Next Generation Software Defined Services and the Global Research Platform: A Software Defined Distributed Environment For High Performance Large Scale Data Intensive Science

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, PI StarLight IRNC SDX,Co-PI Chameleon, PI-iGENI, PI-OMNINet (www.startap.net/starlight)

Supercomputing Asia
Singapore
March 26-29, 2018

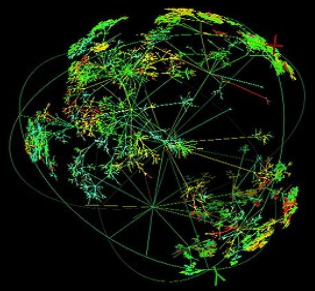

iCAIR

STARLIGHT℠

The Official Launch of the SINGAREN - vBNS LINK
7 November 1997 Washington D.C.

Hosted by
The USA National Science Foundation
The Singapore National Science and Technology board
The Telecommunication Authority of Singapore

Introduction to iCAIR:

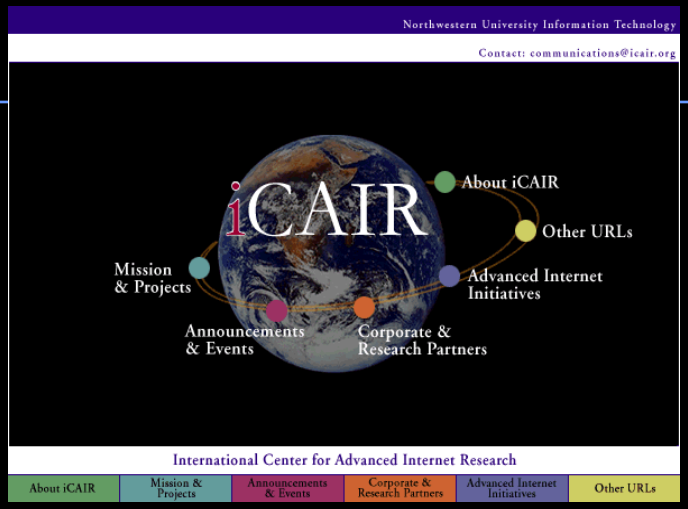

Accelerating Leading Edge Innovation and Enhanced Global Communications through Advanced Internet Technologies, in Partnership with the Global Community

- **Creation and Early Implementation of Advanced Networking Technologies - The Next Generation Internet All Optical Networks, Terascale Networks, Networks for Petascale and Exascale Science**
- **Advanced Applications, Middleware, Large-Scale Infrastructure, NG Optical Networks and Testbeds, Public Policy Studies and Forums Related to Optical Fiber and Next Generation Networks**
- **Three Major Areas of Activity: a) Basic Research b) Design and Implementation of Prototypes and Research Testbeds, c) Operations of Specialized Communication Facilities (e.g., StarLight, Specialized Science Networks)**

iCAIR Undertakes Basic Research In These Areas of Network Science

- Transition From Legacy Networks To Networks That Take Full Advantage of IT Architecture and Technology
- Extremely Large Capacity (Multi-Tbps Streams)
- Specialized Network Services, Architecture and Technologies for Data Intensive Science
- High Degrees of Communication Services Customization
- Highly Programmable Networks
- Network Facilities As Enabling Platforms for Any Type of Service
- Network Virtualization
- Tenet Networks
- Network Virtualization
- Network Programming Languages (e.g., P4) API (e.g., Jupyter)
- Disaggregation
- Orchestrators
- Highly Distributed Signaling Processes
- Network Operations Automation (Including Through AI/Machine Learning)
- SDN/SDX/SDI/OCX/SDC/SDE

Large Scale Data Intensive Science Motivates the Creation of Next Generation Communications

- Large Scale, Data (and Compute) Intensive Sciences Encounter Technology Challenges Many Years Before Other Domains
- Resolving These Issues Creates Solutions That Later Migrate To Other Domains
- 30+ Year History of Communication Innovations Has Been Driven Primarily By Data and Compute Intensive Sciences
- Best Window To the Future = Examining Requirements of Data and Compute Intensive Science Research
- Science Is Transitioning From Using Only Two Classic Building Blocks, Theory and Experimentation To Also Utilizing a Third – Modeling and Simulation – With Massive Amounts of Data
- Petabytes, Exabytes, Zettabytes
- For Communications, Data Volume Capacity Not Only Issue, But a Major Issue

Context Issues

- Today, Almost All Networks Provide Only Generic "One-Size-Fits-All" Services
- This Model Suboptimal For Many Types of Emerging and Anticipated Applications And Services
- E.g, With Today's Networks, Even R&E Networks, It Is Difficult To Transport Extremely Large Files and Collections of Many Files Over WANs, Especially Over Multi-Domains
- Future Networks Will Provide Differentiated Services, e.g., Using Software Defined Networking (SDN) and Software Defined Exchange (SDX) Resource Slicing
- These Capabilities Provide Opportunities To Address The Special Requirements of Global Data Intensive Science
- Goal: Convergence of A) Segmented Research Platforms, e.g., Science DMZ, National Science Foundation's Campus Cyberinfrastructure, Pacific Research Platform (PRP), Potential National Research Platform, Global Research Platform, GLIF, GLIF GOLES, and Related Specialized Environments and B) Software Defined Infrastructure (SDI)

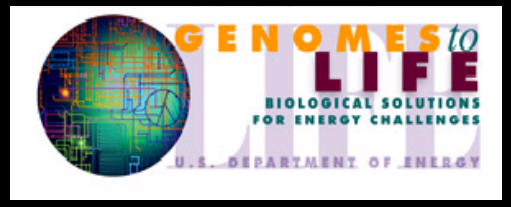

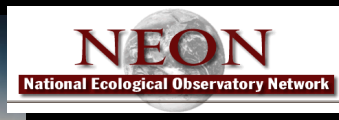

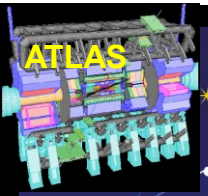

ALMA: Atacama Large Millimeter Array

ANDRILL: Antarctic Geological Drilling
www.andrill.org

BIRN: Biomedical Informatics Research Network
www.nbirn.net

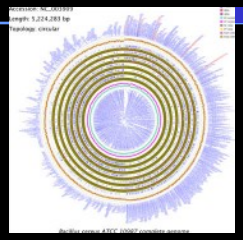

CAMERA metagenomics
camera.calit2.net

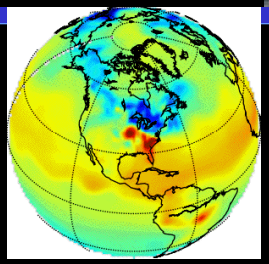

Carbon Tracker
www.esrl.noaa.gov/gmd/ccgg/carbontracker

CineGrid
www.cinegrid.org

LHCONE
www.lhcone.net

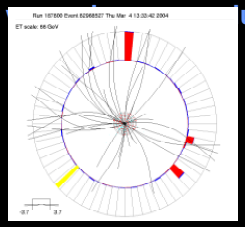

DØ (DZero)
www-d0.fnal.gov

GEON: Geosciences Network
www.geongrid.org

GLEON: Global Lake Ecological Observatory Network

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

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

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

IVOA: International Virtual Observatory
www.ivoa.net

OSG
www.opensciencegrid.org

Globus Alliance
www.globus.org

SKA
www.skatelescope.org

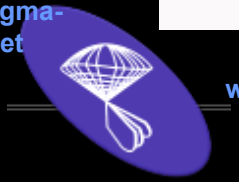

Sloan Digital Sky Survey
www.sdss.org

TeraGrid
www.teragrid.org

XSEDE
www.xsede.org

Compilation By Maxine Brown

STARLIGHT SM

Petascale Computational Science

For Decades, Computational Science
Has Driven Network Innovation
Today –
Petascale Computational Science

BLUE WATERS
SUSTAINED PETASCALE COMPUTING

National Center for Supercomputing Applications, UIUC

STARLIGHT℠

XSEDE

- Extreme Science and Engineering Discovery Environment (XSEDE)
- Goal: Create a Distributed Computational Science Infrastructure to Enable Distributed Data Sharing and High-Speed Computing for Data Analysis and Numerical Simulations
- Builds on Prior Distributed TeraGrid

Open Science Grid: Selected Investigations

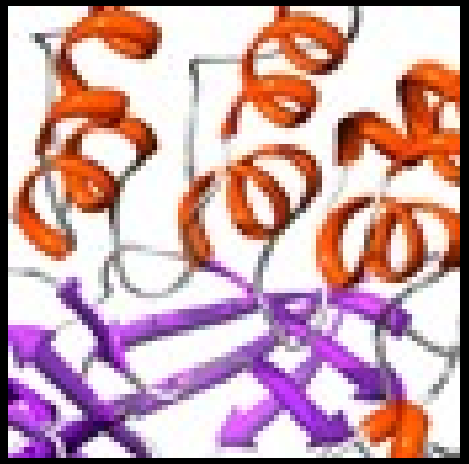

DNA Modeling

Gravity Wave Modeling

Nutrino Studies

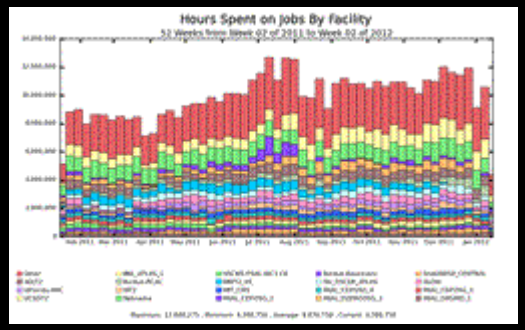

Usage

Open Science Grid

This Distributed Facility
Supports Many Sciences

The Open Science Data Cloud (OSDC) is an **open-source**, **cloud-based** infrastructure that allows scientists to manage, share, and analyze medium to large size scientific datasets.

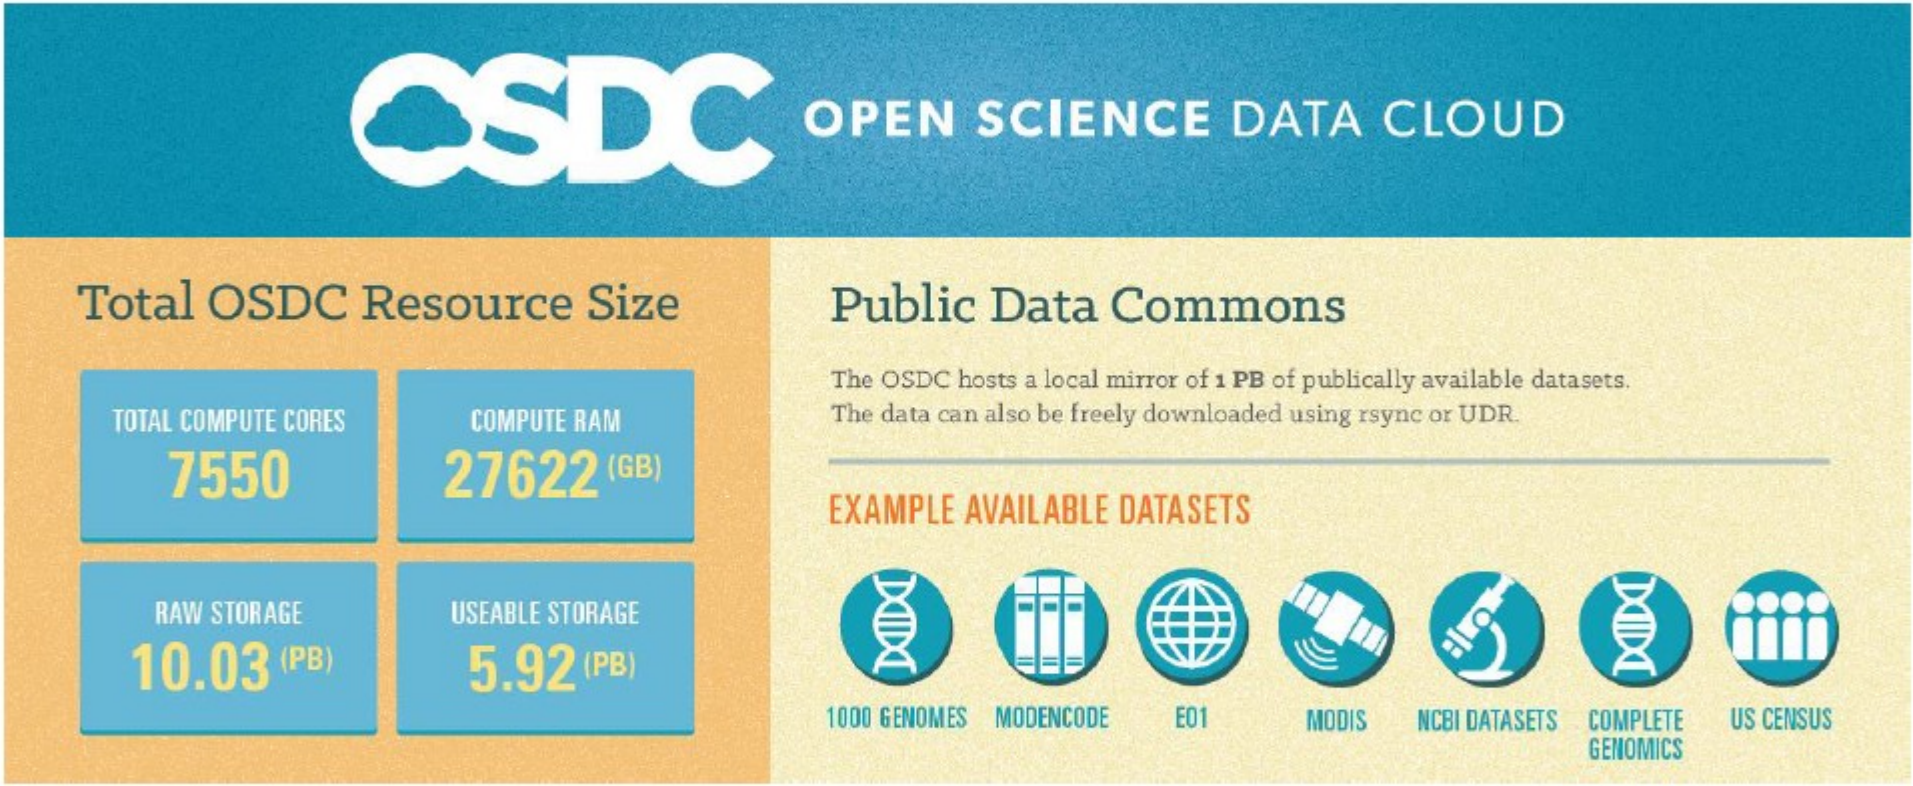

Application for resources available to anyone doing scientific research:

Open Commons Consortium

www.opensciencedatacloud.org

Maria Patterson (mtpatter@uchicago.edu) *Center for Data Intensive Science, University of Chicago*

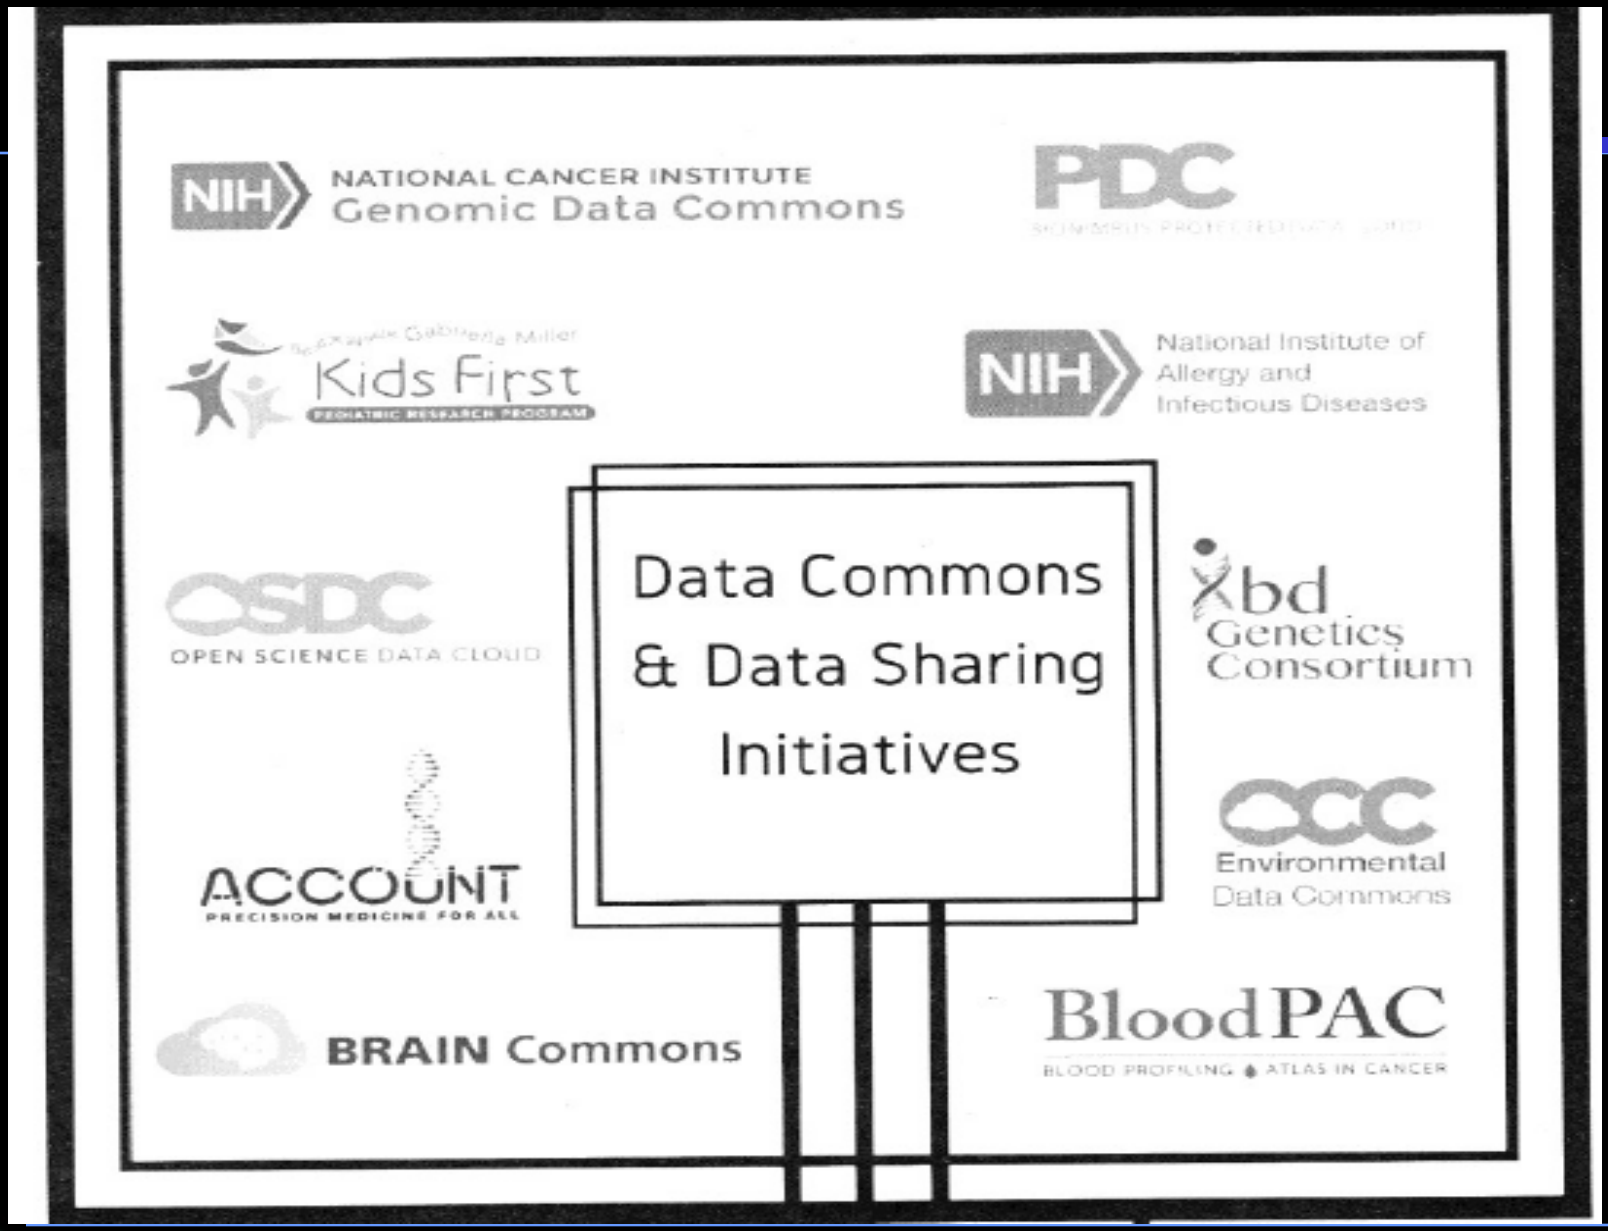
NIH
NATIONAL CANCER INSTITUTE
Genomic Data Commons
PDC
Kids First
PEDIATRIC RESEARCH PROGRAM
NIH
National Institute of
Allergy and
Infectious Diseases
Data Commons
& Data Sharing
Initiatives
OSDC
OPEN SCIENCE DATA CLOUD
bd
Genetics
Consortium
ACCOUNT
PRECISION MEDICINE FOR ALL
CCC
Environmental
Data Commons
BRAIN Commons
BloodPAC
BLOOD PROFILING ATLAS IN CANCER

