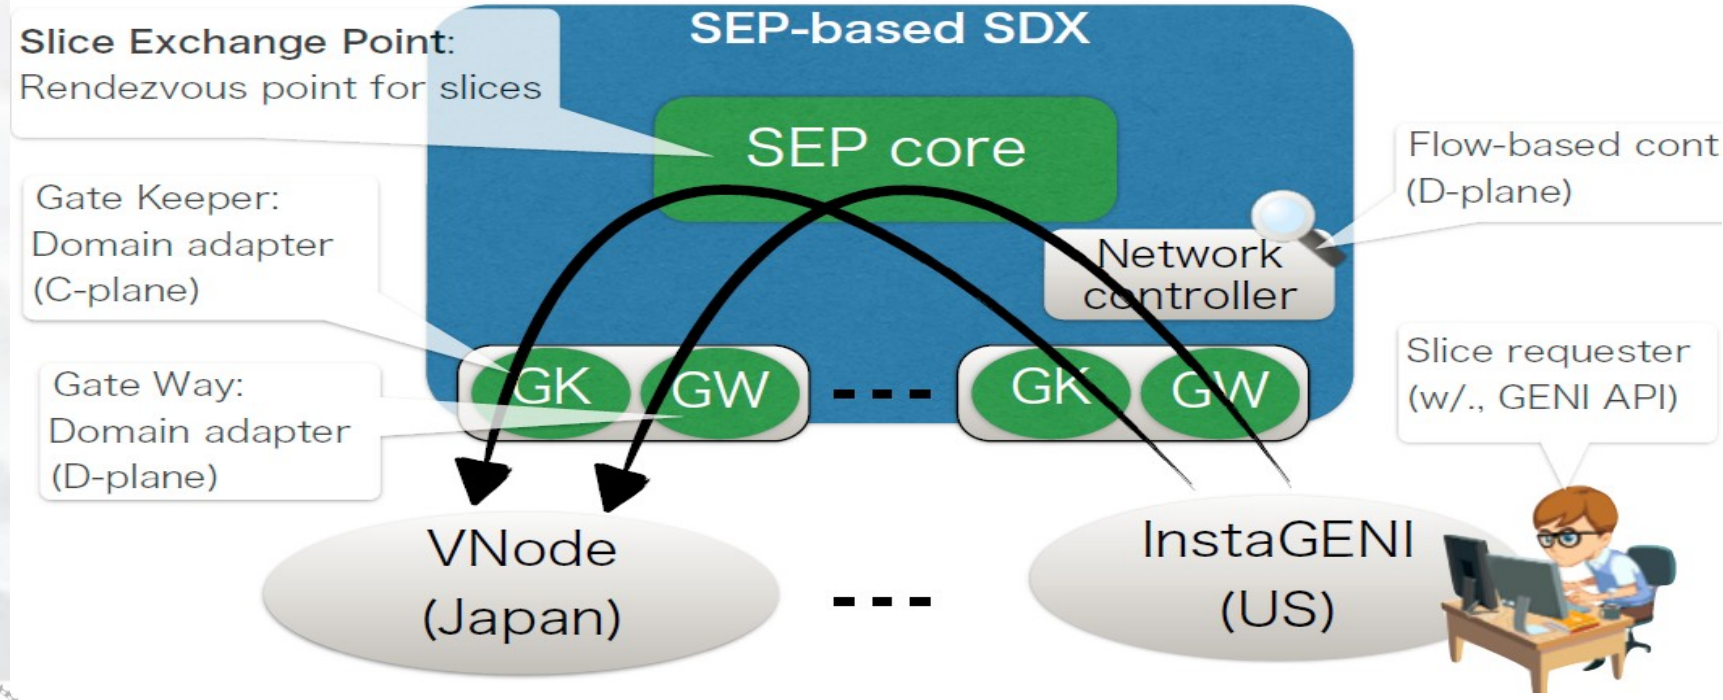

Inter-SDX federation for GEC21

Multi-architecture Federation

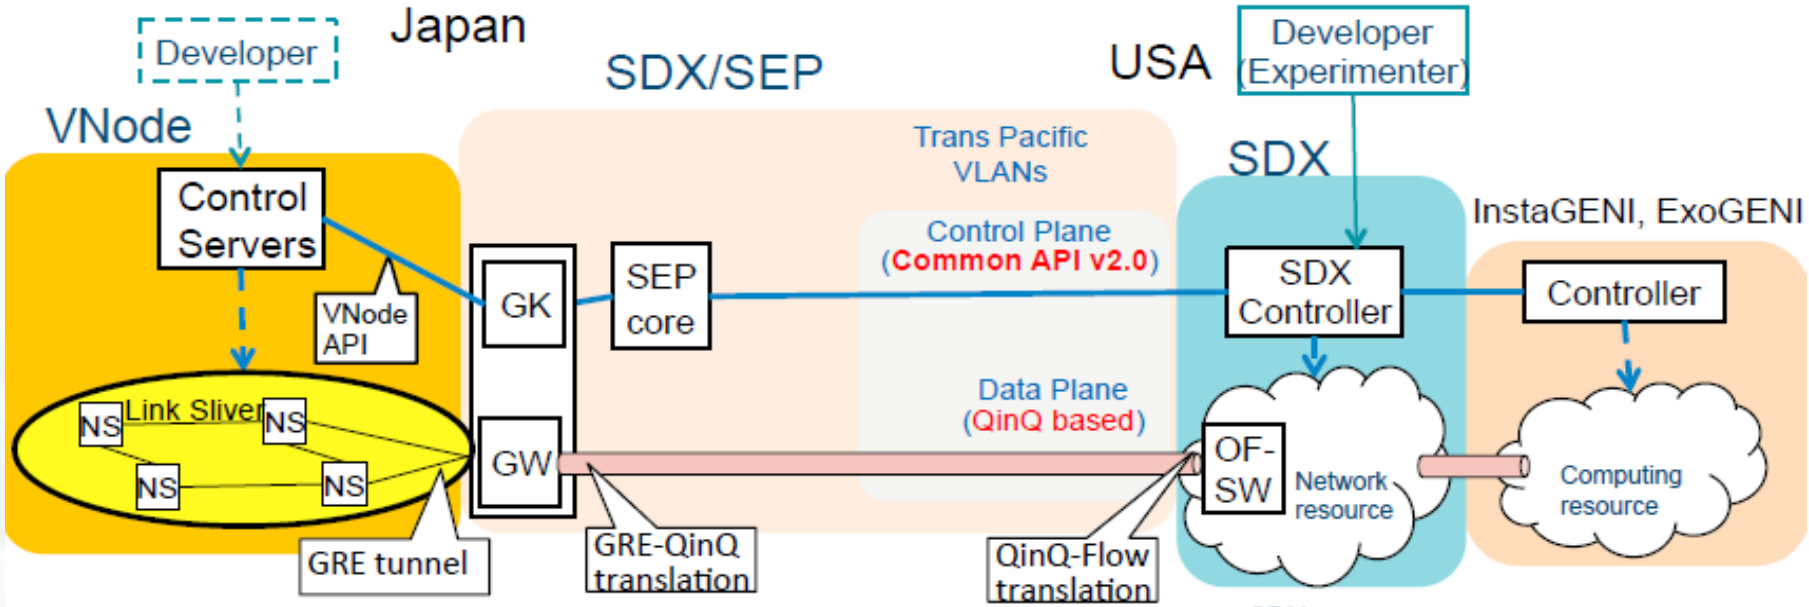

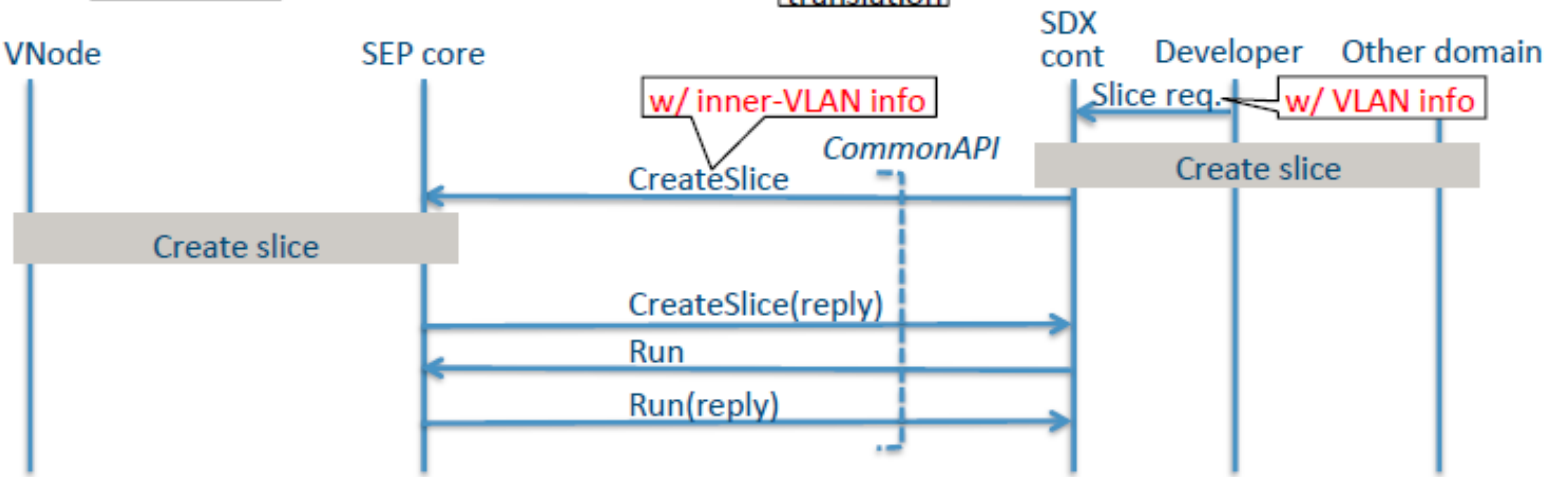

Cross-Domain Information Exchange Through NSI

- **Modified LLDP** (also **NDP** and **BDDP**) Was Used To Exchange Neighbor Controller Information Among SDN Domains Automatically
  - However, These LLDP-like Packets May Be Filtered When Traversing Domain Boundaries
- For That Case, We Utilize **NSI** To Share Information Among NSAs
  - A New **sendmsg** Service Is Implemented
  - In Addition To Topology Discovery, Service Announcements Could Also Be Delivered

