Software-Defined Network Exchanges (SDXs): Architecture, Services, Capabilities, and Foundation 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)

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

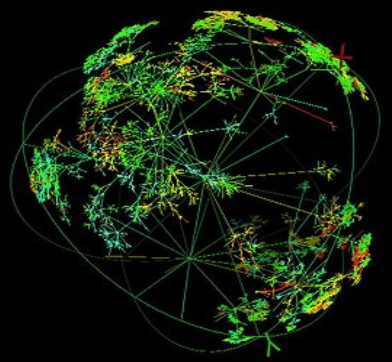

iCAIR

STARLIGHT℠

The Global Lambda Integrated Facility (GLIF) Is Based On GOLEs (GLIF Optical Lambda Exchanges)

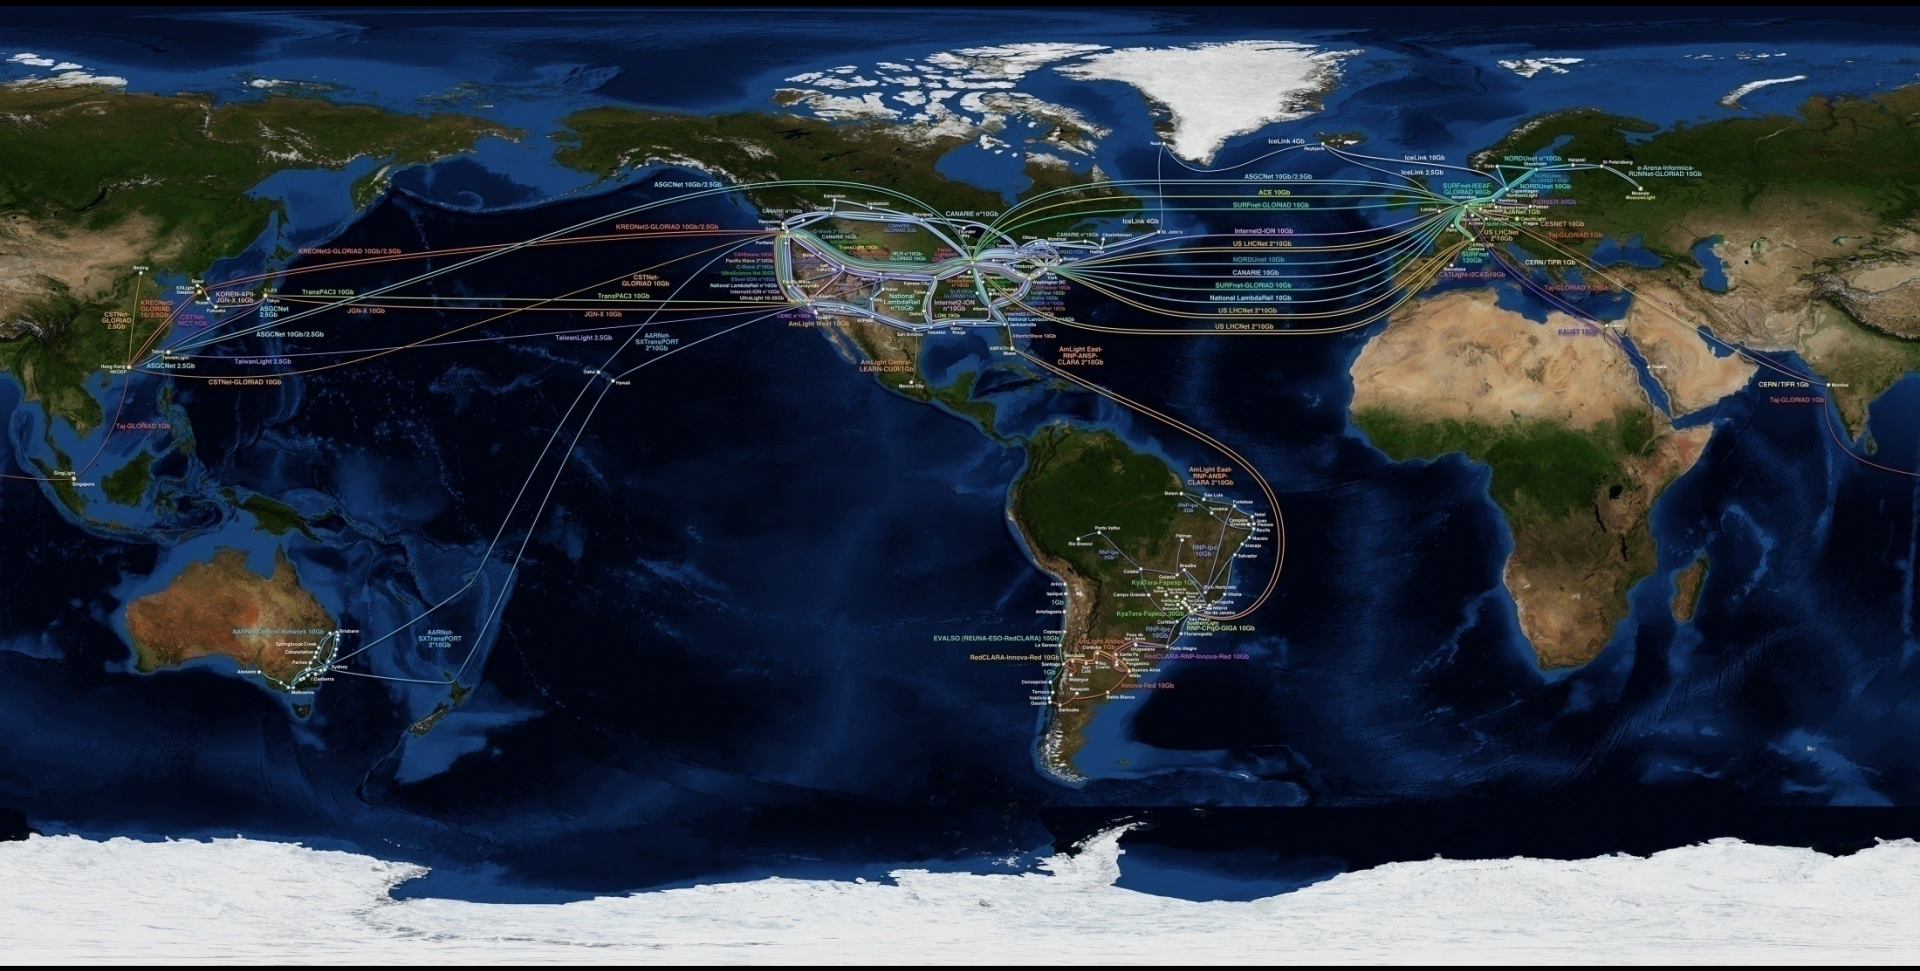

Many Years Ago, the GLIF Community Began Developing The Network Service Interface (NSI) Architecture With the Open Grid Forum (OGF) To Enable Path Resource Sharing Among GOLEs Which Have Multiple Control Frameworks