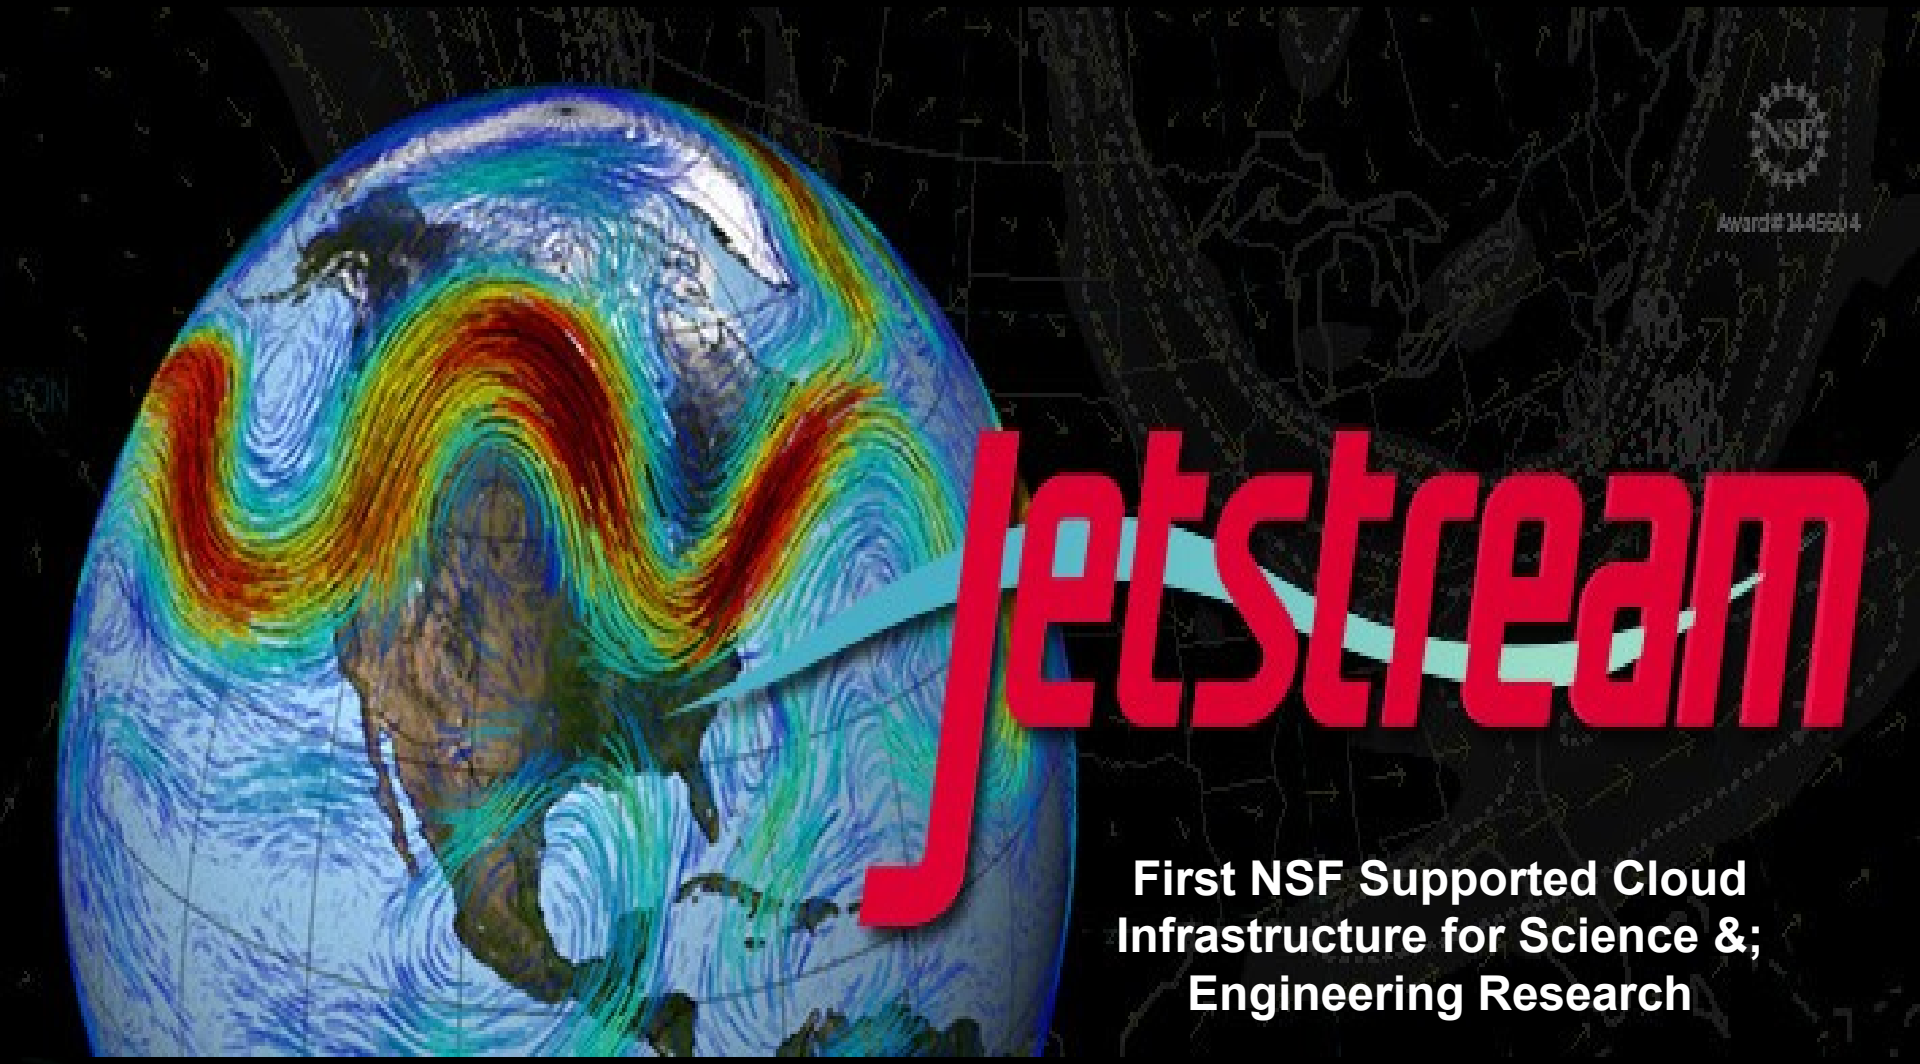
Award# 1445604
Jetstream
First NSF Supported Cloud Infrastructure for Science &; Engineering Research

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

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

HEP = Staggering Amounts of Data

BaBar 0.3 PetaByte/year (2001)

CDF or D0 Run II

0.5 PetaByte/year (2003)

LHC Mock Data Challenge
1 PetaByte/year (~2005)

CMS or ATLAS

2 PetaBytes/year (~2008)

KTeV 50 TeraBytes /year (1999)

SLD 3 TB /year (1998)

Run I (CDF or D0) 20 TB /year (1995)

L3 5 TB /year (1993)

E791 50 TB /year (1991)

EMC 400GB /year (1981)

In 1977 the Upsilon (bottom quark) was discovered at Fermilab by experiment E288 led by now Nobel laureate Leon Lederman

The experiment took about 1 million events and recorded the raw data on ~ 300 magnetic tapes for about 6 GB of raw data

Source: Fermi Lab

Large Hadron Collider at CERN

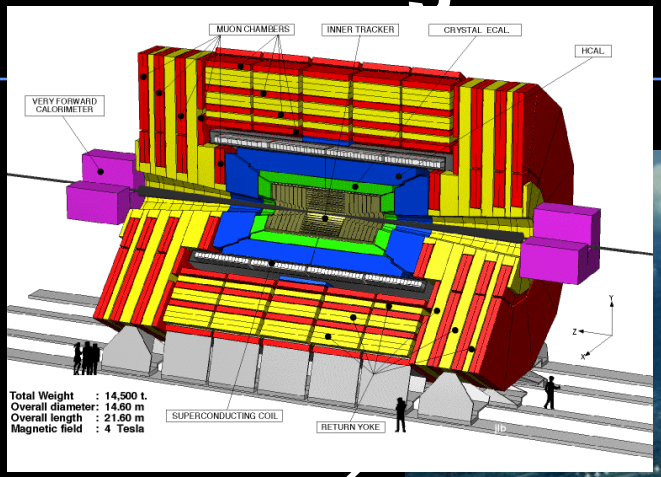

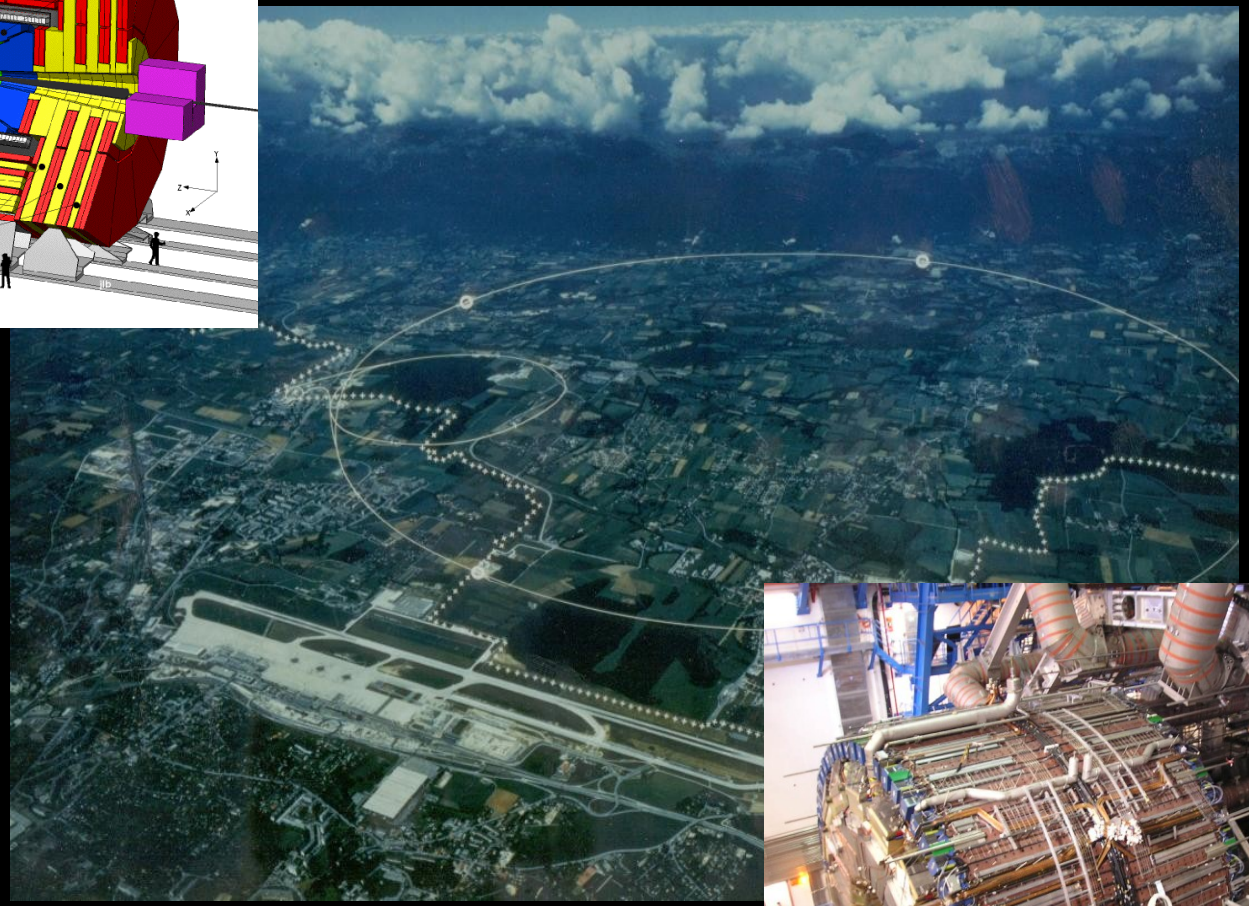

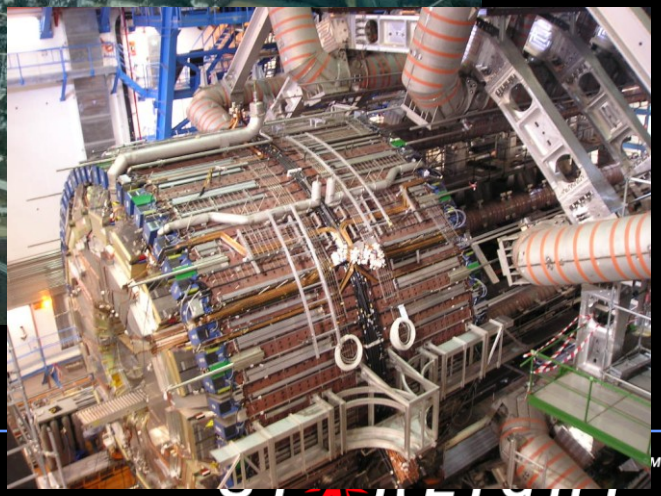

iCAIR

Network-Centric View of Large Hadron Collider (@CERN)

Q1: where does "discovery" occur?

CERN →T1	miles	kms
France	350	565
Italy	570	920
UK	625	1000
Netherlands	625	1000
Germany	700	1185
Spain	850	1400
Nordic	1300	2100
USA – New York	3900	6300
USA - Chicago	4400	7100
Canada – BC	5200	8400
Taiwan	6100	9850

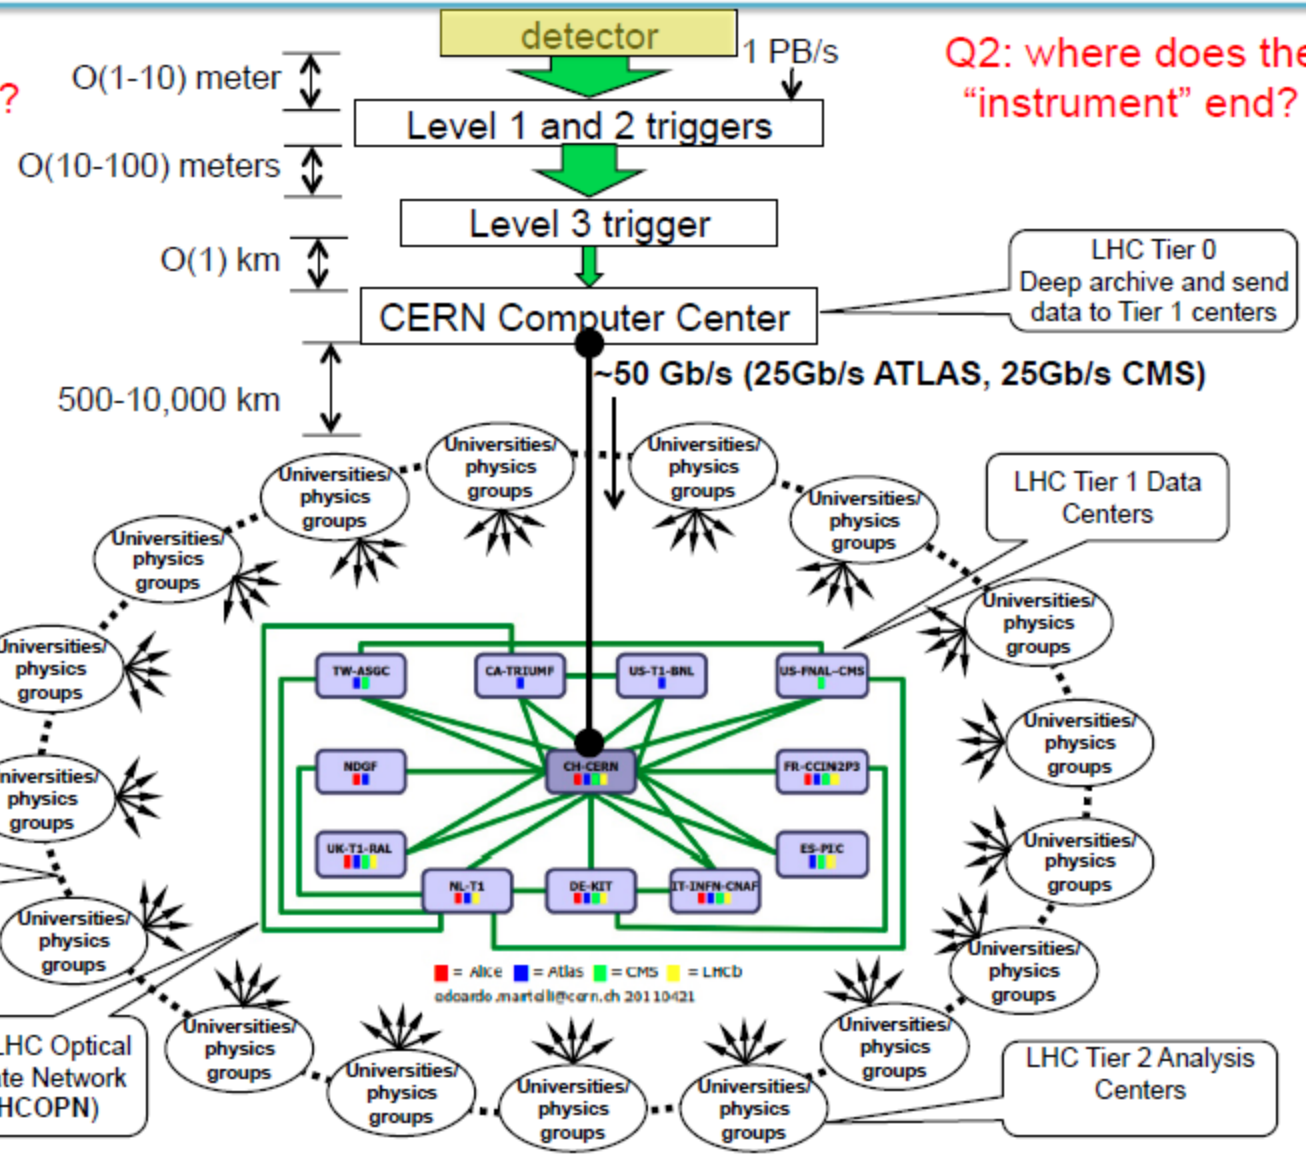

Q2: where does the "instrument" end?