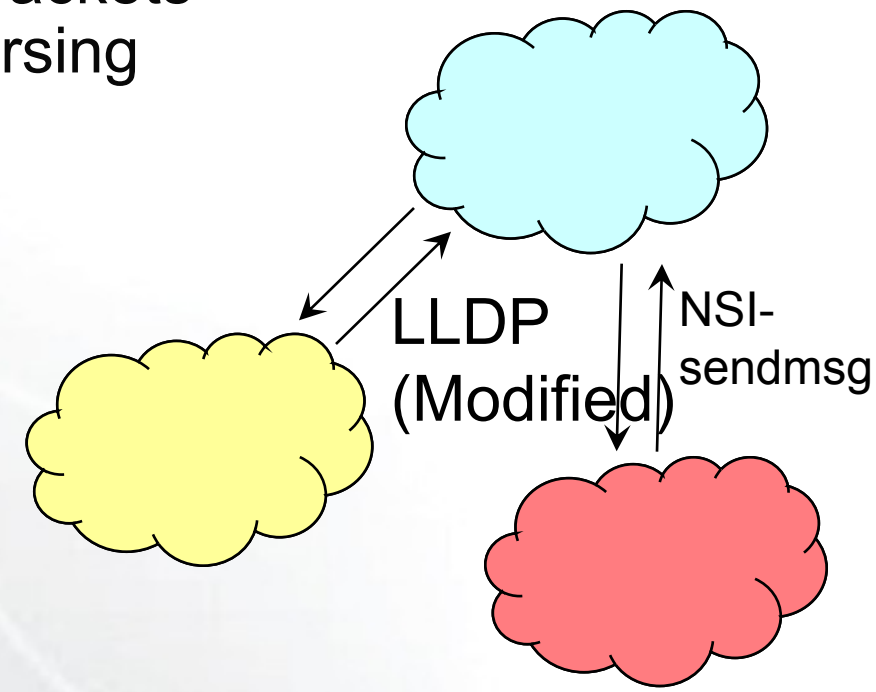

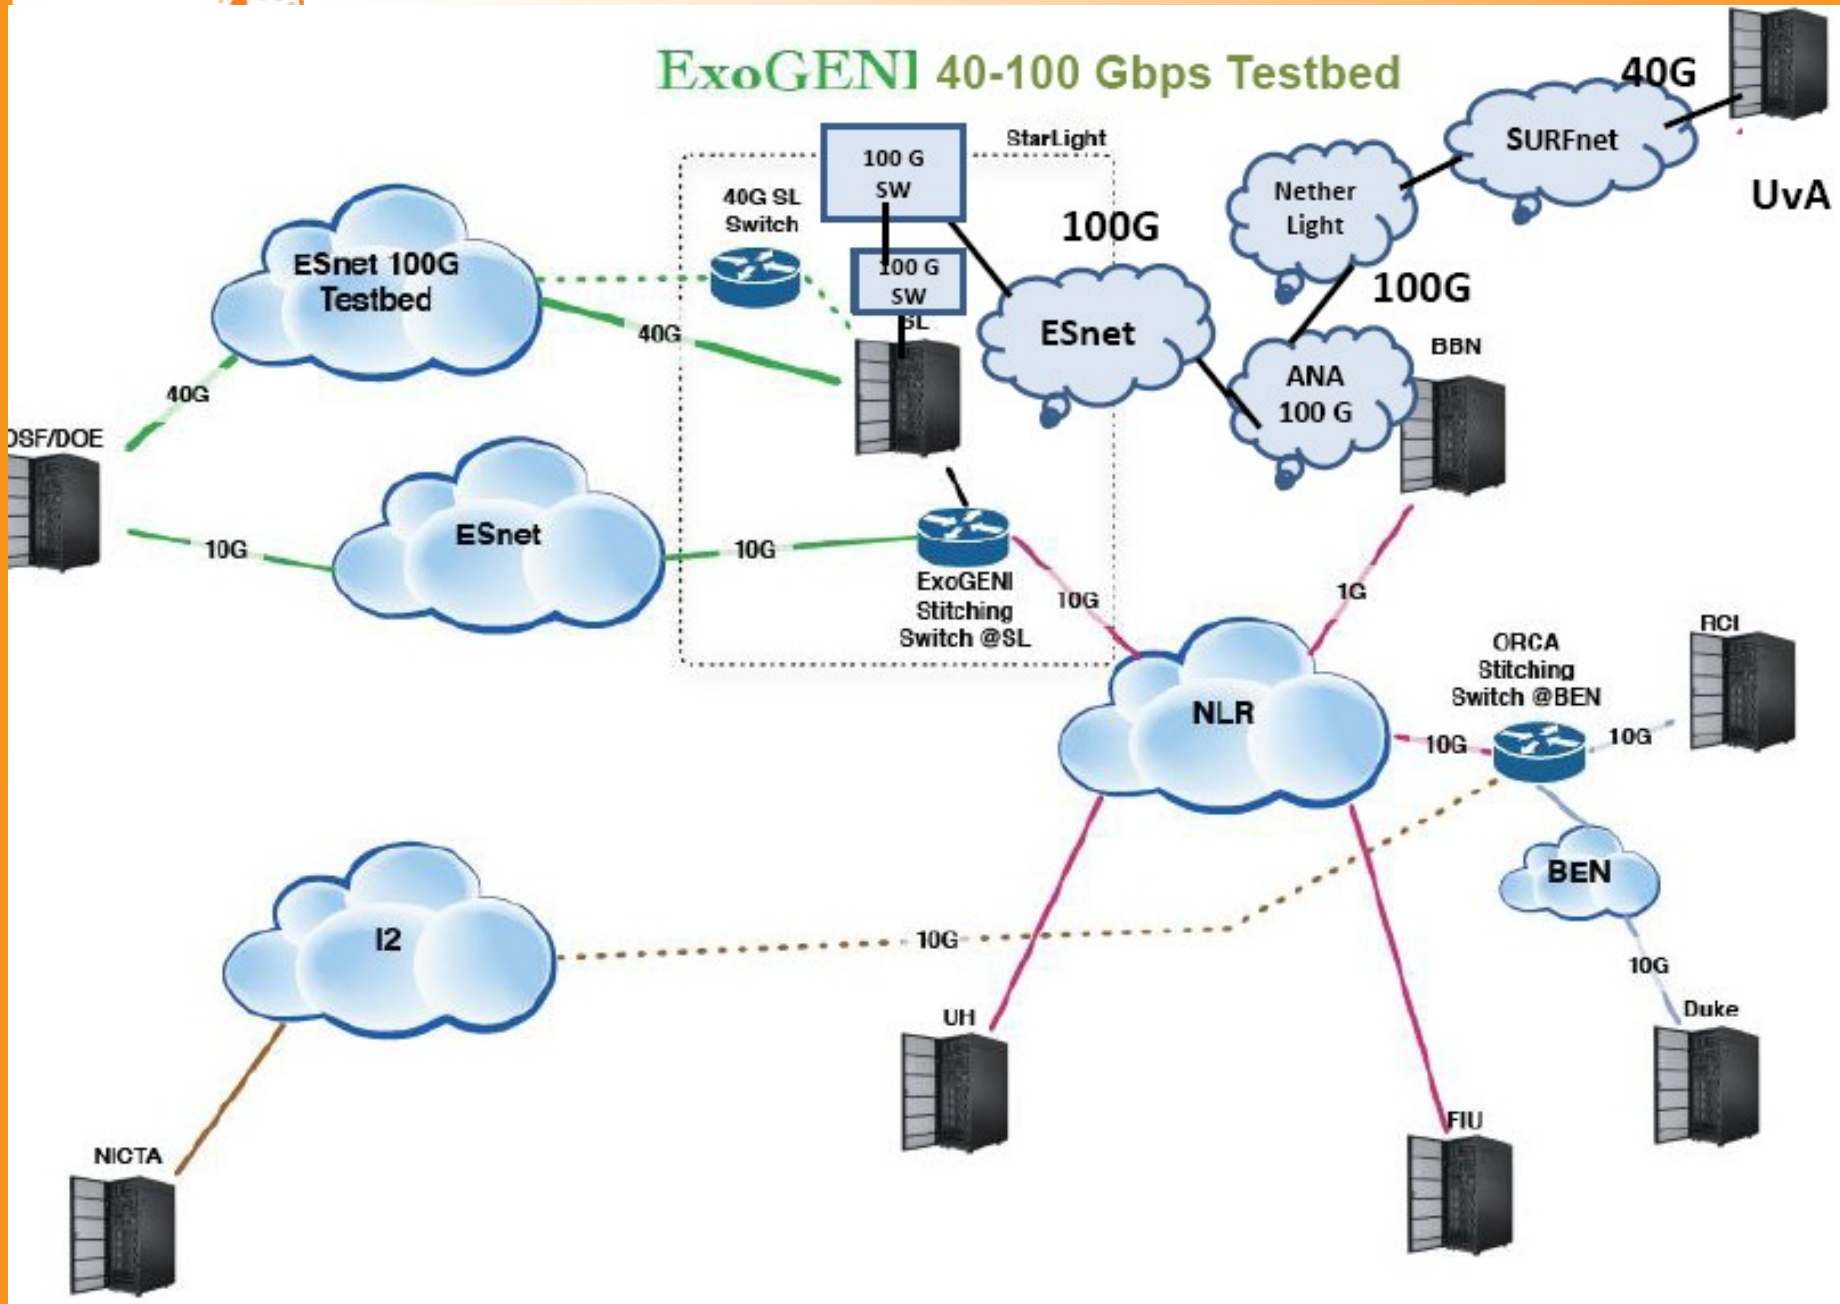
ExoGENI 40-100 Gbps Testbed
40G
SURFnet
UvA
StarLight
100 G SW
40G SL Switch
100G
Nether Light
100G
ESnet 100G Testbed
40G
100 G SW
SL
ESnet
ANA 100 G
BBN
40G
OSF/DOE
ESnet
10G
10G
ExoGENI Stitching Switch @SL
10G
1G
NLR
ORCA Stitching Switch @BEN
10G
10G
RCI
BEN
I2
10G
10G
Duke
UH
FIU
NICTA