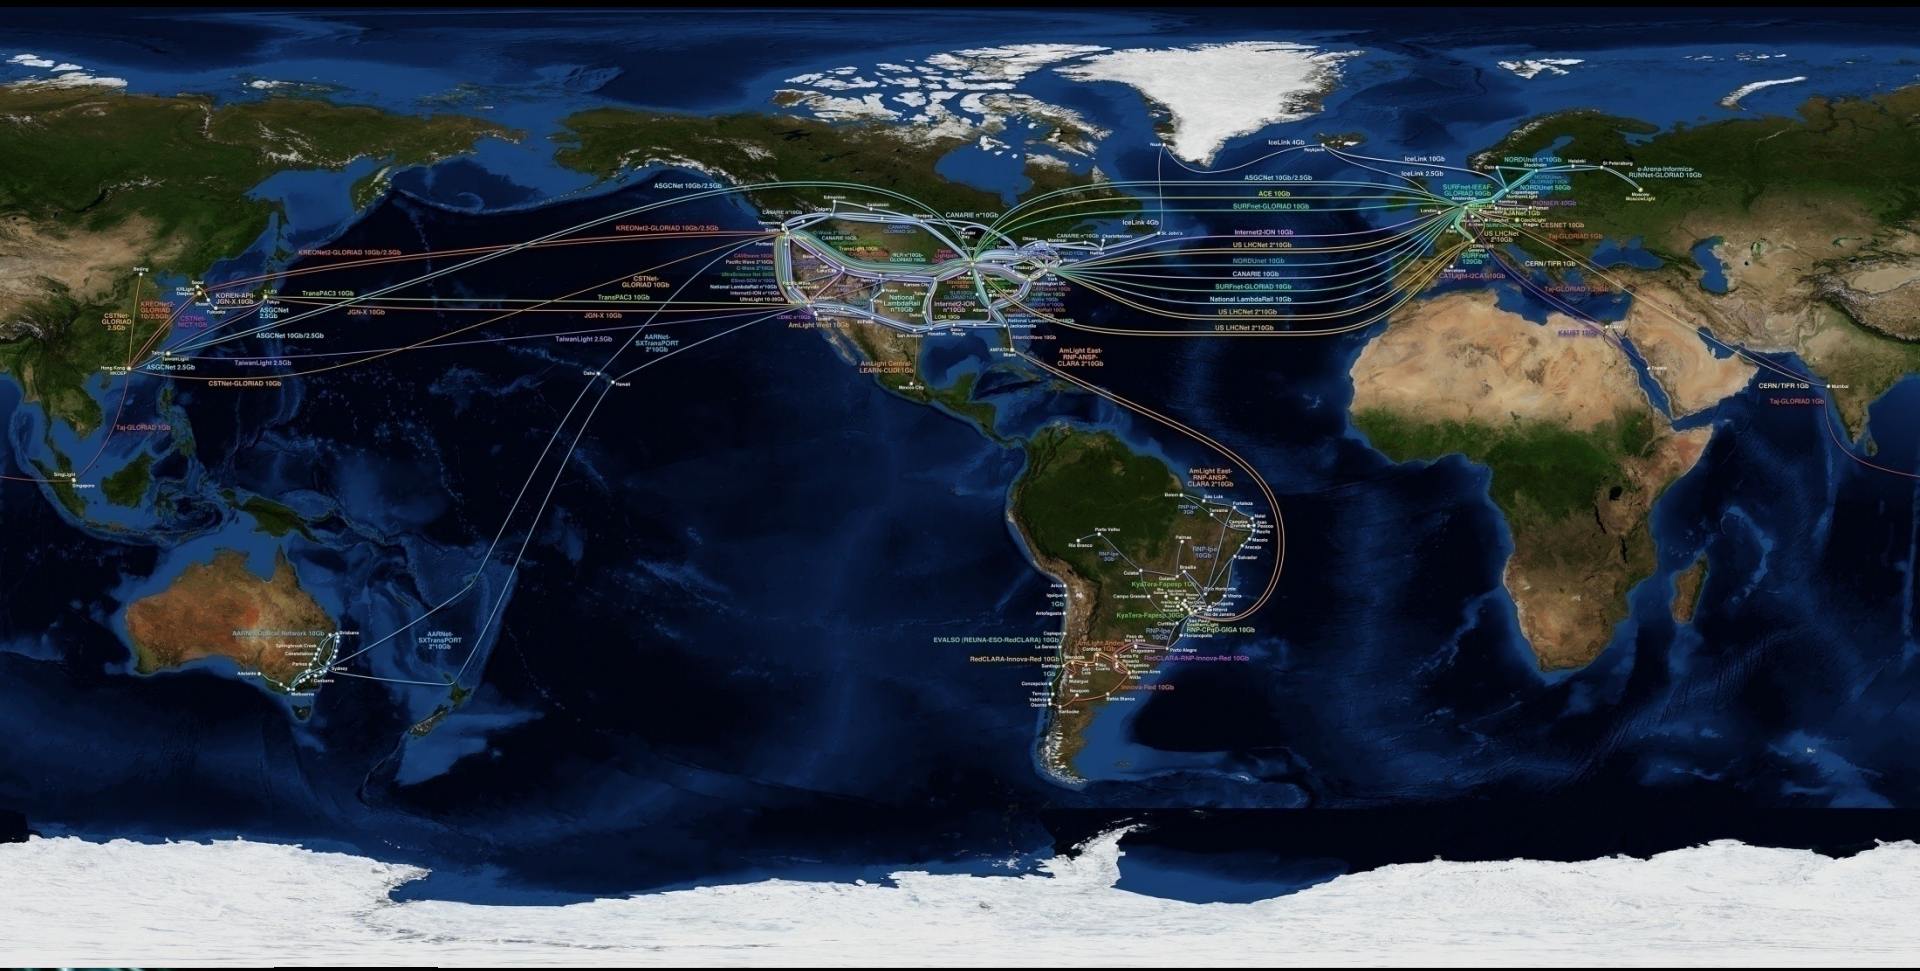

Automated GOLE Fabric

Source: GLIF Auto GOLE Group

GLIF AutoGOLE Initiative Oct 2013

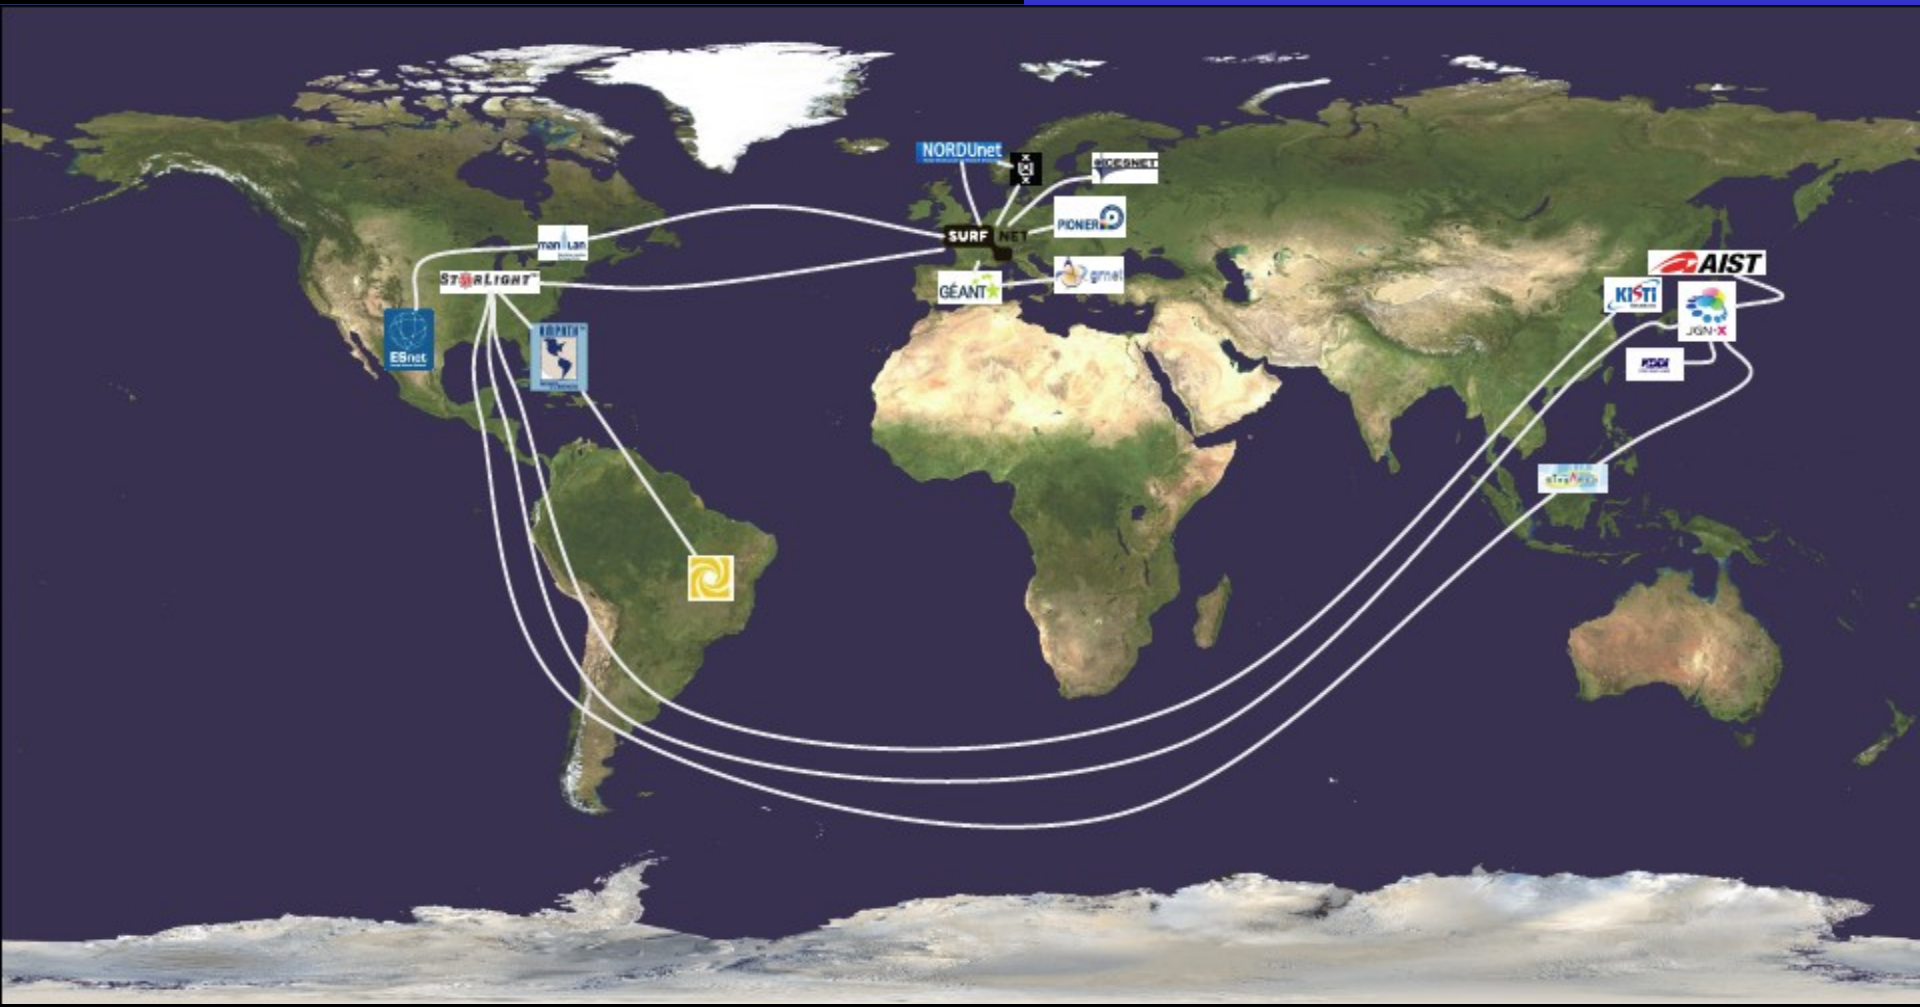

Source: GLIF Auto GOLE Group

STARLIGHT℠

Building an AutomatedGOLE backbone

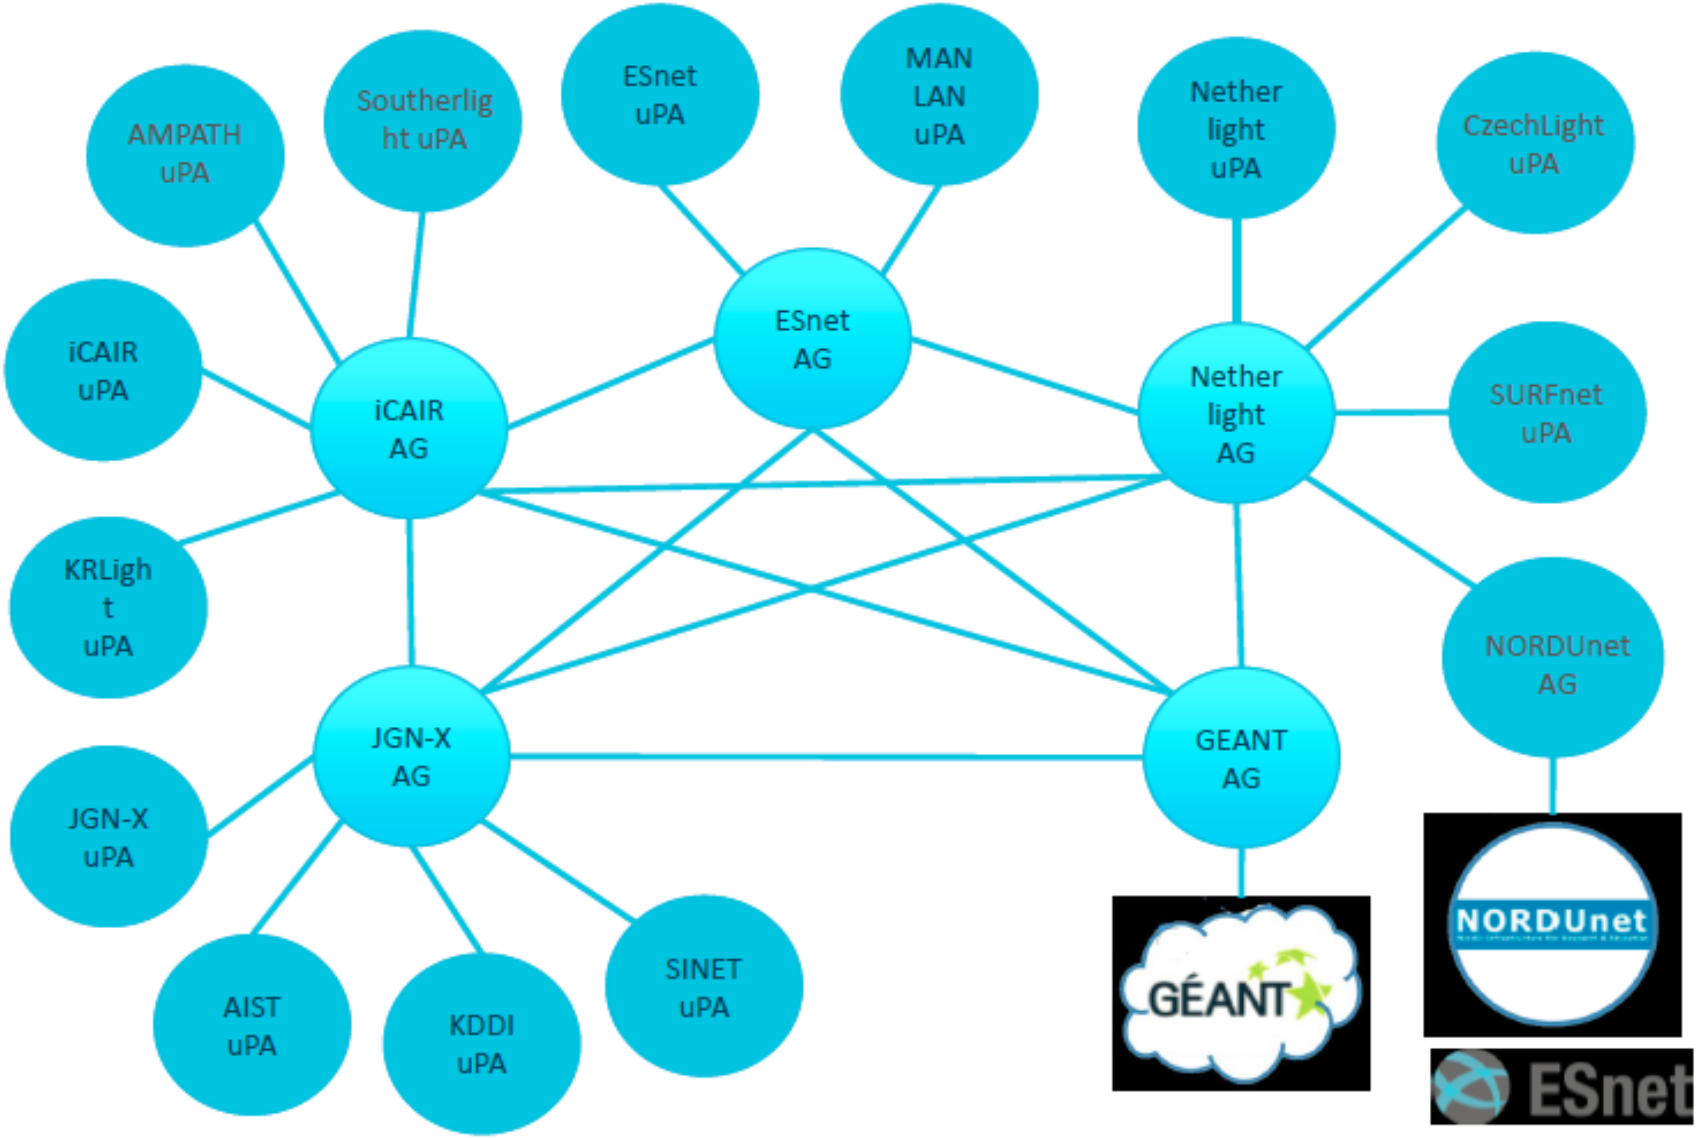

AutoGOLE Topology SC'13

Source: John MacAuley

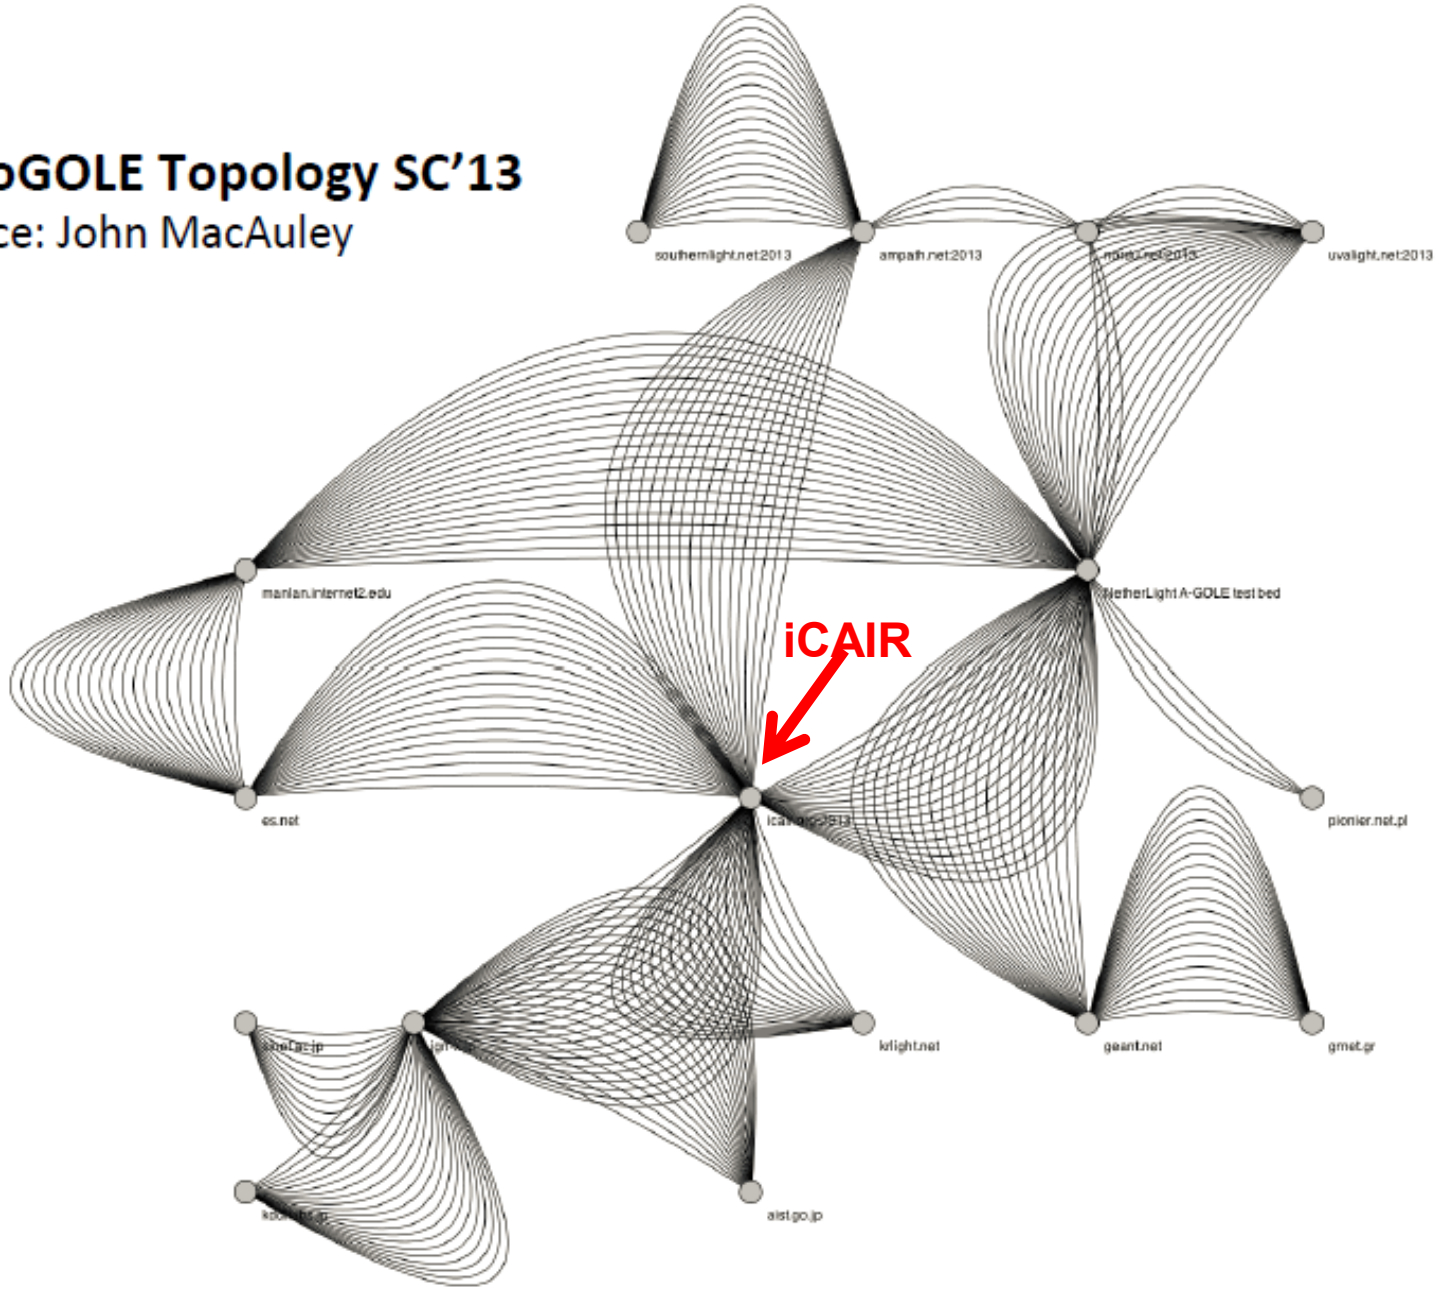

Tasks/Goals For 2014

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

The iGENI SDN/OpenFlow Consortium Uses The Global Lambda Integrated Facility

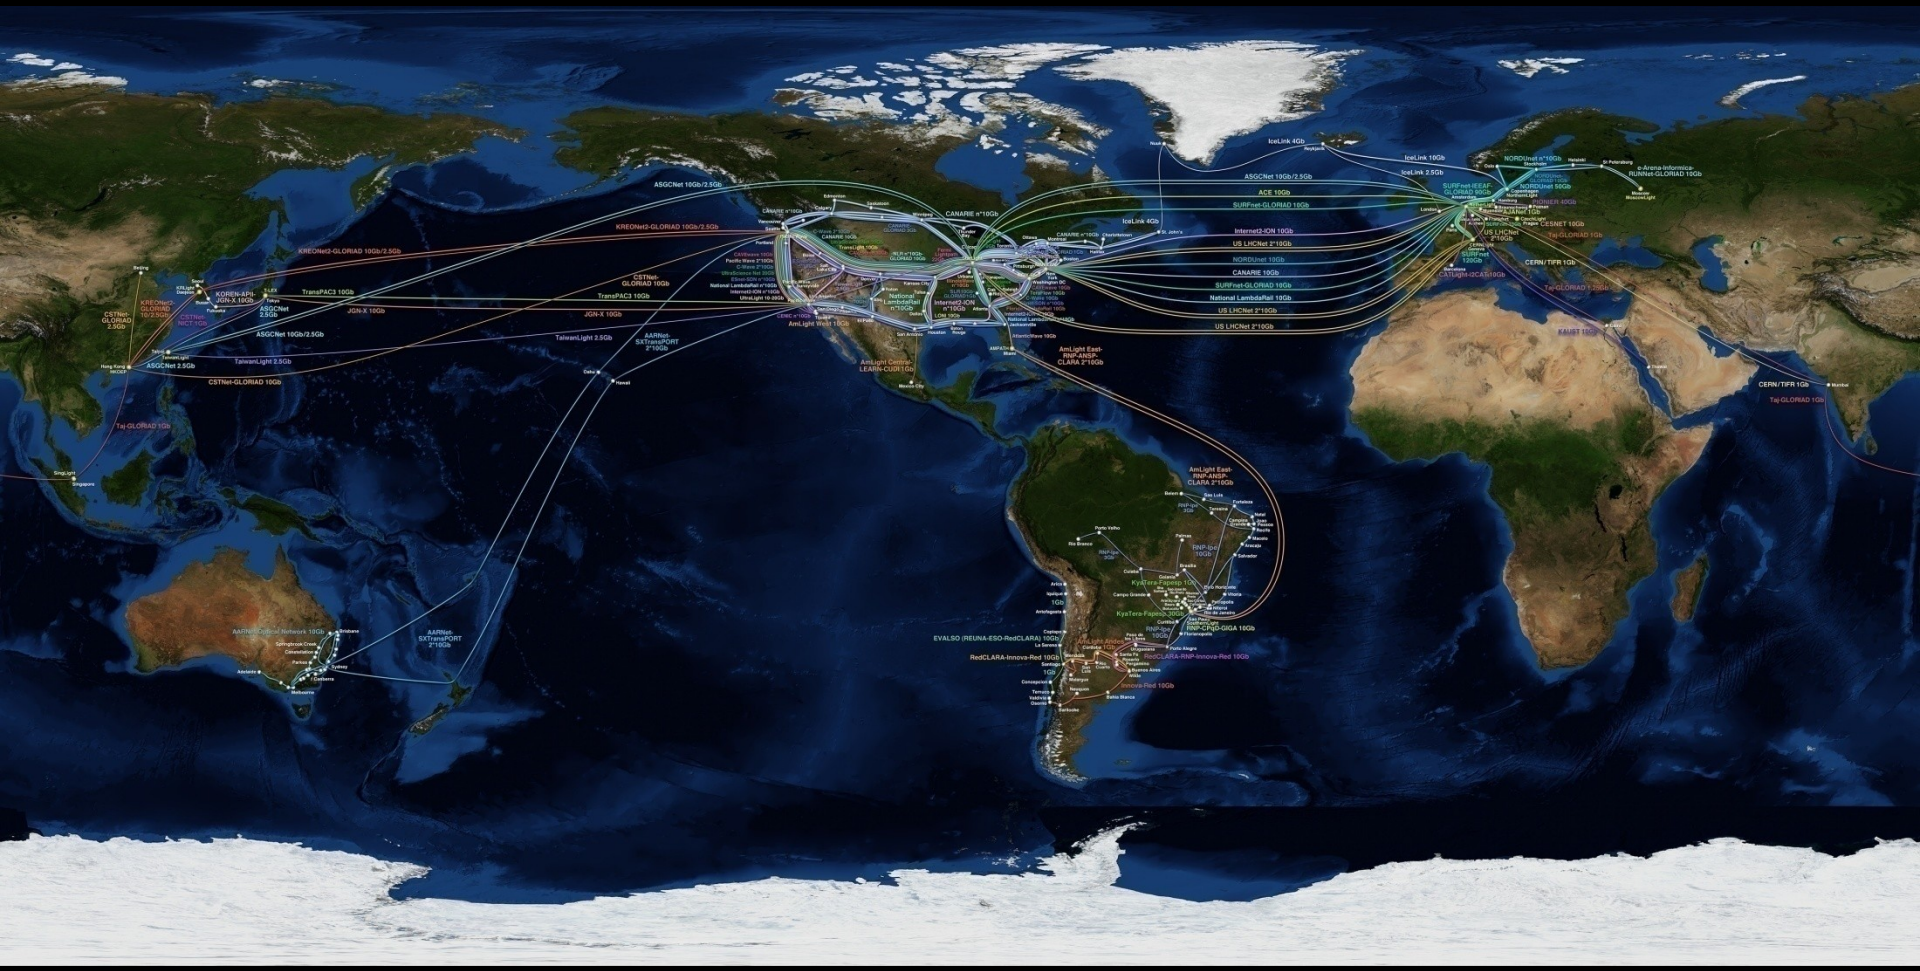


KUAS
NetFPGA
NOX
OpenFlow
NCKU
NetFPGA
NOX
OpenFlow
ETRI
NOX
OpenFlow
NetFPGA
KiSTi
University of Victoria
TOMATO
OpenFlow
InstaGENI
CRC
NOX
OpenFlow
NetFPGA
cenic
InstaGENI
OpenFlow
JGN-X
NSI
東京大学
THE UNIVERSITY OF TOKYO
VNode
PACIFIC NORTHWEST GIGAPOP
aarnet
NICTA
ExoGENI
OpenFlow
NCHC
NetFPGA
NOX
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

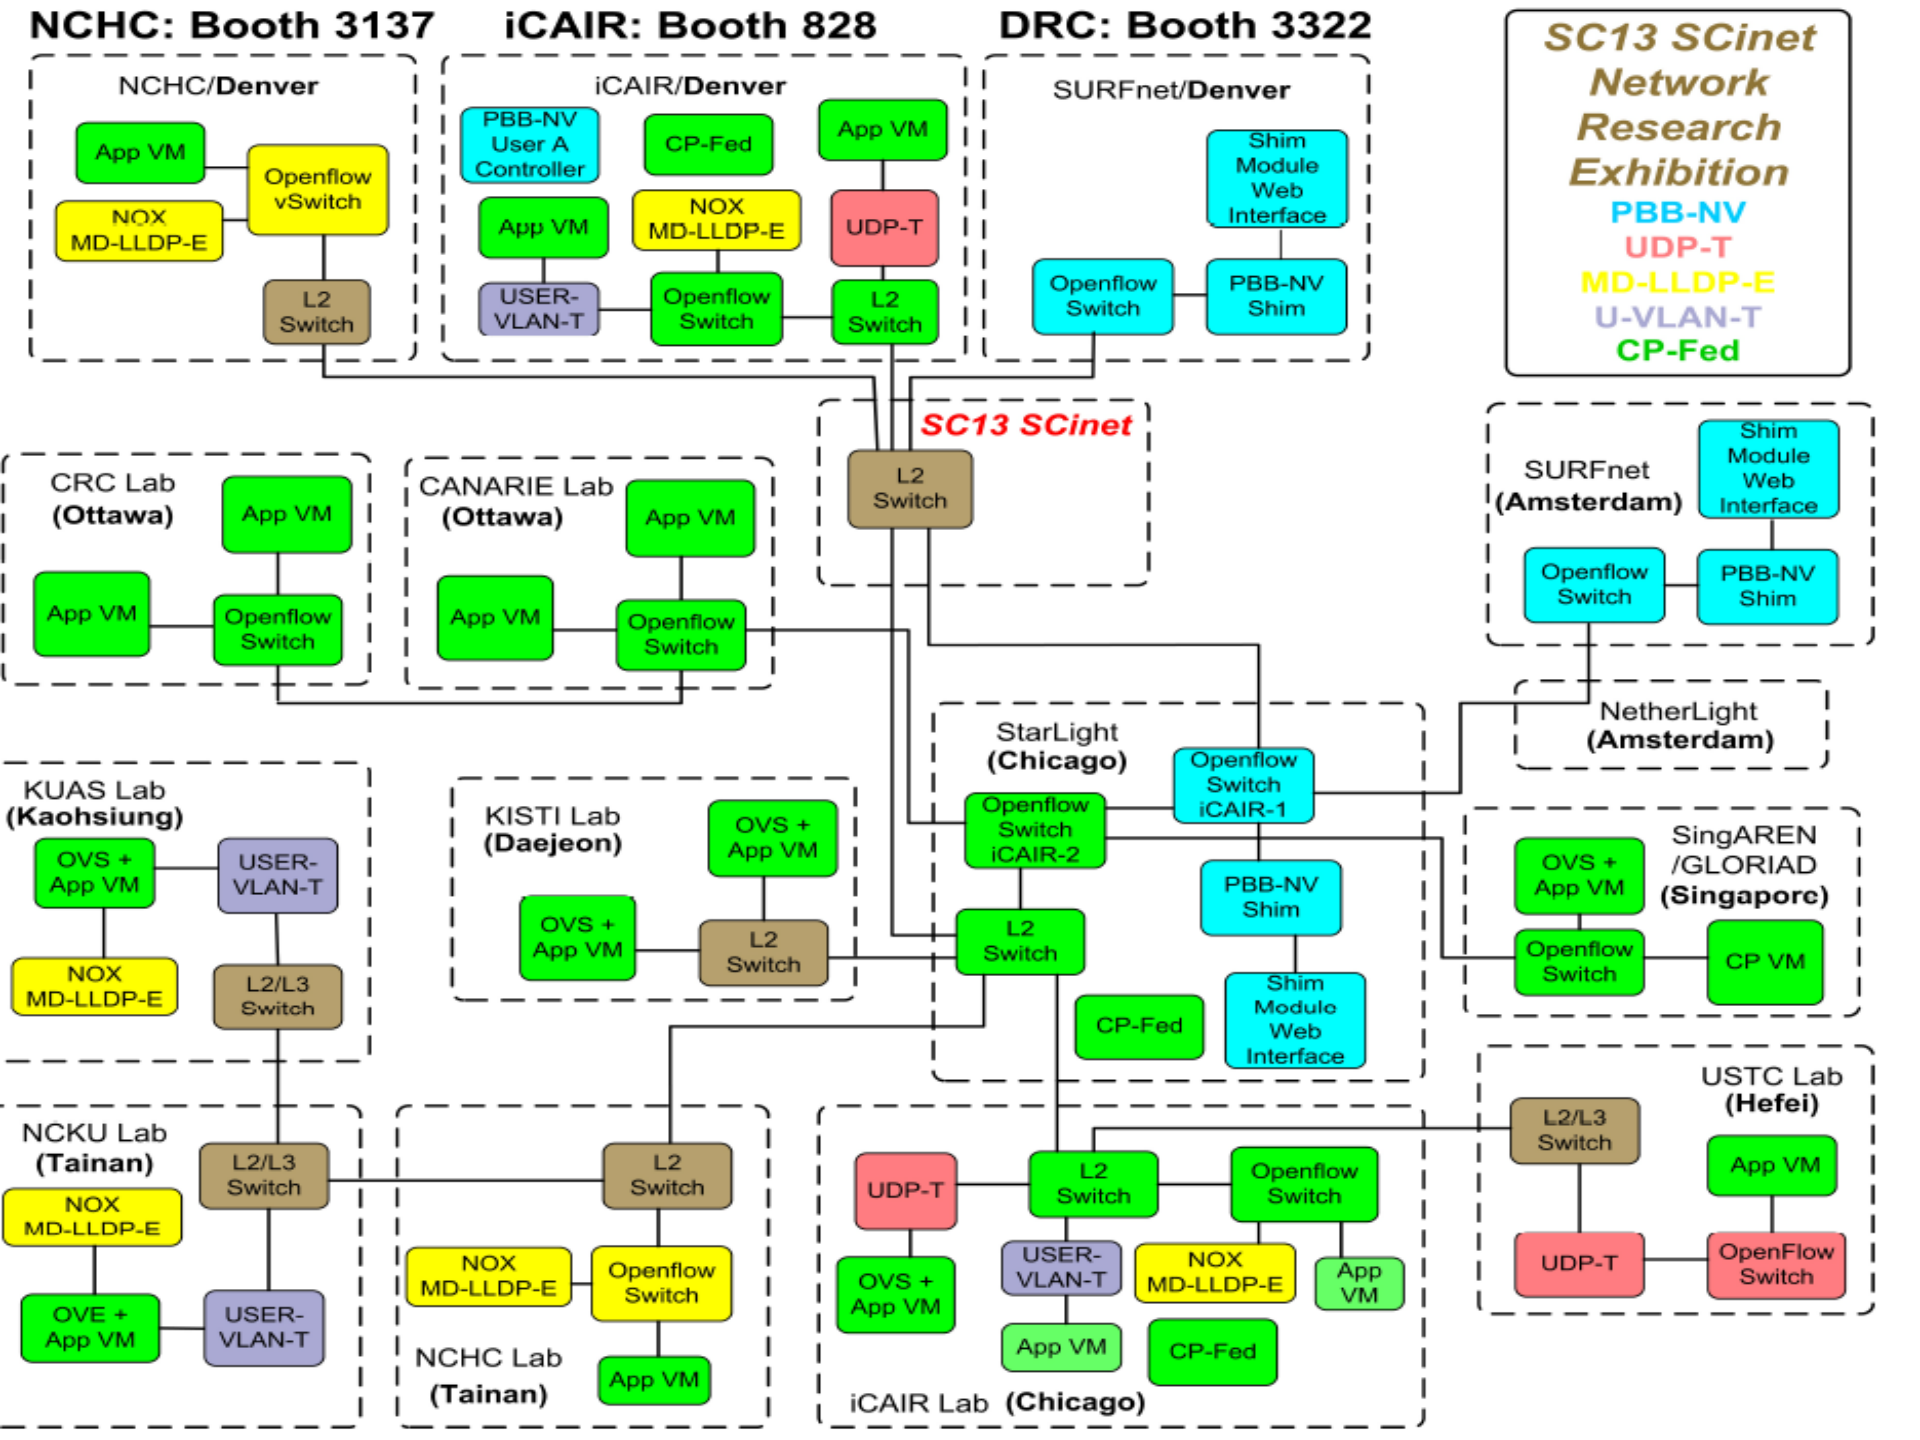
NCHC: Booth 3137
iCAIR: Booth 828
DRC: Booth 3322
SC13 SCinet Network Research Exhibition
PBB-NV
UDP-T
MD-LLDP-E
U-VLAN-T
CP-Fed
NCHC/Denver
App VM
Openflow vSwitch
NOX MD-LLDP-E
L2 Switch
iCAIR/Denver
PBB-NV User A Controller
CP-Fed
App VM
NOX MD-LLDP-E
UDP-T
App VM
USER-VLAN-T
Openflow Switch
L2 Switch
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)