Source: Bill Johnston

LHCOPN map

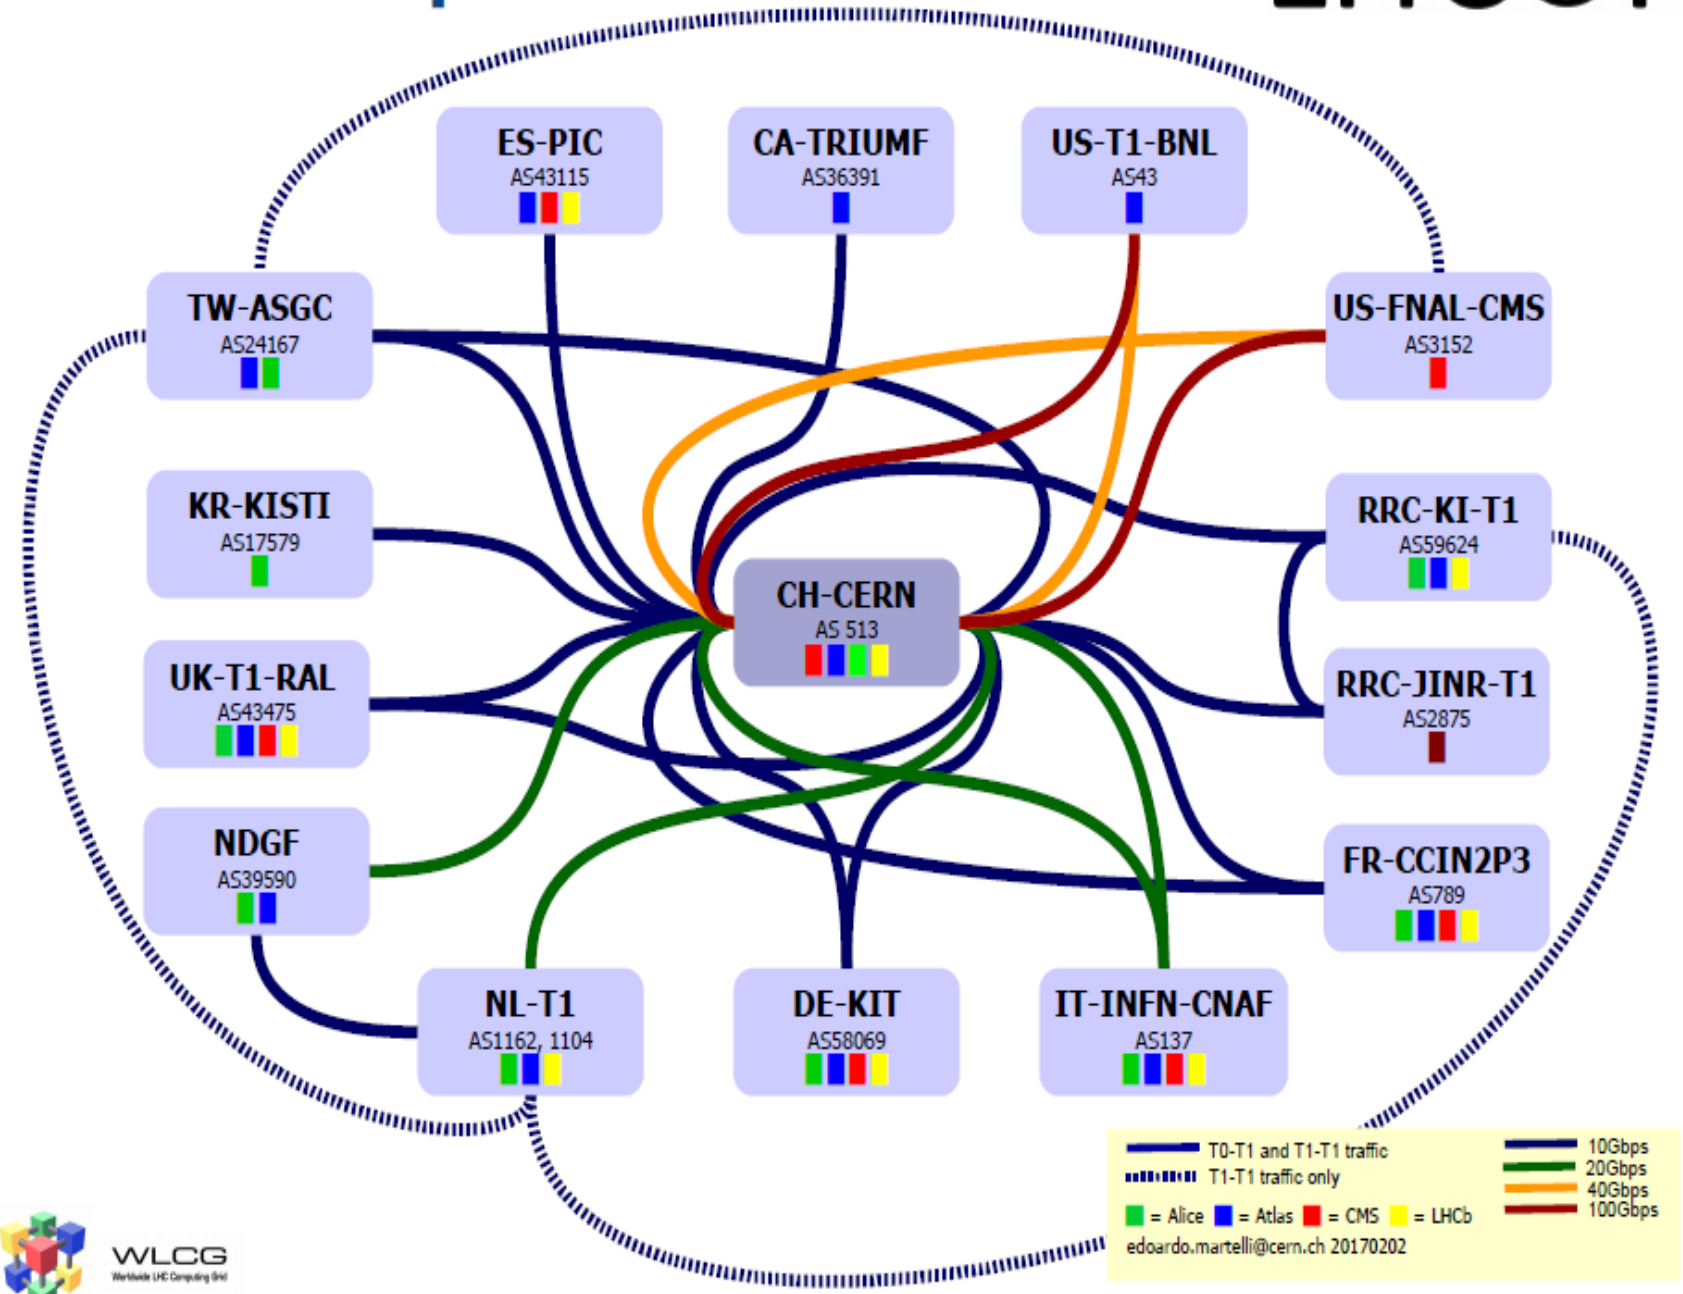
ES-PIC
AS43115
CA-TRIUMF
AS36391
US-T1-BNL
AS43
TW-ASGC
AS24167
US-FNAL-CMS
AS3152
KR-KISTI
AS17579
RRC-KI-T1
AS59624
CH-CERN
AS 513
UK-T1-RAL
AS43475
RRC-JINR-T1
AS2875
NDGF
AS39590
FR-CCIN2P3
AS789
NL-T1
AS1162, 1104
DE-KIT
AS58069
IT-INFN-CNAF
AS137
T0-T1 and T1-T1 traffic
T1-T1 traffic only
= Alice = Atlas = CMS = LHCb
edoardo.martelli@cern.ch 20170202
10Gbps
20Gbps
40Gbps
100Gbps

LHCONE: A global infrastructure for the High Energy Physics (LHC and Belle II) data management

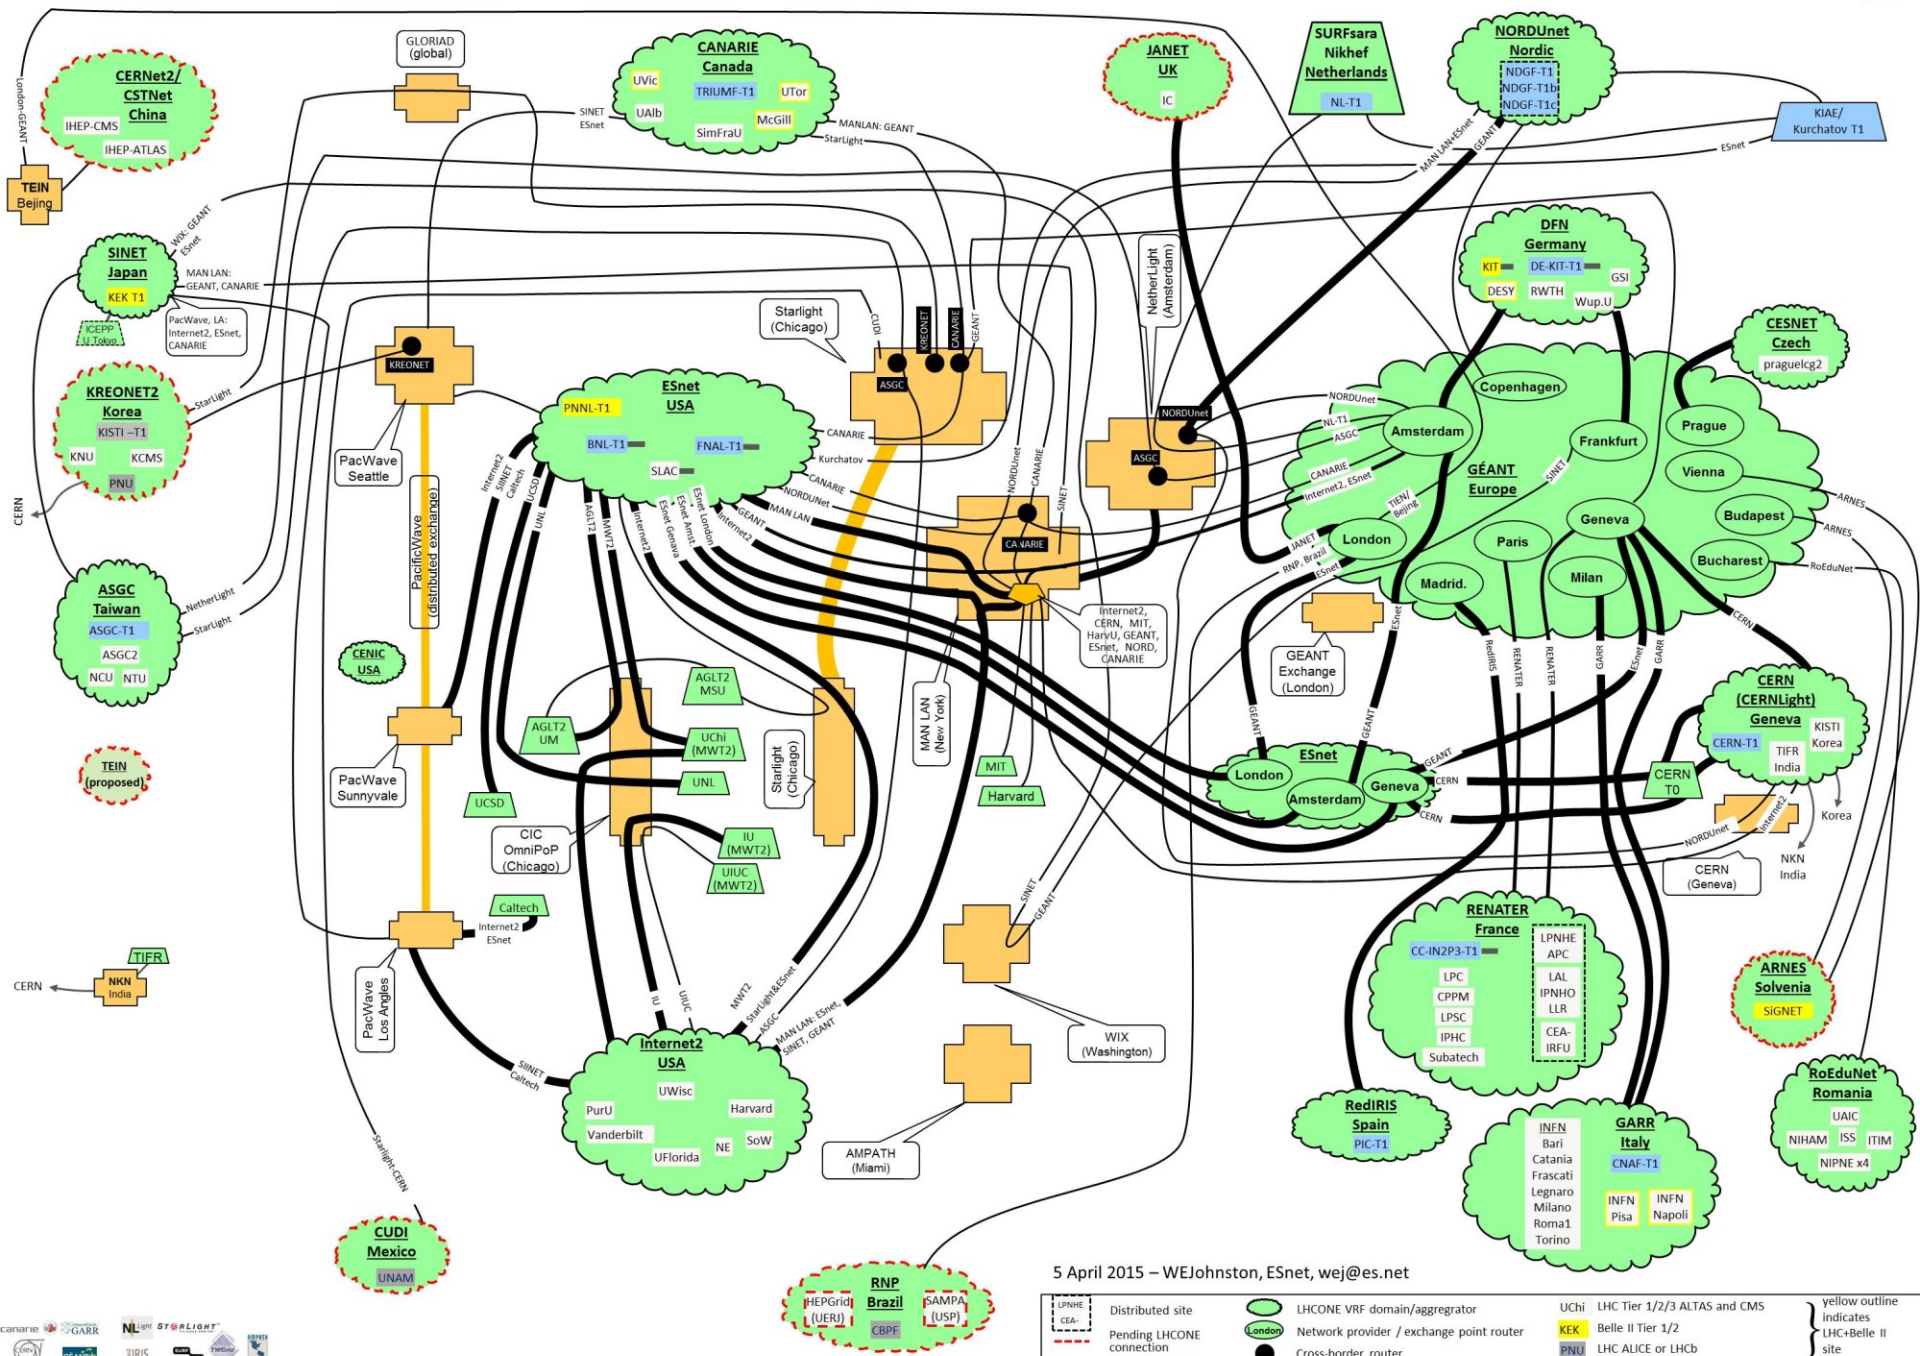

5 April 2015 – WEJohnston, ESnet, wej@es.net

Also see http://lhcone.net for details.