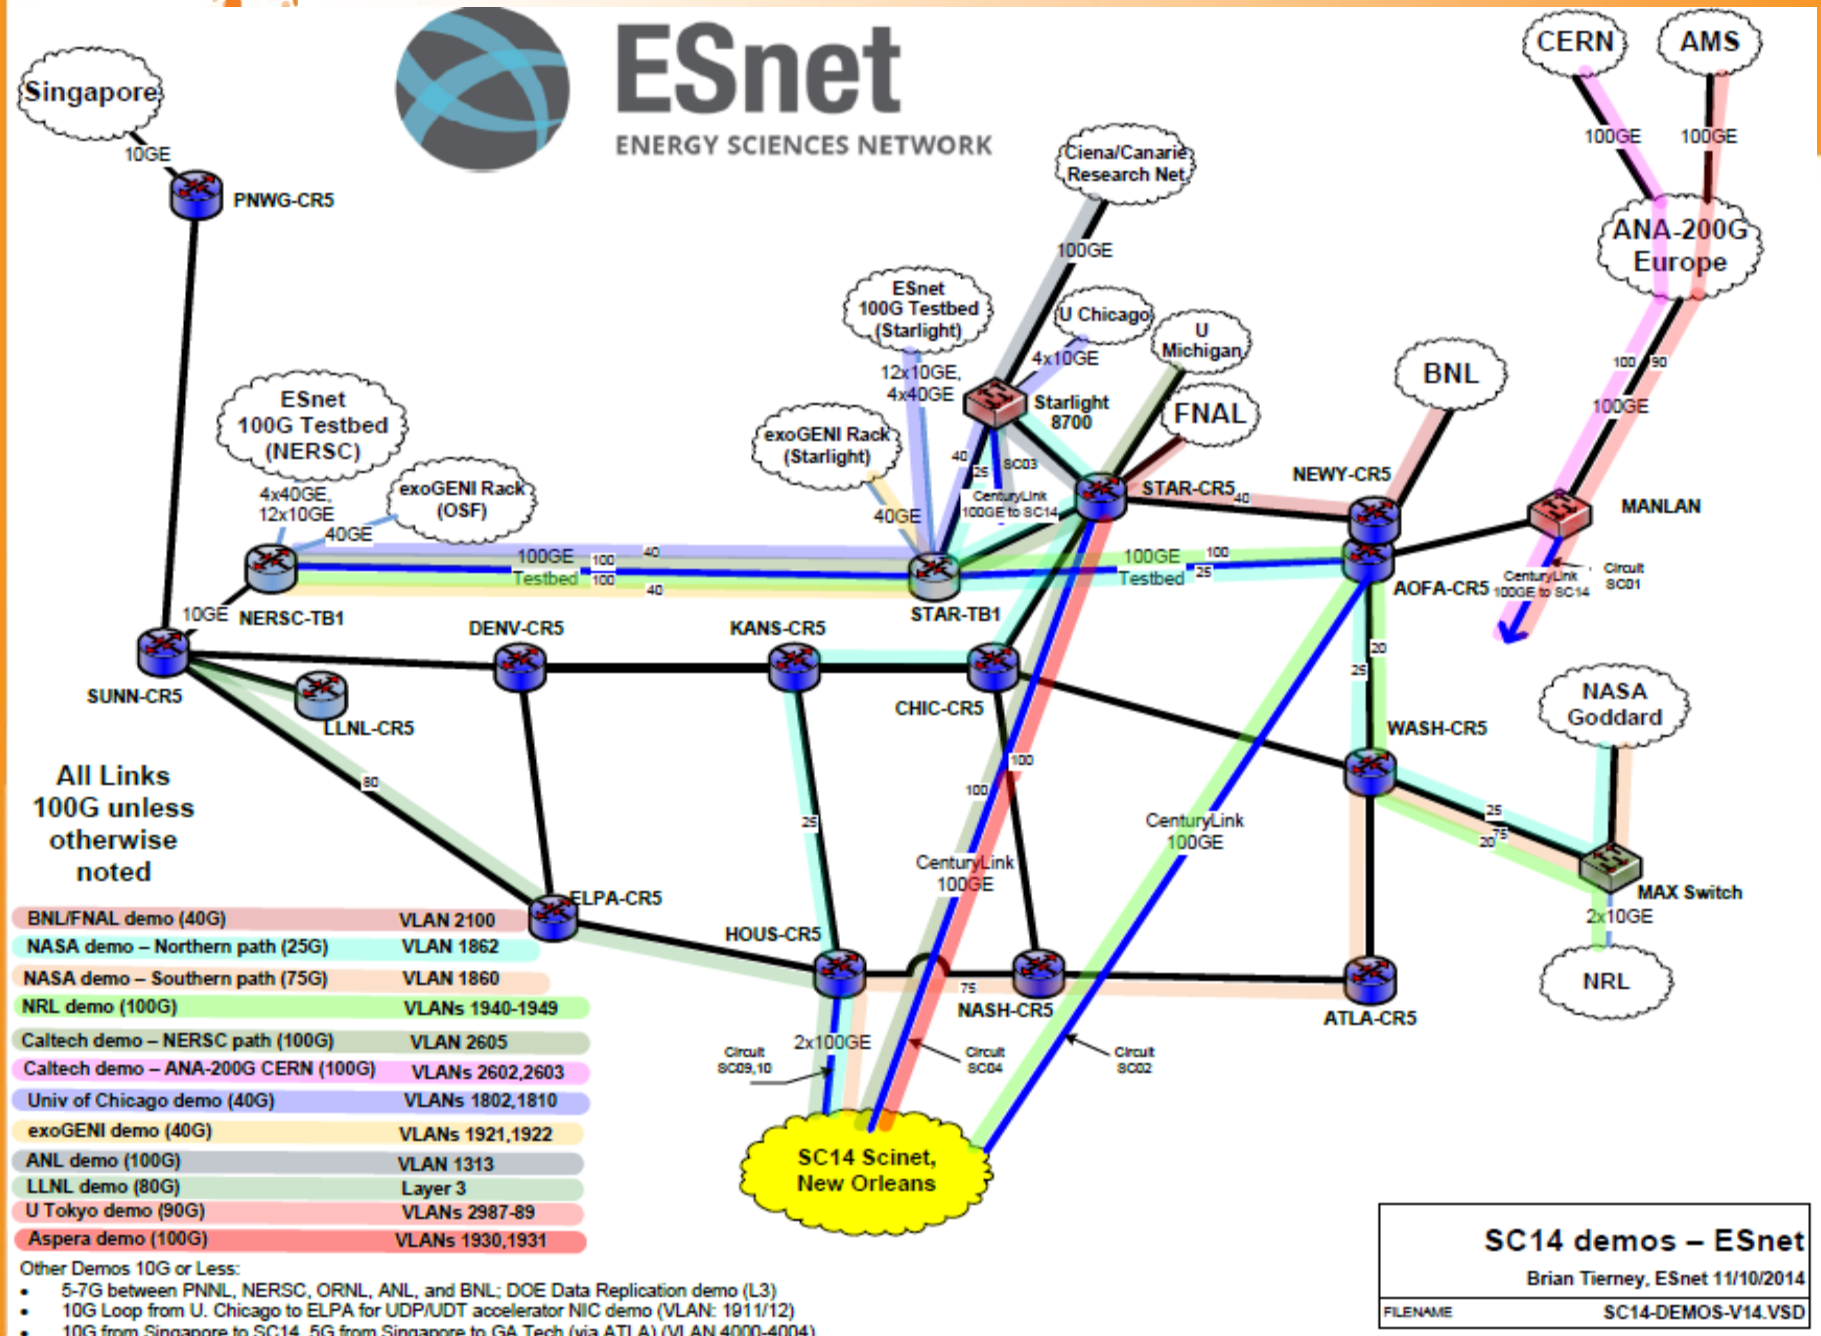
ESnet
ENERGY SCIENCES NETWORK
Singapore
10GE
PNWG-CR5
CERN
AMS
100GE
100GE
ANA-200G
Europe
Ciena/Canarie
Research Net
100GE
ESnet
100G Testbed
(Starlight)
U Chicago
U
Michigan
4x10GE
12x10GE,
4x40GE
Starlight
8700
BNL
FNAL
ESnet
100G Testbed
(NERSC)
exoGENI Rack
(Starlight)
exoGENI Rack
(OSF)
4x40GE,
12x10GE
40GE
40GE
SC03
CenturyLink
100GE to SC14
STAR-CR5
NEWY-CR5
MANLAN
100GE
Testbed
100GE
Testbed
NERSC-TB1
STAR-TB1
AOFA-CR5
CenturyLink
100GE to SC14
Circuit
SC01
10GE
SUNN-CR5
LLNL-CR5
DENV-CR5
KANS-CR5
CHIC-CR5
WASH-CR5
NASA
Goddard
All Links
100G unless
otherwise
noted
CenturyLink
100GE
CenturyLink
100GE
ELPA-CR5
HOUS-CR5
NASH-CR5
ATLA-CR5
MAX Switch
2x10GE
NRL
2x100GE
Circuit
SC09,10
Circuit
SC04
Circuit
SC02
SC14 Scinet,
New Orleans
BNL/FNAL demo (40G) VLAN 2100
NASA demo – Northern path (25G) VLAN 1862
NASA demo – Southern path (75G) VLAN 1860
NRL demo (100G) VLANs 1940-1949
Caltech demo – NERSC path (100G) VLAN 2605
Caltech demo – ANA-200G CERN (100G) VLANs 2602,2603
Univ of Chicago demo (40G) VLANs 1802,1810
exoGENI demo (40G) VLANs 1921,1922
ANL demo (100G) VLAN 1313
LLNL demo (80G) Layer 3
U Tokyo demo (90G) VLANs 2987-89
Aspera demo (100G) VLANs 1930,1931
Other Demos 10G or Less:
5-7G between PNNL, NERSC, ORNL, ANL, and BNL; DOE Data Replication demo (L3)
10G Loop from U. Chicago to ELPA for UDP/UDT accelerator NIC demo (VLAN: 1911/12)
10G from Singapore to SC14, 5G from Singapore to GA Tech (via ATLA) (VLAN 4000-4004)
SC14 demos – ESnet
Brian Tierney, ESnet 11/10/2014
FILENAME SC14-DEMOS-V14.VSD

PetaTrans: Peta Byte Science Data Transfer

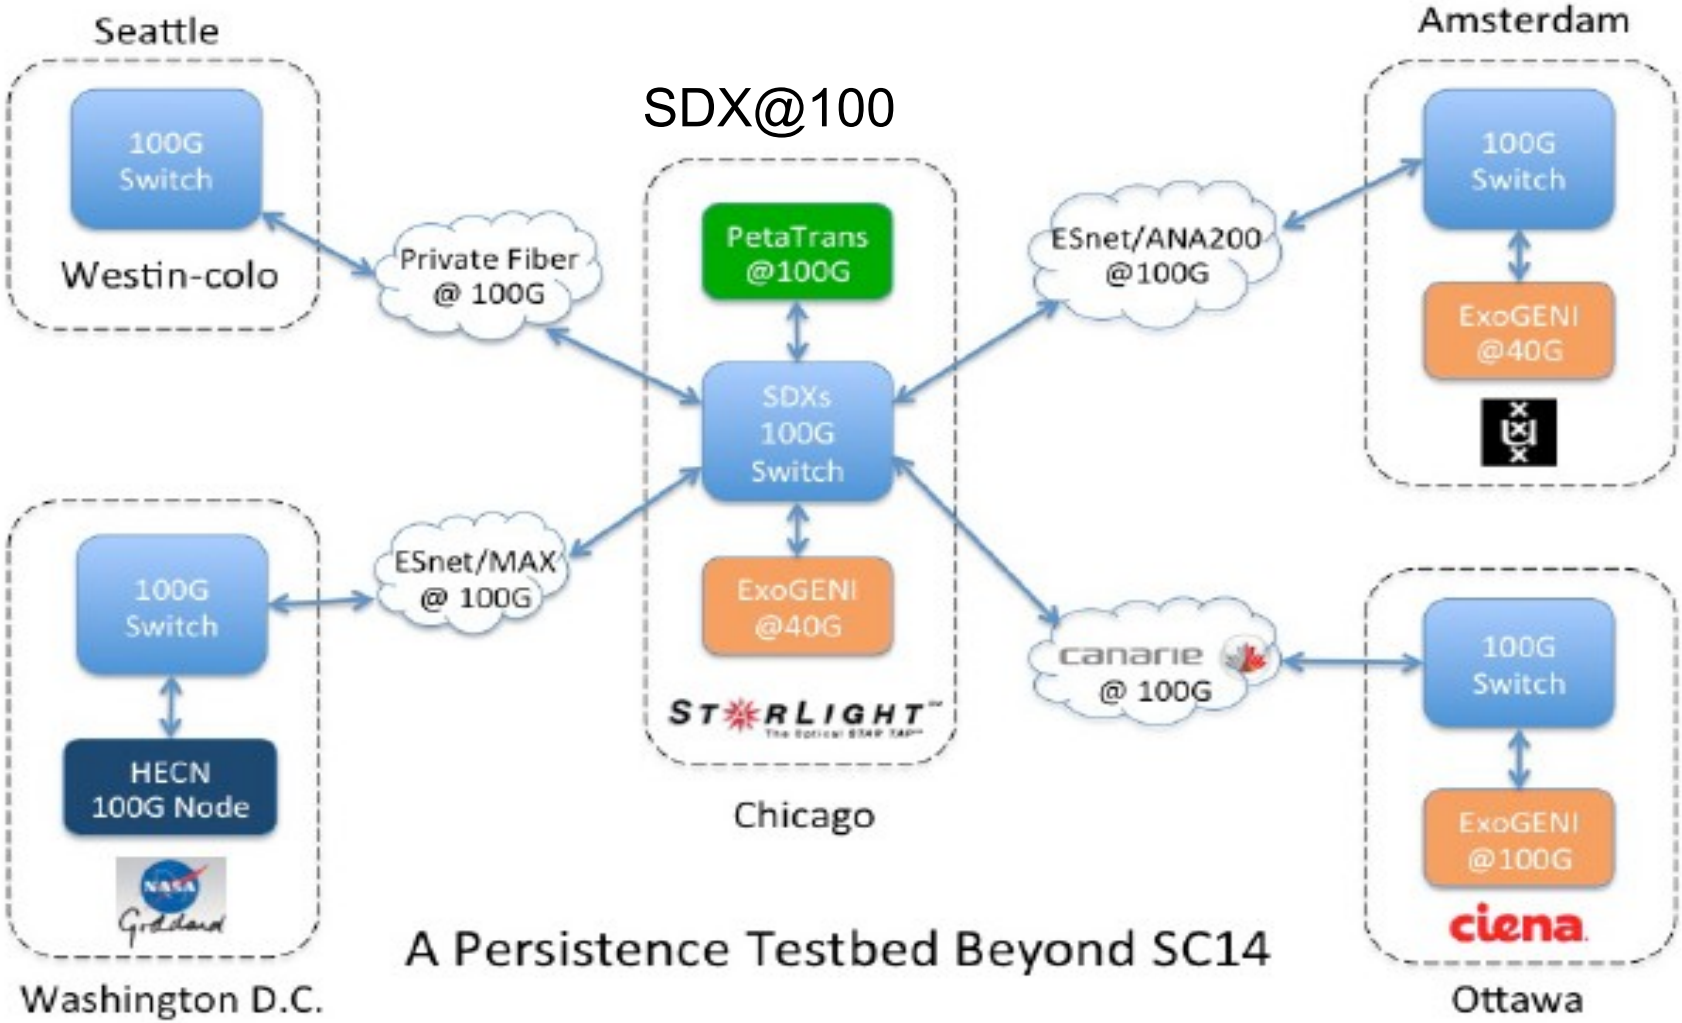

Global Software-Defined Dynamic Circuits for Data Intensive Science
(PhEDEx – ANSE – PANDA - OpenDayLight)