GENI Mesoscale Network: InstaGENI

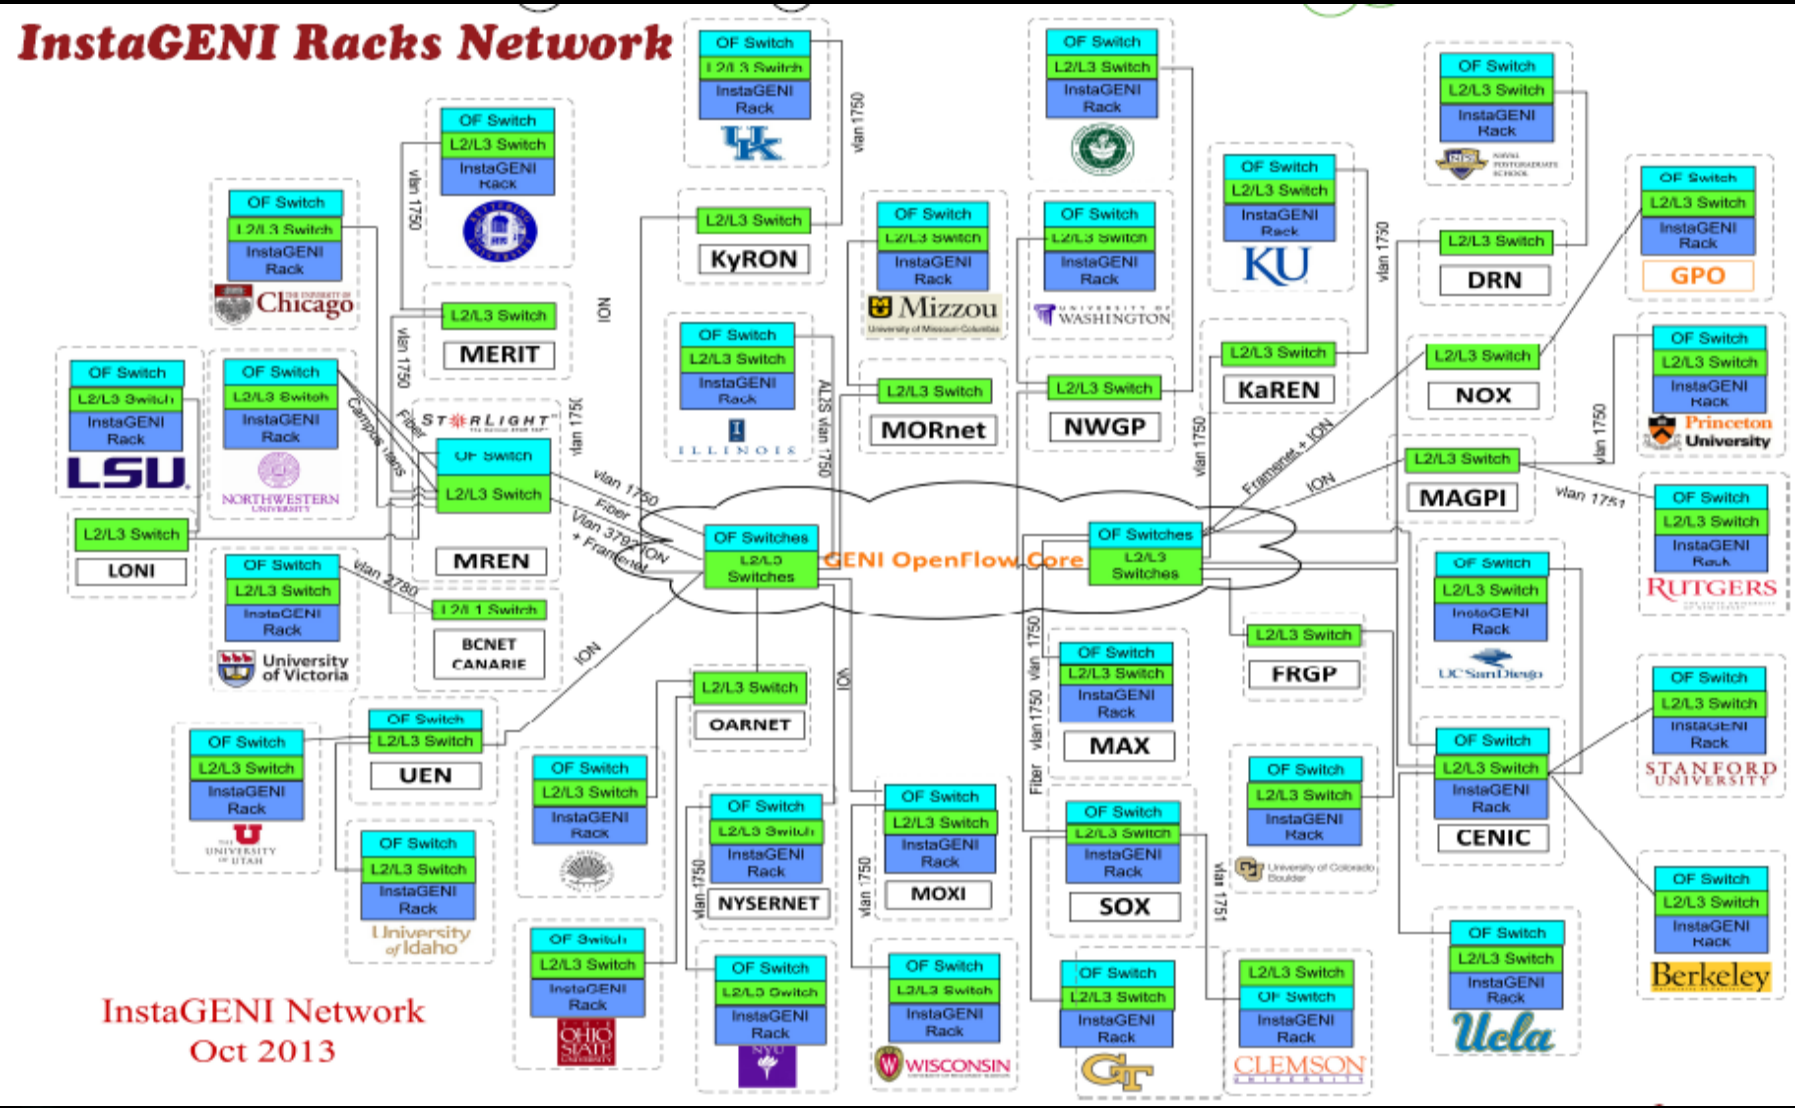

Software Defined Networking Exchanges (SDXs)

- 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
- Required Capabilities for Multi-Domain Distributed SDN Resource Discovery, Signaling Provisioning, Operations, and Fault Detection and Recovery Are Fairly Challenging.
- Nonetheless – Many Motivations Exist for SDXs

Motivations for SDXs

- WH Office of Science and Technology Policy – Large Scale Science Instrumentation
- Large Scale Ultra High Resolution Digital Media Services
- Multi-Domain Networks Interconnecting Data Centers (SDN Is Already in Production Within Large Scale Data Centers)
- Multi-Domain SDN Services
- Providing Capabilities for Edge Control
- Democratization Of Exchange Facilities
- Network Function Virtualization (NFV)
- And Much, Much More!
- ***NB: In the Future ALL Exchanges Will Be SDXs***

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

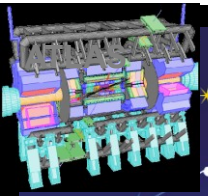

ALMA: Atacama Large Millimeter Array

ANDRILL: Antarctic Geological Drilling
www.andrill.org

BIRN

BIRN: Biomedical Informatics Research Network
www.nbirn.net

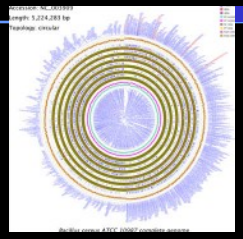

CAMERA metagenomics
camera.calit2.net

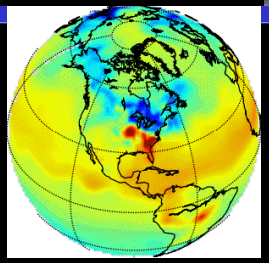

Carbon Tracker
www.esrl.noaa.gov/gmd/ccgg/carbontracker

CineGrid™

CineGrid
www.cinegrid.org

LHCONE
www.lhcone.net

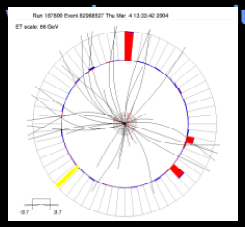

DØ (DZero)
www-d0.fnal.gov

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

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

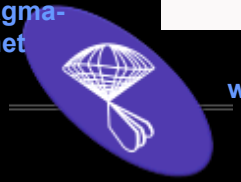

Sloan Digital Sky Survey
www.sdss.org

XSEDE

XSEDE
www.xsede.org

Compilation By Maxine Brown

STARLIGHT℠

Google's Software Defined WAN

Google

Examples from NZ #2: SDX

12 November 2013

CityLink and NoviFlow sign agreement to create the first SDN Based Internet Exchange in the World

CityLink and NoviFlow today announced their agreement for the supply of Software Defined Networking (SDN) equipment, which will be deployed in New Zealand's Internet Exchange Points (IXPs) around the country.

Group Chief Technology Officer Jamie Baddeley said that "...the agreement is the conclusion of an exhaustive set of tests of key SDN switching vendors that has taken place at CityLink over the last 12 months. We have very exacting requirements when it comes to SDN and the IXPs. NoviFlow clearly demonstrated their capability and commitment to very high performance Software Defined Networking and because of that we're innovating the future architecture of Internet Exchange Points with them."

Software Defined Networking Exchanges (SDXs)

- Today, No Production SDX Exists
- However, Currently the International Center for Advanced Internet Research (iCAIR) and Its Research Partners Are Designing and Implementing a Prototype SDX at the StarLight International/National Communications Exchange Facility
- Georgia Tech and SOX Are Prototyping a SDX In Atlanta
- With Support from the National Science Foundation's Global Environment for Network Innovations (GENI) Program
- Others are Being Developed in Many Places, Including the Netherlands, Japan, Canada, the Republic of Korea – and New Zealand
- Ref: GLIF Multi-Country SDX Demonstration
- The StarLight SDX Is a Multi-Domain Service Enabling Federated Controllers To Exchange Signaling and Provisioning Information

Selected SDX Architectural Attributes

- Control and Network Resource APIs
- Multi Domain Integrated Path Controller
- Controller Signaling, Including Edge Signaling
- SDN/OF Multi Layer Traffic Exchange
- Multi Domain Resource Advertisement/Discovery
- Topology Exchange
- Multiple Service Levels At All Layers
- Granulated Resource Access (Policy Based), Including Through Edge Processes
- 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

StarLight – "By Researchers For Researchers"

StarLight is an experimental optical 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

iCAIR

STARLIGHT℠

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 Demonstrations

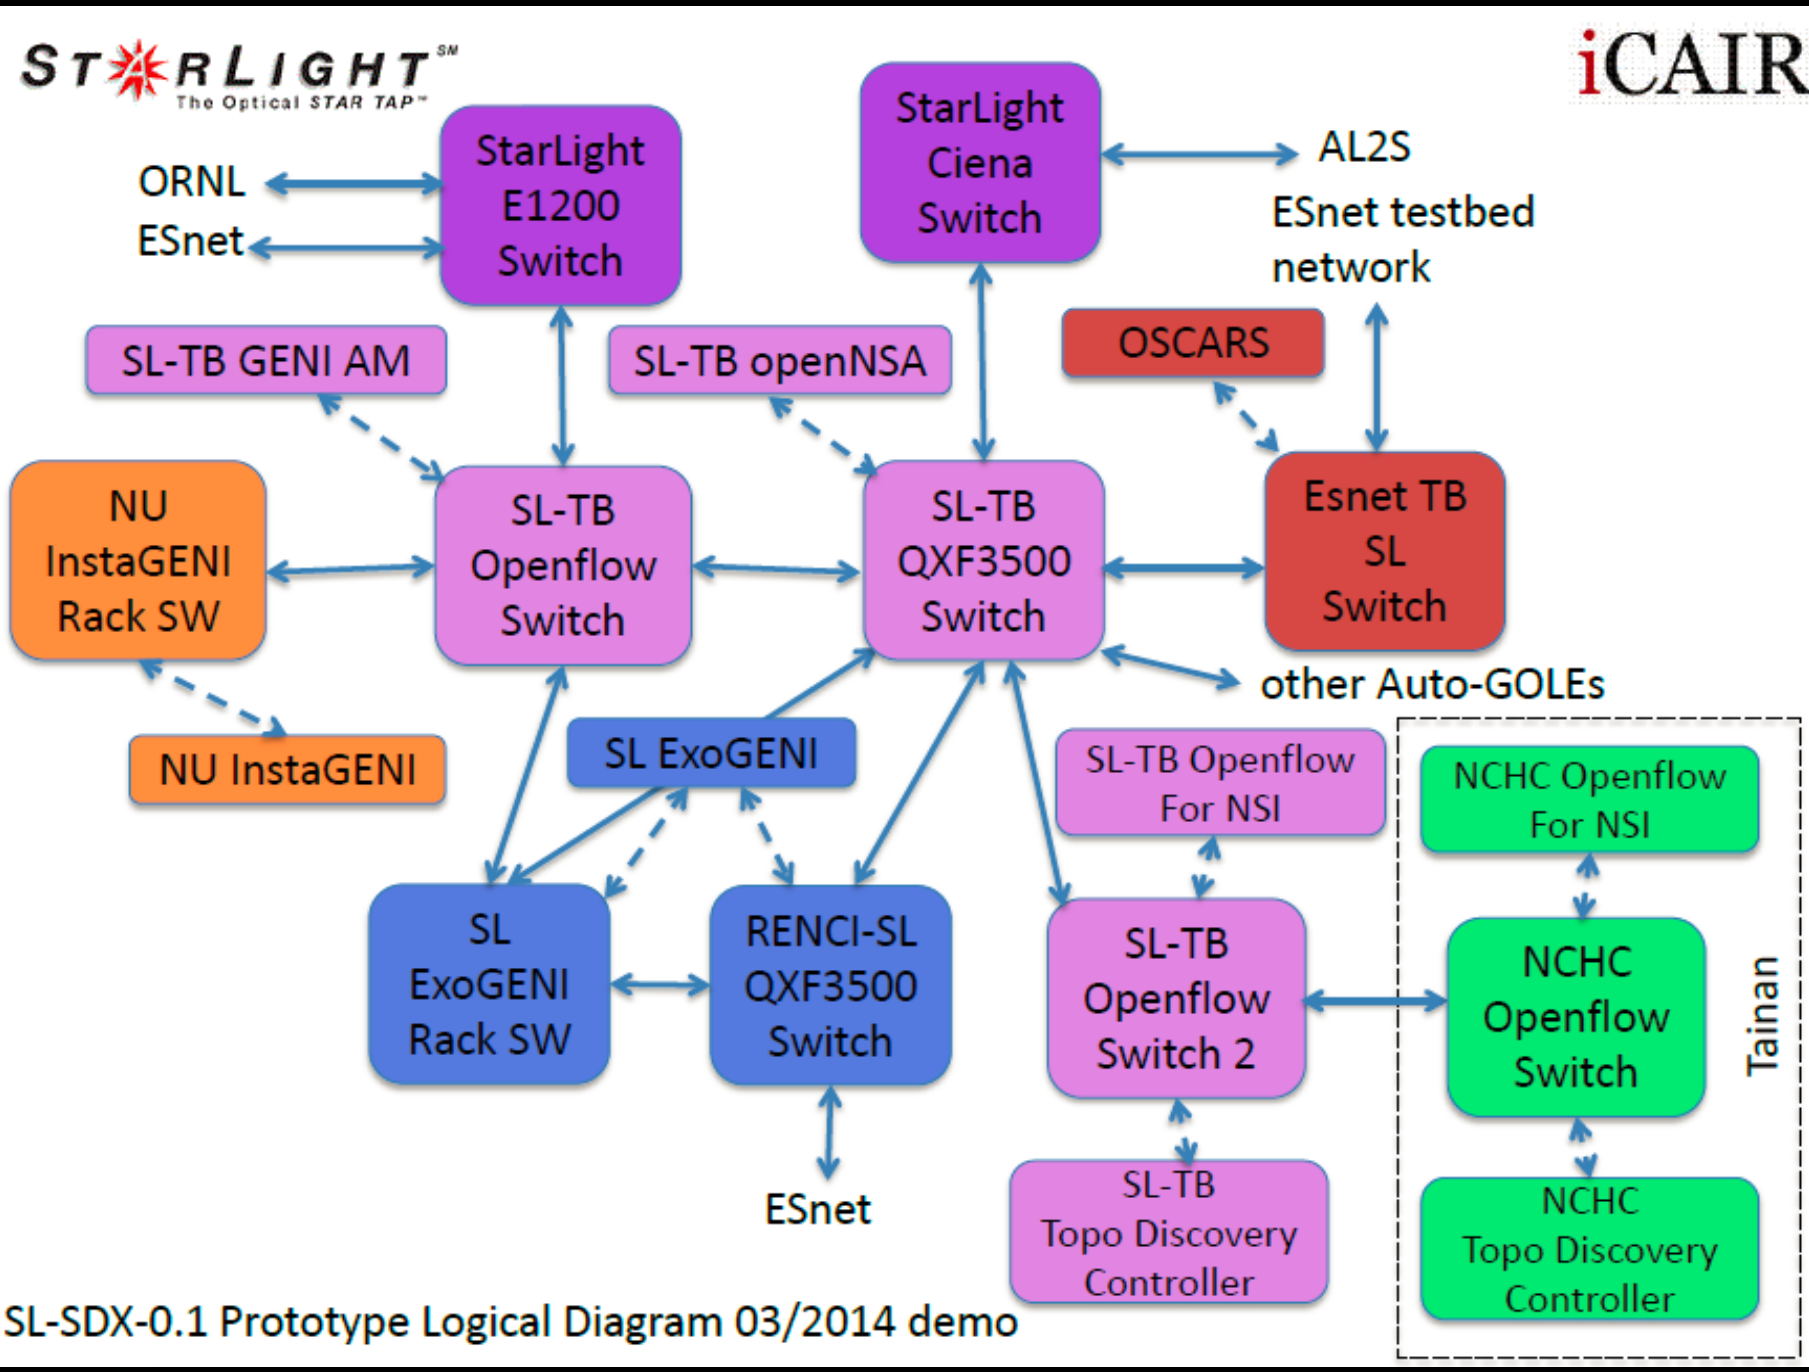
STARLIGHT
The Optical STAR TAP
iCAIR
ORNL
ESnet
StarLight E1200 Switch
StarLight Ciena Switch
AL2S
ESnet testbed network
SL-TB GENI AM
SL-TB openNSA
OSCARS
NU InstaGENI Rack SW
SL-TB Openflow Switch
SL-TB QXF3500 Switch
Esnet TB SL Switch
other Auto-GOLEs
NU InstaGENI
SL ExoGENI
SL-TB Openflow For NSI
NCHC Openflow For NSI
SL ExoGENI Rack SW
RENCI-SL QXF3500 Switch
SL-TB Openflow Switch 2
NCHC Openflow Switch
Tainan
ESnet
SL-TB Topo Discovery Controller
NCHC Topo Discovery Controller
SL-SDX-0.1 Prototype Logical Diagram 03/2014 demo