Distributed site
Pending LHCONE connection
Sites connected at 40G-100G
Broadcast VLAN
LHCONE VRF domain/aggregator
Network provider / exchange point router
Cross-border router
Regional R&E communication nexus w/ switch providing VLAN connections
UChi LHC Tier 1/2/3 ALTAS and CMS
KEK Belle II Tier 1/2
PNU LHC ALICE or LHCb
UNL Sites that are standalone VRFs,
yellow outline indicates LHC+Belle II site
Communication links: 1/10, 20/30/40, and 100Gb/s


Science
21 December 2012 | $10
BREAKTHROUGH
of the YEAR
The HIGGS
BOSON
AAAS

New Science Communities Using LHCONE

- Belle II Experiment, Particle Physics Experiment Designed To Study Properties of B Mesons (Heavy Particles Containing a Bottom Quark).
- Pierre Auger Observatory, Studying Ultra-High Energy Cosmic Rays, the Most Energetic and Rarest of Particles In the Universe.
- In August 2017 the PAO, LIGO and Virgo Collaboration Measured a Gravitational Wave Originating From a Binary Neutron Star Merger.
- The NOvA Experiment Is Designed To Answer Fundamental questions in neutrino Physics.
- The XENON Dark Matter Project Is a Global Collaboration Investing Fundamental Properties of Dark Matter, Largest Component Of The Universe.
- ProtoNUMA/NUMA – Collaborative Research On Nutrinos

Magnetic Fusion Energy

New Sources
Of Power

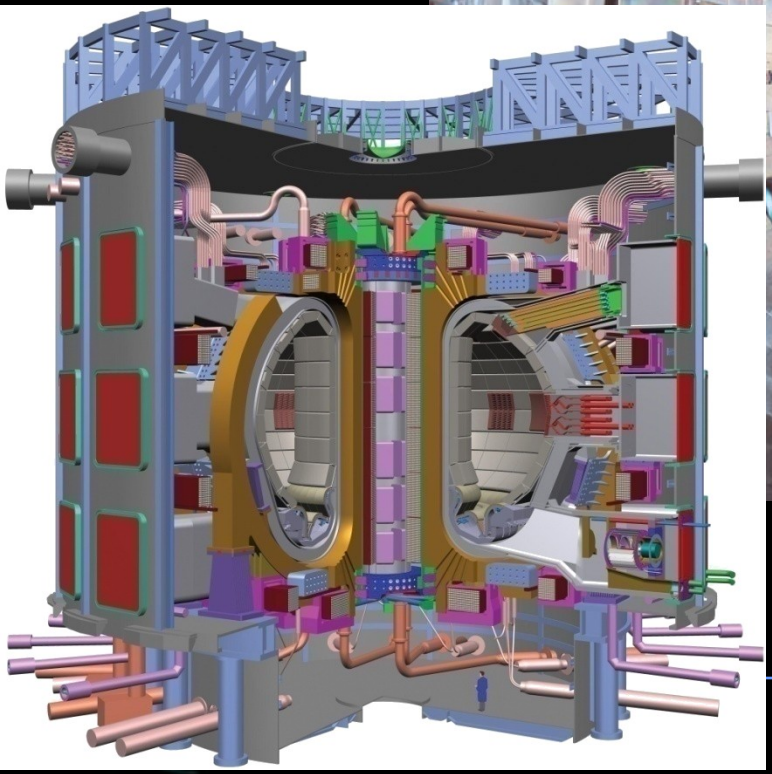

Source: DOE

Source: DOE

STARLIGHT SM

ITER (Formally- International Thermonuclear Experimental Reactor)

- ITER Is One of the World's Largest and Most Ambitious International Science Project Extremely Data Intensive

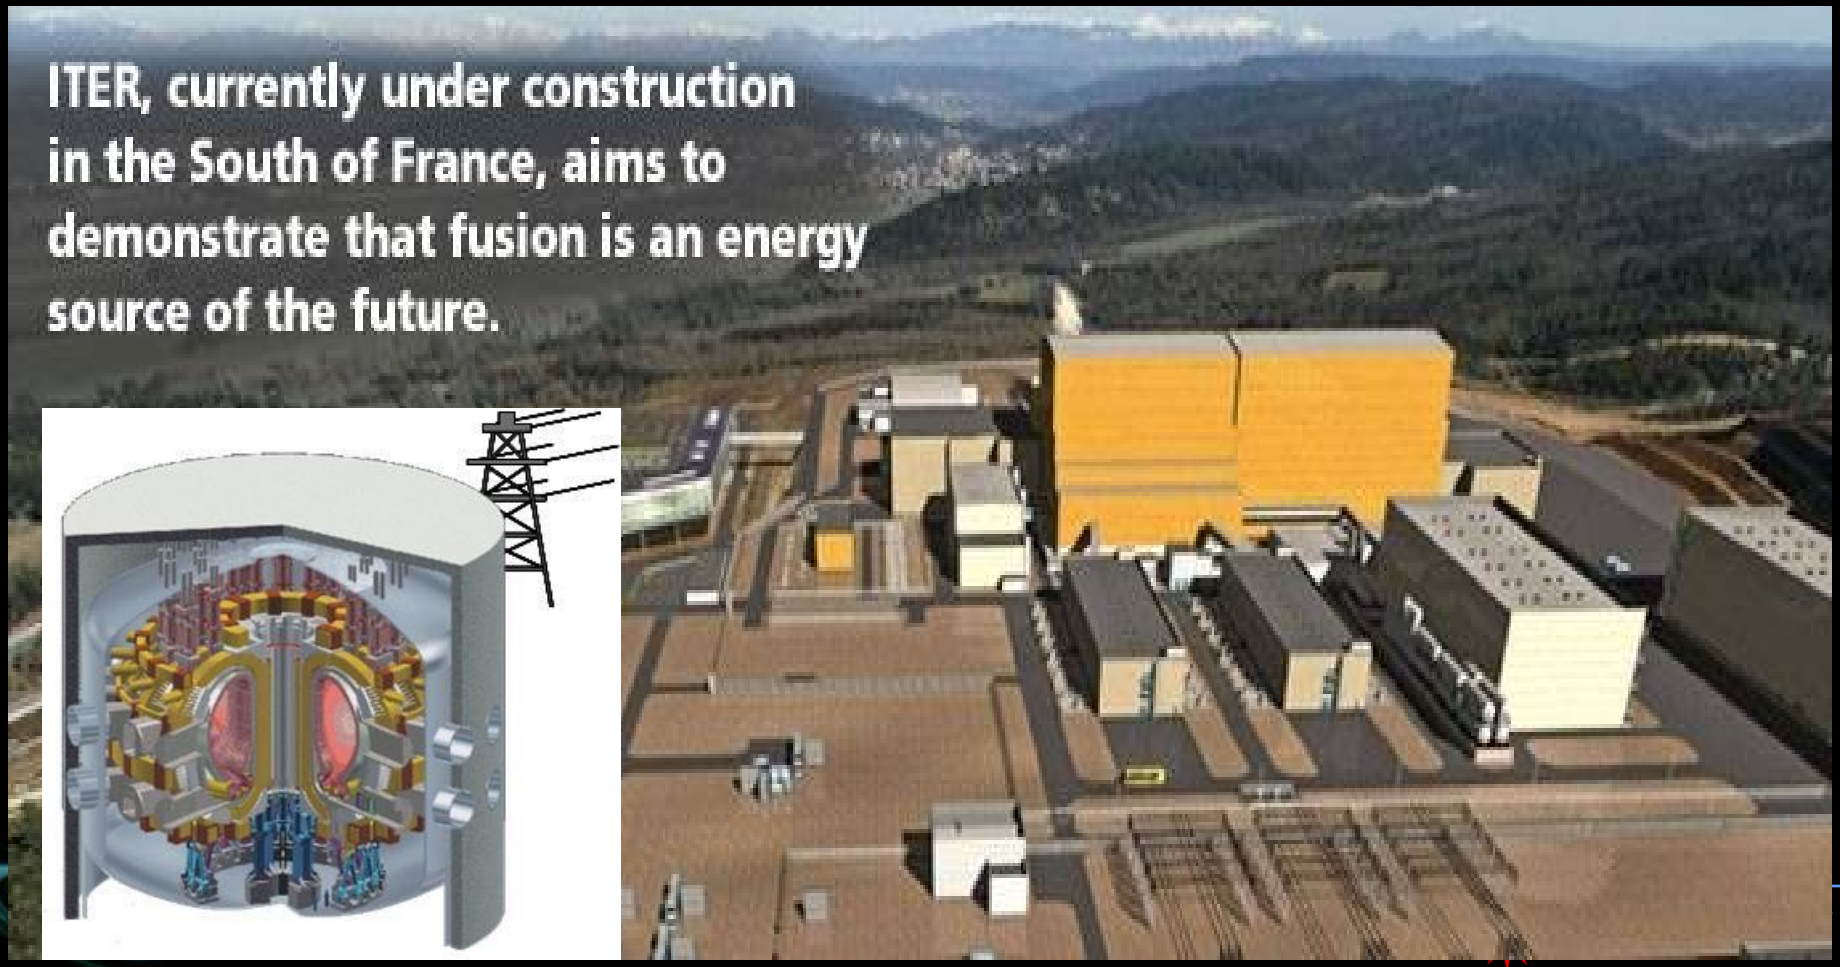

Fusion Energy Research

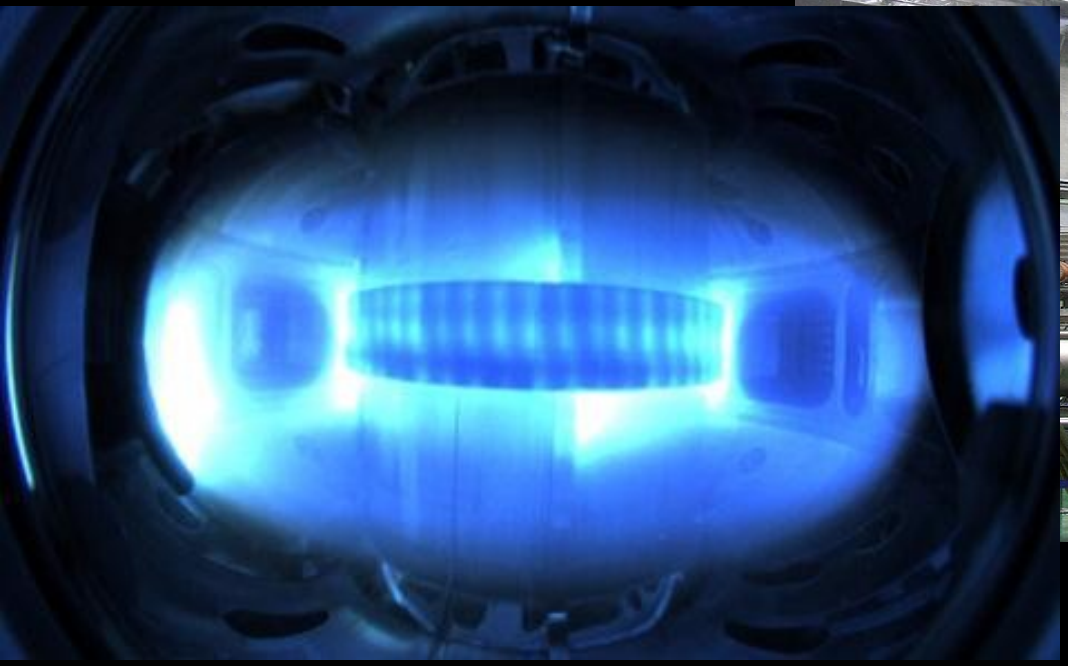

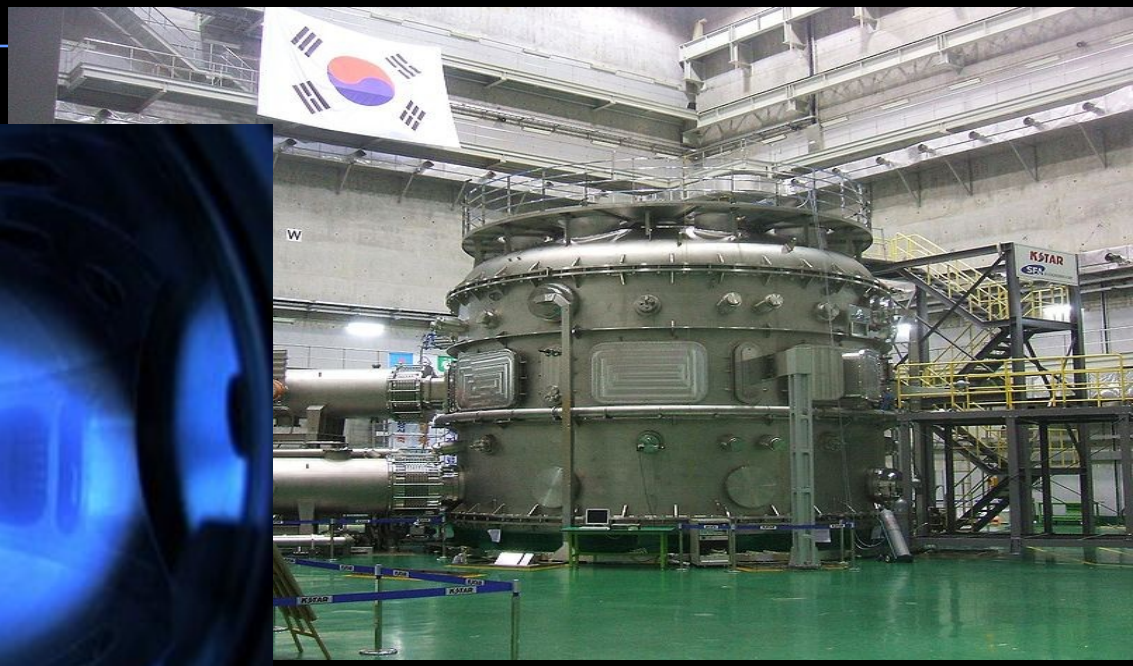

KSTAR, or Korea Superconducting Tokamak Advanced Research: Magnetic Fusion Device At the National Fusion Research Institute in Daejon, South Korea. KSTAR Is Providing Major Contributions To ITER.

Spallation Neutron Source (SNS) at ORNL

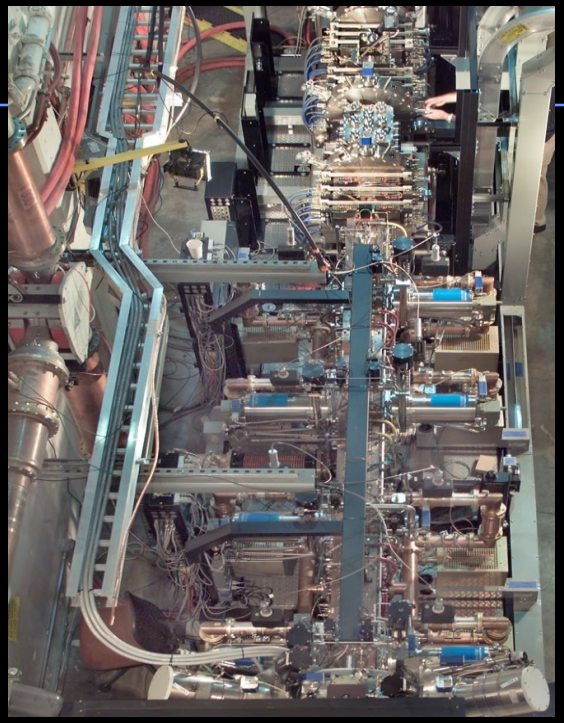

Neutron Beams Are Directed At Different Types of Materials To Investigate Their Atomic Properties, Including Structures

iCAIR

Source: DOE

Argonne National Laboratory Advanced Photon Source

STARLIGHT℠

Real-Time Global e-Very Long Baseline Interferometry

DRAGON (Dynamic Resource Allocation via GMPLS Optical Networks)

- Mid Atlantic Crossroads (MAX) GigaPoP, USA
- Information Sciences Institute, USA
- Westford Observatory, MIT Haystack, USA
- Goddard Geophysical and Atmospheric Observatory, NASA, USA
- Kashima, NiCT, Japan
- Onsala, Sweden
- Jodrell Bank, UK
- JIVE, The Netherlands
- Westerbork, Observatory/ ASTRON, The Netherlands

Real-time e-VLBI data correlation from telescopes in USA, Sweden, the Netherlands, UK and Japan

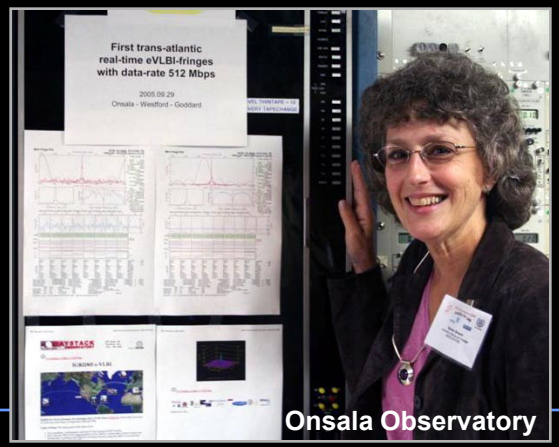

Onsala Observatory

http://dragon.maxgigapop.net

STARLIGHT SM

Square Kilometer Array

LSST Data Movement
Upcoming challenges for Astronomy

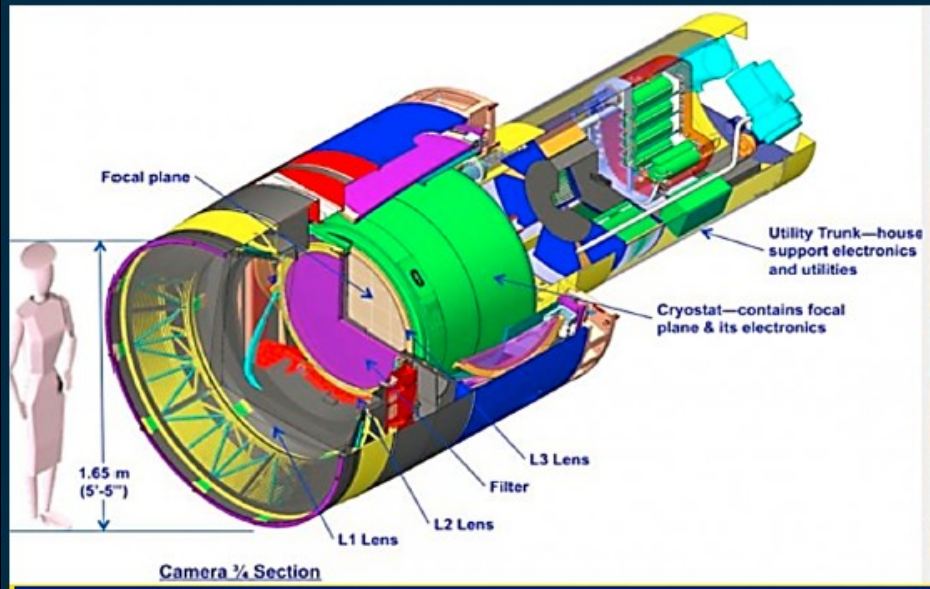

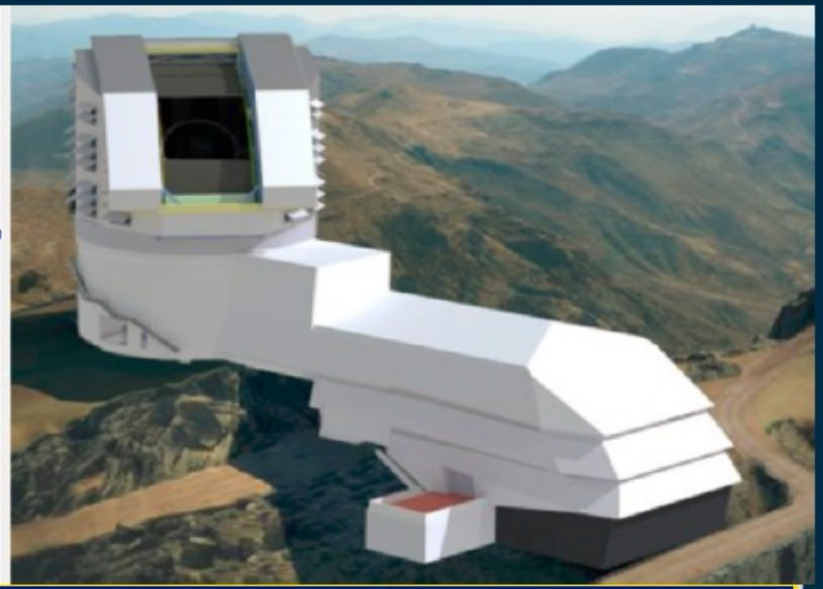

- 3.2 Gigapixel Camera with calibrated exposures at (10 Bytes / pixel)
- Planned Networks: Dedicated 100G for image data, Second 100G for other traffic, and 40G for diverse path
- Lossless compressed Image size = 2.7GB (~5 images transferred in parallel over a 100 Gbps link)
- UDP based custom image transfer protocols

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℠

iCAIR: Founding Partner of the Global Lambda Integrated Facility
Available Advanced Network Resources

Visualization courtesy of Bob Patterson, NCSA; data compilation by Maxine Brown, UIC.

www.glif.is

iCAIR

STARLIGHT℠

IRNC: RXP: StarLight SDX A Software Defined Networking 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
Jim Chen, Associate Director, International Center for Advanced Internet Research, Northwestern University

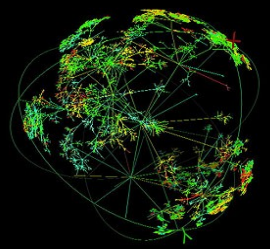

National Science Foundation
International Research Network Connections Program
Workshop
Chicago, Illinois
May 15, 2015