BI Factoids

- Haploid Human Genome (23 Chromosomes) = Approximately **3.2** B**illion** Bases Long
- Contains 20,000–25,000 Distinct Protein-Coding Genes.
- 6 Billion Base Pairs Per Diploid Cell.
- Kilobase (kb): Molecular Biology Measurement Unit Equal To 1000 Base Pairs of Deoxyribonucleic Acid (DNA) or Ribonucleic acid (RNA).
- Information Flow : From DNA Through RNA To Proteins
- Soon: 85 Petabytes of Data
- 1st Sequence: 13 Years
- Today: < 30 Minutes

Big Science Data Generators

Sequencers

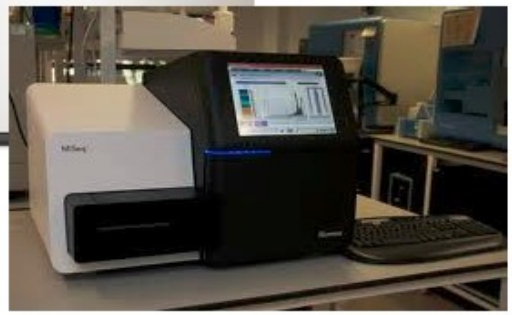

Source: Don Pruess

Cost Effective Sequencing = Massive Data

NATURE|Vol 464|1 April 2010 HUMAN GENOME AT TEN NEWS FEATURE

HUMAN GENOME

THE SEQUENCE EXPLOSION

At the time of the announcement of the first drafts of the human genome in 2000, there were 8 billion base pairs of sequence in the three main databases for 'finished' sequence: GenBank, run by the US National Center for Biotechnology Information; the DNA Databank of Japan; and the European Molecular Biology Laboratory (EMBL) Nucleotide Sequence Database. The databases share their data regularly as part of the International Nucleotide Sequence Database Collaboration (INSDC). In the subsequent first post-genome decade, they have added another 270 billion bases to the collection of finished sequence, doubling the size of the database roughly every 18 months. But this number is dwarfed by the amount of raw sequence that has been created and stored by researchers around the world in the Trace archive and Sequence Read Archive (SRA).

See Editorial, page 649, and human genome special at www.nature.com/humangenome

DNA SEQUENCES BY TAXONOMY

International Nucleotide Sequence Database Collaboration: The main repositories of 'finished' sequence span a wide range of organisms, representing the many priorities of scientists worldwide.

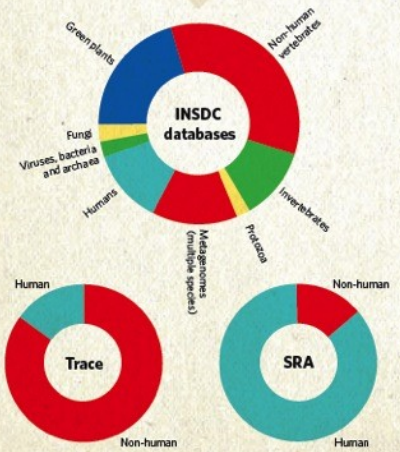

Trace Archive: Developed to house the raw output of high-throughput sequencers built in the late 1990s, the trace archive spans a wide range of taxa.

Sequence Read Archive: Houses raw data from next-generation sequencers. Dominated by human sequence, including multiple coverage for more than 170 people.

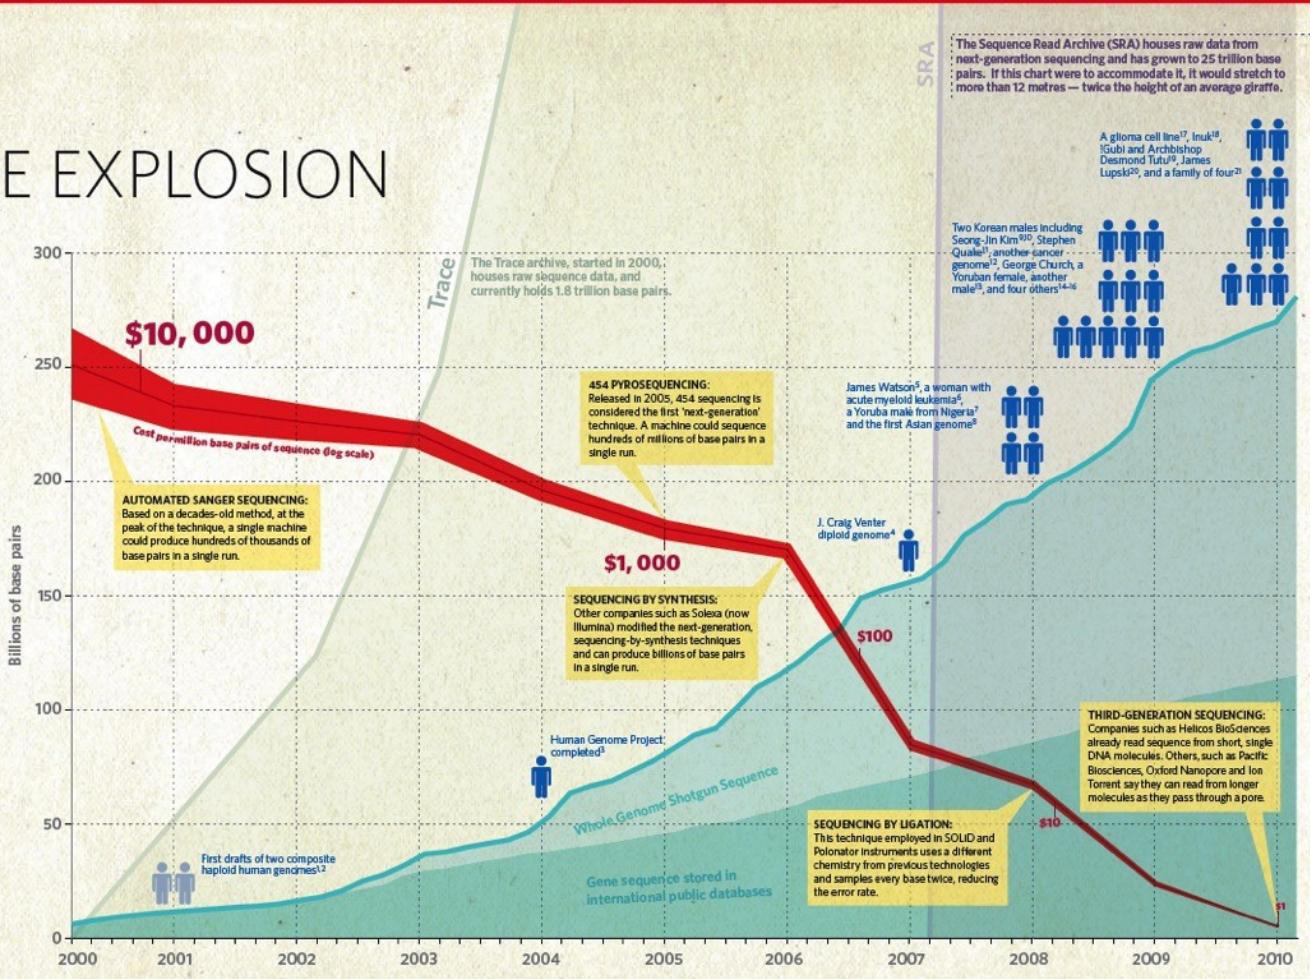

HOW MANY HUMAN GENOMES?

The graphic shows all published, fully sequenced human genomes since 2000, including nine from the first quarter of 2010. Some are resequencing efforts on the same person and the list does not include unpublished completed genomes.

1. Venter, J. C. *et al. Science* **291**, 1304–1351 (2001).
2. International Human Genome Sequencing Consortium *Nature* **409**, 860–921 (2001).
3. International Human Genome Sequencing Consortium *Nature* **431**, 931–945 (2004).
4. Levy, S. *et al. PLoS Biol.* **5**, e254 (2007).
5. Wheeler, D. A. *et al. Nature* **452**, 872–876 (2008).
6. Ley, T. J. *et al. Nature* **456**, 66–72 (2008).
7. Bentley, D. R. *et al. Nature* **456**, 53–59 (2008).
8. Wang, J. *et al. Nature* **456**, 60–65 (2008).
9. Ahn, S.-M. *et al. Genome Res.* **19**, 1622–1629 (2009).
10. Kim, J.-I. *et al. Nature* **460**, 1011–1015 (2009).
11. Pushkarev, D., Neff, N. F. & Quake, S. R. *Nature Biotechnol.* **27**, 847–850 (2009).
12. Wheeler, D. A. *et al. Nature* **452**, 872–876 (2008).
13. Ormanac, R. *et al. Science* **327**, 78–81 (2009).
14. McKernan, K. J. *et al. Genome Res.* **19**, 1527–1541 (2009).
15. Pleasance, E. D. *et al. Nature* **463**, 191–196 (2010).
16. Pleasance, E. D. *et al. Nature* **463**, 184–190 (2010).
17. Clark, M. J. *et al. PLoS Genet.* **6**, e1000832 (2010).
18. Rasmussen, M. *et al. Nature* **463**, 757–762 (2010).
19. Schuster, S. C. *et al. Nature* **463**, 943–947 (2010).
20. Lupski, J. R. *et al. N. Engl. J. Med.* doi:10.1056/NEJMoa0908094 (2010).
21. Roach, J. C. *et al. Science* doi:10.1126/science.1186802 (2010).