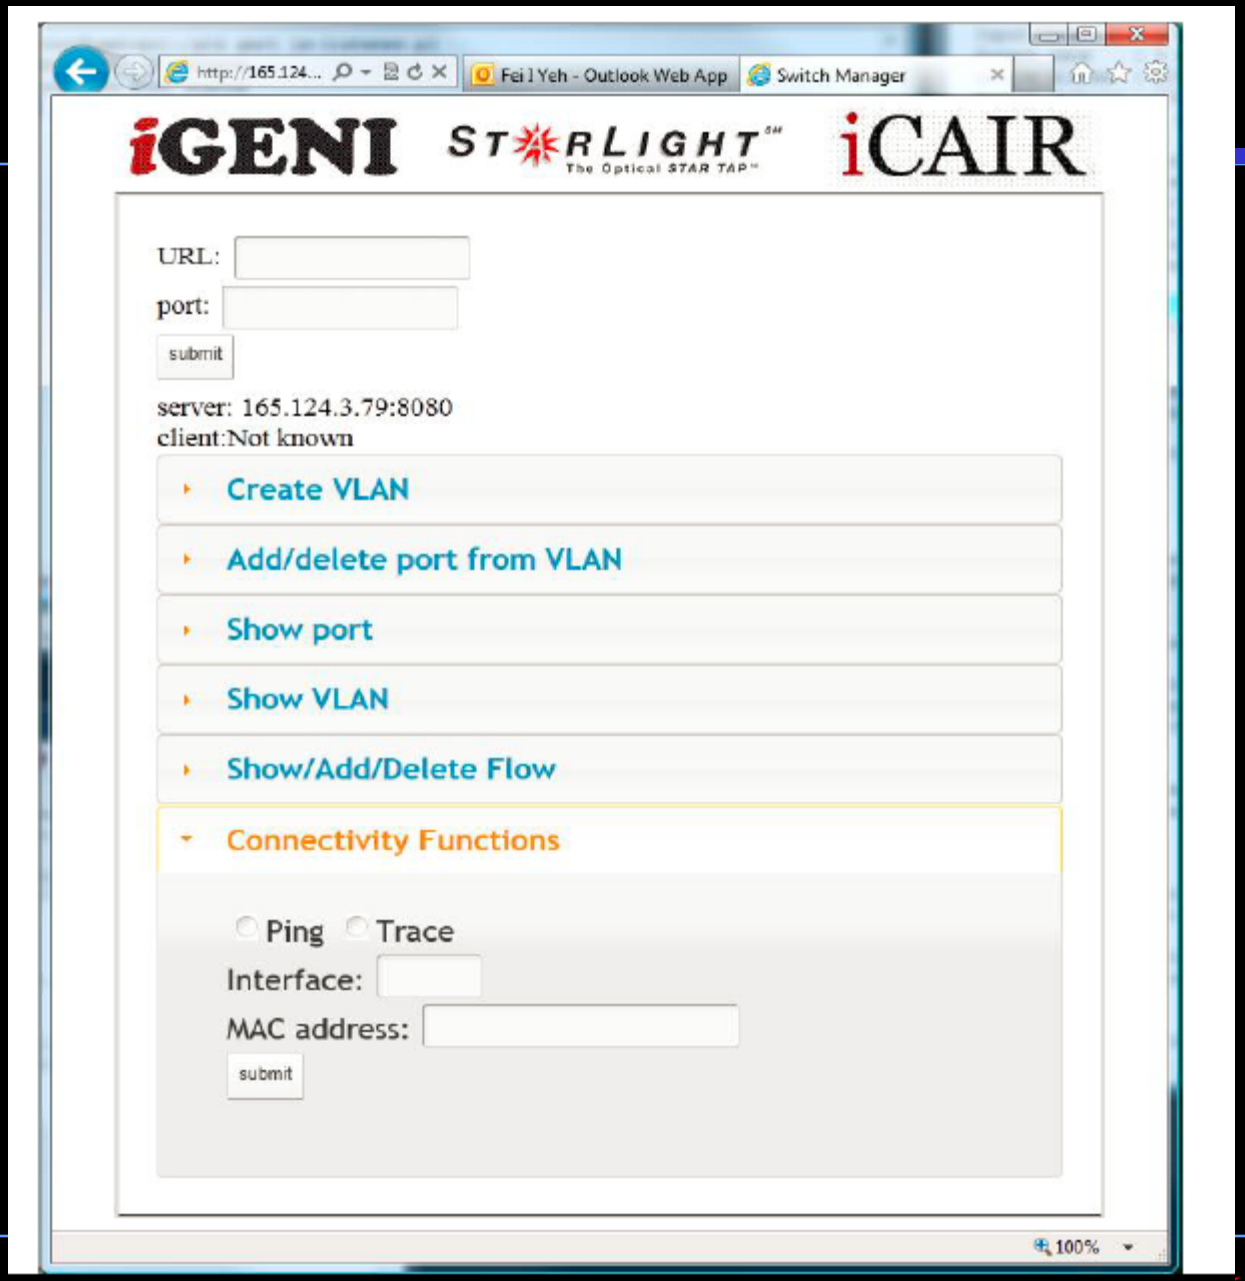
http://165.124...
Fei Yeh - Outlook Web App
Switch Manager
iGENI
STARLIGHT
The Optical STAR TAP
iCAIR
URL:
port:
submit
server: 165.124.3.79:8080
client:Not known
Create VLAN
Add/delete port from VLAN
Show port
Show VLAN
Show/Add/Delete Flow
Connectivity Functions
Ping
Trace
Interface:
MAC address:
submit
100%

GENI Engineering Conference (GEC 19) SDX Demonstration Atlanta March 18-20, 2014

- Initial SDX Capability Between GENI Sites (StarLight and SOX) Was Demonstrated
- Motivation: To Share a Vision of Interconnected US Nationwide SDN Infrastructure, With Multiple SDN Capable Networks and Domains
- SDX Benefits Was Showcased Through a Compelling Application – Nowcast – Developed By Mike Zink et al at University of Massachusetts, Amherst
- (GEC 19 Is Co-Located With the GLIF Tech Workshop, March 19-20)

SDX Demonstration Application: Nowcast

- Radar Data Based Approach For Short-term Weather Prediction
- Short-term: 1-15 Minutes In The Future
- Current Forecasts:
  - Assimilate Data From Many Sensors: Radar, Satellite, Balloons, etc.
  - Usually For Large Regions
  - Takes Super Computers to Calculate