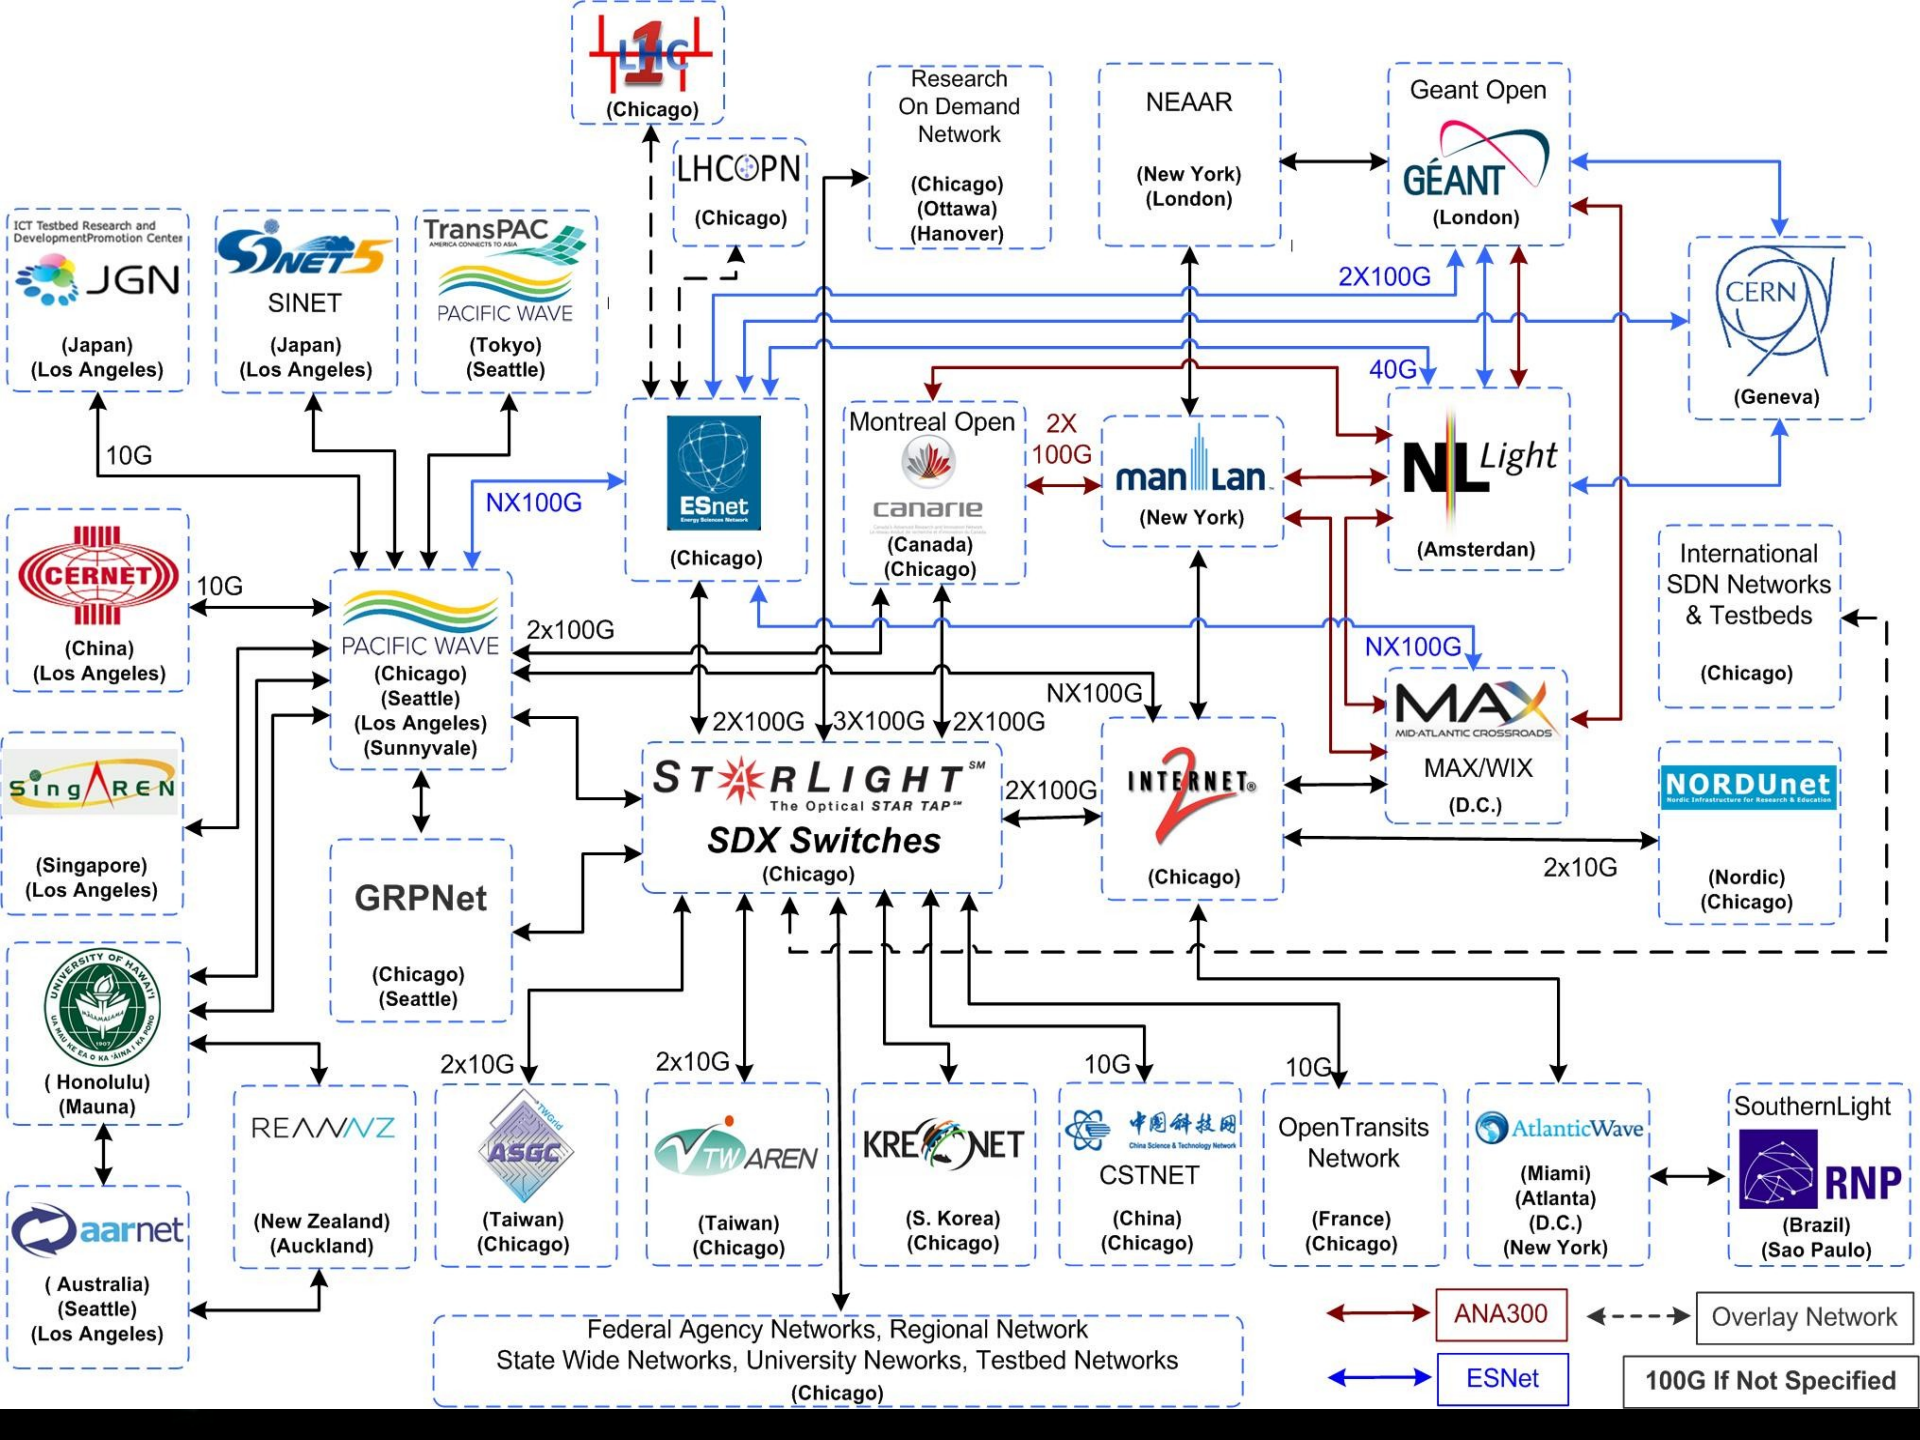
LHC
(Chicago)
LHCOPN
(Chicago)
Research On Demand Network
(Chicago)
(Ottawa)
(Hanover)
NEAAR
(New York)
(London)
Geant Open
GÉANT
(London)
CERN
(Geneva)
2X100G
40G
ICT Testbed Research and DevelopmentPromotion Center
JGN
(Japan)
(Los Angeles)
SINET
(Japan)
(Los Angeles)
TransPAC
PACIFIC WAVE
(Tokyo)
(Seattle)
10G
NX100G
ESnet
(Chicago)
Montreal Open
canarie
(Canada)
(Chicago)
2X 100G
manLan
(New York)
NL Light
(Amsterdan)
CERNET
(China)
(Los Angeles)
10G
PACIFIC WAVE
(Chicago)
(Seattle)
(Los Angeles)
(Sunnyvale)
2x100G
International SDN Networks & Testbeds
(Chicago)
NX100G
MAX
MID-ATLANTIC CROSSROADS
MAX/WIX
(D.C.)
2X100G
3X100G
2X100G
NX100G
SingAREN
(Singapore)
(Los Angeles)
STARLIGHT
The Optical STAR TAP
SDX Switches
(Chicago)
2X100G
INTERNET2
(Chicago)
2x10G
NORDUnet
(Nordic)
(Chicago)
GRPNet
(Chicago)
(Seattle)
UNIVERSITY OF HAWAI'I
(Honolulu)
(Mauna)
2x10G
2x10G
10G
10G
REANNZ
(New Zealand)
(Auckland)
ASGC
(Taiwan)
(Chicago)
TWAREN
(Taiwan)
(Chicago)
KREONET
(S. Korea)
(Chicago)
CSTNET
(China)
(Chicago)
OpenTransits Network
(France)
(Chicago)
AtlanticWave
(Miami)
(Atlanta)
(D.C.)
(New York)
SouthernLight
RNP
(Brazil)
(Sao Paulo)
aarnet
(Australia)
(Seattle)
(Los Angeles)
Federal Agency Networks, Regional Network
State Wide Networks, University Neworks, Testbed Networks
(Chicago)
ANA300
Overlay Network
ESnet
100G If Not Specified

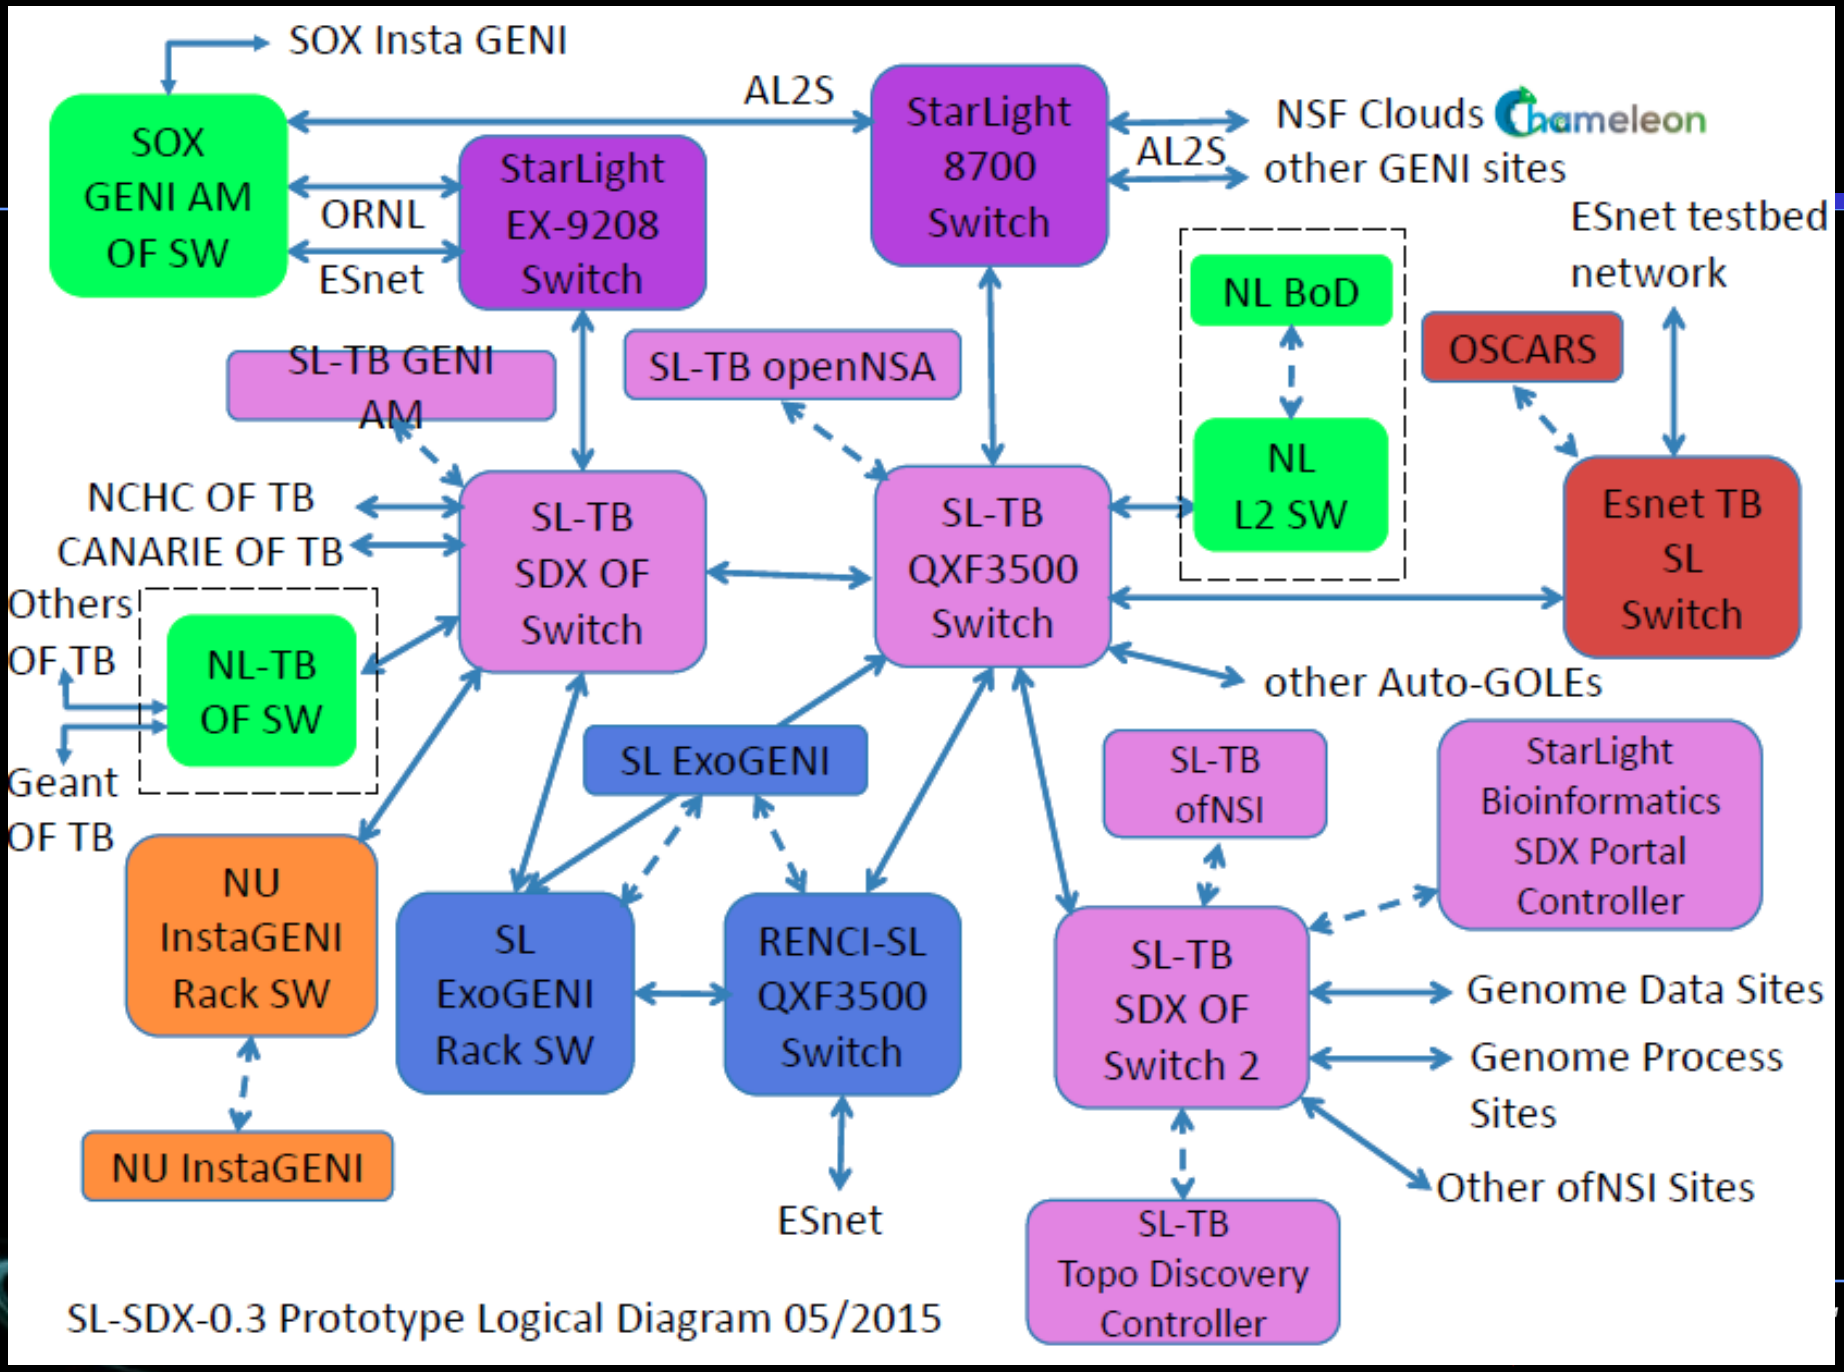
SOX Insta GENI
SOX GENI AM OF SW
AL2S
StarLight 8700 Switch
NSF Clouds
Chameleon
AL2S
other GENI sites
ORNL
ESnet
StarLight EX-9208 Switch
ESnet testbed network
NL BoD
OSCARS
SL-TB GENI AM
SL-TB openNSA
NL L2 SW
NCHC OF TB
CANARIE OF TB
SL-TB SDX OF Switch
SL-TB QXF3500 Switch
Esnet TB SL Switch
Others 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

Emerging Capabilities/Technologies

- Tenant Service Workflows Integrated With Foundation Infrastructure Workflows (Orchestration)
- Built-In Preconfigured Examples/Templates To Establish Infrastructure Foundation Workflows
- Zero-Touch “Playbooks” For Different Segments of Infrastructure Foundation Workflows After Implementing Initial Suites (e.g., Using Jupyter)
- Interactive Control Over Running Workflows
- Portability for Different Infrastructure Foundation Workflows
- Options/Capabilities for Specialized Customization
- Options For Real Time Visualization Of Individual Workflows, At Highly Granulated Levels

Opportunities For New Distributed Environments for Computational Science, Including Computer Science