Page size by comparison

Beyond Today's Internet
Experiencing a Smart Future

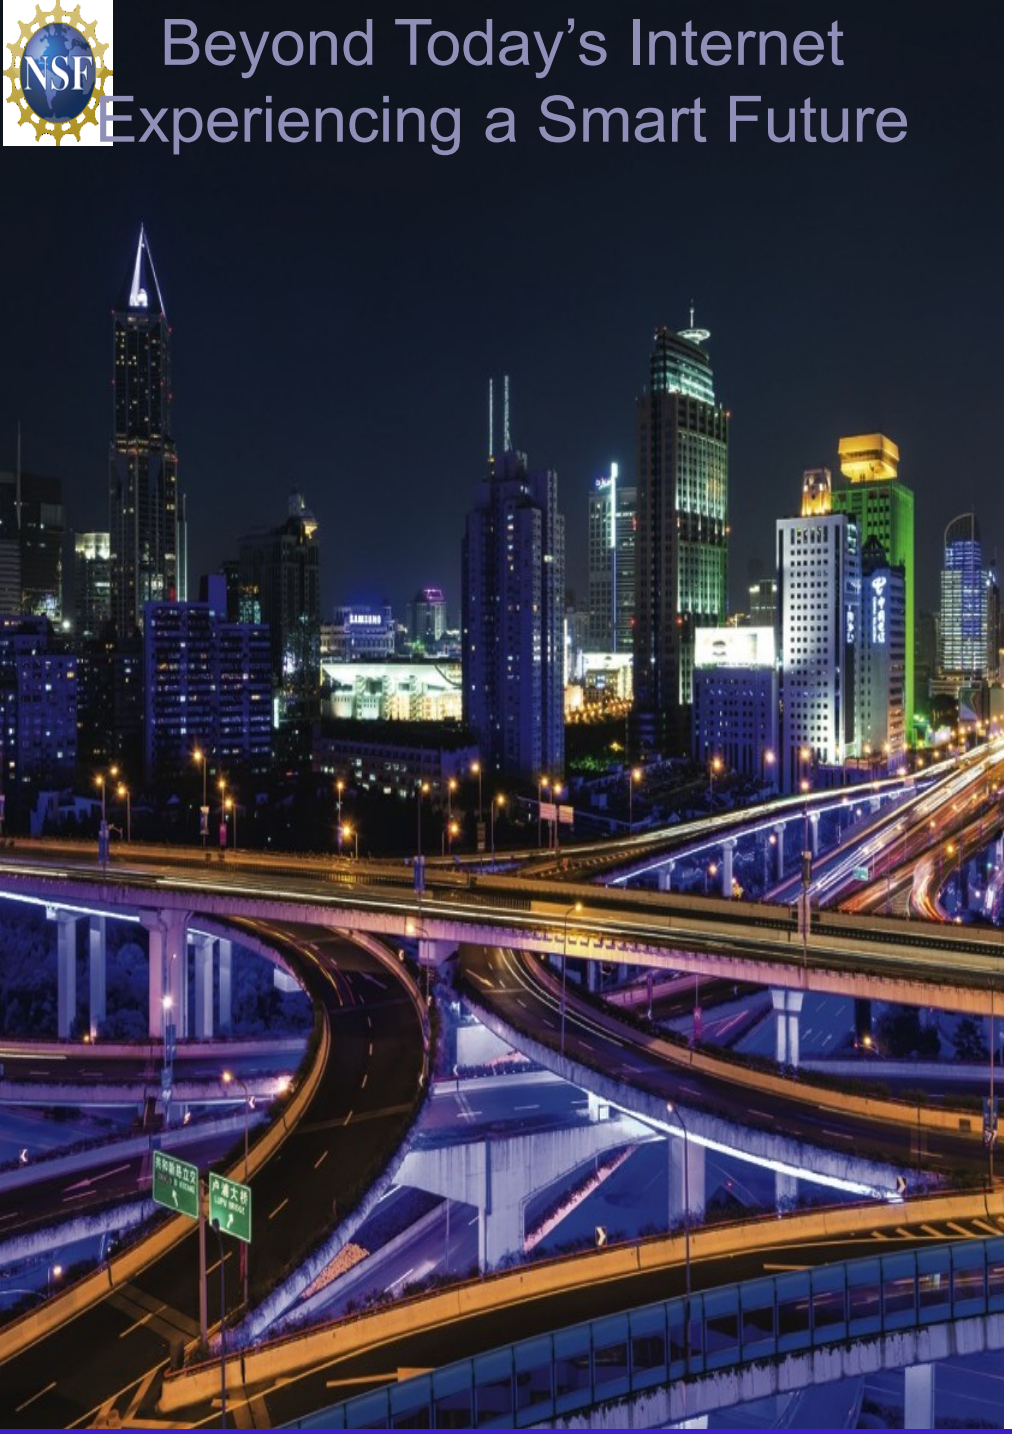

Prototype SDX Bioinformatics Exchange: Demonstrating an Essential Use-Case for Personalized Medicine

Robert Grossman – University of Chicago
Joe Mambretti – Northwestern University
Piers Nash – University of Chicago
Jim Chen – Northwestern University
Allison Heath – University of Chicago

THE UNIVERSITY OF CHICAGO MEDICINE

NORTHWESTERN UNIVERSITY

Precision Medicine Enabled By Precision Networks

- Precisely match treatments to patients and their specific disease
- Genomic data promises optimal matching.
- 1.7 million cancer cases diagnosed in America each year.
- A single RNA-seq file is 10-20 GB, Whole genome raw data files are > 100 GB.
- Analysis has become the bottleneck and data size is an issue.
  - 2,000,000 genomes ≈ 1 Exabyte (1,000,000,000,000 MB)
  - Cost to sequence 1 genome less than $5,000 and falling fast.
  - Cost to analyze 1 genome is approx. $100,000 and rising.
- A key step towards Algorithm-assisted Personalized medicine is building Data Commons/Cloud analytics and the *Programmable* Networks & Communication Exchanges (SDXs) for high performance, flexible data transport.

Future Vision: A Nationwide Virtual Comprehensive Cancer Center

Cloud Computation
Genomic Data Commons

Hospitals,
Doctors

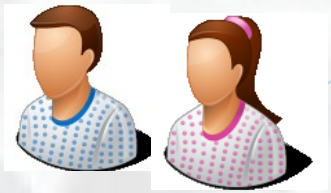

Patients

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

Bionimbus Protected Data Cloud

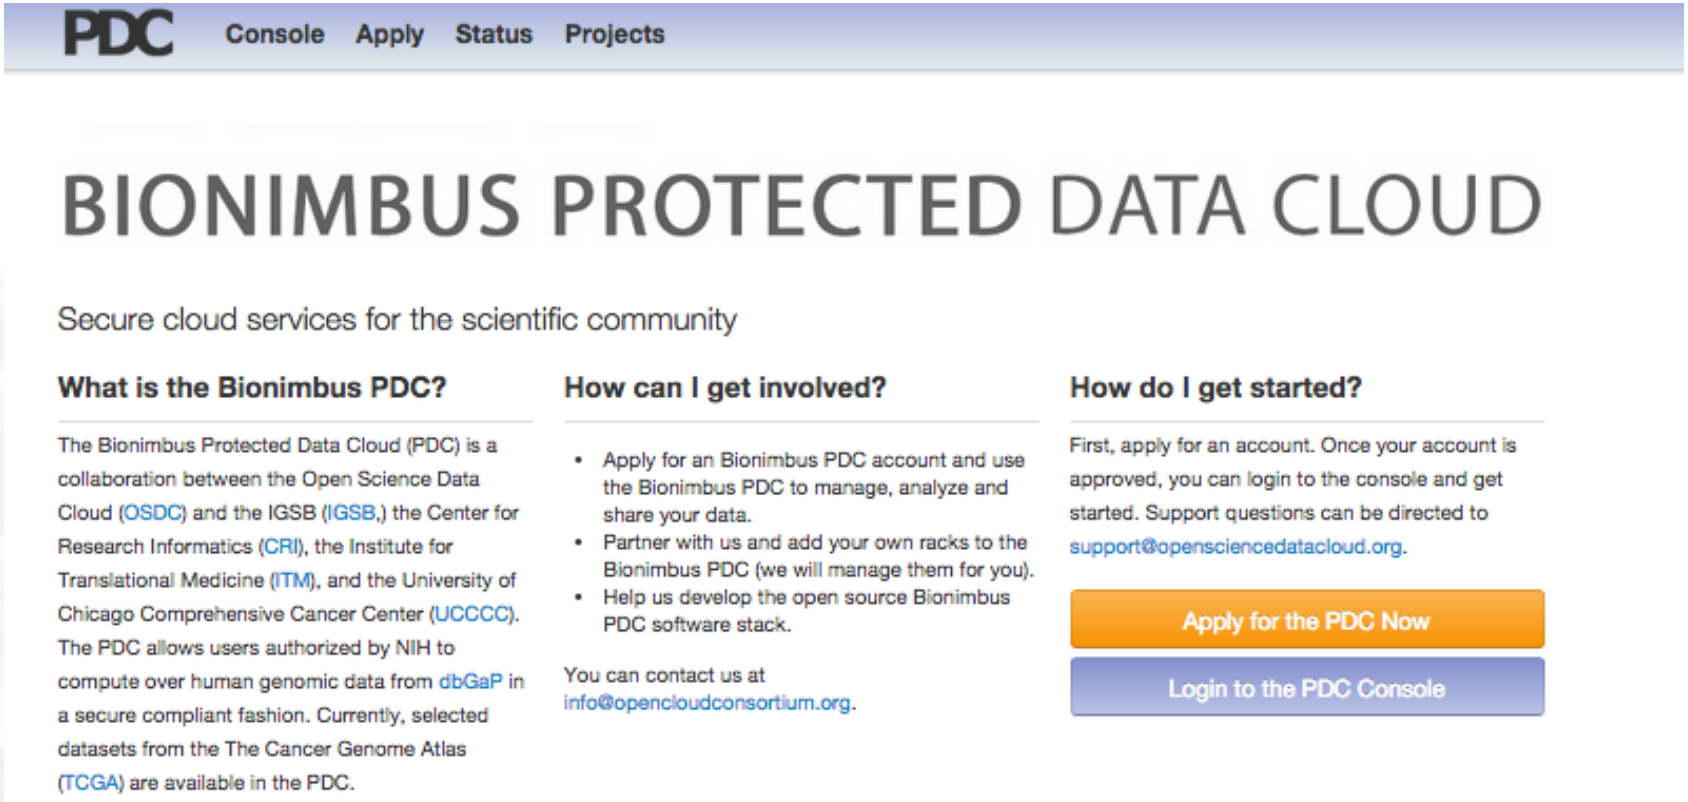

- Petabyte-scale, secure compliant biomedical cloud that interoperates with dbGaP controlled access data at NIH.

Opportunity: Close Integration of Research Workflows and Foundation Networks

- Opportunity: Using GENI To Develop Innovative Techniques for Extremely Close Integration of Research WorkFlows and Dynamic Programmable Network Resources, Enabling Precision Networking
- Network Foundation Architecture: GENI + Innovative Customized Software Defined Networking Exchange (SDX)
- For This Demonstration: Specifically To Meet The Requirements of Bioinformatic Workflows