Source: Mike Zink, UMass Amherst

Nowcast – Processing

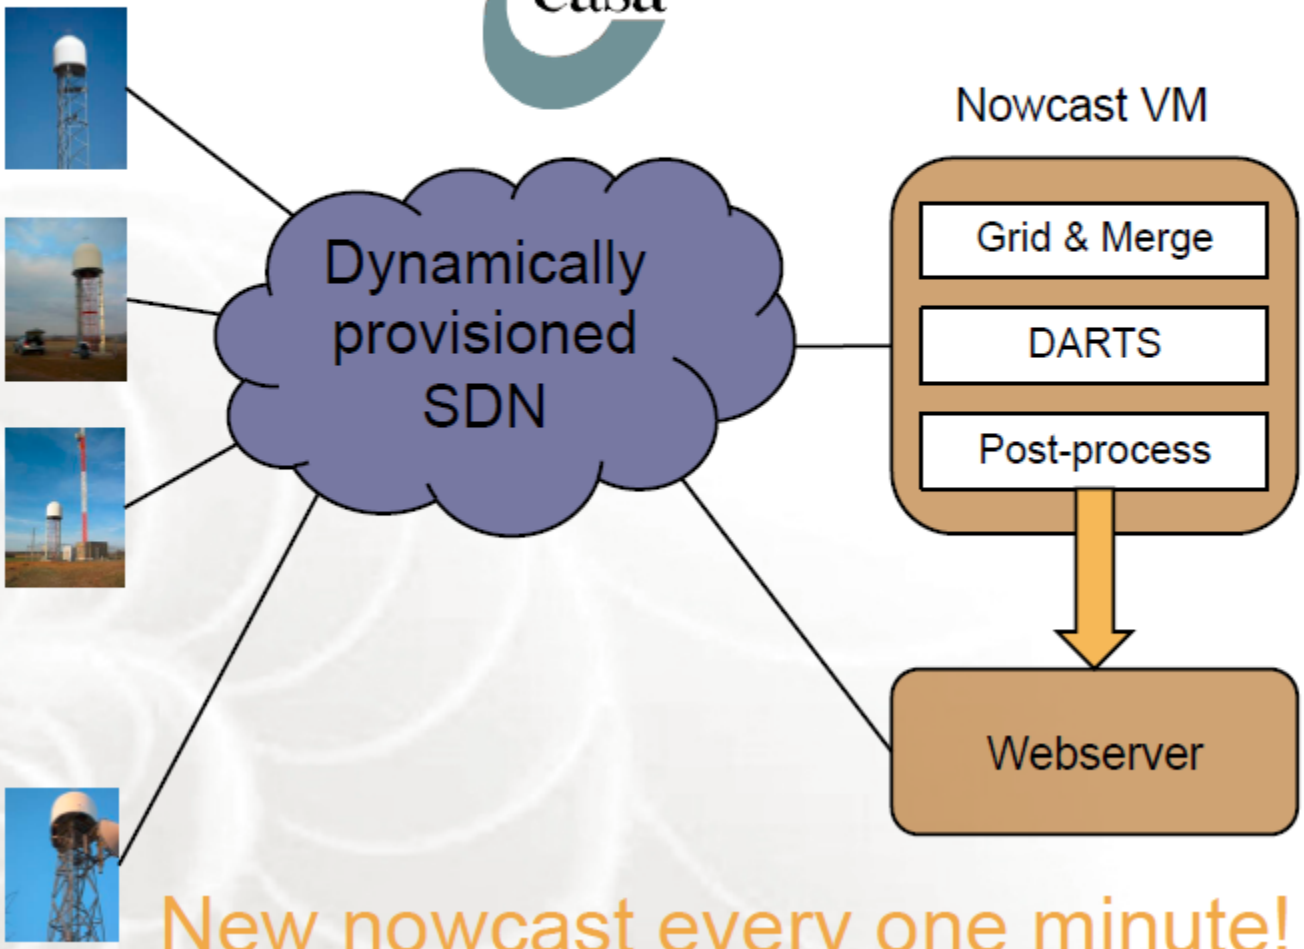

Source: Mike Zink

Nowcast Example

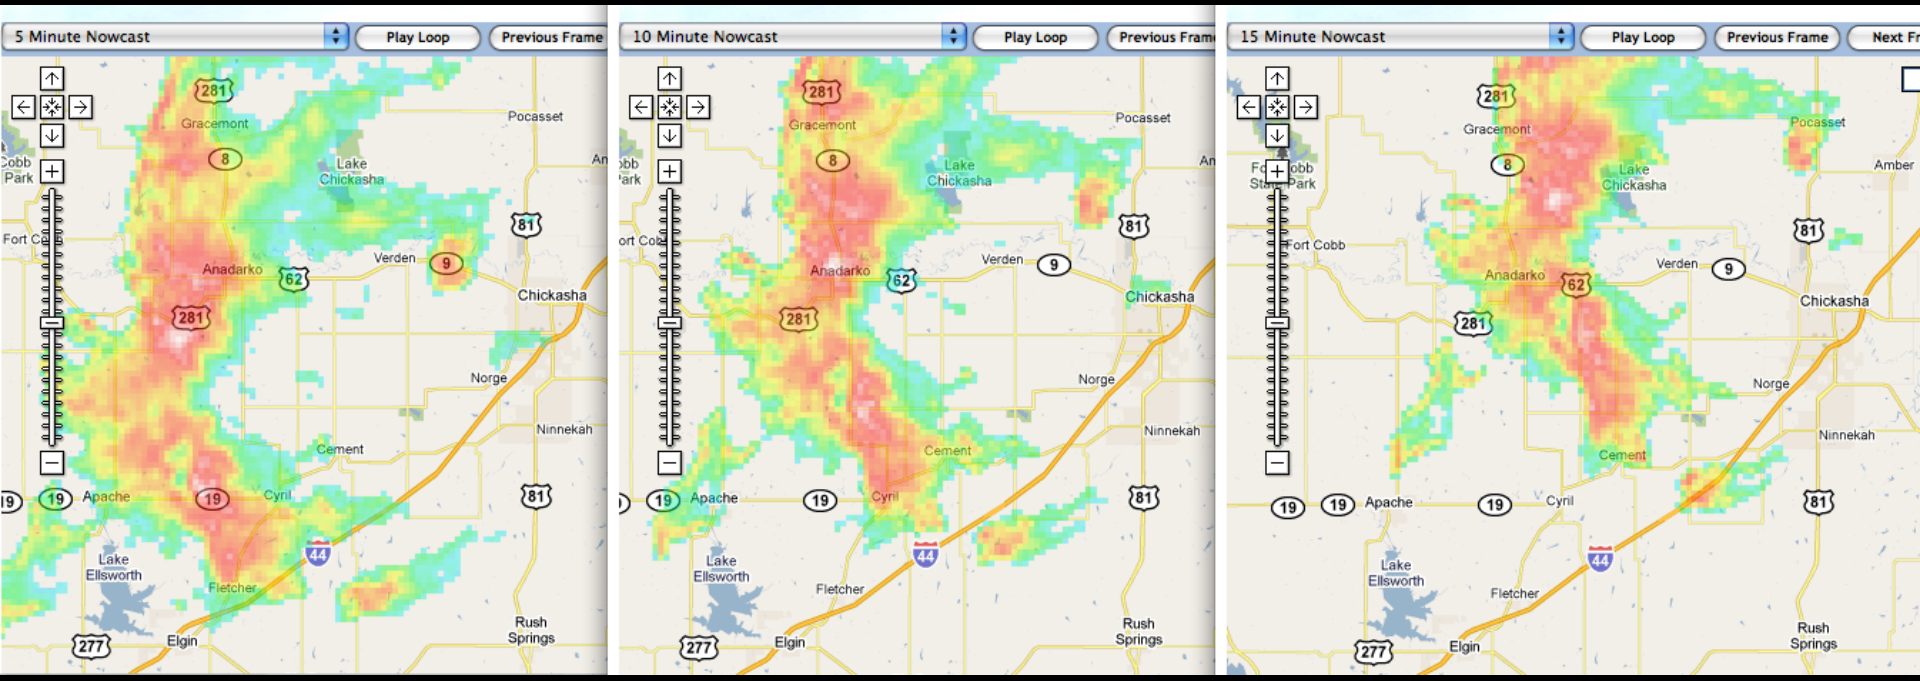

Source: Mike Zink, UMass Amherst

Comparison With Existing System

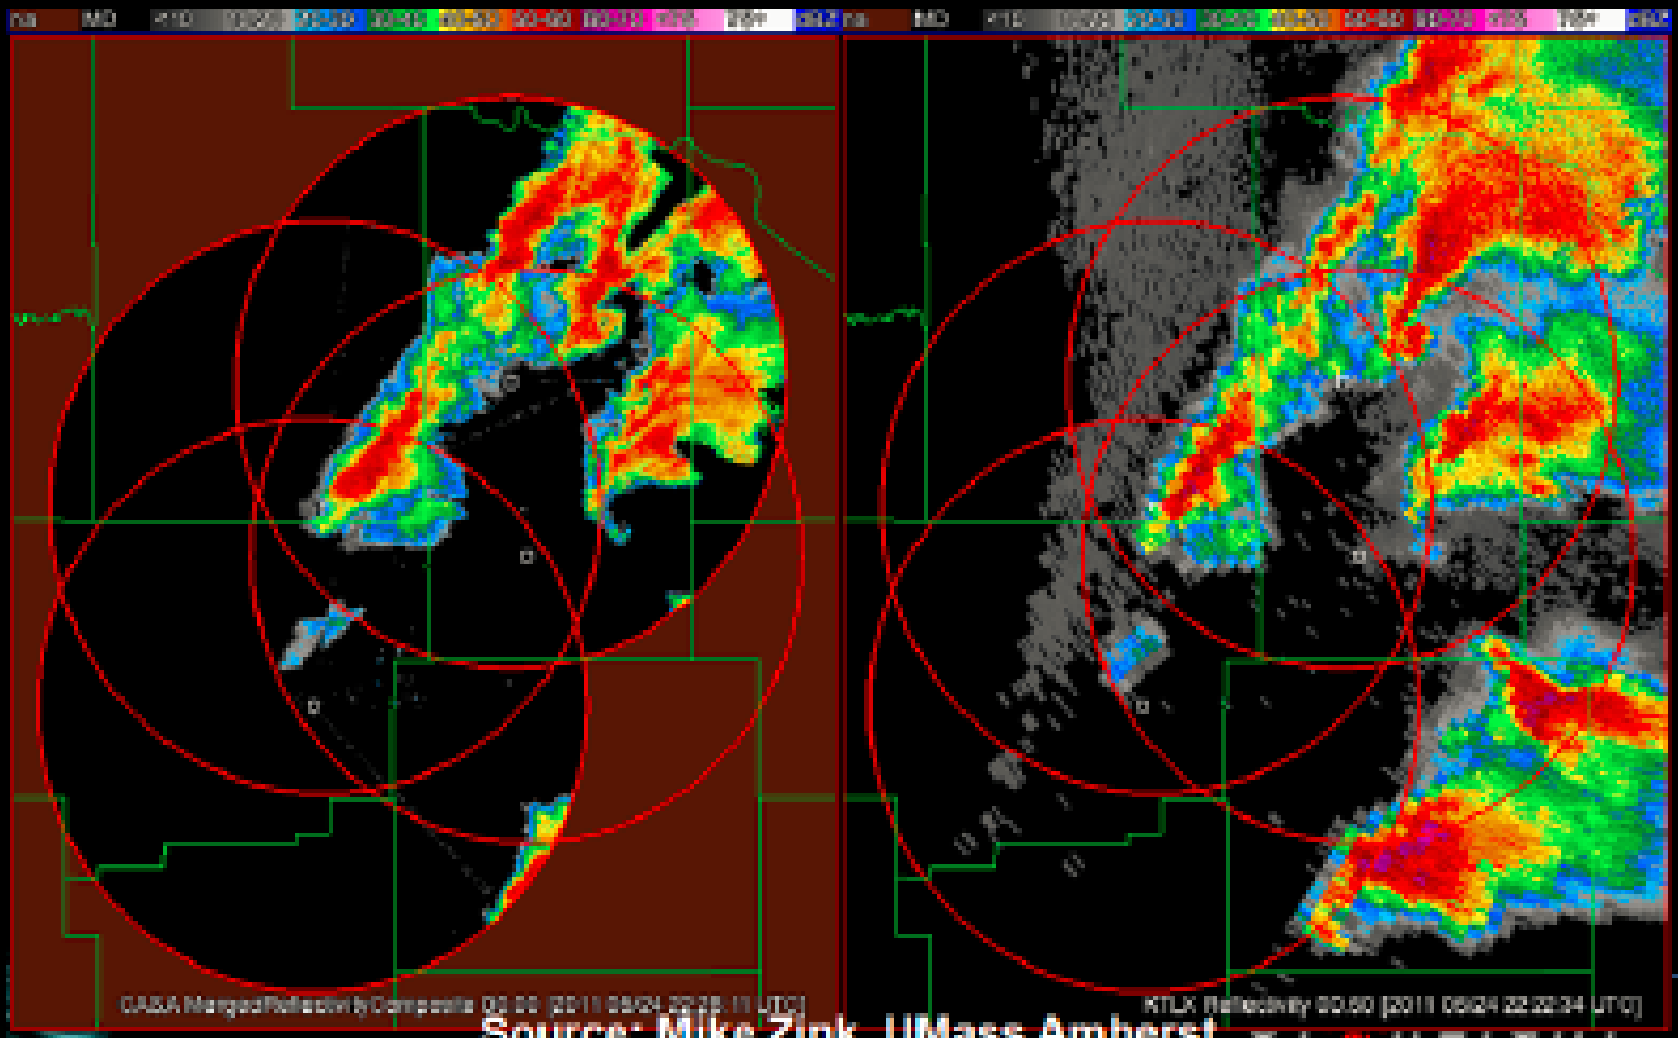

Source: Mike Zink, UMass Amherst

STORMLIGHT℠

Potential

Source: Mike Zink, UMass Amherst

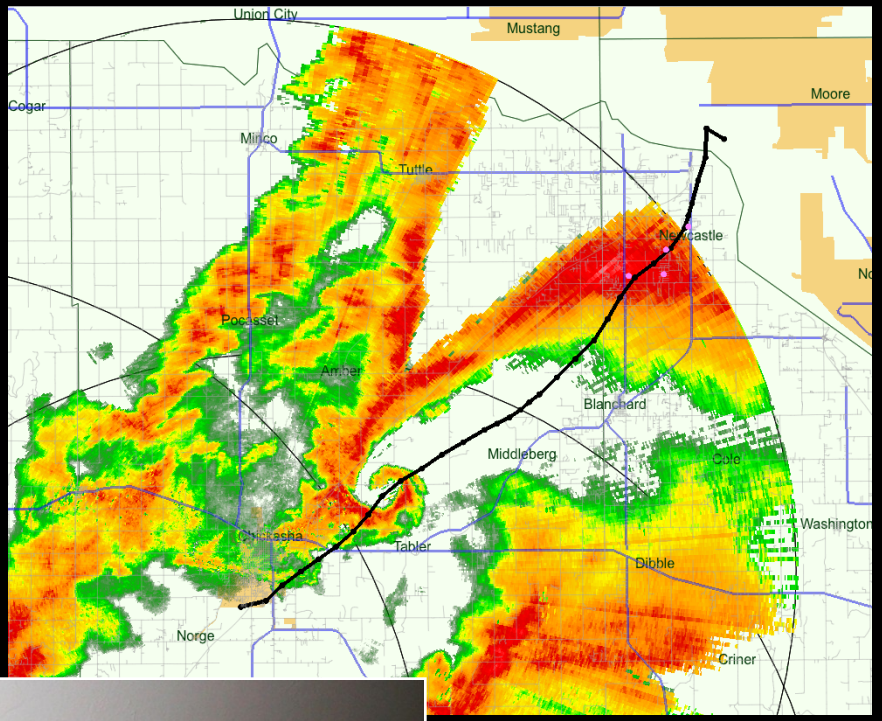

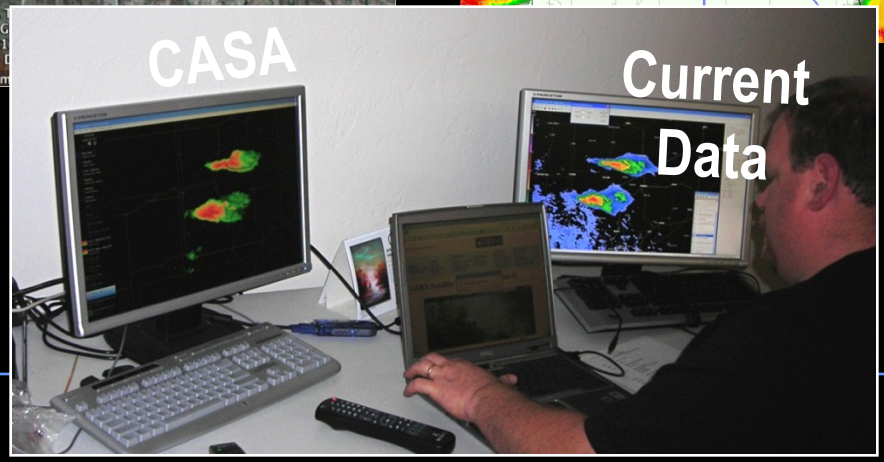

CASA Data, EM Decision-Making Protects First Responders and Public

STORMLIGHT℠

GENI SDX Demo Scenario 1

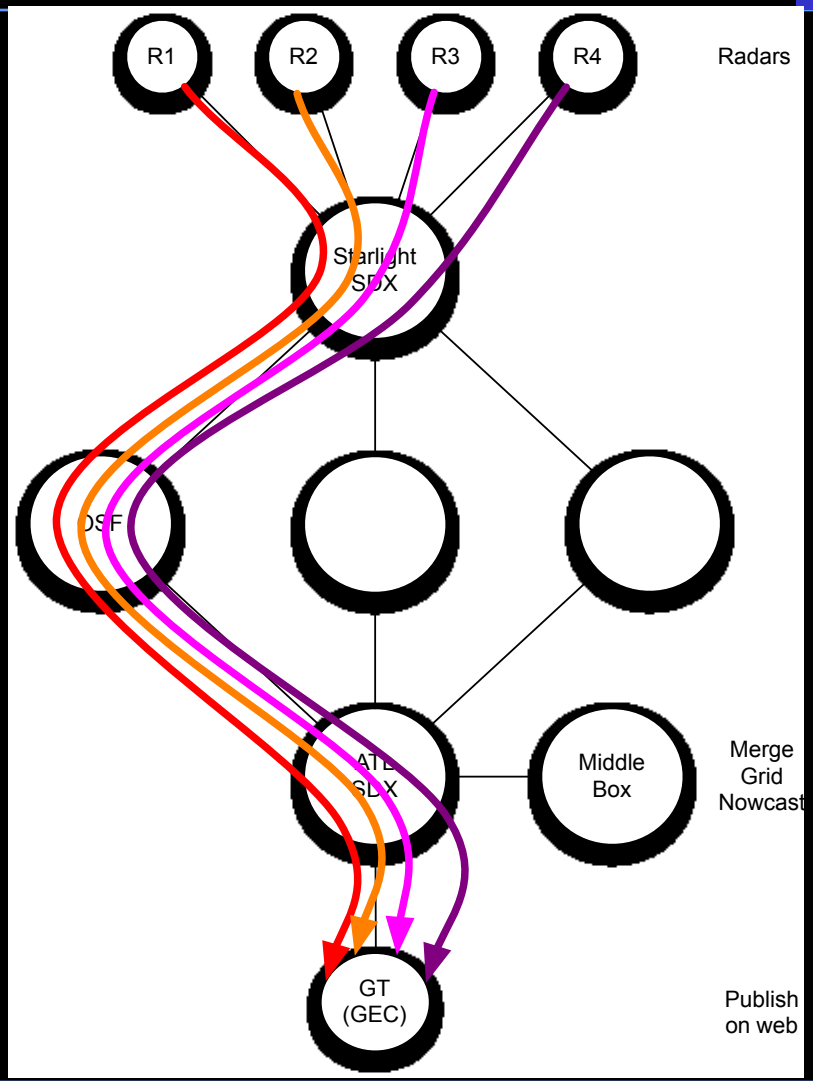

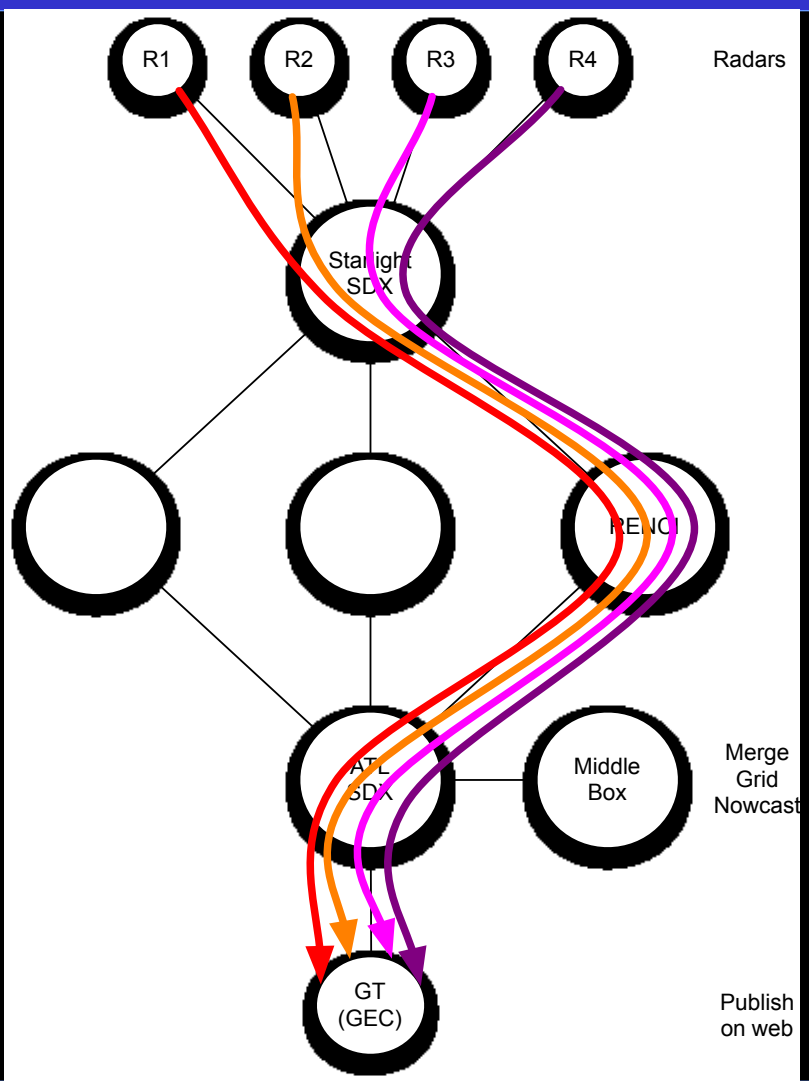

Slide by Mike Zink, UMass Amherst

GENI SDX Demo Scenario 2

Simulated Radar (4)

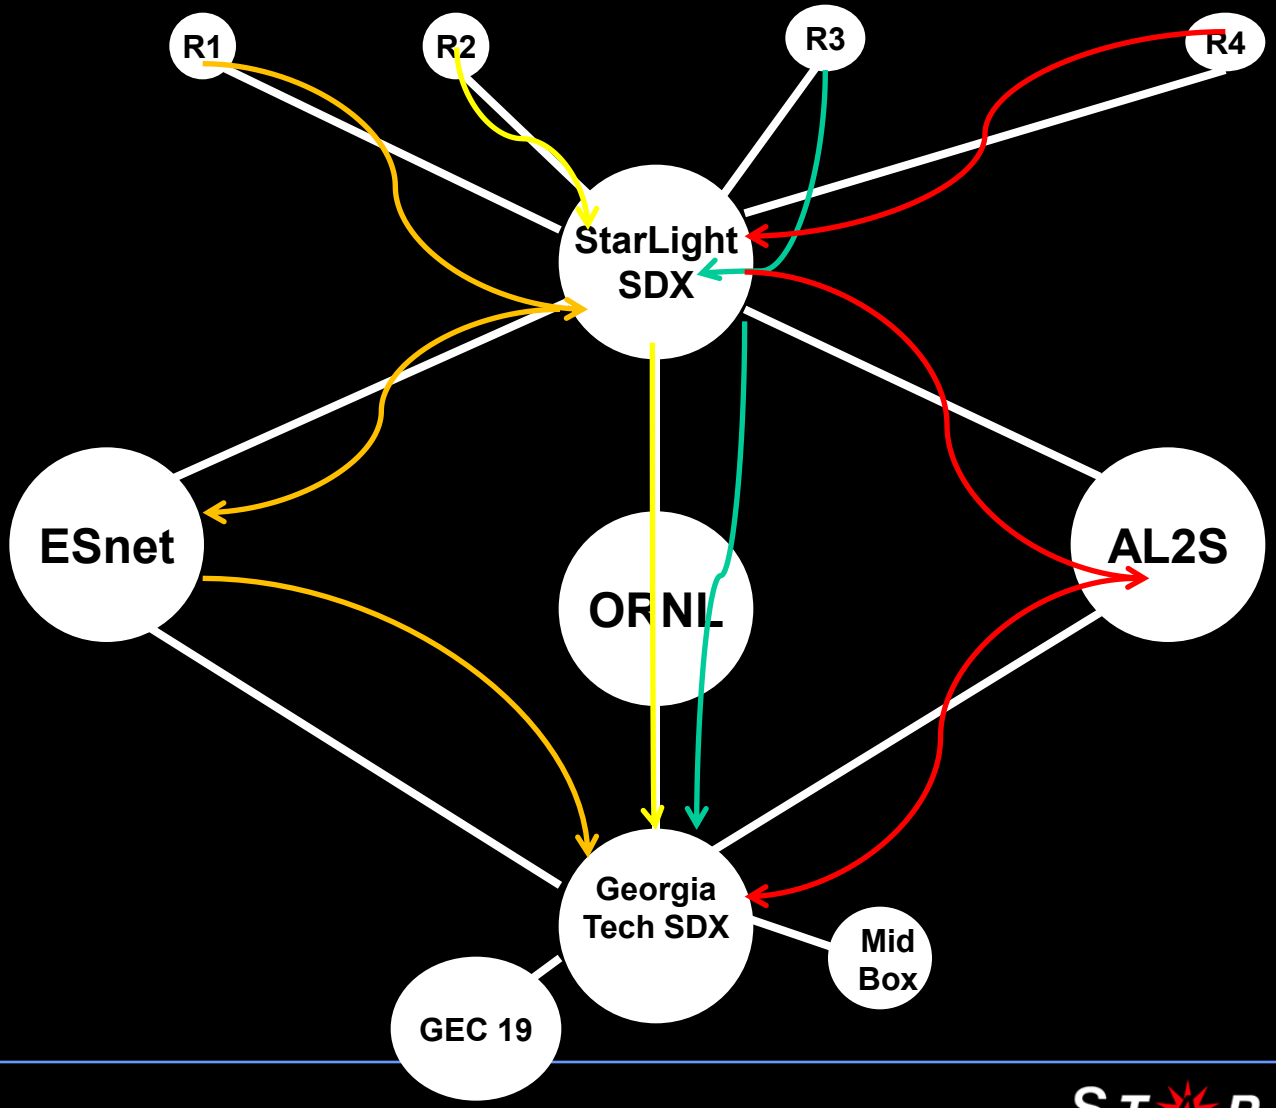

SDX StarLight⇔NetherLight

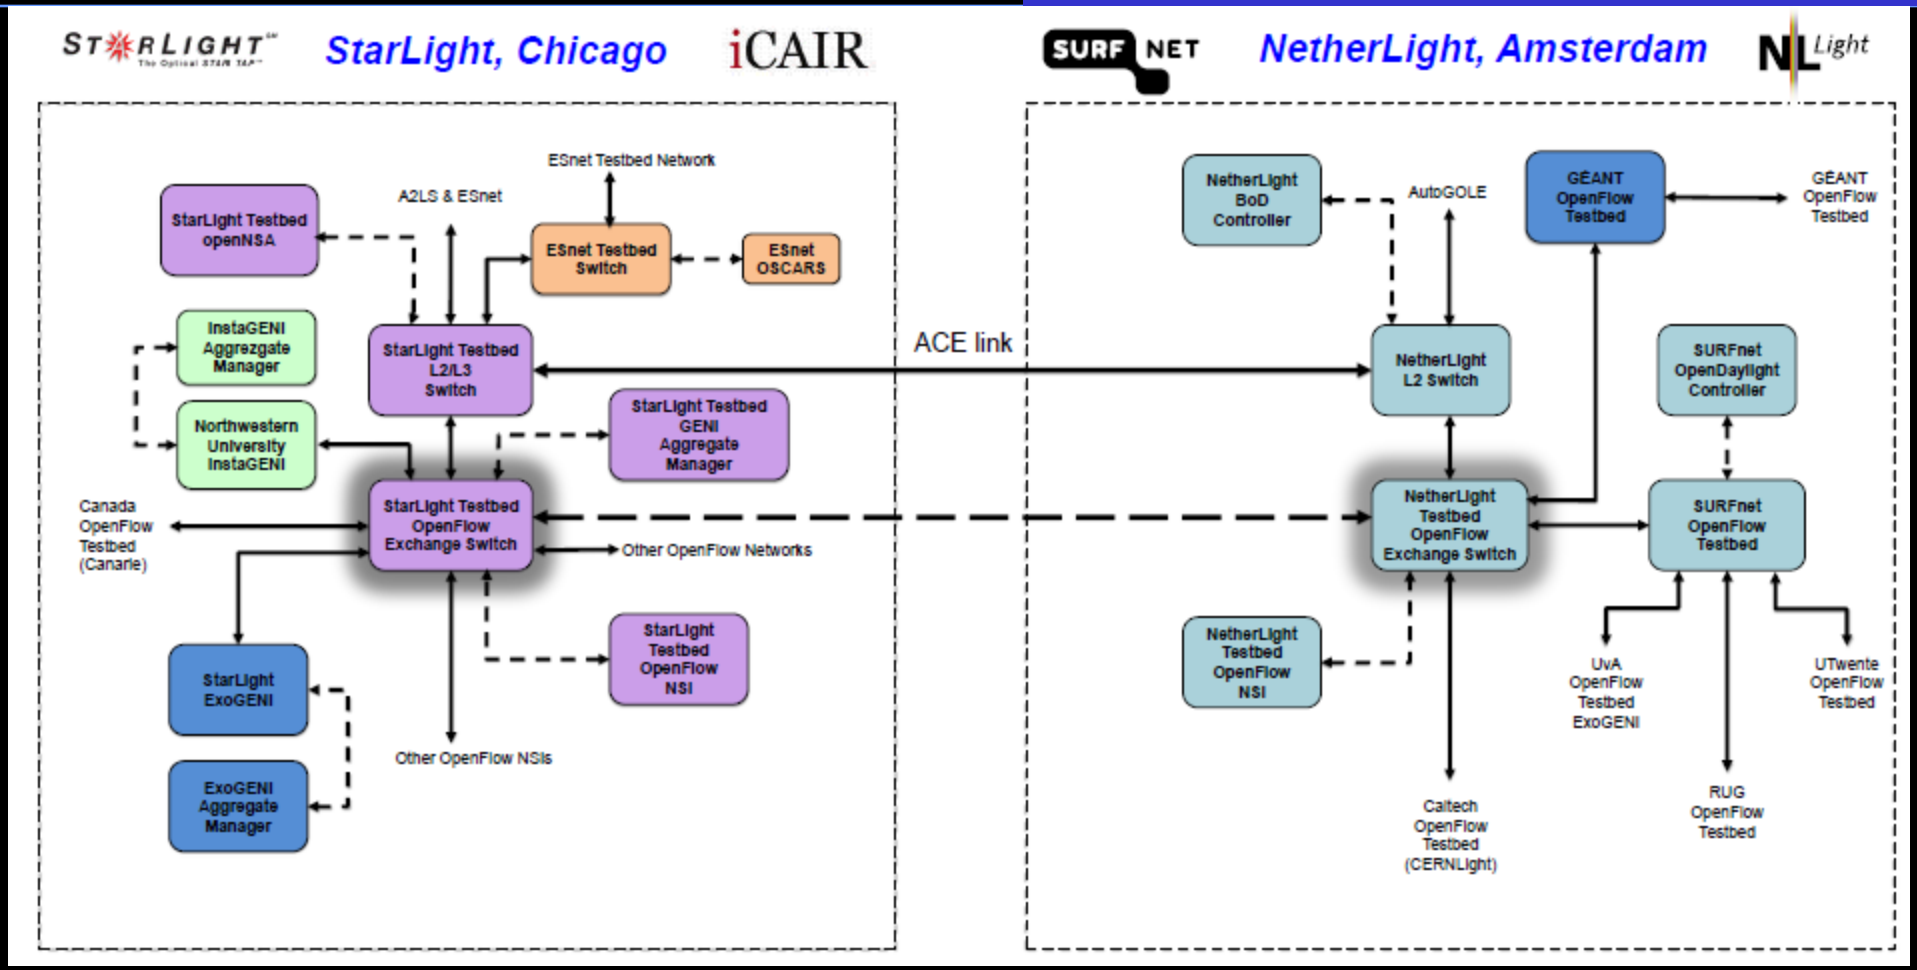

Ronald van der Pol, Joe Mambretti, Jim Chen, John Shillington

STARLIGHT SM

iSDXs @ Global LambdaGrid Workshop

- International SDX (iSDX) Demonstrations Showcase a World-wide Prototype Environment That Could Be Used for Modeling Major Weather Systems, Including The Depiction/Prediction of Severe Weather Patterns.
- Application Based on Nowcast System Being Developed by the NSF Engineering Research Center for Collaborative Adaptive Sensing of the Atmosphere (Led By University of Massachusetts at Amherst), Which Is Being Designed for Next Generation Weather Depiction/Prediction and Visualization Systems.

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℠

International SDX Demonstrations

- Demonstrations of A Prototype Implemented Across the Globe Using the GLIF To Interconnect Sites World-Wide, Supported by Interoperable International SDXs, Including Sites at
  - a) the StarLight International/National Communications Exchange Facility, Designed by iCAIR
  - b) New Zealand, Designed by REANNZ and Google
  - c) Republic of Korea, KISTI/KREONET
  - d) Taiwan, Designed by High Performance Computing Center/TWAREN
  - e) Amsterdam, SDX prototype at NetherLight Designed by SURFnet,
  - f) Ottawa, Designed by CANARIE and Cybera,

    g) Tokyo, Designed by the University of Tokyo

    h) Poznan, Designed By the Poznan Supercomputing and Networking Center

    i) Atlanta, SDX Prototype at SOX, Designed by Georgia Tech
- SDN/OpenFlow Controllers, Federated To Enable Cross-Domain Interoperability.
- Controllers Use Distributed Control Plane To Directly Address and Dynamically Manage Multiple Paths Among Sites Via Distributed Data Plane - Transporting Nowcast Instrumentation Traffic Among Sites.