- National Science Foundation's Ten Year Transition Point
- New Architecture, Services, Technologies
- National Research Platform (NRP) Workshop, Montana State University, August 2017.
- Regional Research Platforms?
- Asia Research Platform Being Discussed
- European Research Platform?
- Emerging Next Gen Distributed SD Infrastructure
- Example: NSF Concept of A Distributed Fabrics For 'Plug-In" Testbeds, e.g., Global Environment for Network Innovations (GENI), Platform for Advanced Wireless Research (PAWR), US Ignite, NSFCloud Initiative, IoT, Edge Clouds, et al
- Potential for Network Science Research Platform

Global Research Platform (GRP)

- A Emerging International Fabric
- A Specialized Globally Distributed Environment/Platform For Science Discovery and Innovation
- Based On State-Of-the-Art-Clouds, Networks, Storage Systems, Data Repositories, etc
- Interconnected With Computational Grids, Supercomputing Centers, Specialized Instruments, et al
- Also, Based On World-Wide 100 Gbps (Soon 100 G+) Networks
- Leveraging Advanced Architectural Concepts, e.g., SDN/SDX/SDI – Science DMZs
- Ref: 1st Demonstrations @ SC15, Austin Texas November 2015
- Subsequent Demonstrations @ SC16 Salt Lake City Utah, November 2016, Global LambdaGrid Workshop 2016 and 2017,
- Major Demonstrations at SC17 in Denver, Colorado

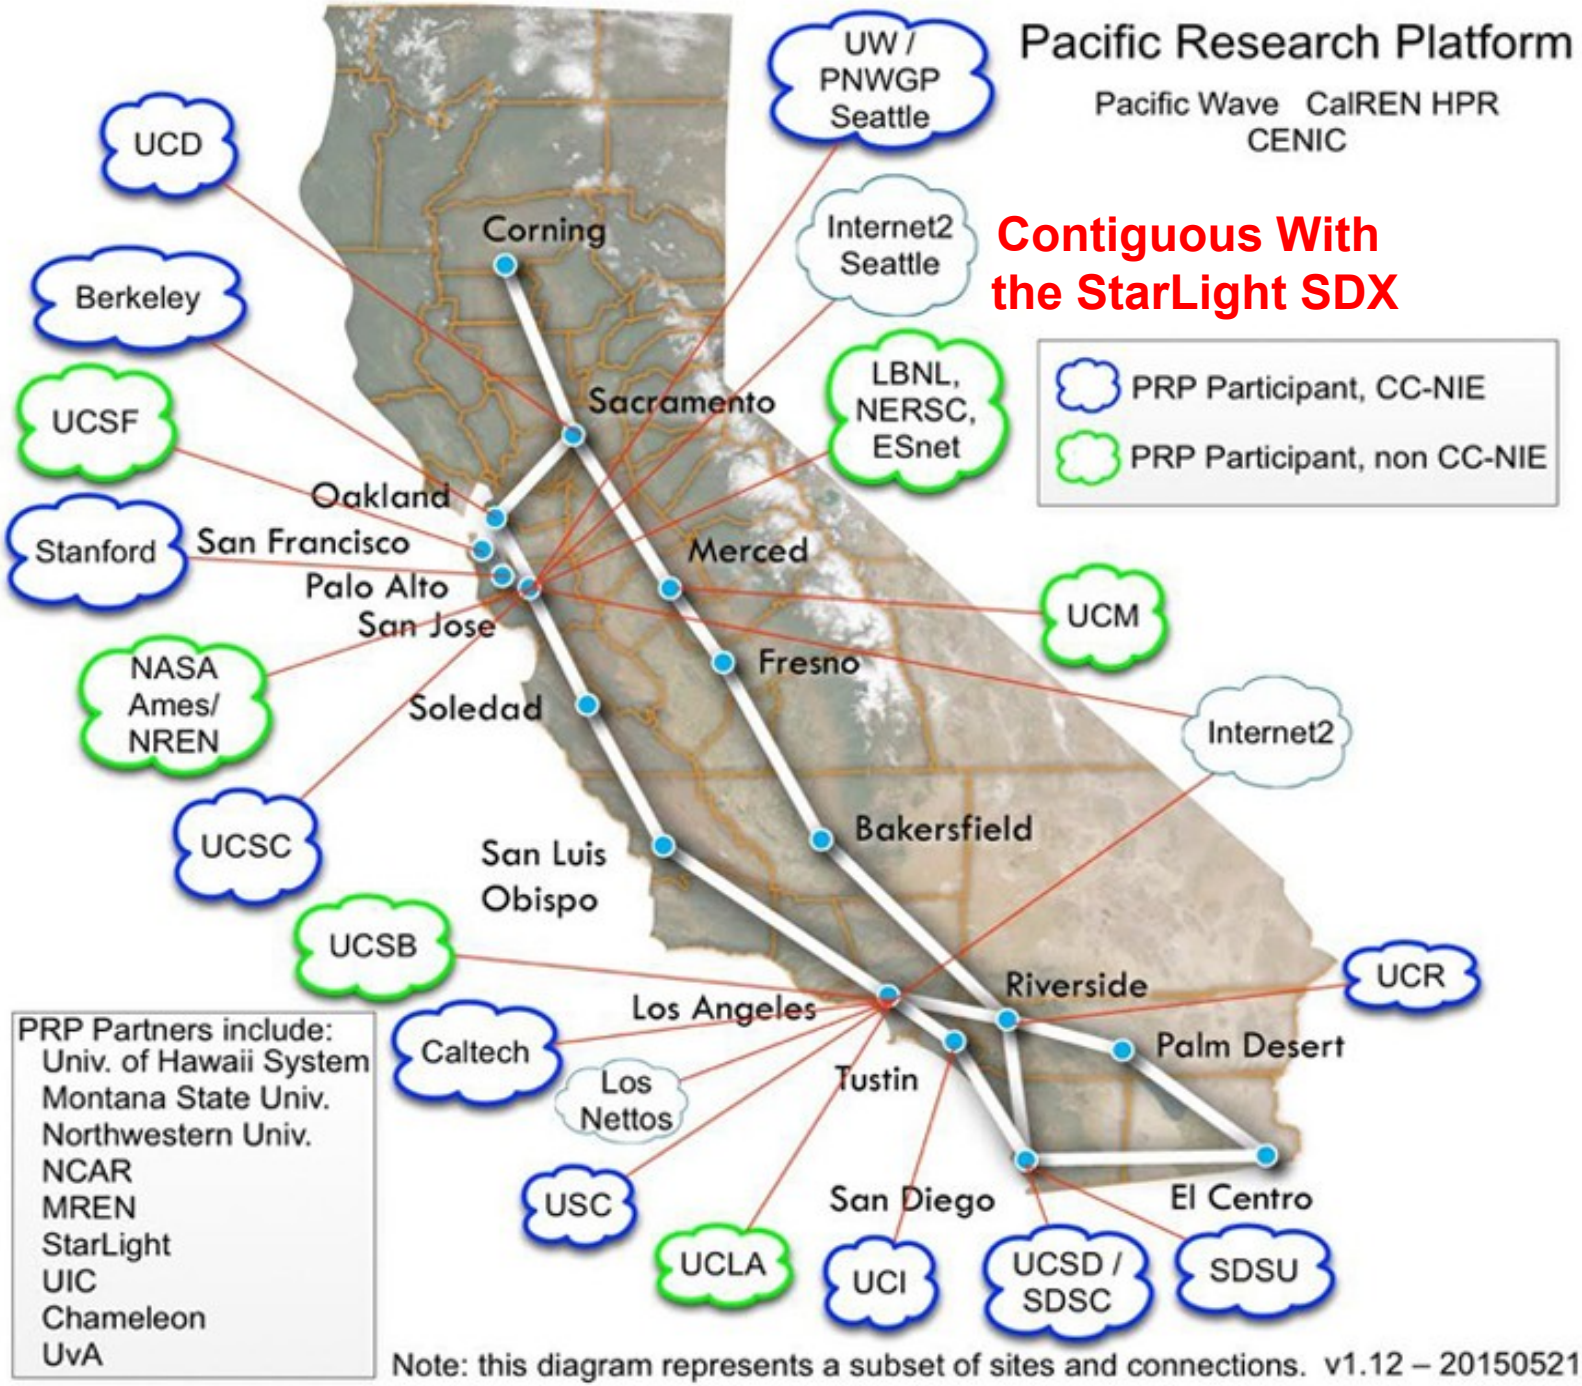
Pacific Research Platform
Pacific Wave CalREN HPR
CENIC
Contiguous With
the StarLight SDX
PRP Participant, CC-NIE
PRP Participant, non CC-NIE
UW / PNWGP Seattle
Internet2 Seattle
LBNL, NERSC, ESnet
UCD
Berkeley
UCSF
Stanford
NASA Ames/ NREN
UCSC
UCSB
Caltech
Los Nettos
USC
UCLA
UCI
UCSD / SDSC
SDSU
UCM
Internet2
UCR
Corning
Sacramento
Oakland
San Francisco
Palo Alto
San Jose
Soledad
Merced
Fresno
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

July 21, 2017

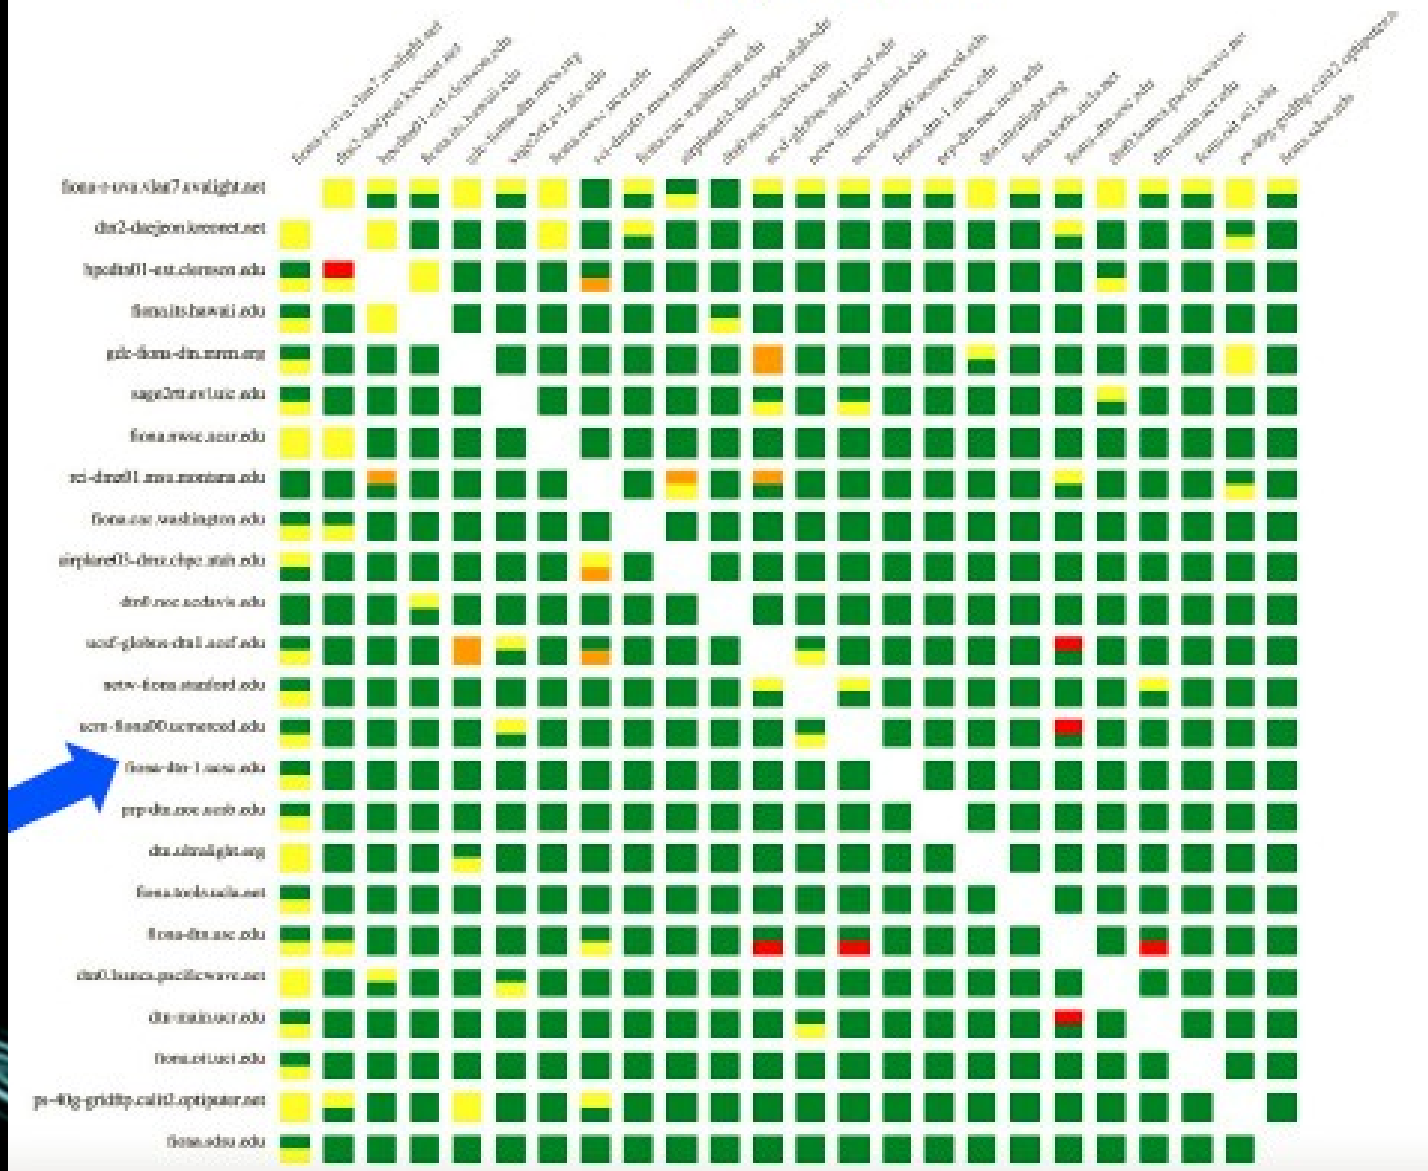
fiona-1-ovs.ultralight.net
dtn2-daejeon.kreonet.net
hpcdtn01-ext.clemson.edu
fiona.its.hawaii.edu
gdc-fiona-dtn.mrnn.org
sage2rtt.evl.uic.edu
fiona.nwsc.ucar.edu
sci-dtn01.nrao.manitara.edu
fiona.cse.washington.edu
airplane03-dtn2.chpc.utah.edu
dtn0.nsrc.ucdavis.edu
ucsf-globus-dtn1.ucsf.edu
netw-fiona.stanford.edu
scm-fiona00.uchicago.edu
fiona-dtn-1.nersc.edu
prp-dtn.nrec.ucsb.edu
dtn.ultralight.org
fiona.tools.sacla.net
fiona-dtn.usc.edu
dtn0.hanea.pacificwave.net
dtn-tmin.ucr.edu
fiona.crt.uci.edu
ps-40g-griddtp.calit2.optiputer.net
fiona.sdsu.edu


NSF


it2

Global Research Platform: Building On CENIC/Pacific Wave, GLIF and GLIF GOLEs (e.g., StarLight et al)

Current International GRP Partners

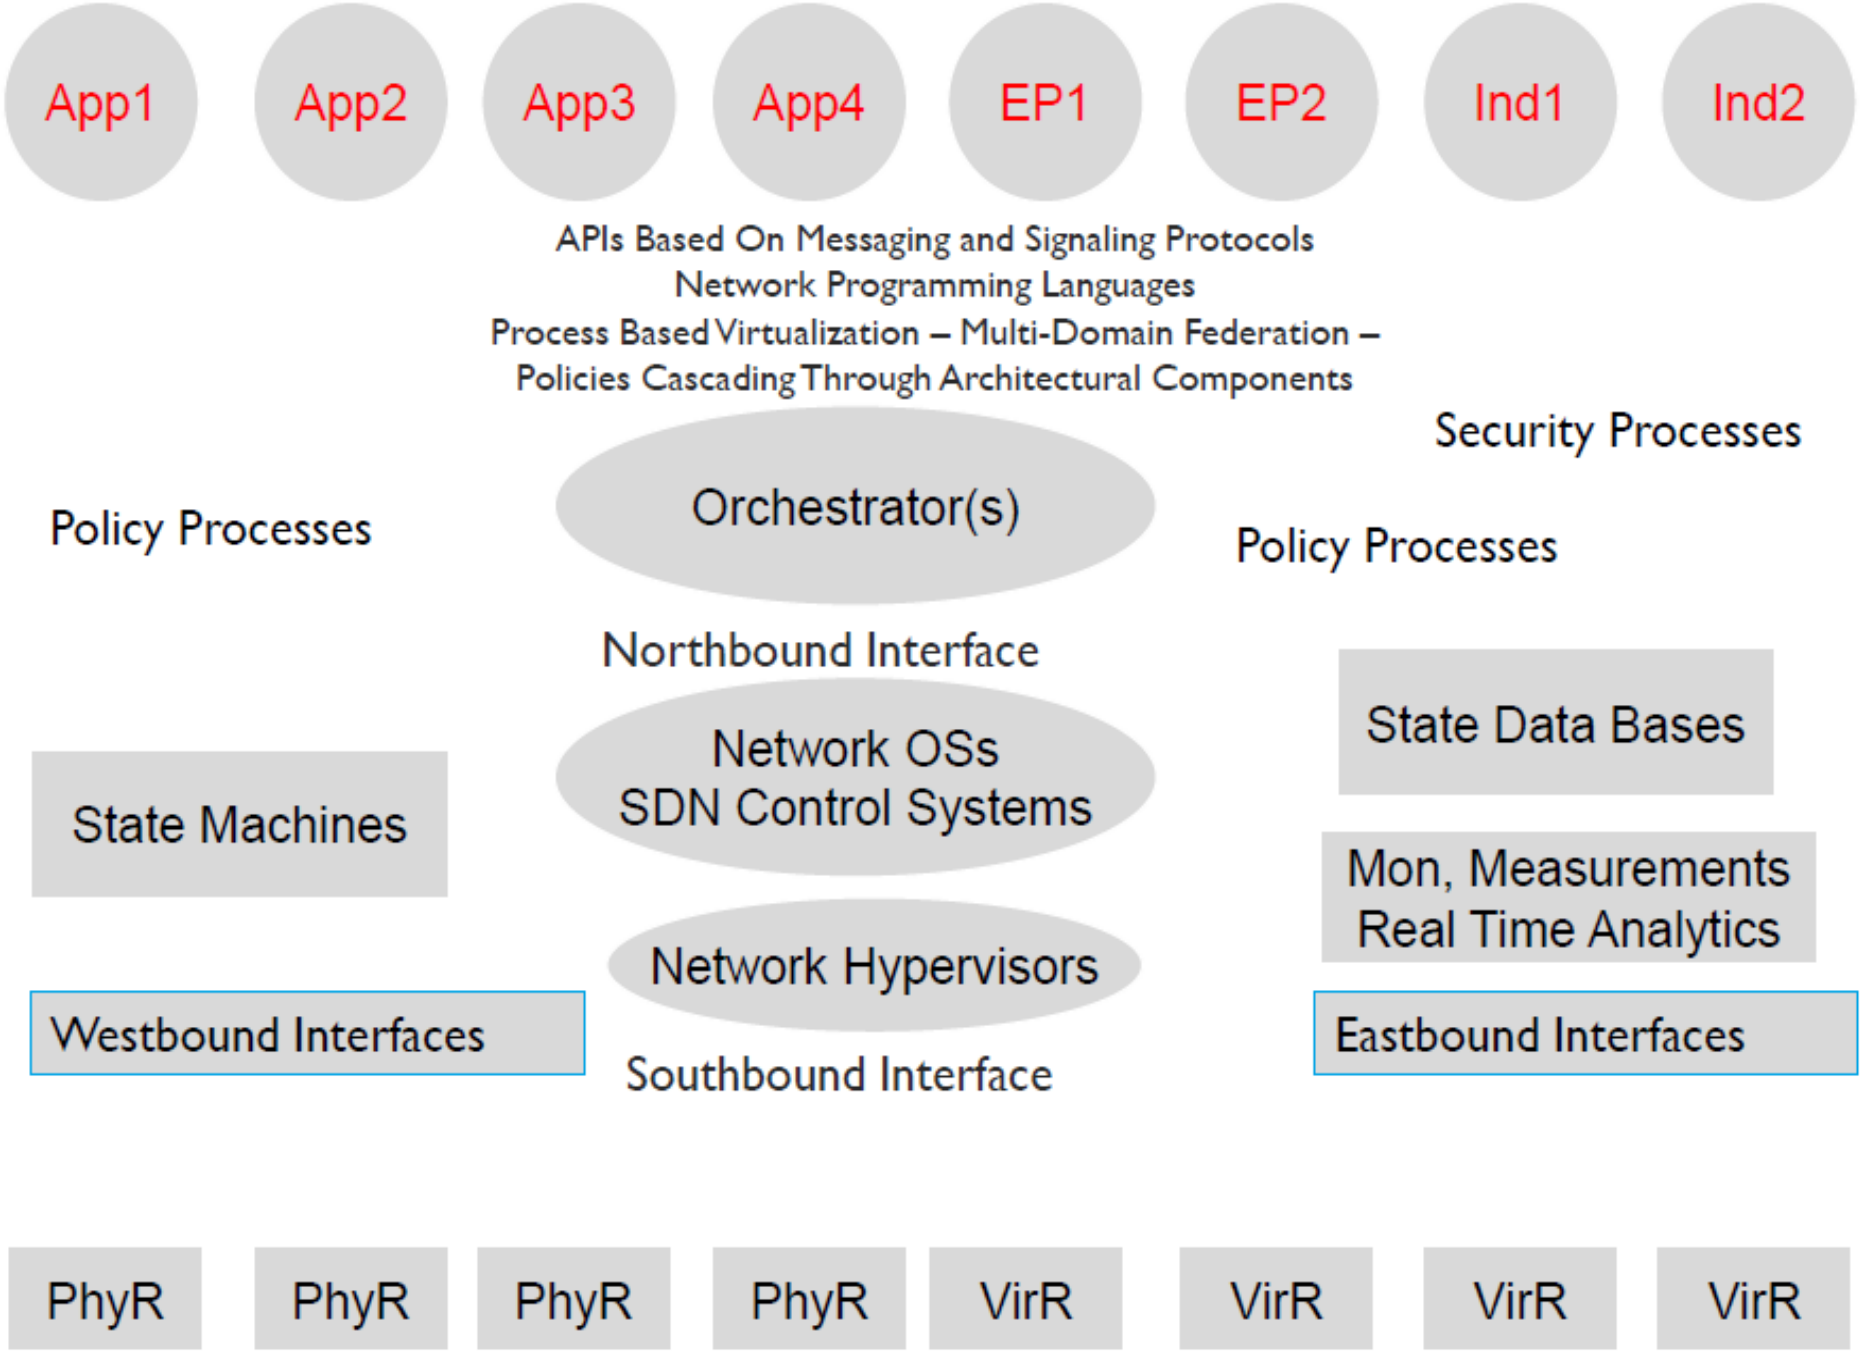
App1
App2
App3
App4
EP1
EP2
Ind1
Ind2
APIs Based On Messaging and Signaling Protocols
Network Programming Languages
Process Based Virtualization – Multi-Domain Federation –
Policies Cascading Through Architectural Components
Security Processes
Orchestrator(s)
Policy Processes
Policy Processes
Northbound 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