Biomedical Data Commons

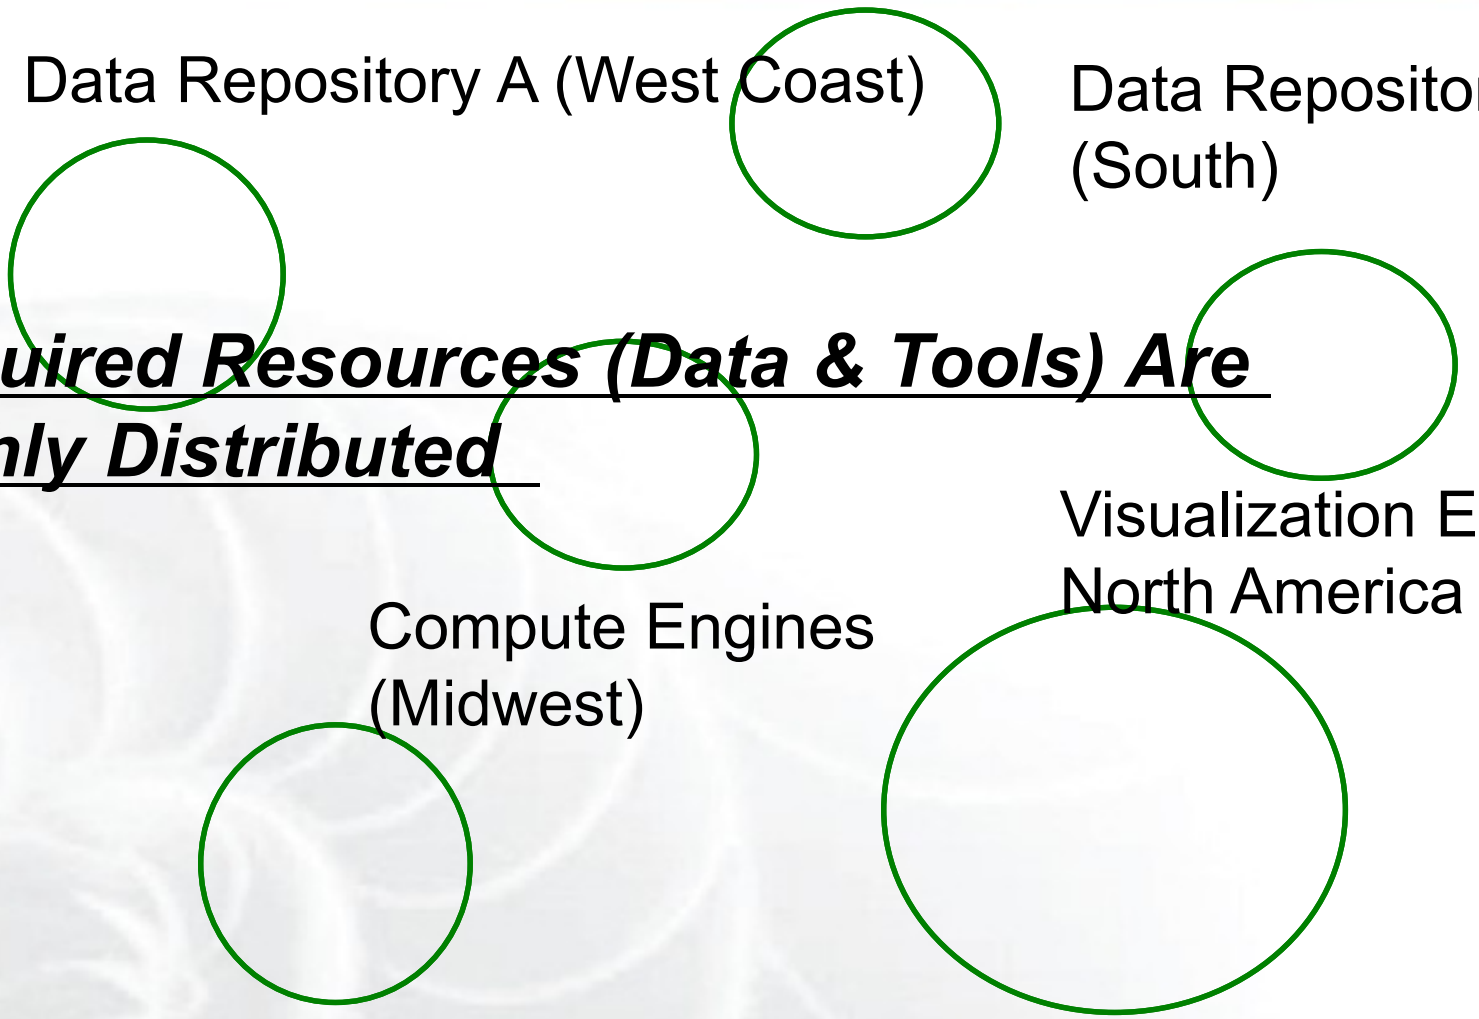

Required Resources (Data & Tools) Are Highly Distributed

Biomedical Data Commons: Flow Orchestration: Data Plane

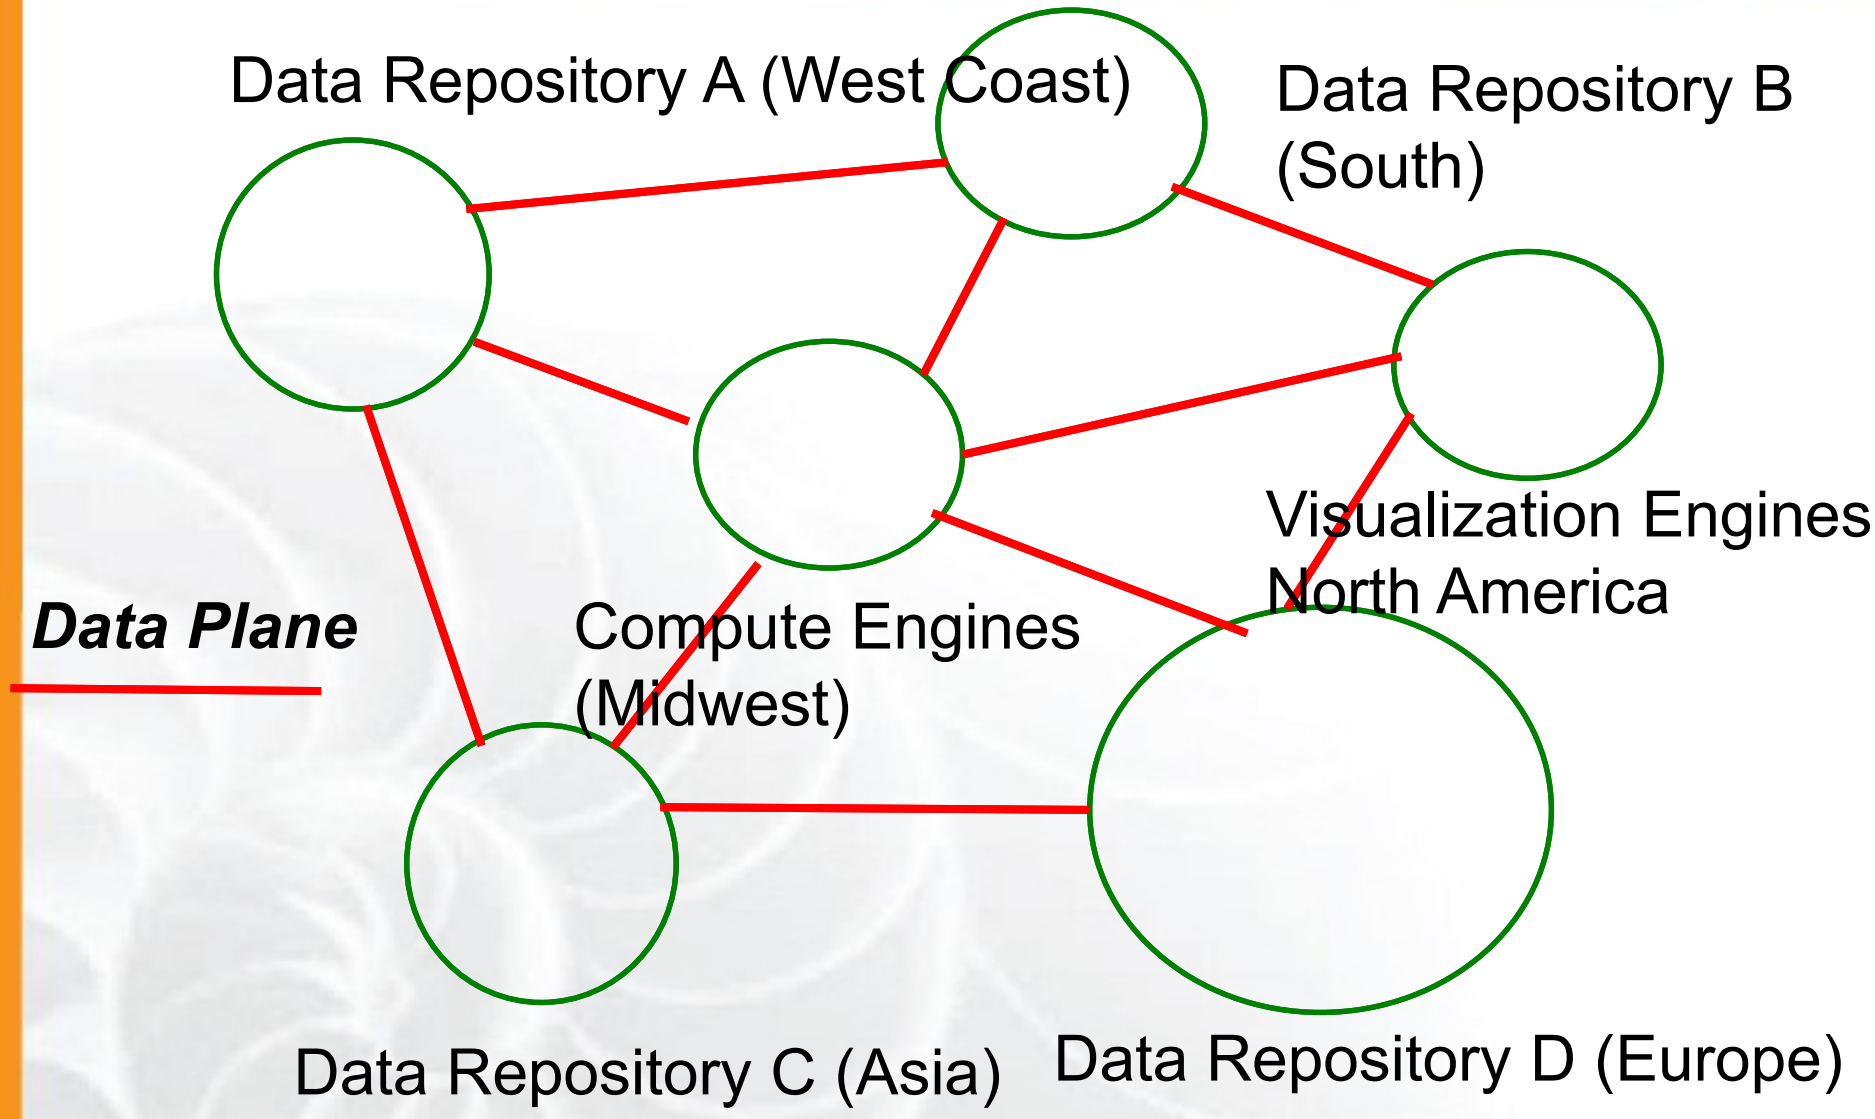

Biomedical Data Commons: Flow Orchestration: Control Plane + Data Plane

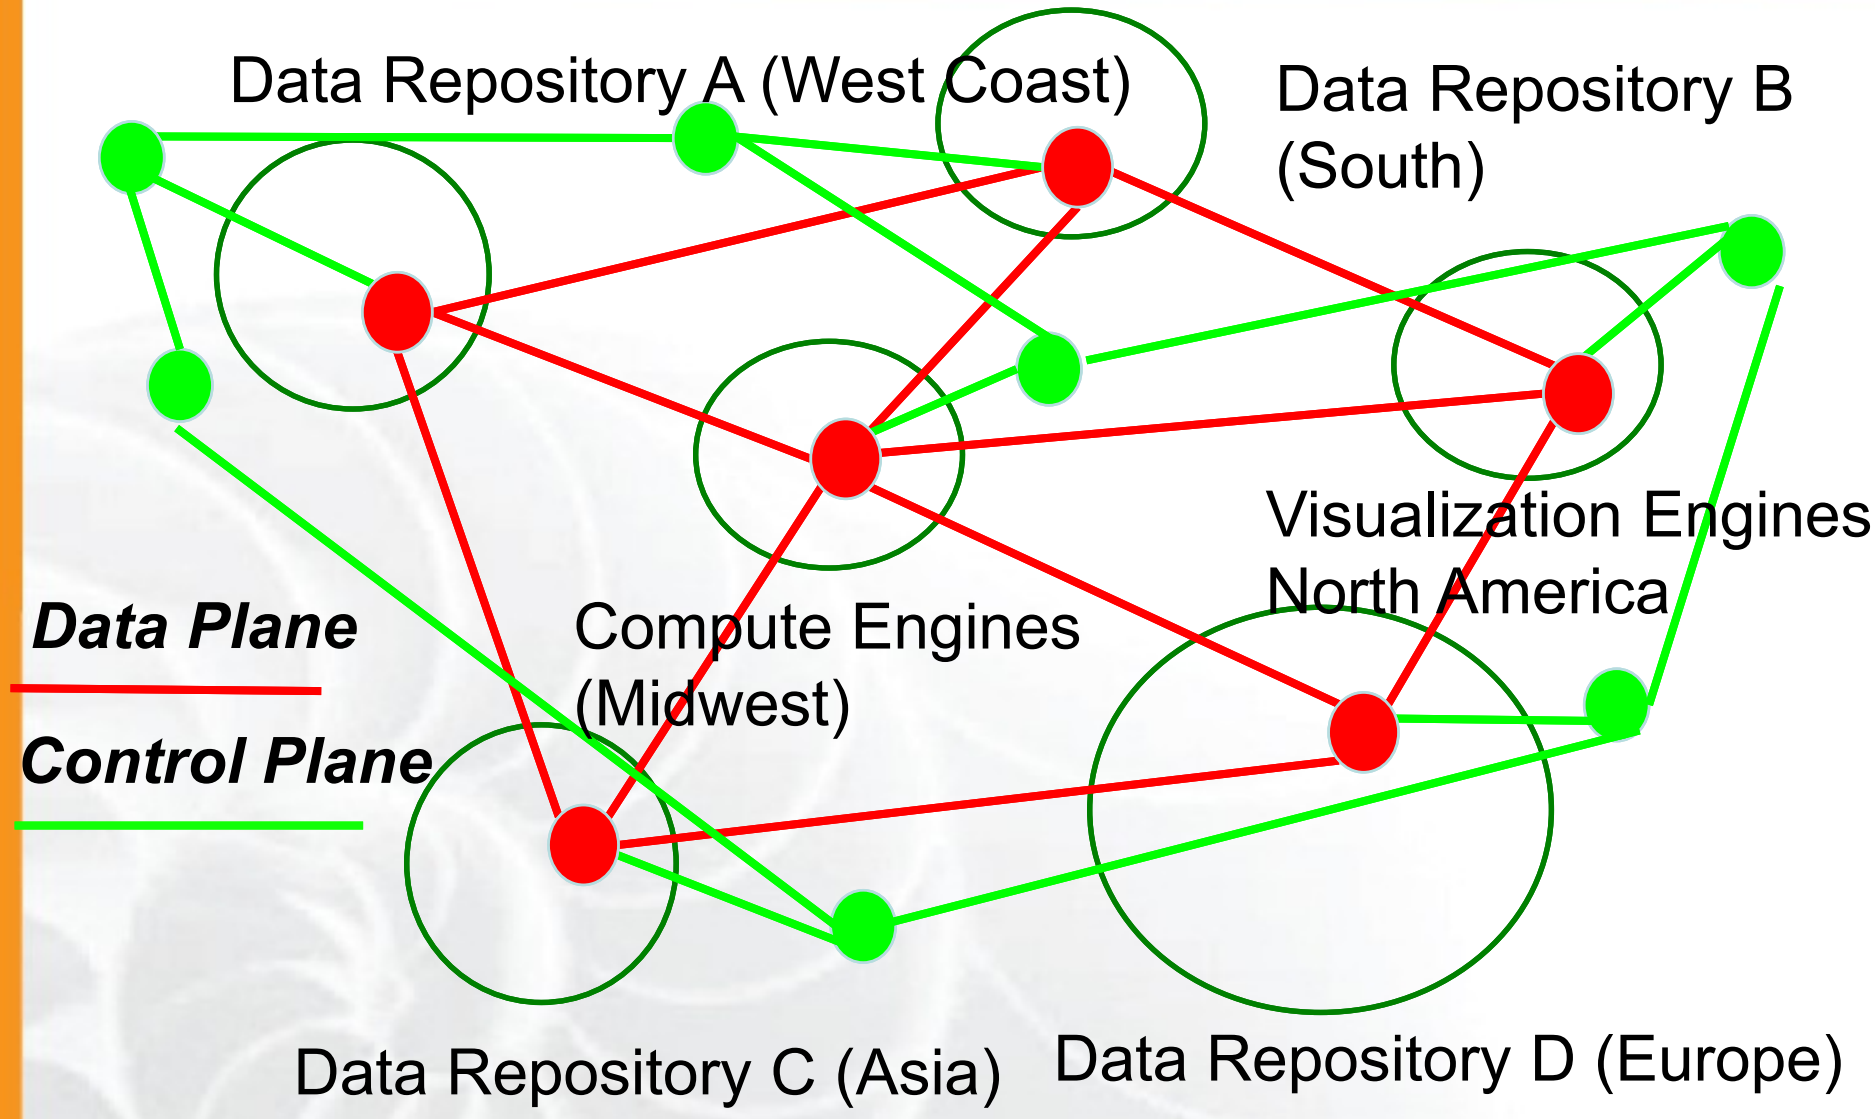

GEC22 Bioinformatics SDXs Demo Network

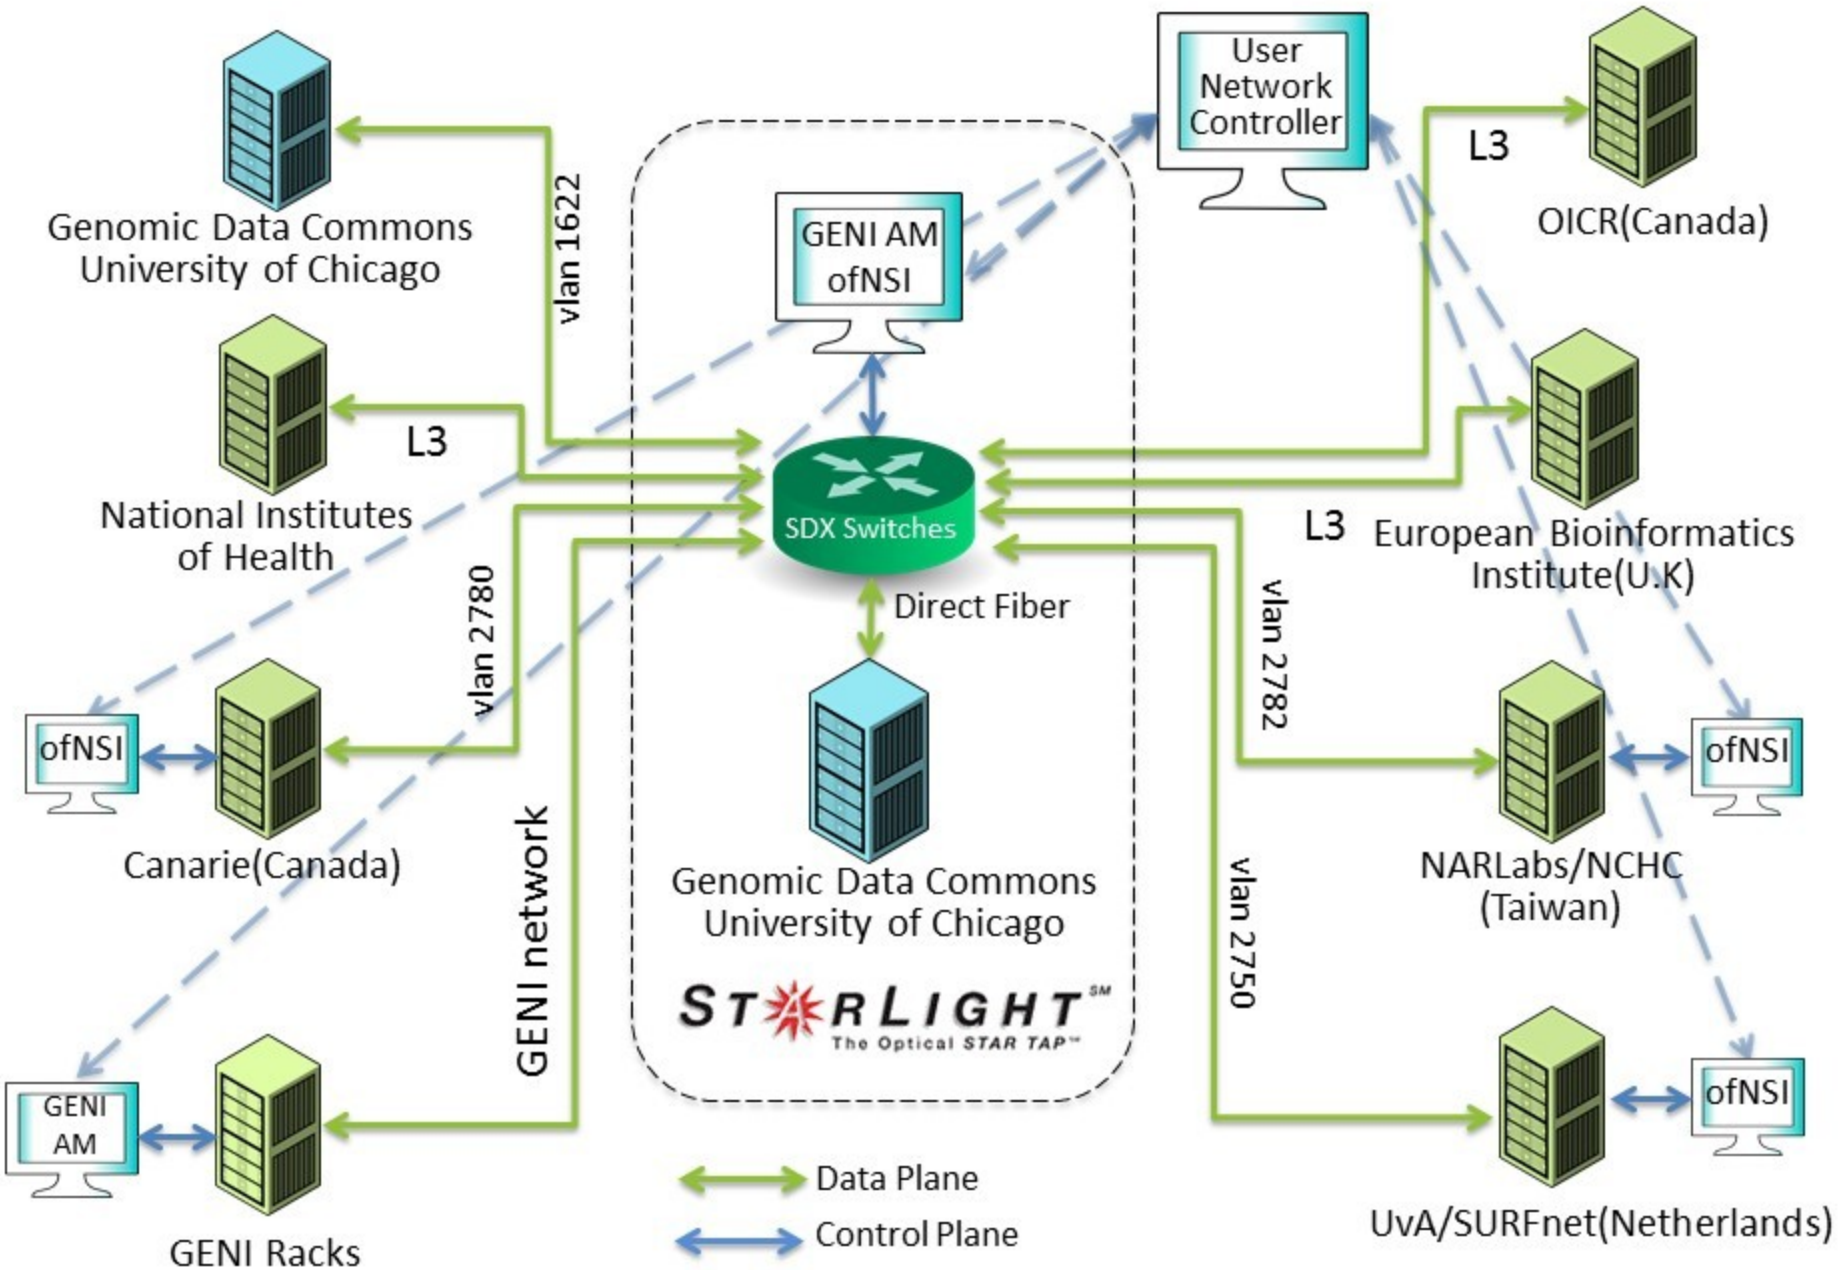

Bioinformatics SDX Demonstration

- A) Dynamically Moving Core Data Files Among Multiple Sites Around the World Via StarLight SDX
- B) Moving RNA-seq Data Files From NCI (Bethesda, MD) and EBI (Hinxton, UK) Through SDX Switch/Routers to The University of Chicago.
  - Analysis By Comparison To Known Data Correlated To Drug Response.
  - Determine Possible Actionable Therapeutic Options.
  - Return Viable Treatment Options To the Originating Site.