STARLIGHT℠

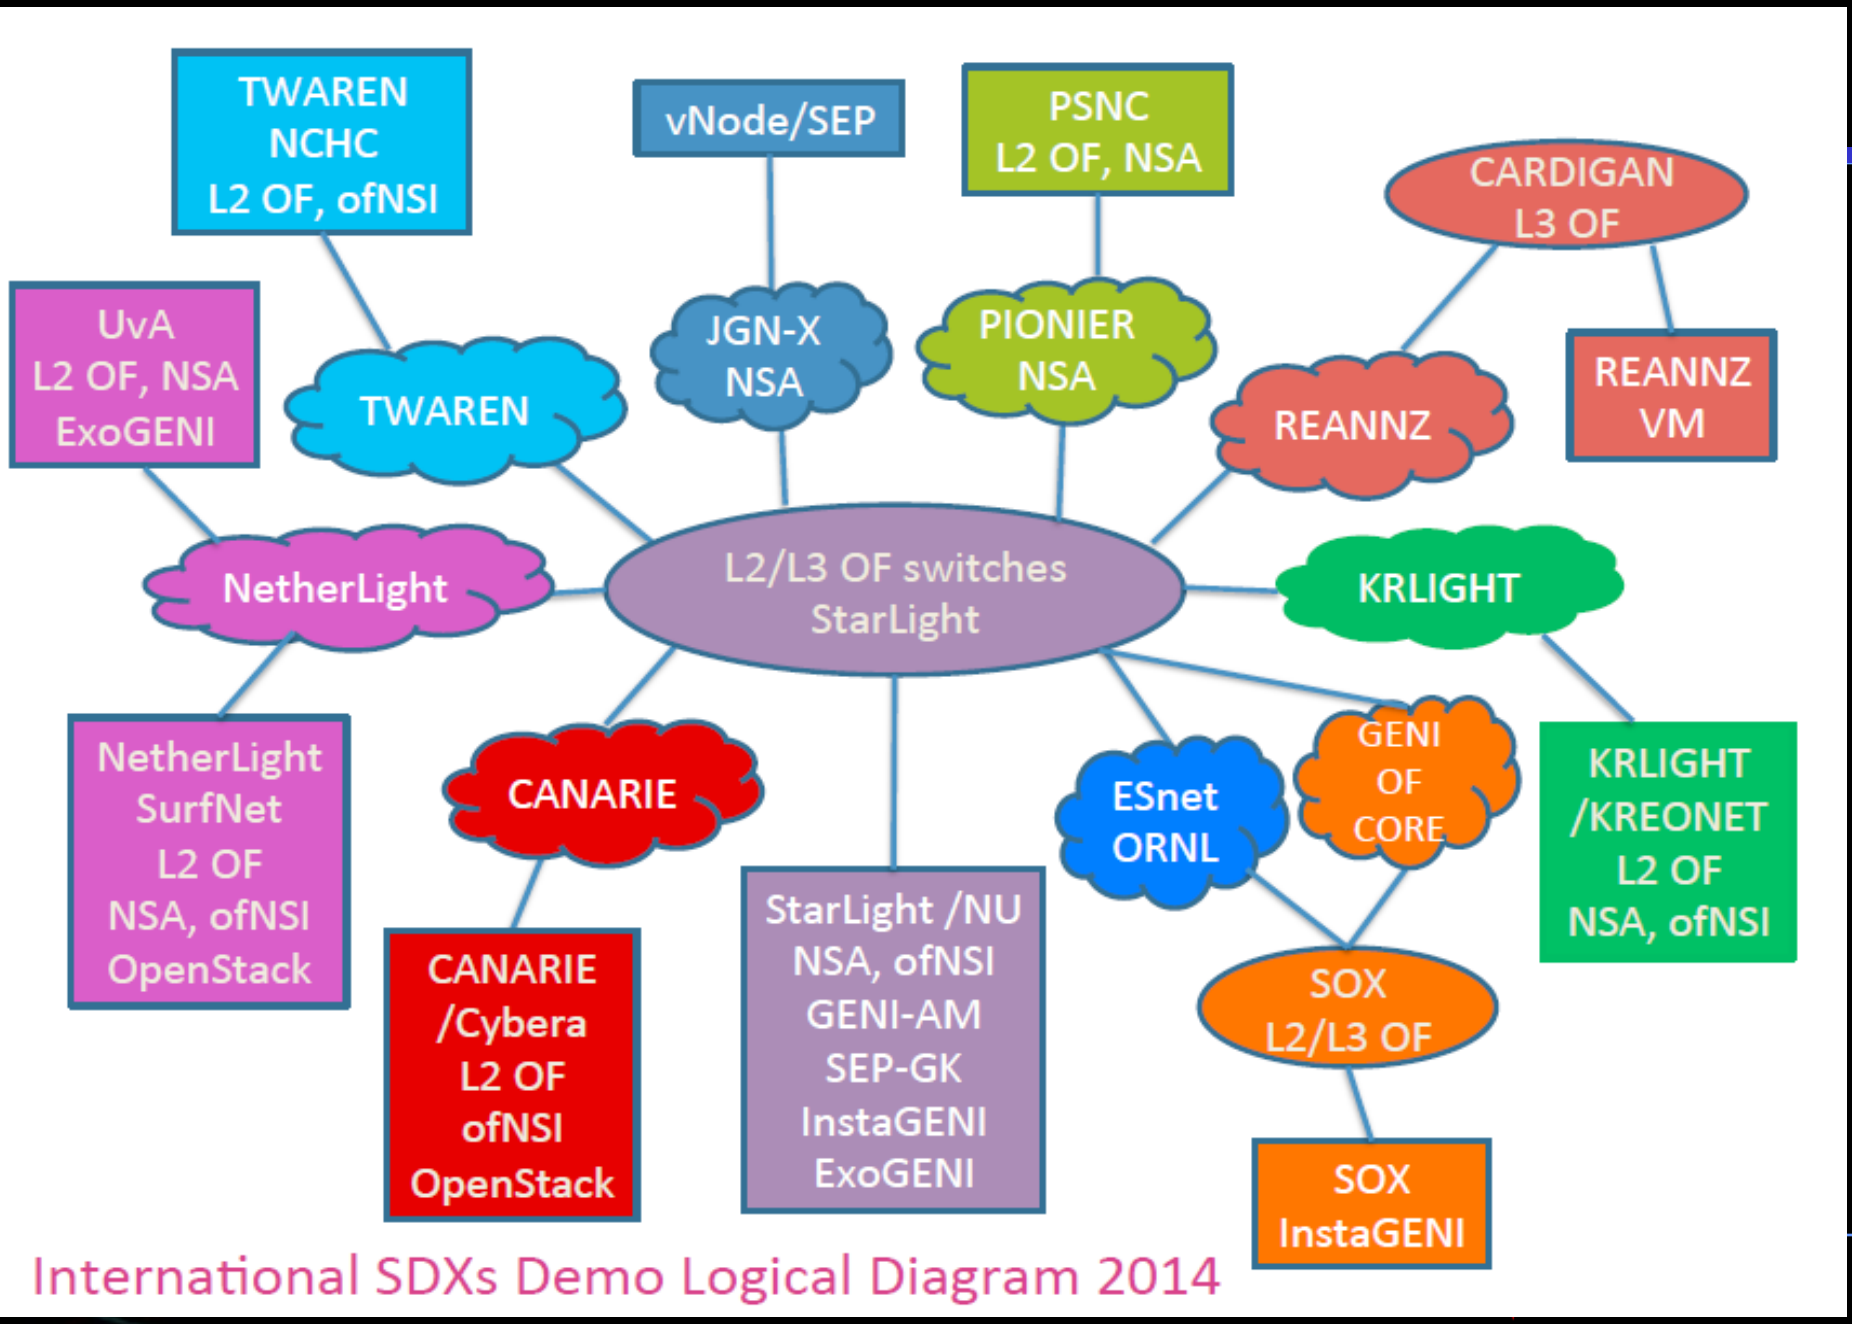

International SDXs Demo Logical Diagram 2014

Inter-SDX federation for GEC21

Multi-architecture Federation

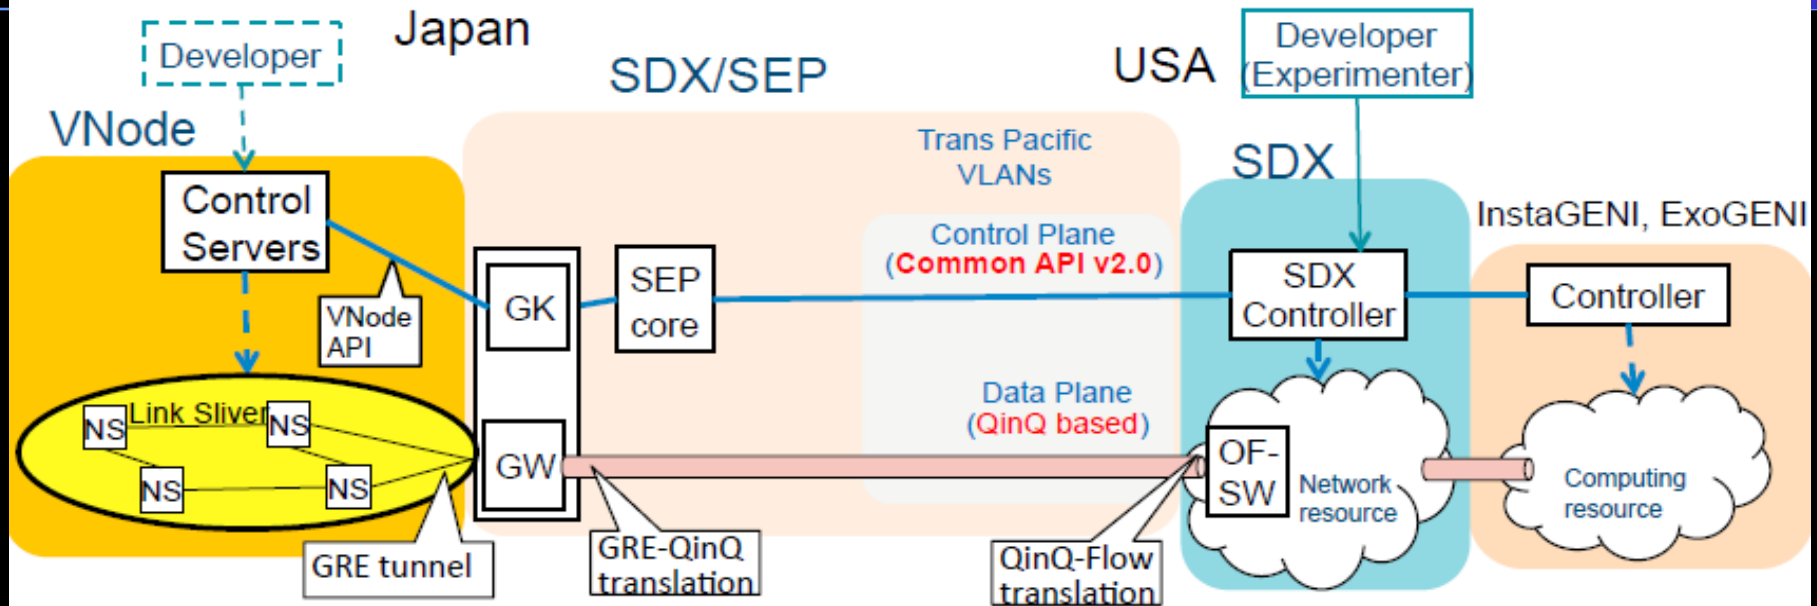

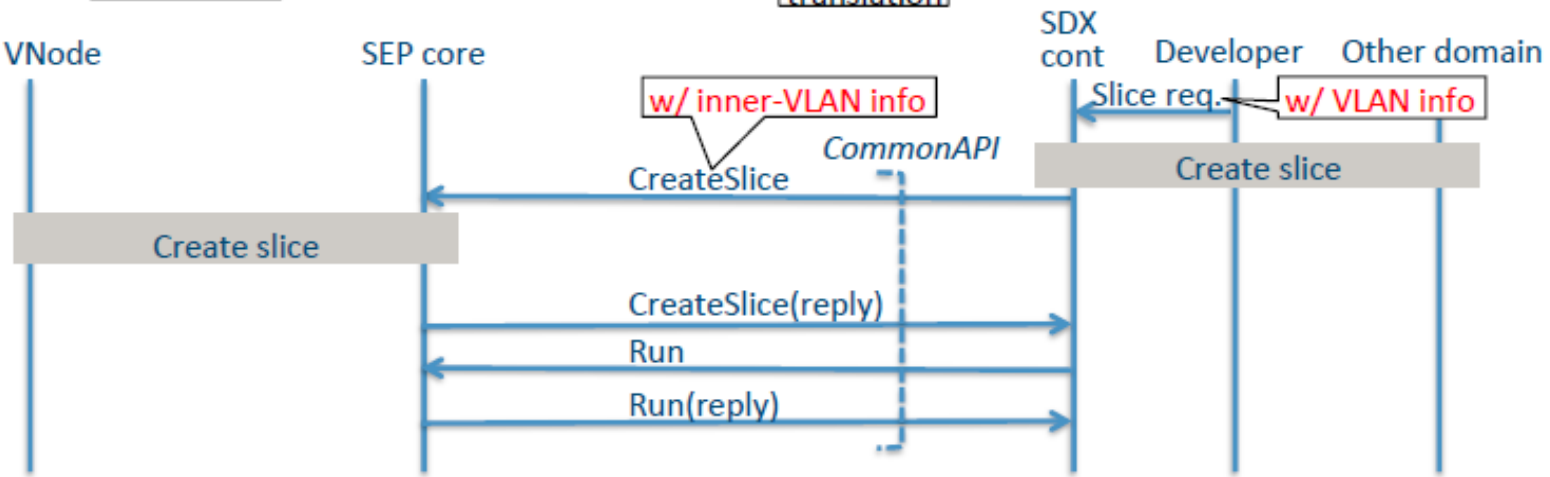

OF@KRLight+KREONET

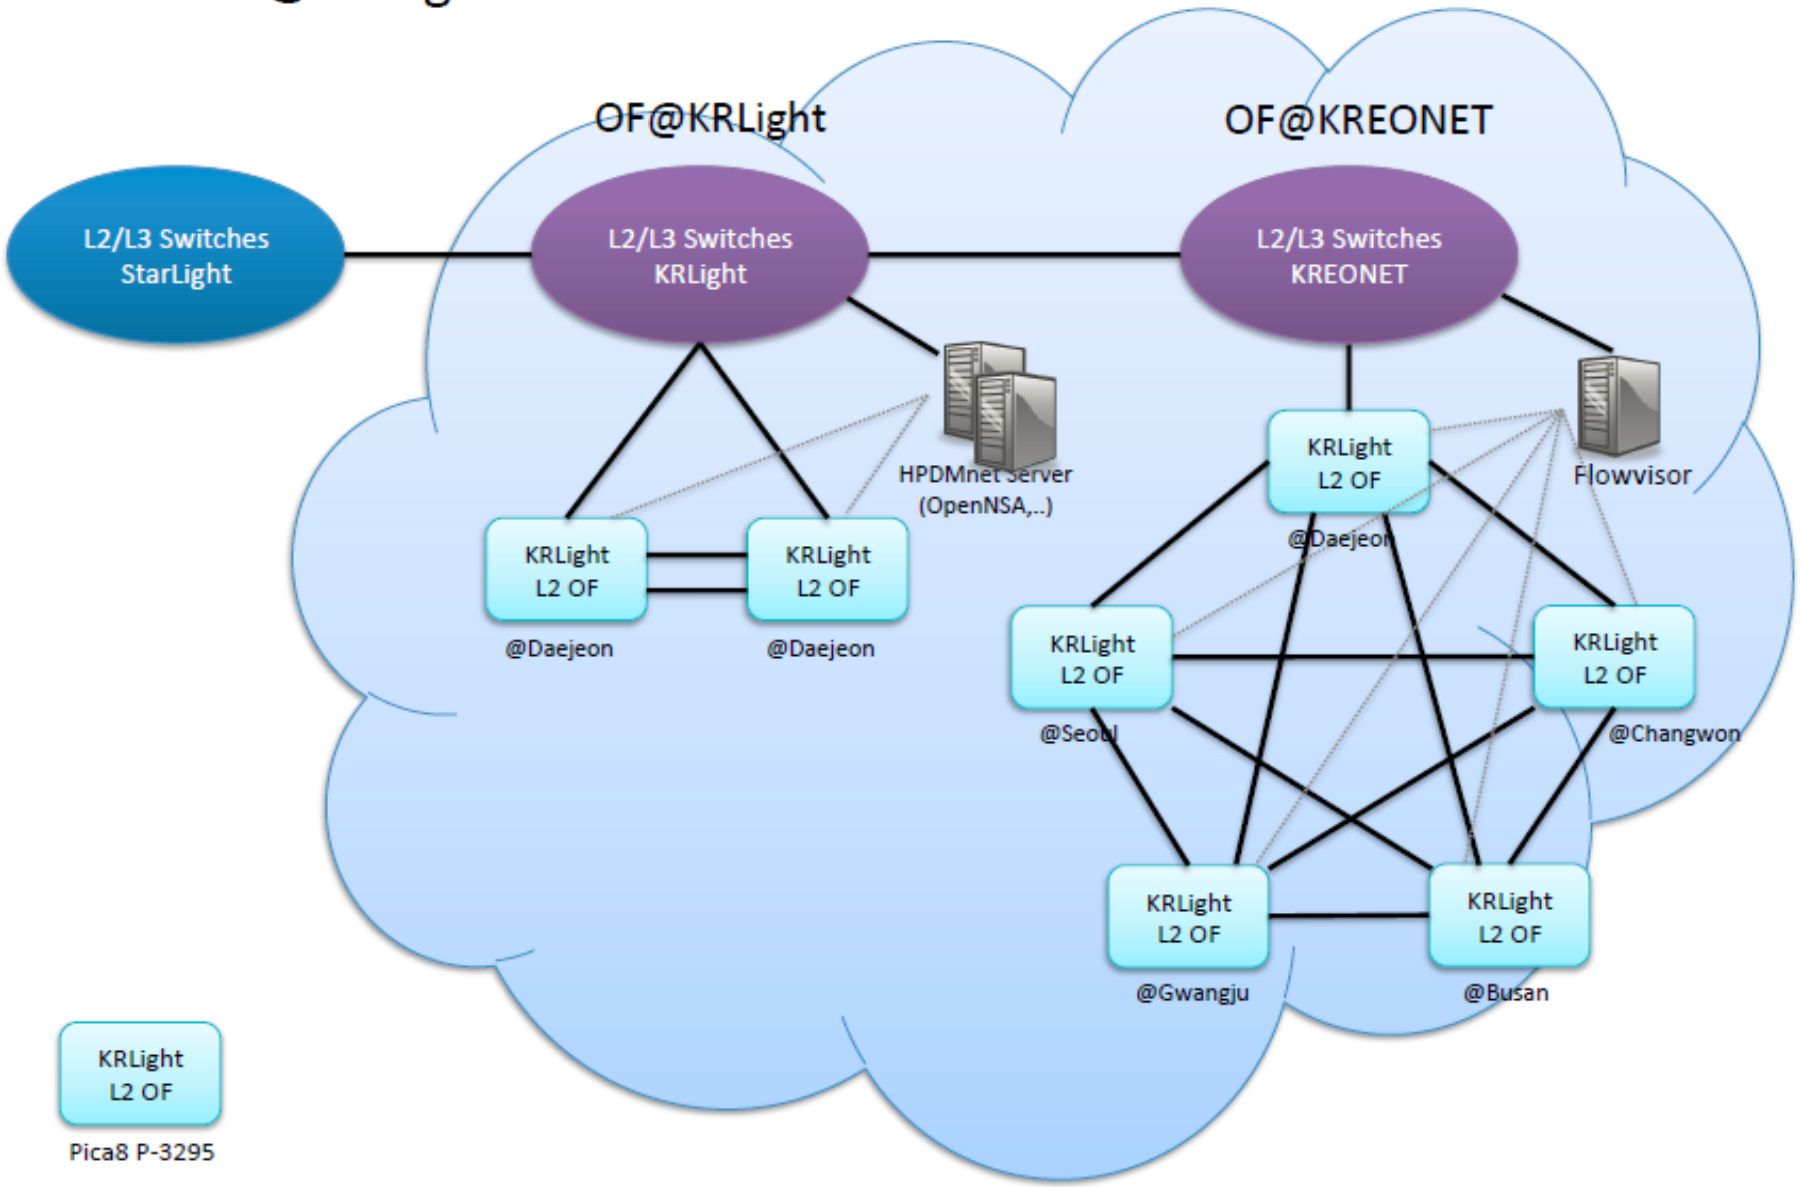

CANARIE/Cybera Openflow testbed

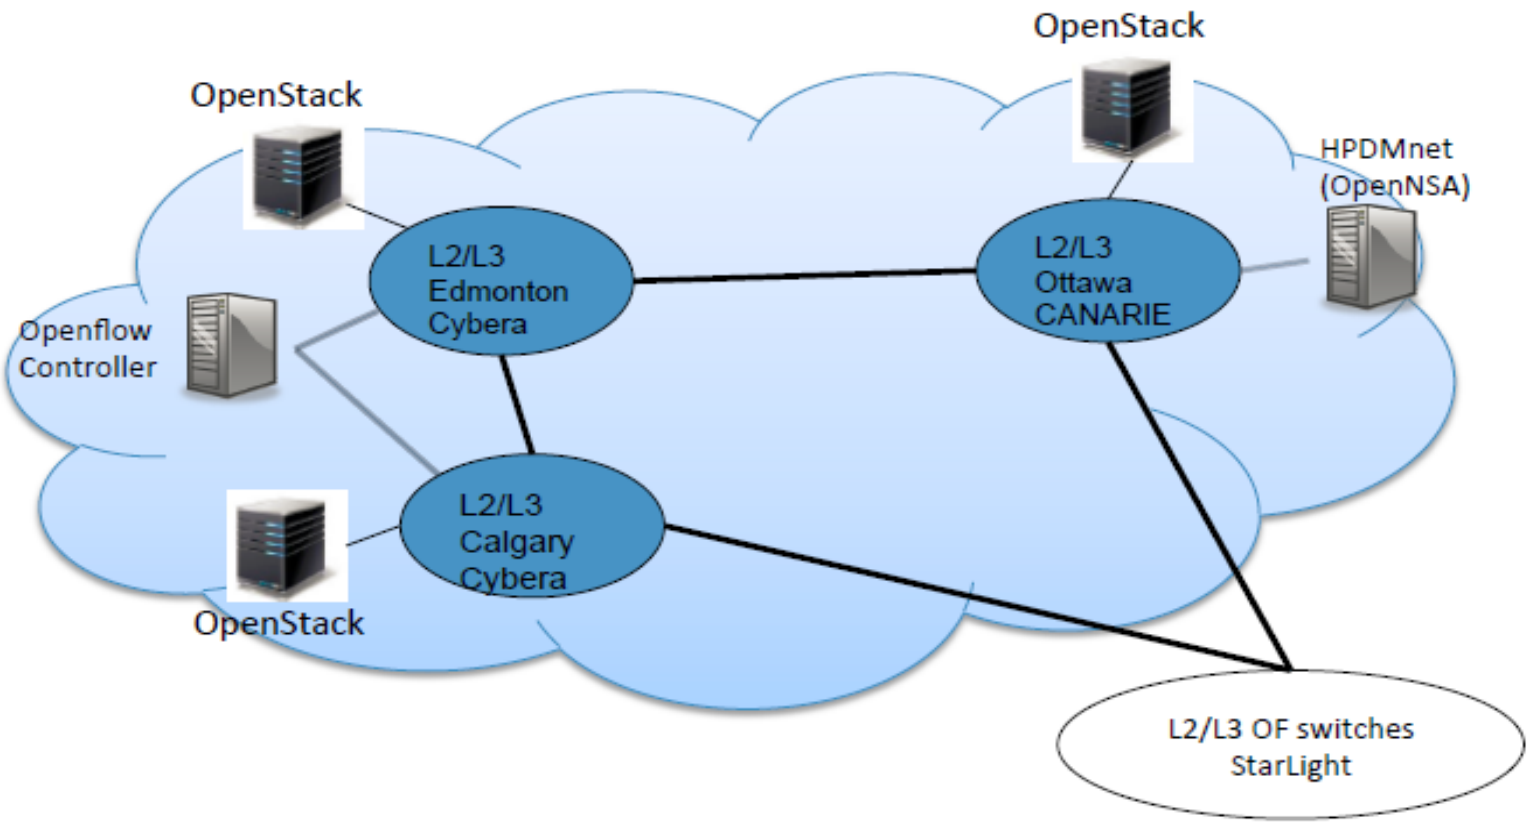

NARLabs

TWAREN OpenFlow Testbed now

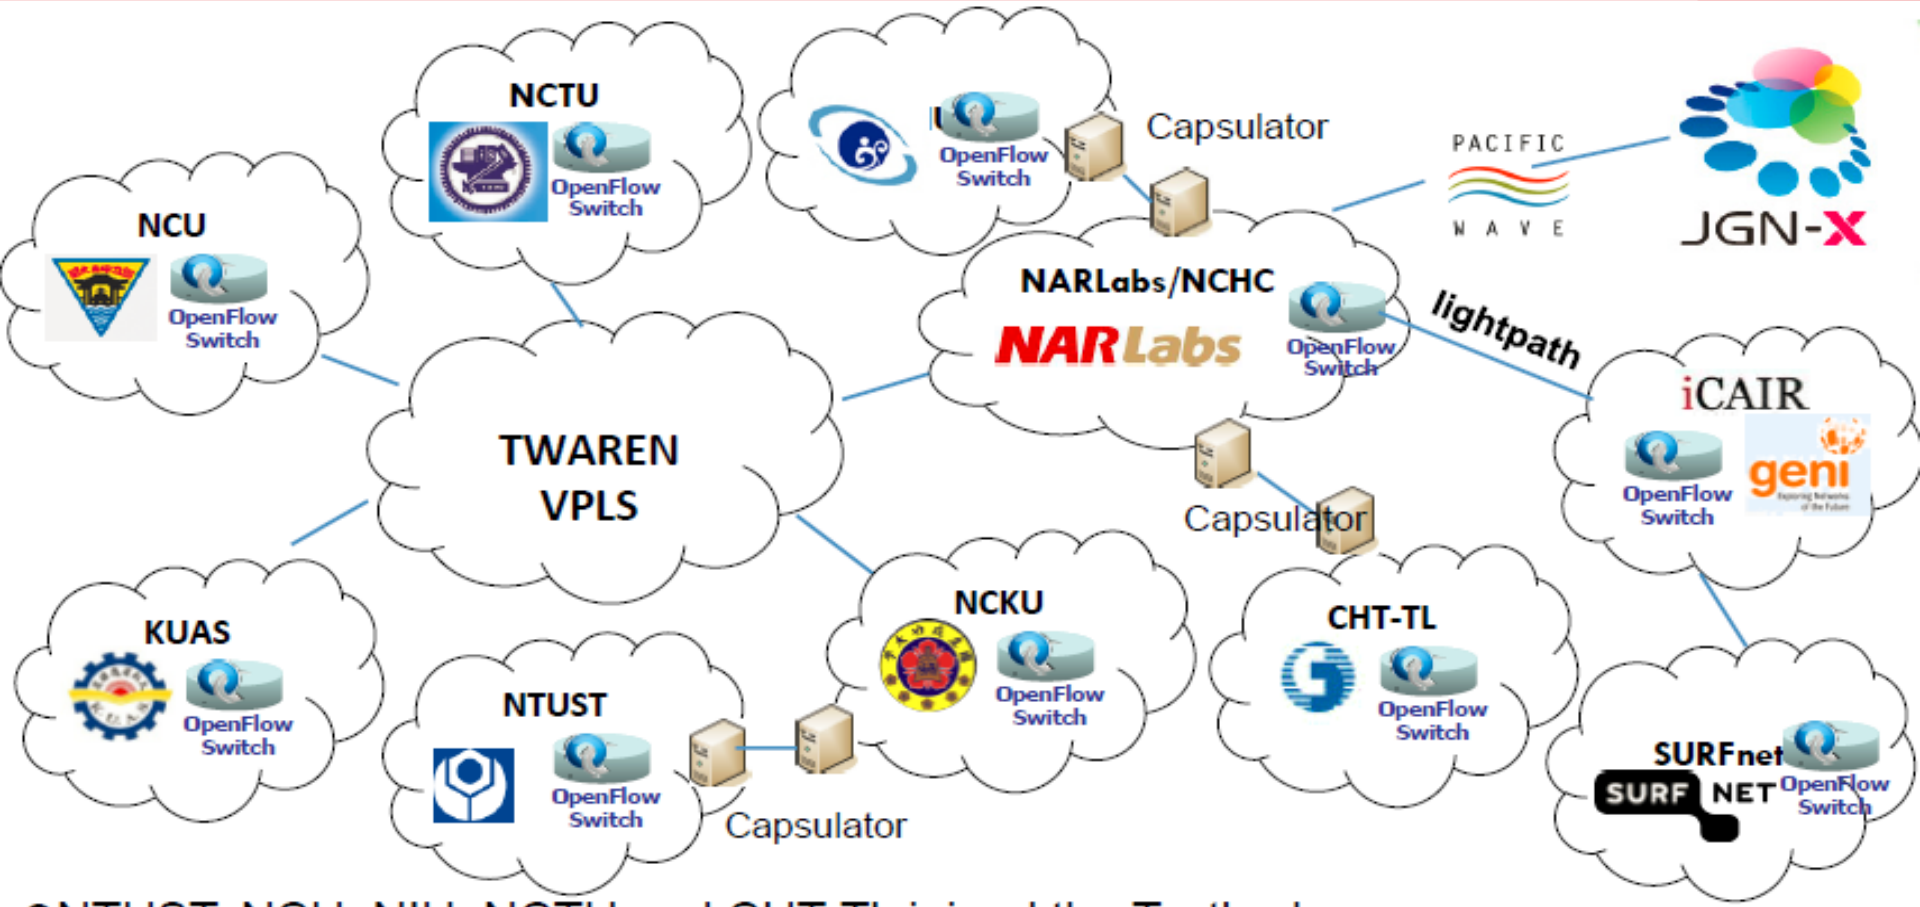

- NTUST, NCU, NIU, NCTU and CHT-TL joined the Testbed.
- For TWAREN connectors (NCTU, NCKU, KUAS and NCU), a dedicated VPLS VLAN is allocated for better transmission performance.
- Extended International connections with JGN-X and SURFnet