National Science Foundation's Global Environment for Network Innovations (GENI)

- GENI: Virtual Laboratory For Networking and Distributed Systems Research and Education.
- GENI Is Being Used To Explore Networks At Scale, Promoting Innovations In Network Science, Services, Security, Operations, And Applications.
- GENI Anticipates Future Communications Infrastructure Design and Technology (e.g., Flexible and Programmable vs Rigid and Static, Decentralized vs Centralized, Customizable vs Predetermined).
- Foundations Consists Of Clouds and Private Networks

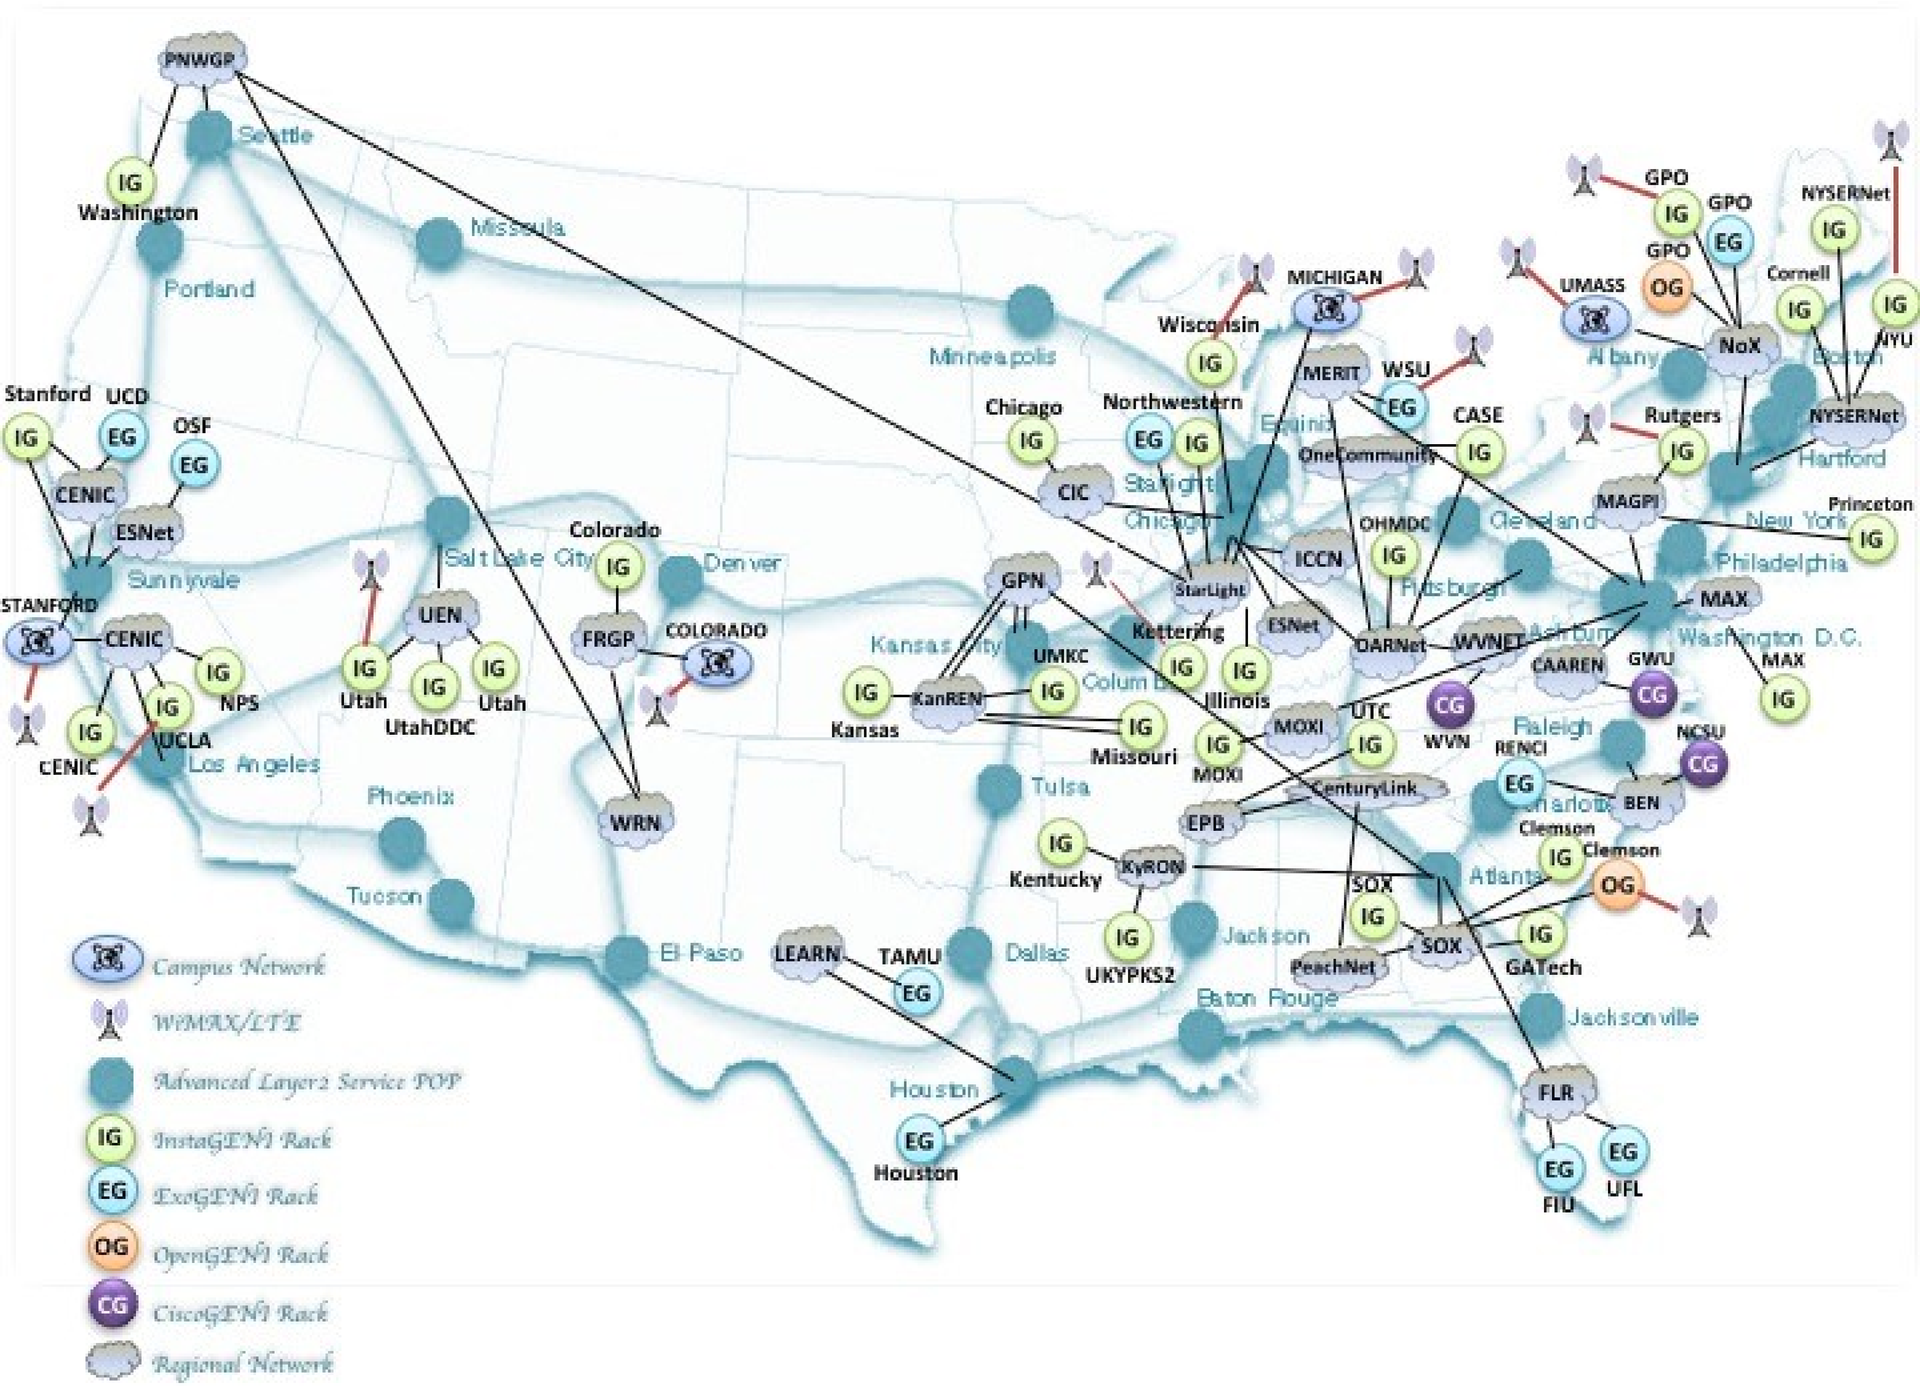
PNWGP
Seattle
IG
Washington
Missoula
Portland
Minneapolis
Stanford
UCD
OSF
CENIC
ESNet
Sunnyvale
STANFORD
CENIC
NPS
cENIC
UCLA
Los Angeles
Utah
UEN
UtahDDC
Salt Lake City
Colorado
FRGP
Denver
COLORADO
WRN
Phoenix
Tucson
El Paso
MICHIGAN
Wisconsin
MERIT
WSU
Chicago
Northwestern
CIC
Starlight
Chicago
Equinix
OneCommunity
CASE
OHMDC
ICCN
StarLight
ESNet
Kettering
Illinois
OARNet
Pittsburgh
Cleveland
WVNET
Ashburn
GPN
Kansas City
UMKC
Columbia
KanREN
Kansas
Missouri
MOXI
UTC
Tulsa
EPB
CenturyLink
Kentucky
KyRON
UKYPKS2
LEARN
TAMU
Dallas
Houston
Jackson
Baton Rouge
PeachNet
SOX
GATech
Atlanta
Clemson
Jacksonville
FLR
FIU
UFL
GPO
NYSERNet
UMASS
Cornell
NoX
NYU
Albany
Boston
Rutgers
Hartford
MAGPI
New York
Princeton
Philadelphia
MAX
Washington D.C.
CAAREN
GWU
WVN
Raleigh
RENCI
NCSU
BEN
Charlotte
Campus Network
WiMAX/LTE
Advanced Layer2 Service POP
InstaGENI Rack
ExoGENI Rack
OpenGENI Rack
CiscoGENI Rack
Regional Network