Bioinformatics Commons Data Transfer

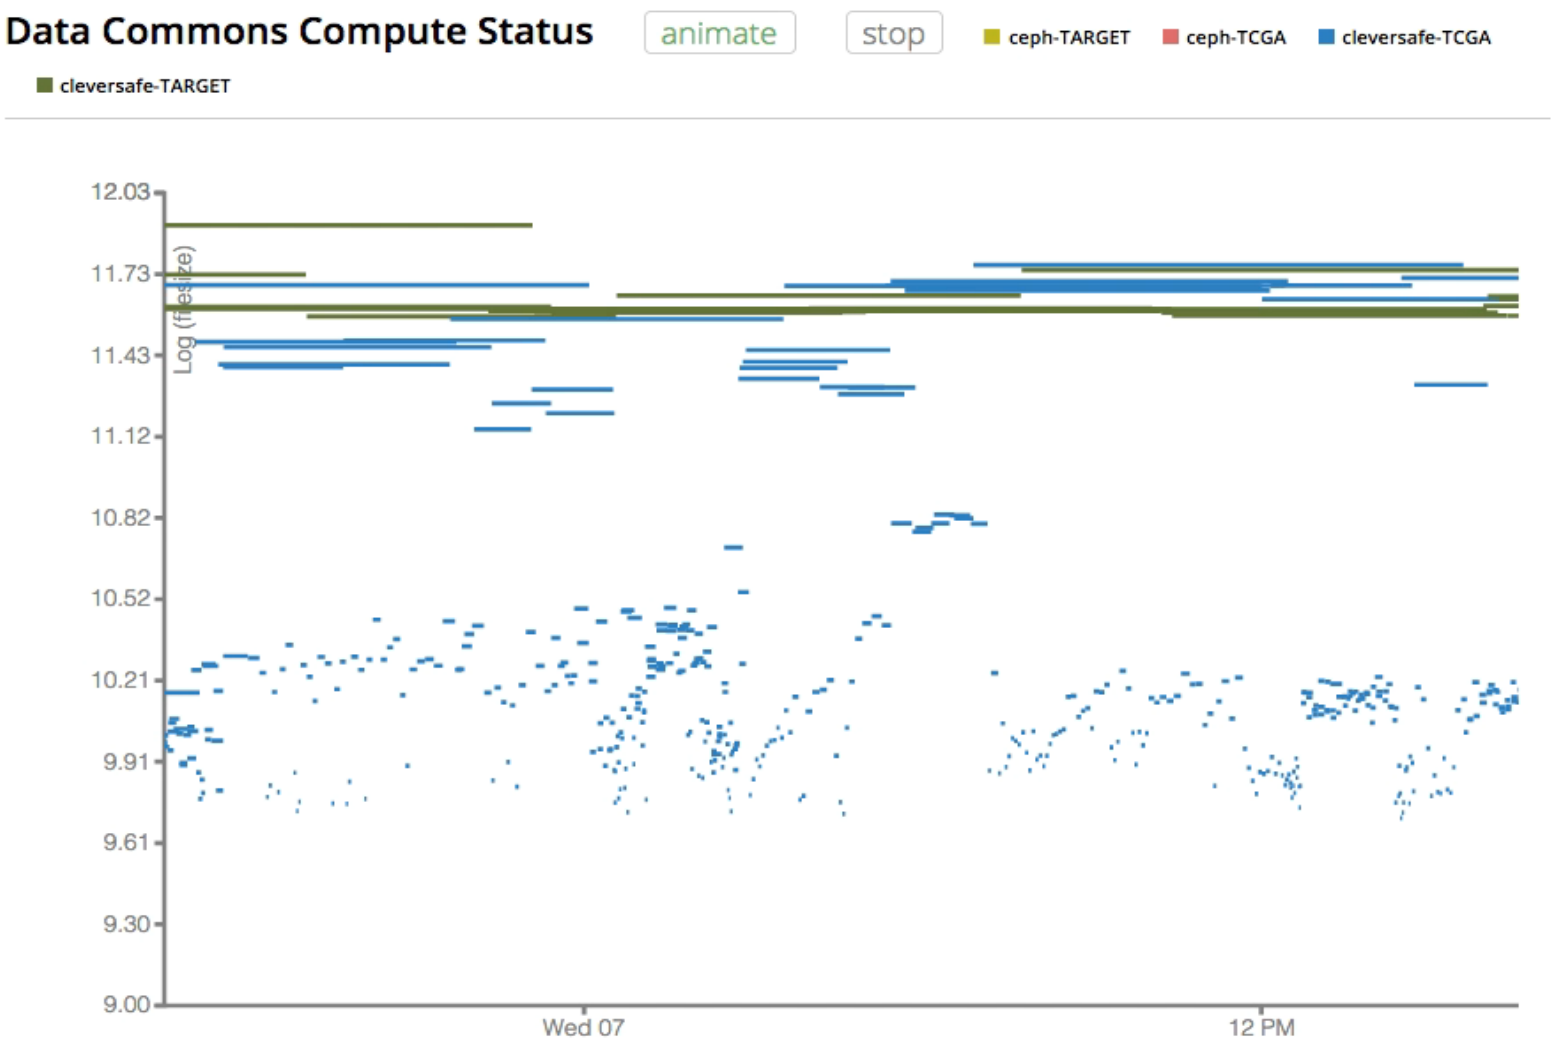

Results

- Precision Medicine Requires Data Commons That Scale To Hundreds of Petabytes, With Programmable Networks and Data Peering To Support Data Sharing.
- Speed Discovery and Support Analytics-Driven Healthcare To Recommend Treatment.
- Large Scale Data Analysis and Dynamic Pipelines For Workflows Are Essential For Determining Optimal Results.

Summary and Future

- An Innovative Approach To Advanced Knowledge Discovery and Medical Treatment**: Precision Medicine Supported By Precision Networking**
- Precision Mapping Of Communication Services To BI Workflow Requirements Across the World Using Advanced Analytics, the Biomedical Data Commons, and a Programmable Dynamic SDX
- Looking Forward:
  - A) Further Development/Refinement of Basic Capabilities
  - B) Transition to ***Actual Production Services***
  - **C) The Biomedical Data Commons and Bionimbus Protected Data Cloud Are Being Developed As a Key Production Knowledge Discovery/Transformational Medical Treatment Facility**

Using GENI To Invent the Future…

Thank You!

NORTHWESTERN
UNIVERSITY

EBI

canarie

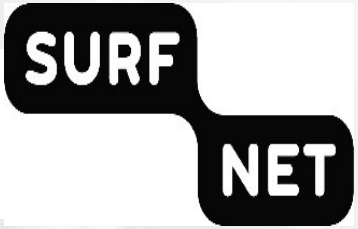

www.chameleoncloud.org

Another SDX Opportunity…

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

THE UNIVERSITY OF CHICAGO · TACC · NORTHWESTERN UNIVERSITY · THE OHIO STATE UNIVERSITY · UTSA · NSF

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

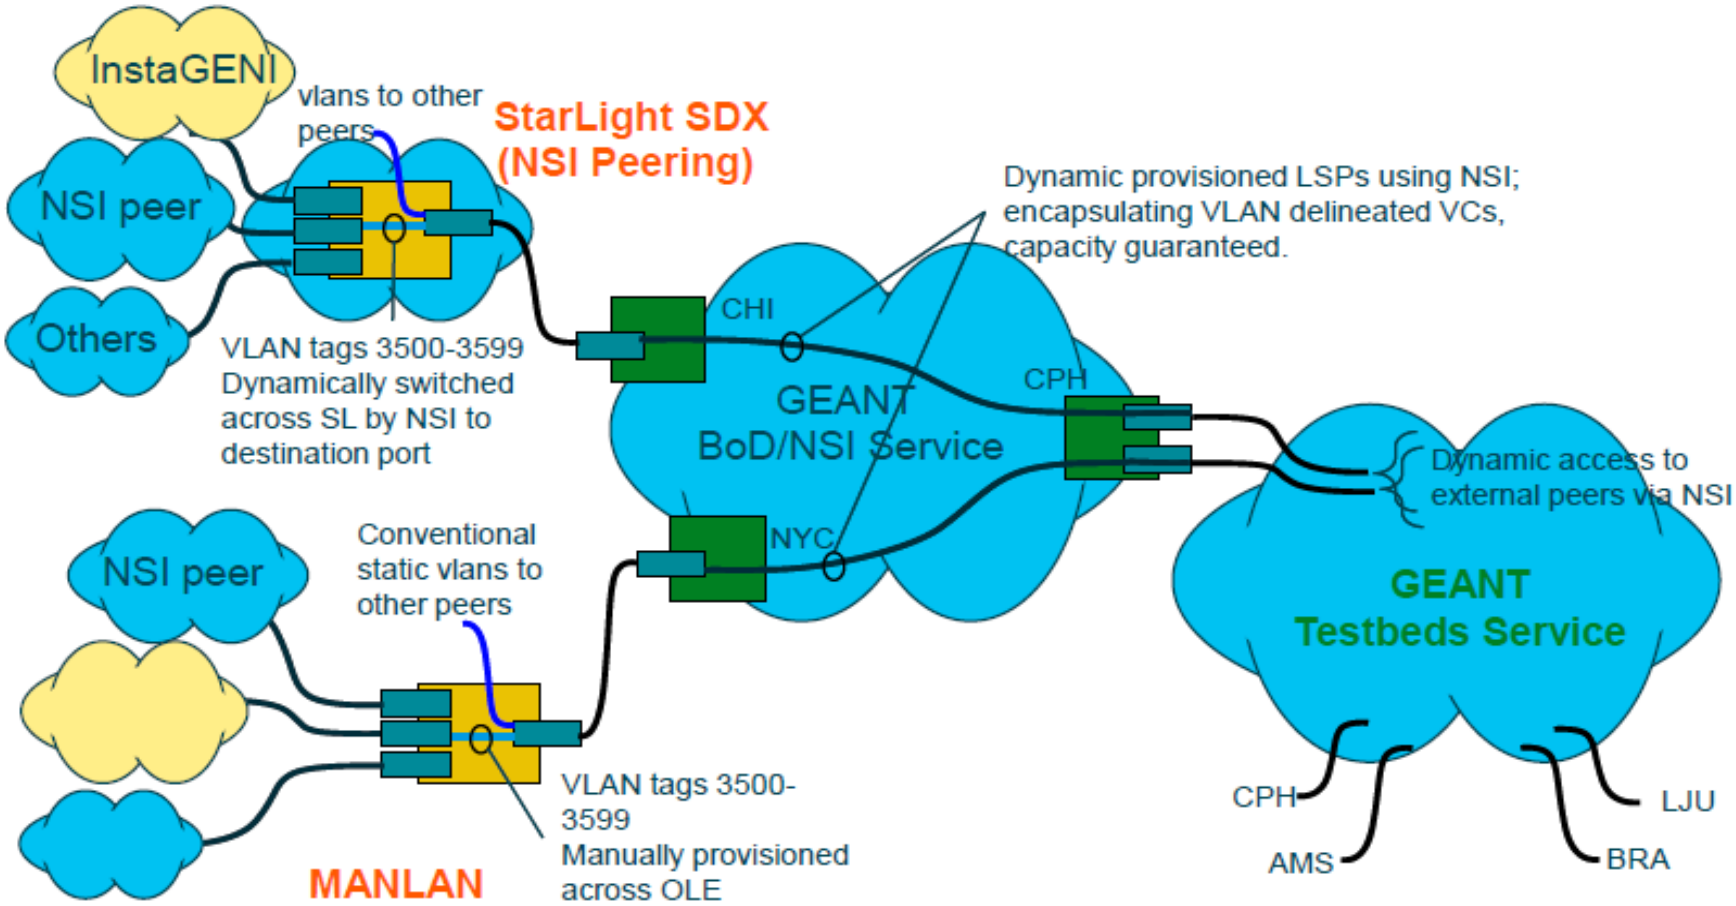

Potential for Major New International SDXs Exists

- Science DMZ Exchange Services
- Emerging Software Defined Infrastructure (SDI)
- Network Function Virtualization (NFV)
- SDX Services Via “App Store”
- Etc…
- Ref: Forthcoming Information On www.startap.net/starlight

New Initiative

- National Science Foundation International Research Network Connections Program: StarLight International SDX: A Software Defined Networking Exchange for Global Science Research and Education

Thanks to the NSF, The GENI GPO,
DOE, DARPA
Universities, National Labs,
International Partners,
and Other Supporters