NARLabs

OFNSI in NCHC

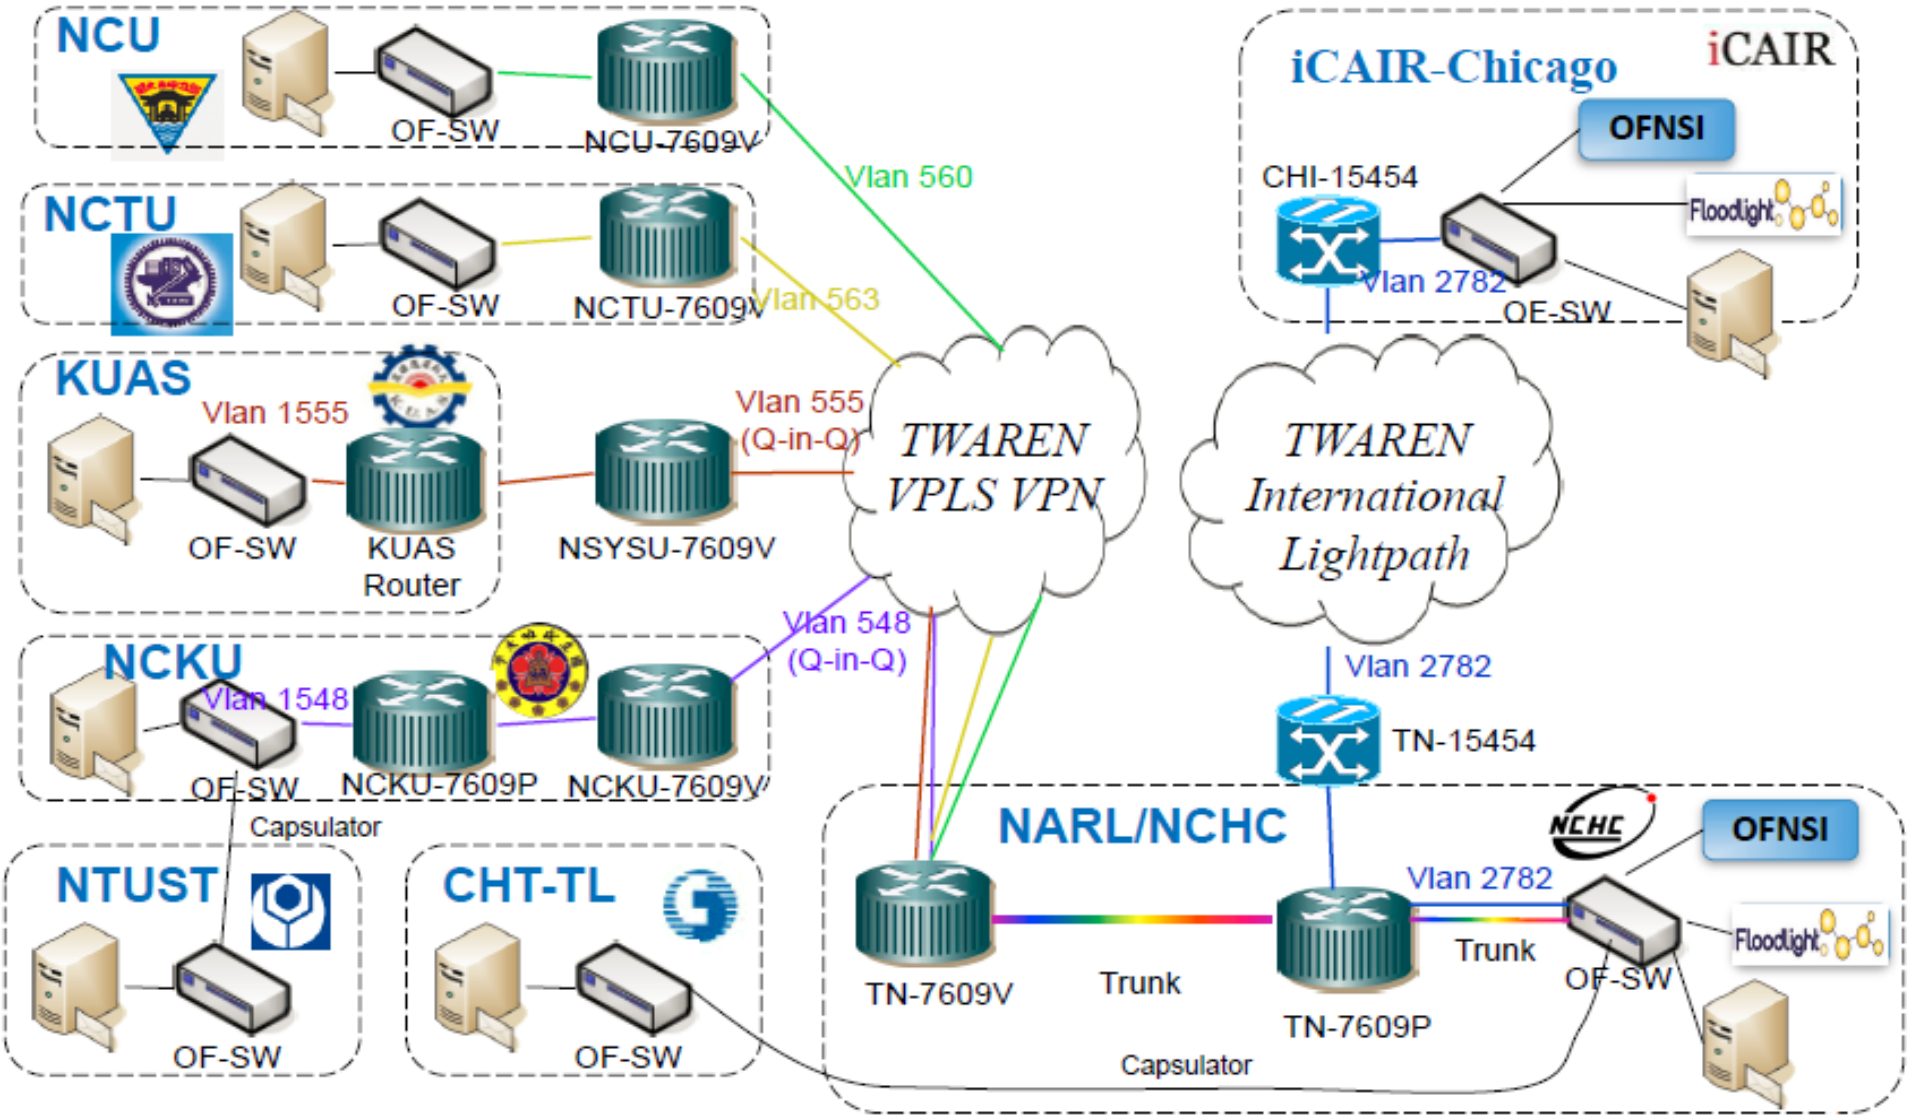

Emerging SDX Initiatives

- Integration of SDXs and Large Scale Data Intensive Science Stream Services
- Specialized SDXs For Specific Applications
- Integration With Testbeds (e.g., GENI, Chameleon, HPDMnet, iGENI, Other International Network and Cloud Testbeds, etc)
- Ref: Future Demonstrations at SC14 Supercomputing Conference in Novemeber in New Orleans

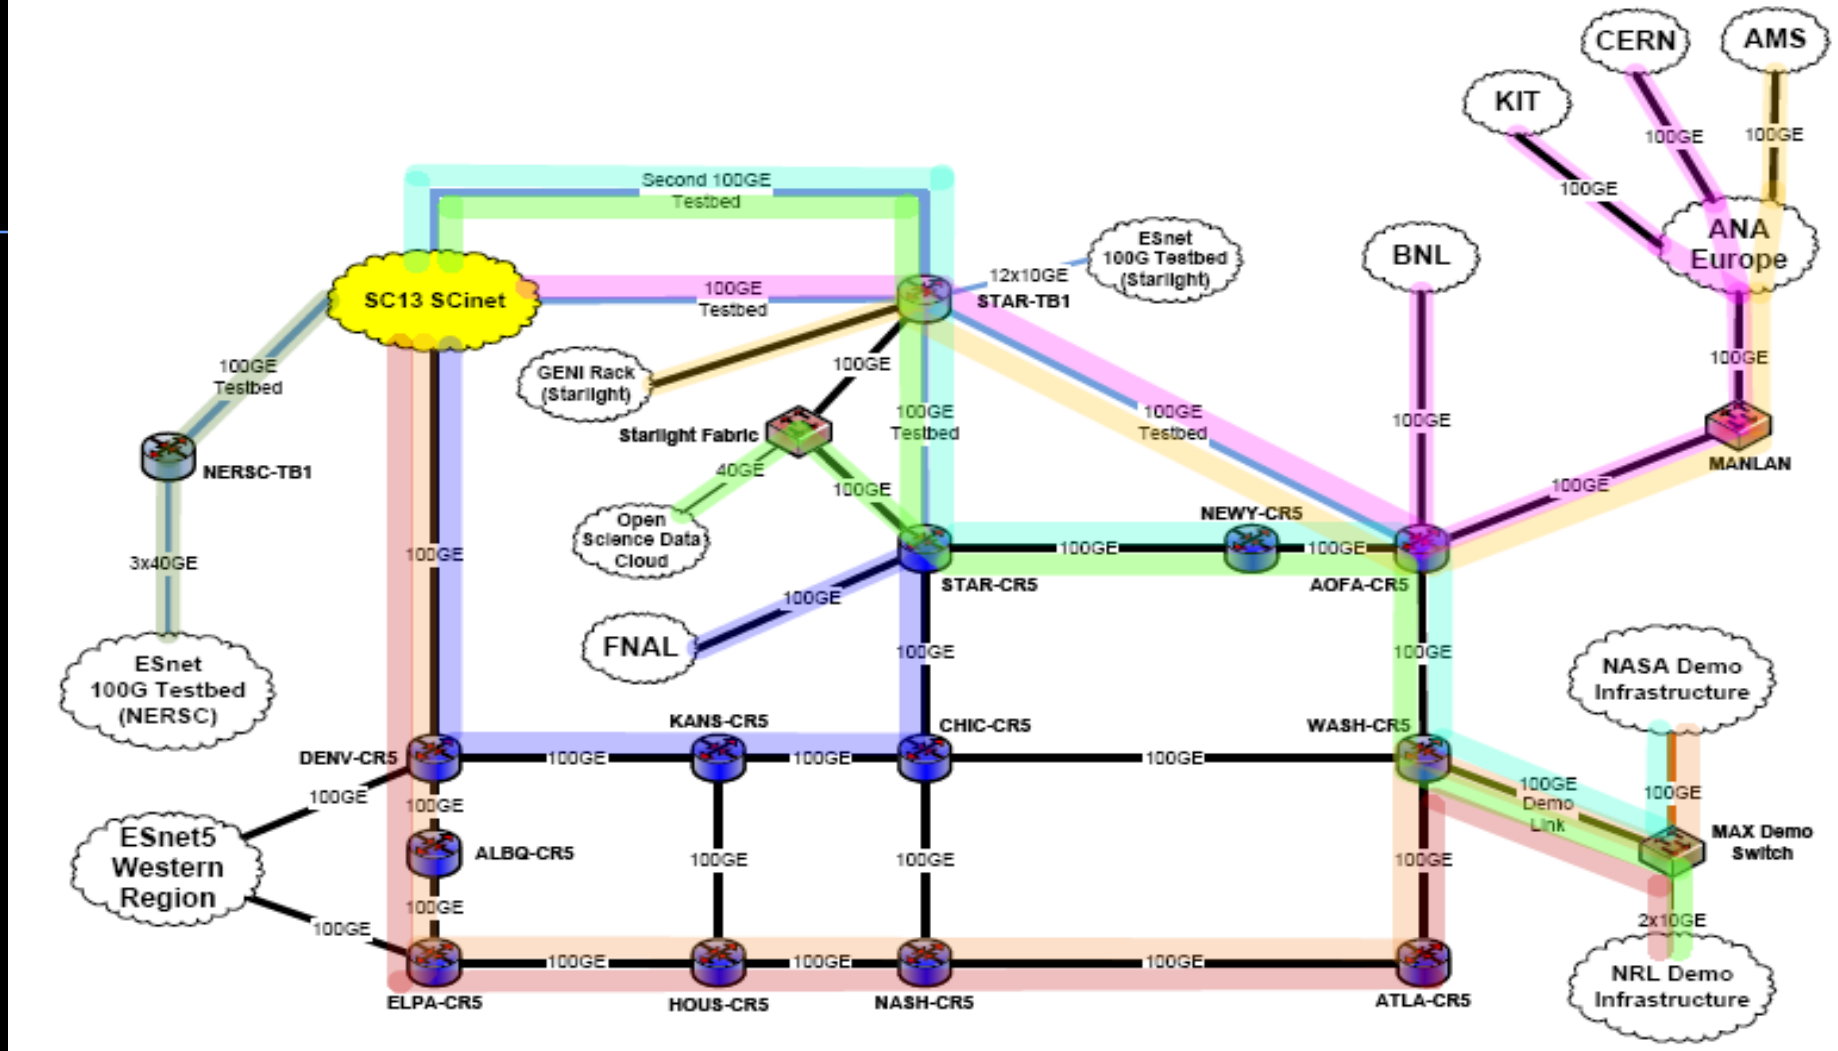

NRL demo Northern path (20G)	VLANs 1900, 1901
NRL demo Southern path (20G)	VLANs 1902, 1903
NASA demo – production path (50G)	VLAN 1801
NASA demo – testbed path (50G)	VLAN 1800
OpenFlow/SDN demo – ANA path (100G)	VLANs 1921-1929
Caltech demo – ANA path (100G)	VLANs 2602, 2603, 2606, 2607
Caltech demo – FNAL path (60G)	VLANs 2600, 2601
Caltech demo – NERSC TB path (100G)	VLANs 2604, 2605

ESnet
Energy Sciences Network

SC13 demos – ESnet5 map

Eli Dart, ESnet 11/14/2013

FILENAME SC13-DEMOS-V24.VSD

STARLIGHT℠

Evaluations/Demonstrations of 100 Gbps Disk-to-Disk WAN File Transfer Performance

Collaborative Initiative Among NASA and Several Partners

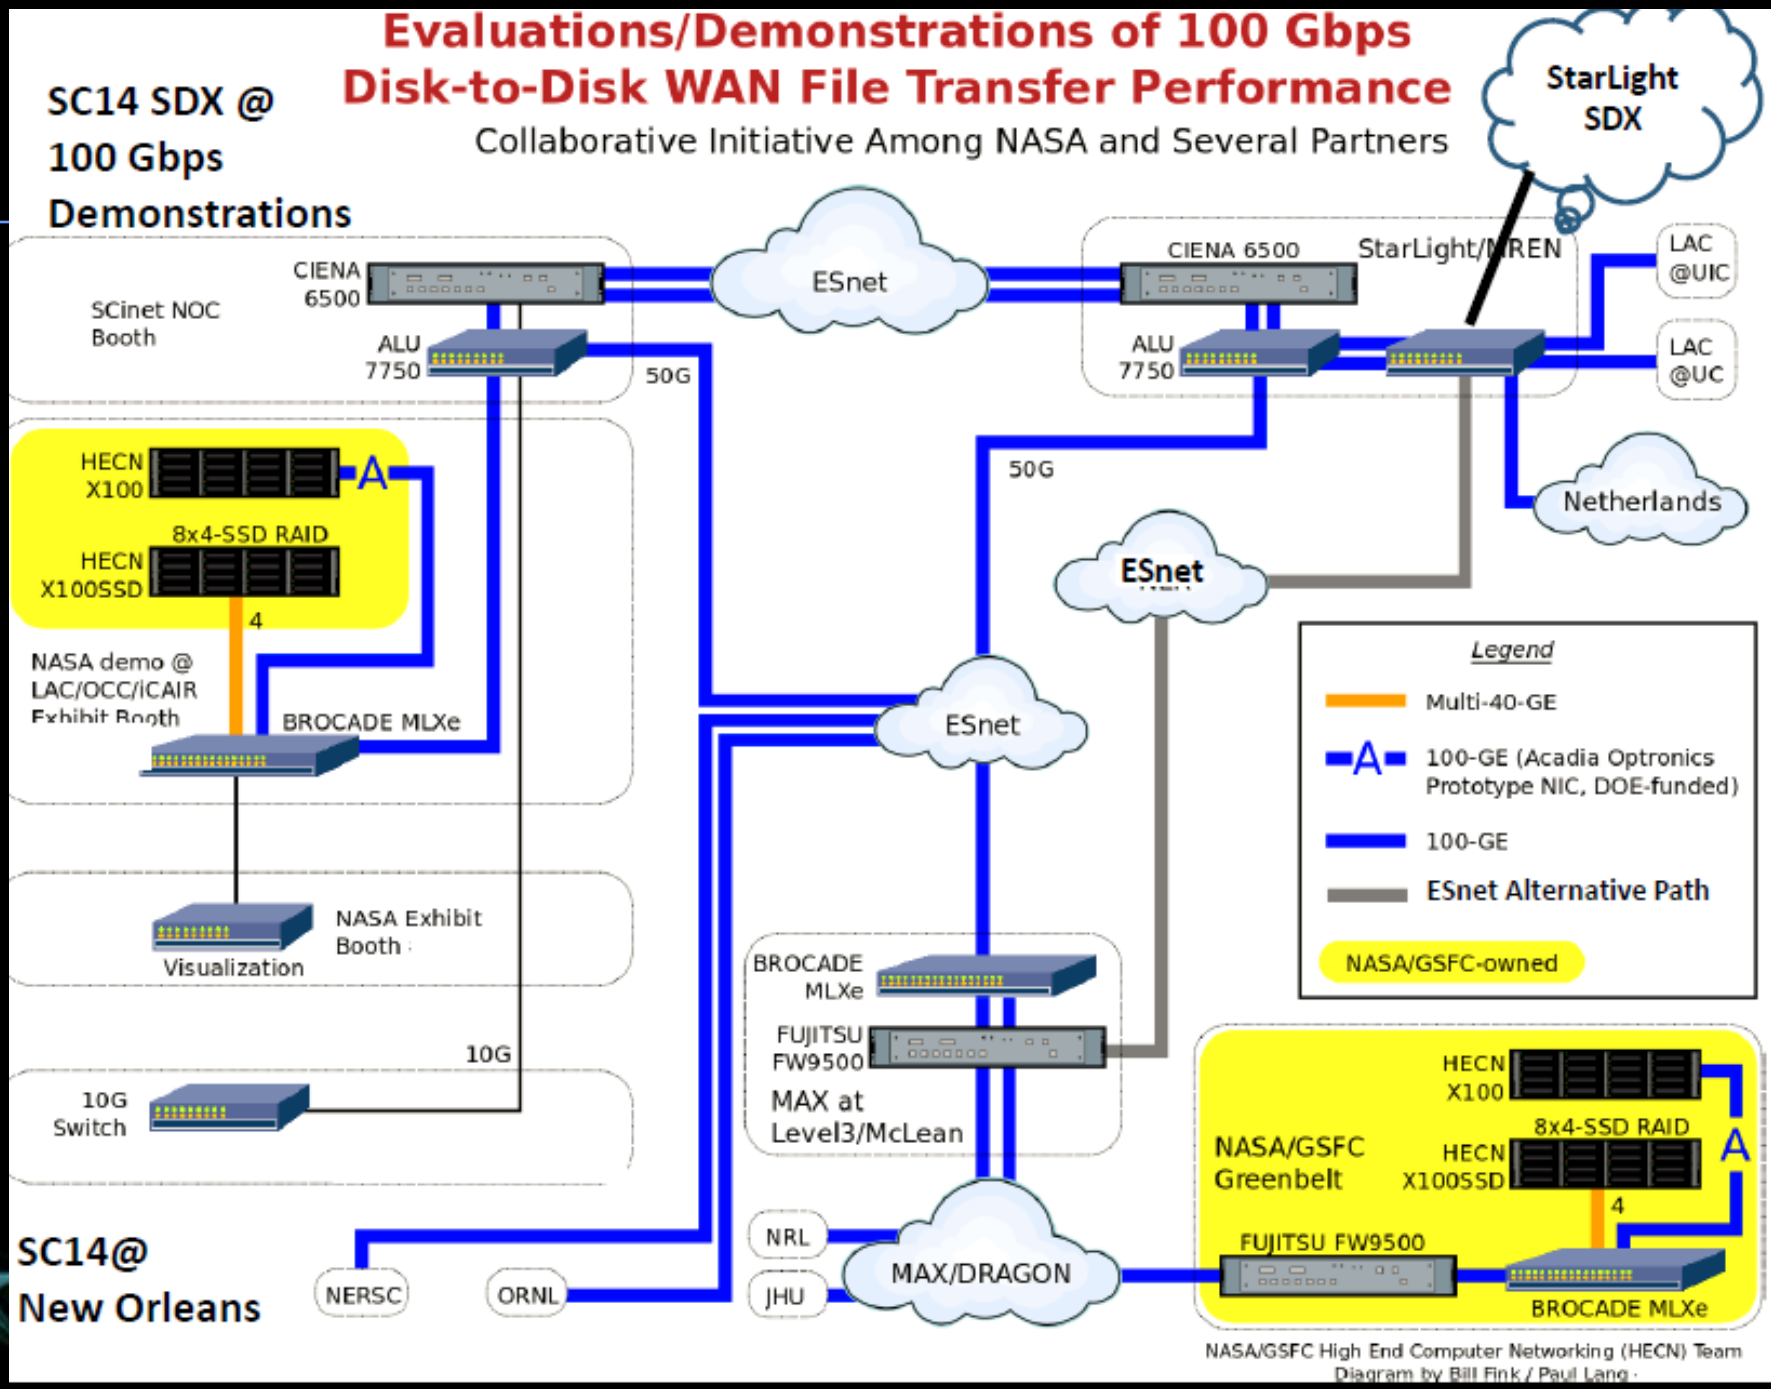

NASA/GSFC High End Computer Networking (HECN) Team
Diagram by Bill Fink / Paul Lang

Computational Genomics @ 100 Gbps at SC13 in Denver Colorado

Computational Genomics at 100gb Speeds

Bob Grossman – University of Chicago
Joe Mambretti – Northwestern University & Starlight
Don Preuss, Chris Cope – NCBI @ National Institutes of Health

NIH-OCC Computational Genomics and Analytics 100 Gbps Testbed

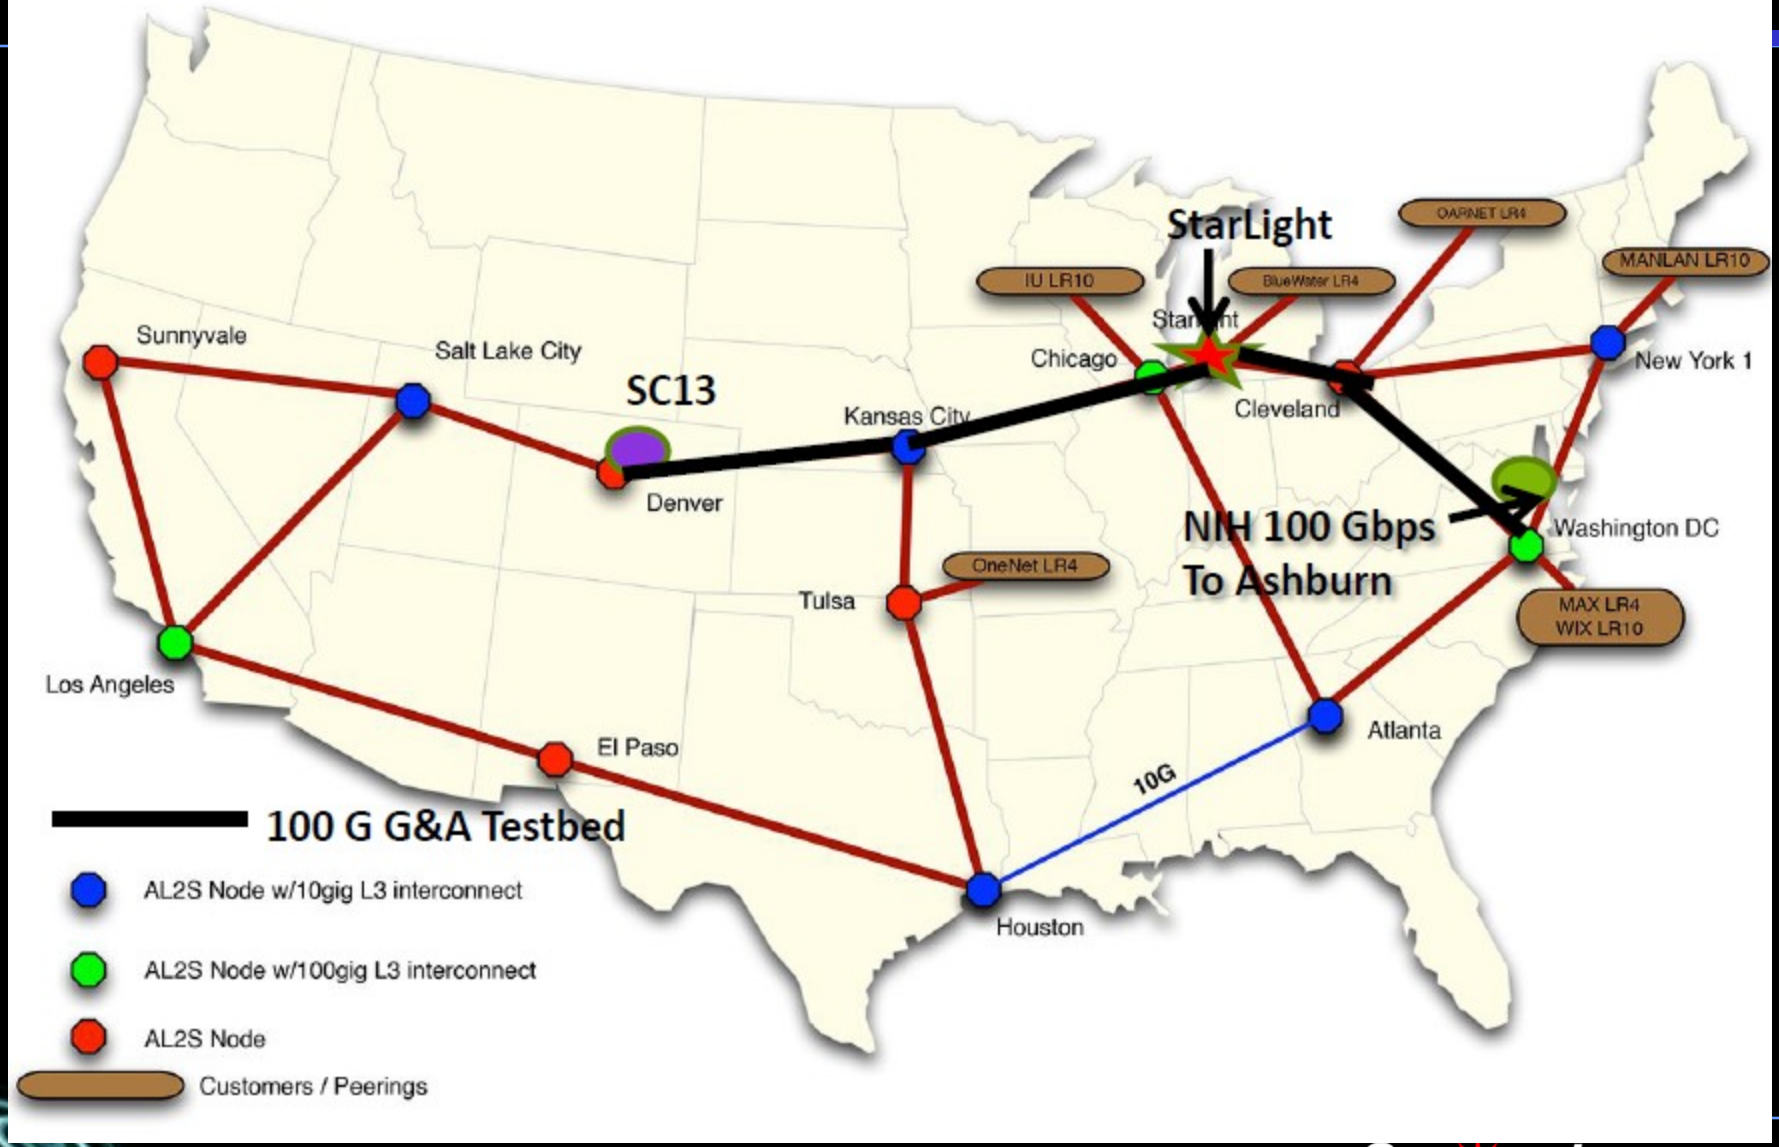

Prototype SDX BioInformatics Exchange

- Demonstration Showcase Designed for SC14 in New Orleans, Louisiana
- Builds On Previous Initiatives Developing Services for Computational BioInformatics and Computational Genomics at 100 Gbps
- Previous Demonstrations At several Conferences, Including SC13
- The SDX BX Is Veing Designed Specifically for BioInformatics/Computational Genomic Worksflows
- Appears As a Private Exchange for BioInformatics Research Communities

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

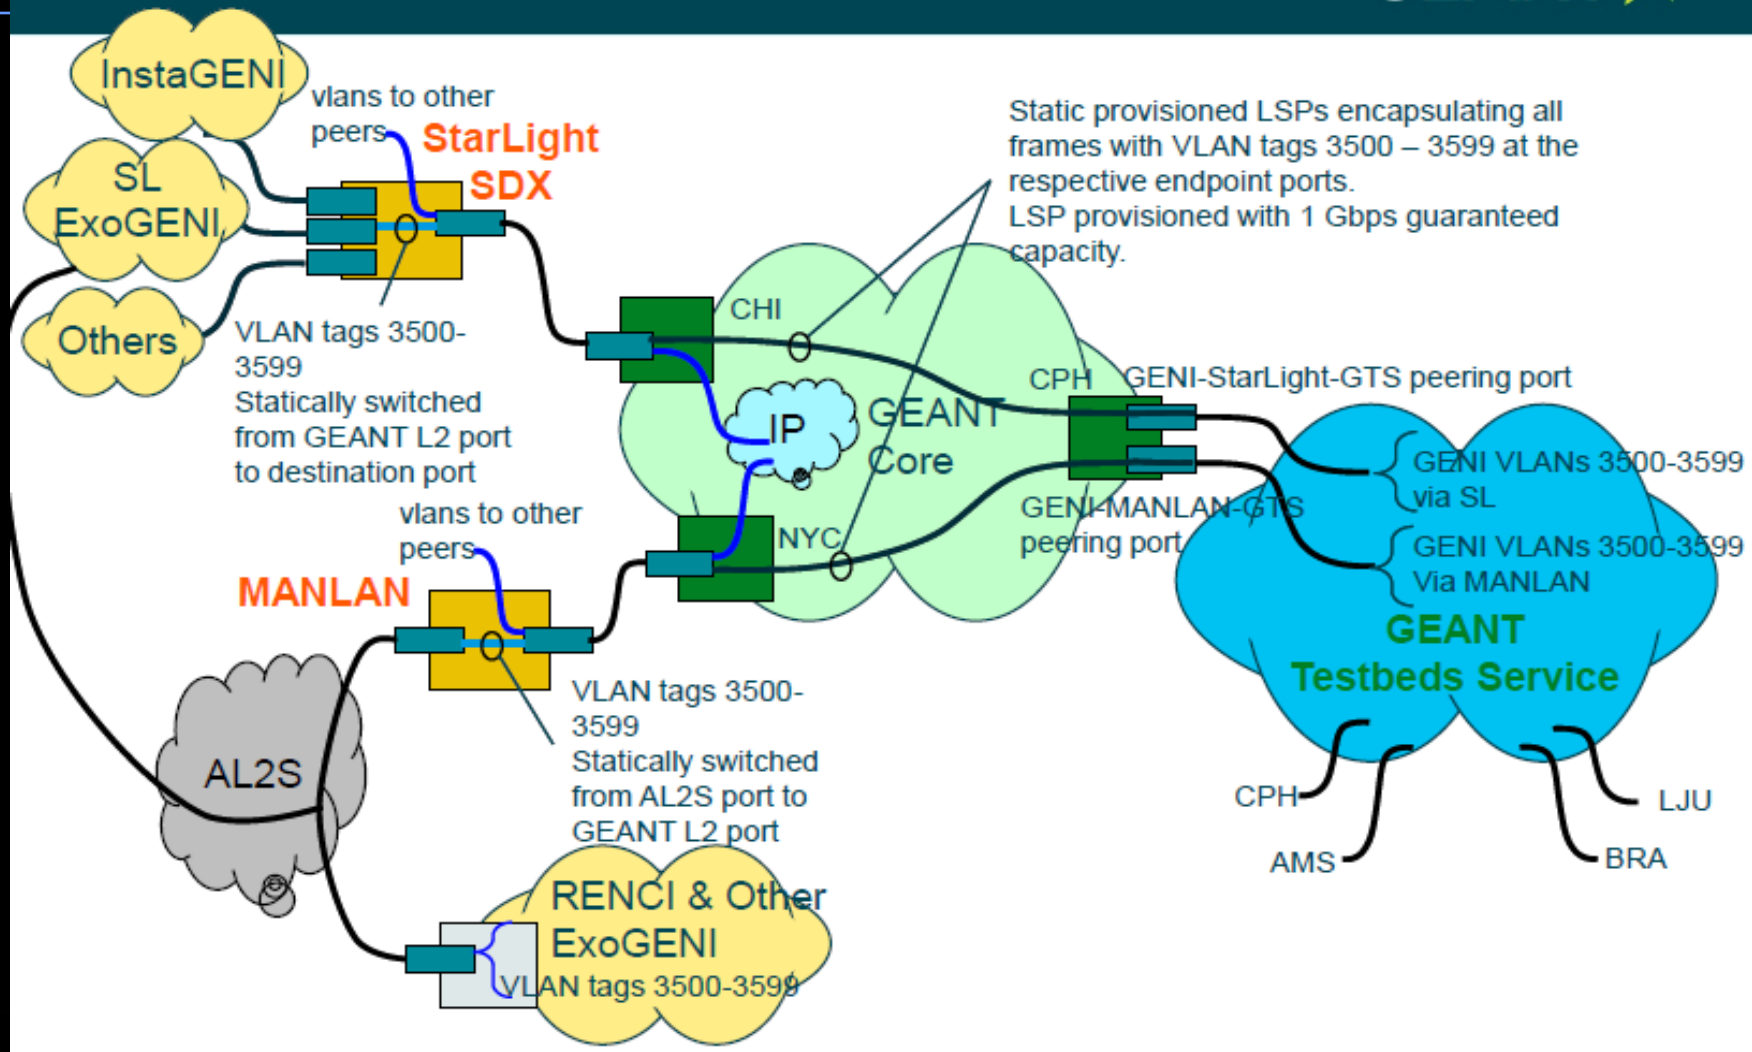

GTS connectivity to GENI Facilities
(step 2: Dynamic transit provisioning)

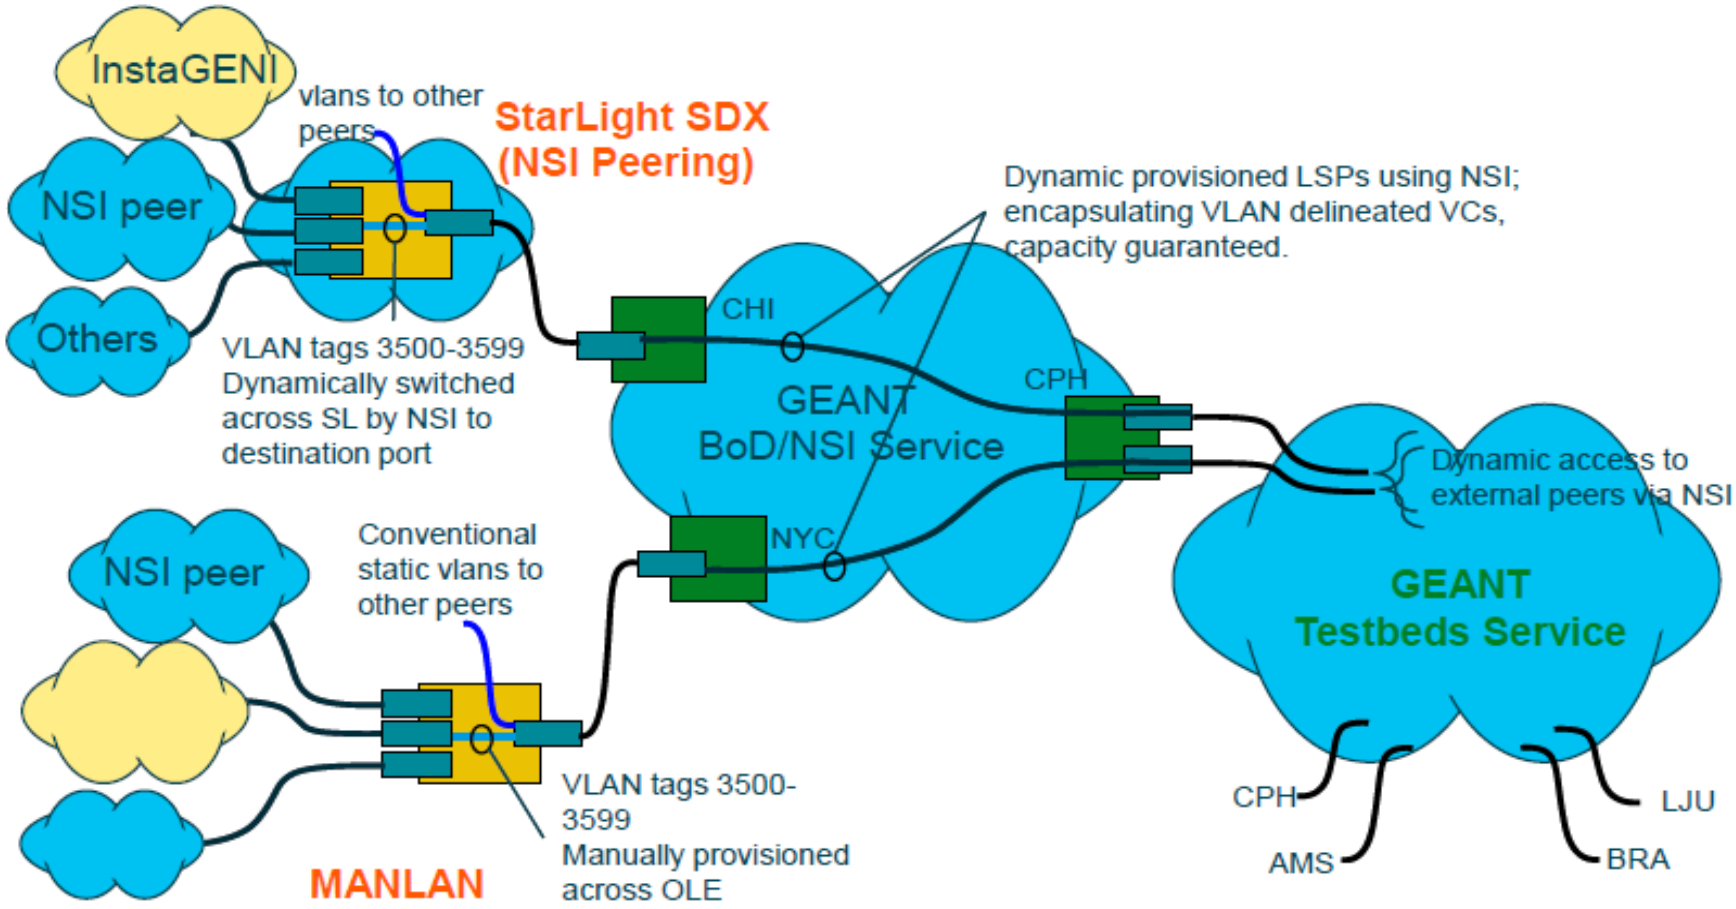

Chameleon

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

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

STARLIGHT℠

SOLVING TODAY'S RESEARCH CHALLENGES

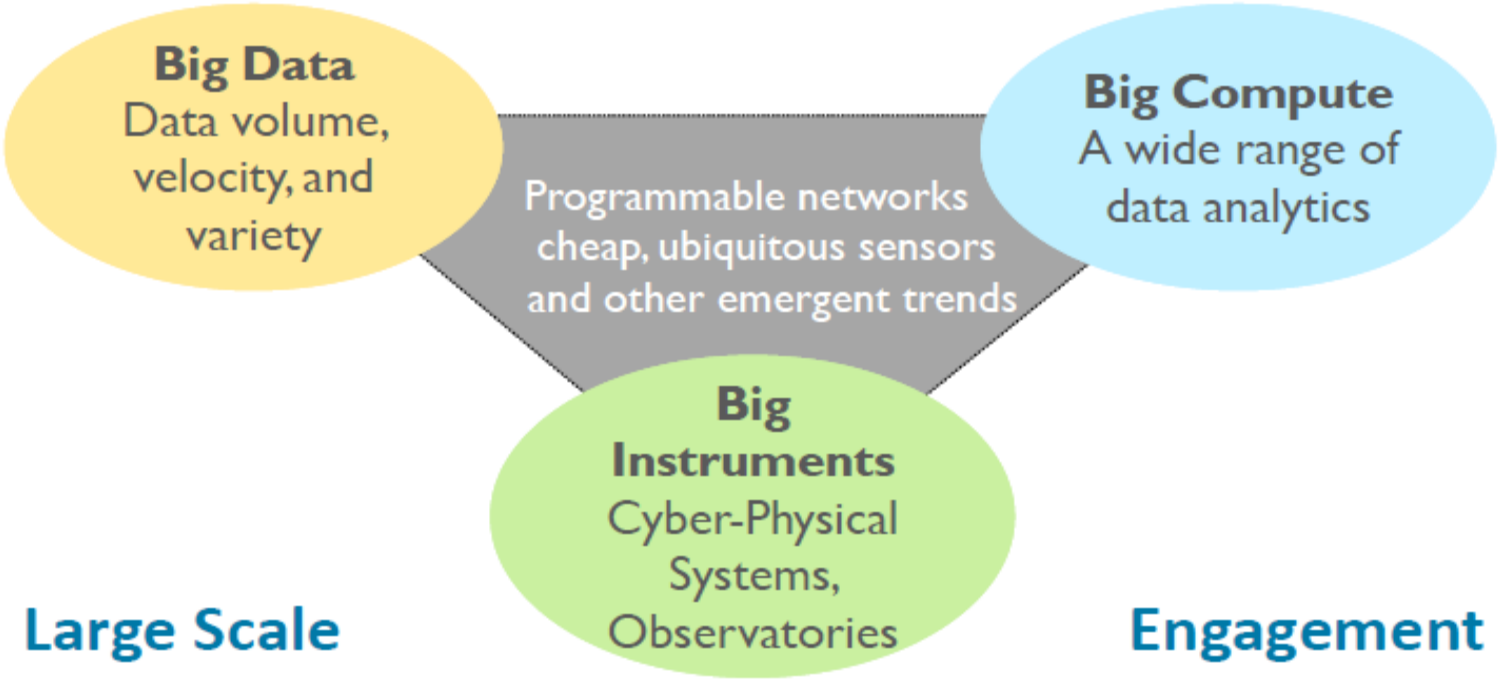

www.startap.net/starlight

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