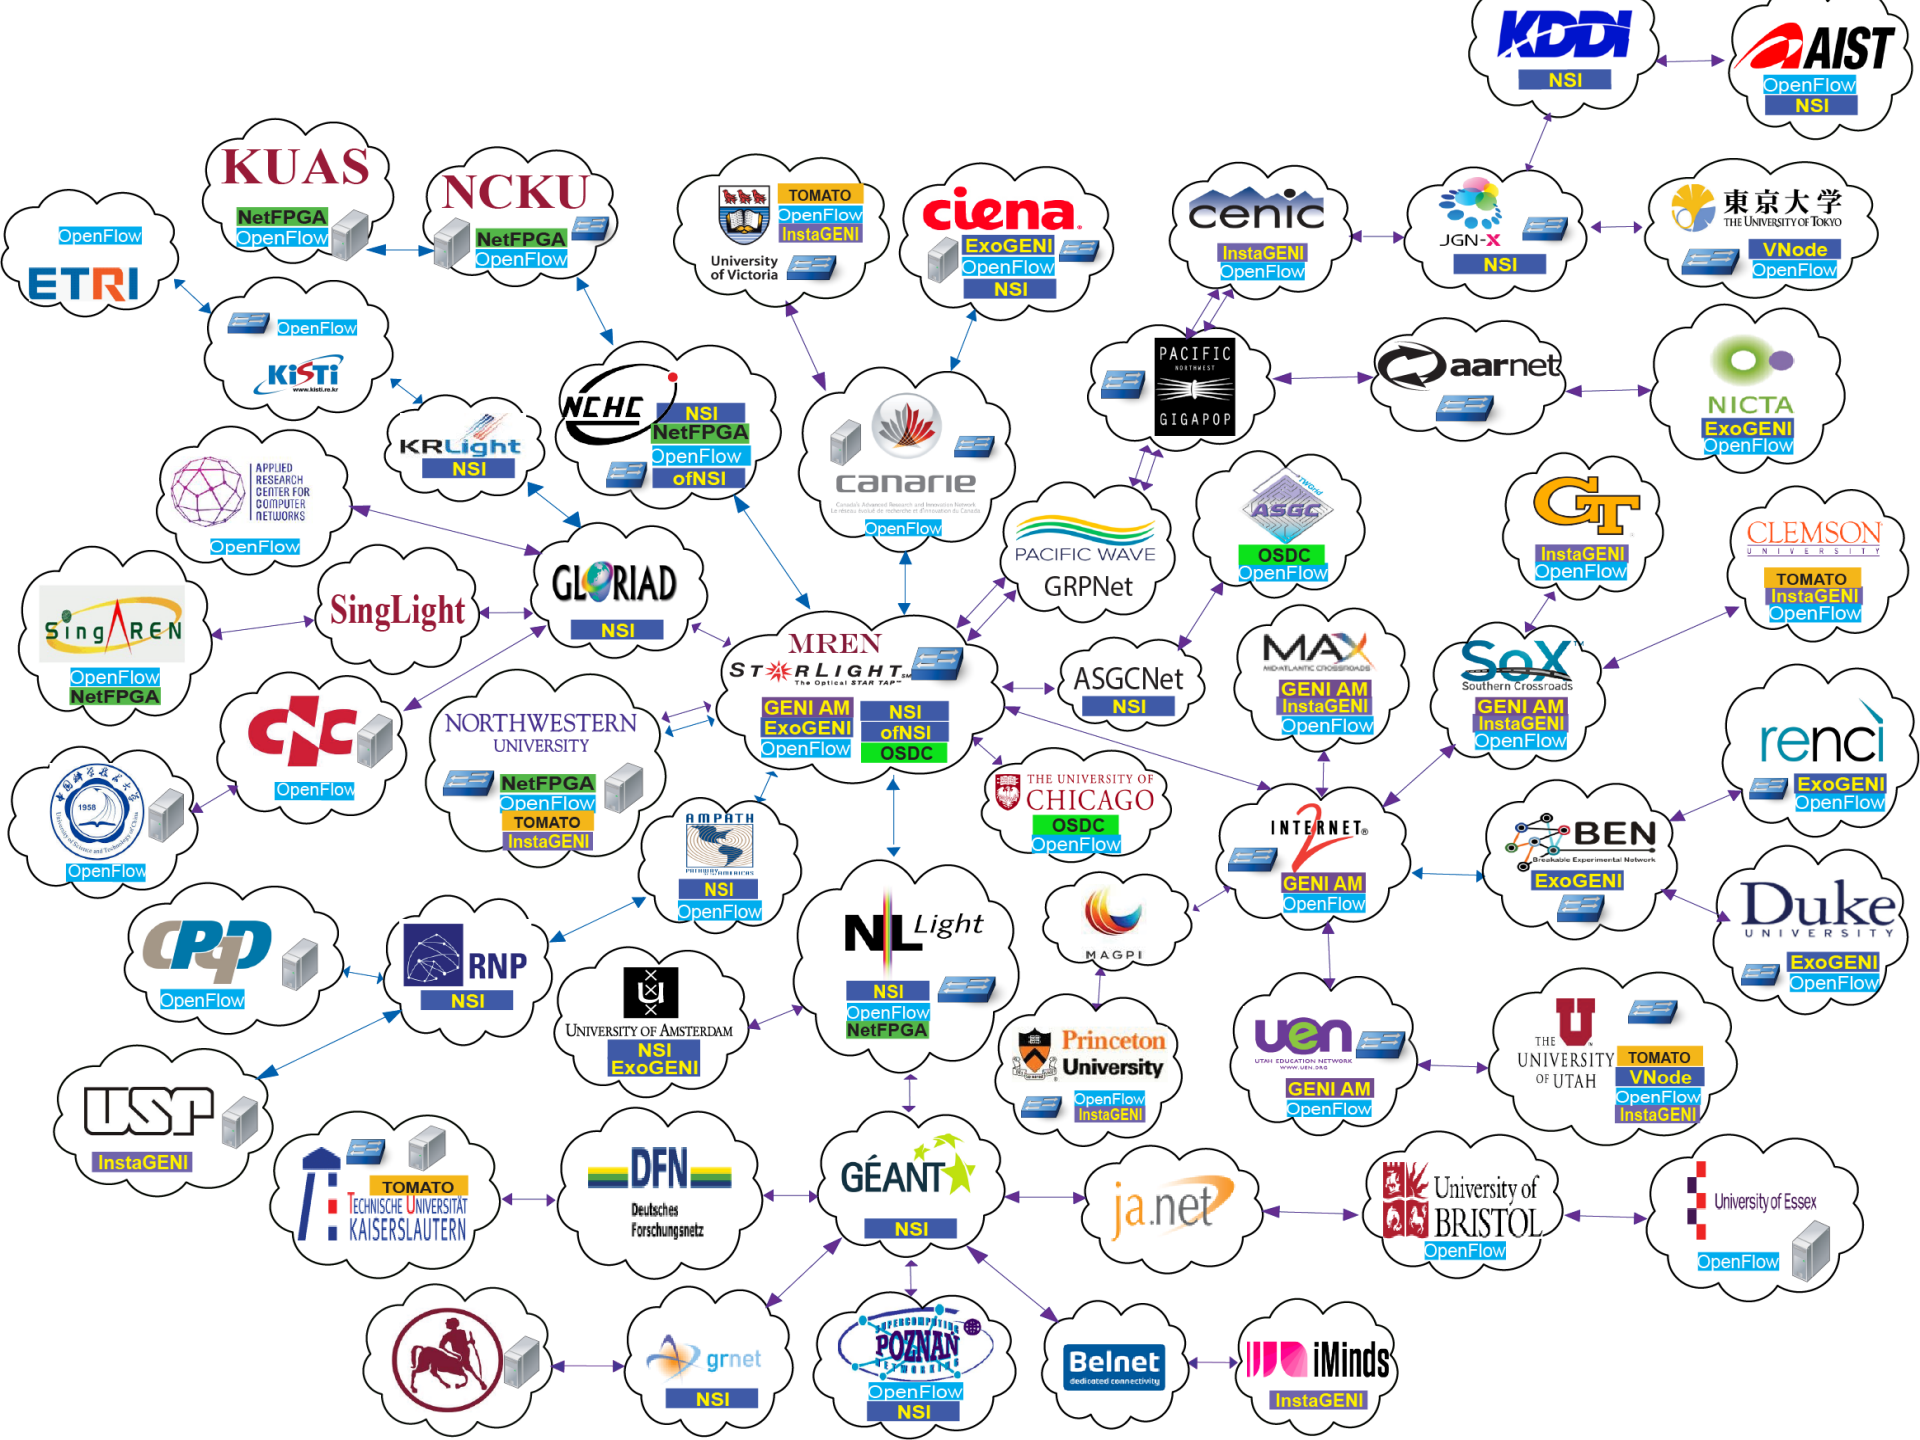
KDDI
NSI
AIST
OpenFlow
NSI
KUAS
NetFPGA
OpenFlow
NCKU
NetFPGA
OpenFlow
TOMATO
OpenFlow
InstaGENI
University of Victoria
ciena
ExoGENI
OpenFlow
NSI
cenic
InstaGENI
OpenFlow
JGN-X
NSI
東京大学
THE UNIVERSITY OF TOKYO
VNode
OpenFlow
OpenFlow
ETRI
OpenFlow
KiSTi
NCHC
NSI
NetFPGA
OpenFlow
ofNSI
KRLight
NSI
canarie
OpenFlow
PACIFIC NORTHWEST GIGAPOP
aarnet
NICTA
ExoGENI
OpenFlow
APPLIED RESEARCH CENTER FOR COMPUTER NETWORKS
OpenFlow
GLORIAD
NSI
PACIFIC WAVE
GRPNet
ASGC
OSDC
OpenFlow
GT
InstaGENI
OpenFlow
CLEMSON UNIVERSITY
TOMATO
InstaGENI
OpenFlow
SingAREN
OpenFlow
NetFPGA
SingLight
MREN
STARLIGHT
The Optical STAR TAP
GENI AM
ExoGENI
OpenFlow
NSI
ofNSI
OSDC
ASGCNet
NSI
MAX
MID-ATLANTIC CROSSROADS
GENI AM
InstaGENI
OpenFlow
SoX
Southern Crossroads
GENI AM
InstaGENI
OpenFlow
renci
ExoGENI
OpenFlow
CNC
OpenFlow
NORTHWESTERN UNIVERSITY
NetFPGA
OpenFlow
TOMATO
InstaGENI
THE UNIVERSITY OF CHICAGO
OSDC
OpenFlow
INTERNET2
GENI AM
OpenFlow
BEN
Boston Educational Network
ExoGENI
Duke UNIVERSITY
ExoGENI
OpenFlow
OpenFlow
AMPATH
NSI
OpenFlow
CPqD
OpenFlow
RNP
NSI
NL Light
NSI
OpenFlow
NetFPGA
MAGPI
UNIVERSITY OF AMSTERDAM
NSI
ExoGENI
Princeton University
OpenFlow
InstaGENI
uen
UTAH EDUCATION NETWORK
GENI AM
OpenFlow
THE UNIVERSITY OF UTAH
TOMATO
VNode
OpenFlow
InstaGENI
USP
InstaGENI
TOMATO
TECHNISCHE UNIVERSITÄT KAISERSLAUTERN
DFN
Deutsches Forschungsnetz
GÉANT
NSI
ja.net
University of BRISTOL
OpenFlow
University of Essex
OpenFlow
grnet
NSI
POZNAN
OpenFlow
NSI
Belnet
dedicated connectivity
iMinds
InstaGENI

Springer

springer.com

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

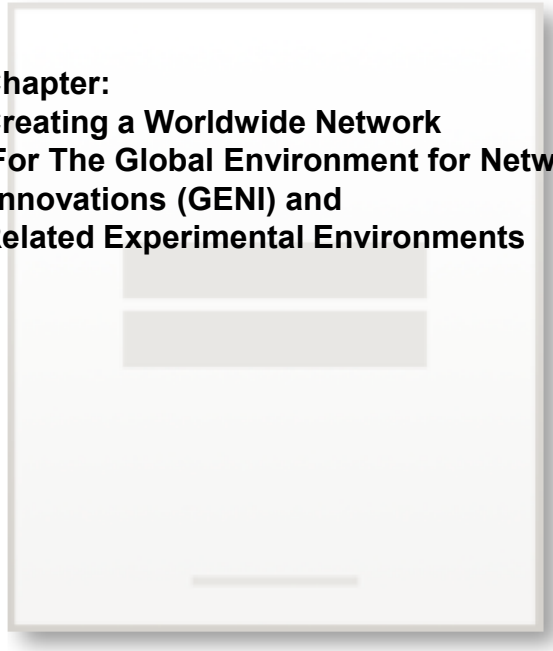

1st ed. 2016, XVIII, 655 p. 216 illus., 183 illus. in color.

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.

Emerging US SDX Interoperable Fabric

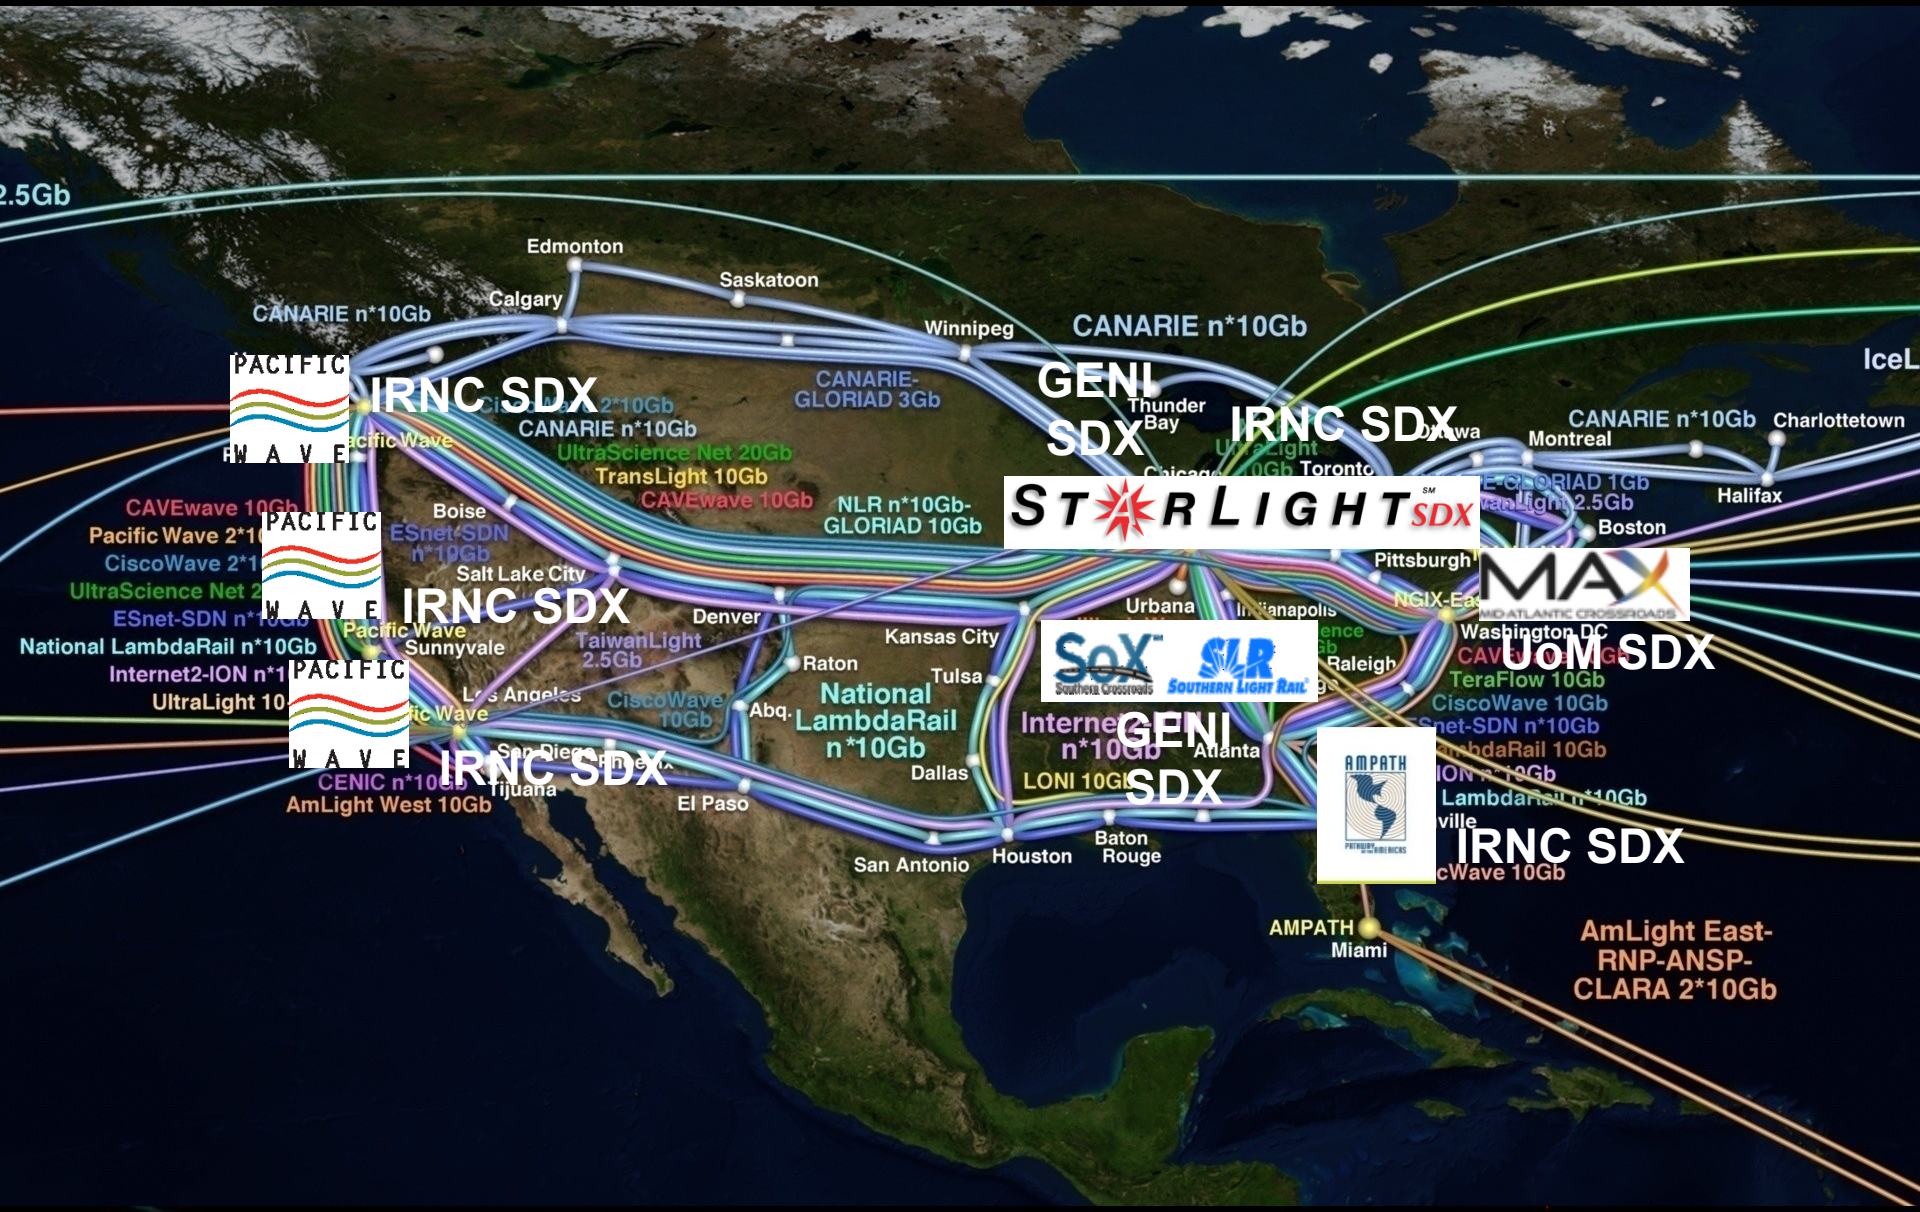

AutoGOLE

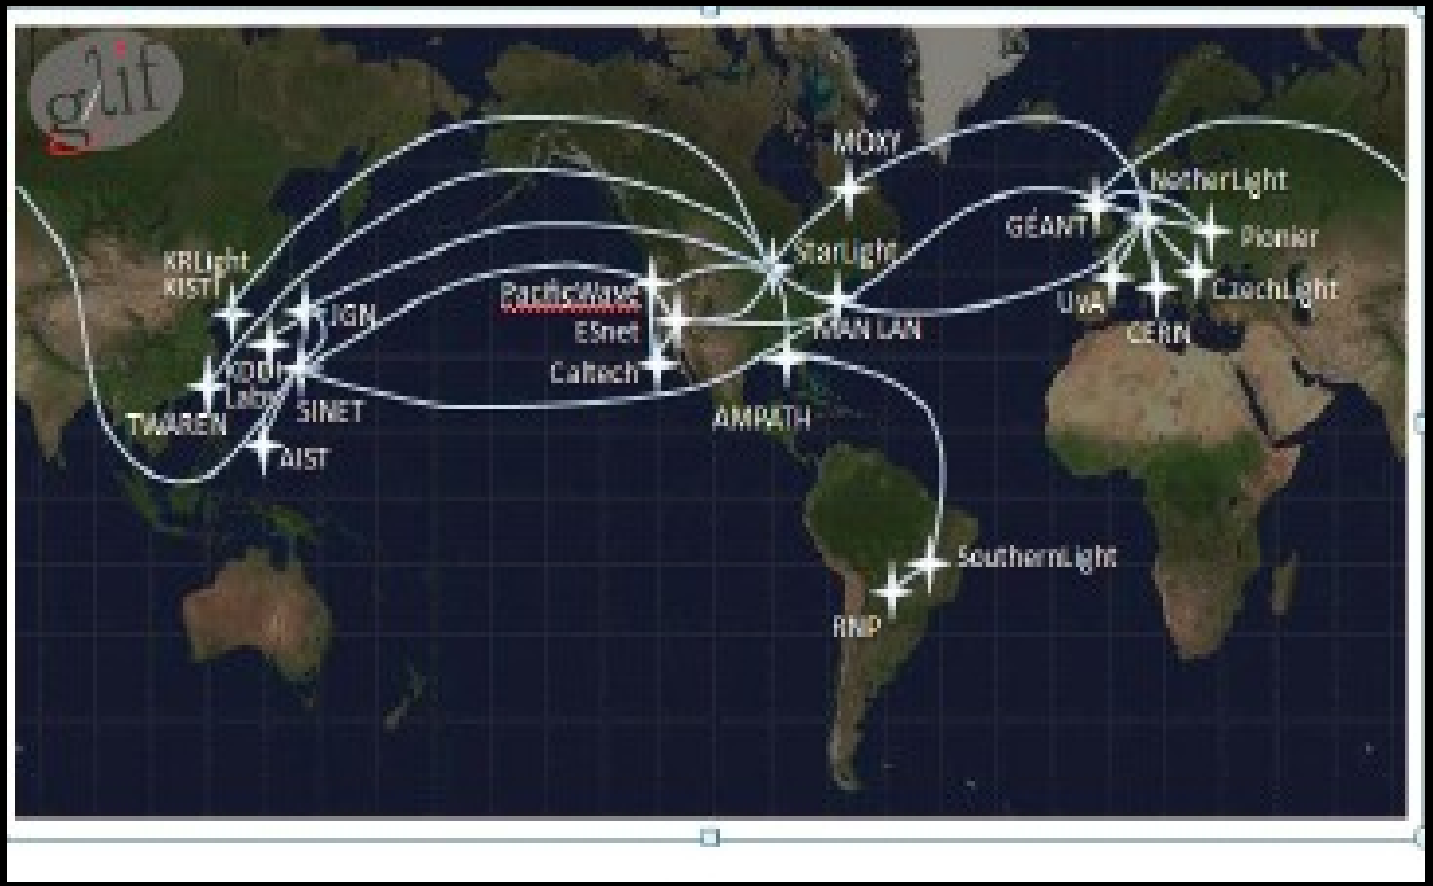
glif
MOXY
NetherLight
GÉANT
Pionier
StarLight
KRLight
KISTI
SURFnet
IGN
PacificWave
ESnet
Caltech
MAN LAN
UvA
CERN
CzechLight
TWAREN
SINET
AIST
AMPATH
SouthernLight
RNP

International Multi-Domain Provisioning Using AutoGOLE Based Network Service Interface (NSI 2.0)

* Network Service Interface (NSI 2.0)
* An Architectural Standard Developed By the *Open Grid Forum (OGF)
* OGF Pioneered Programmable Networking (Initially Termed "Grid Networking")
Techniques That Made Networks 'First Class Citizens" in Grid Environments – Programmable With Grid Middleware
* Currently Being Placed Into Production By R&E Networks Around the World

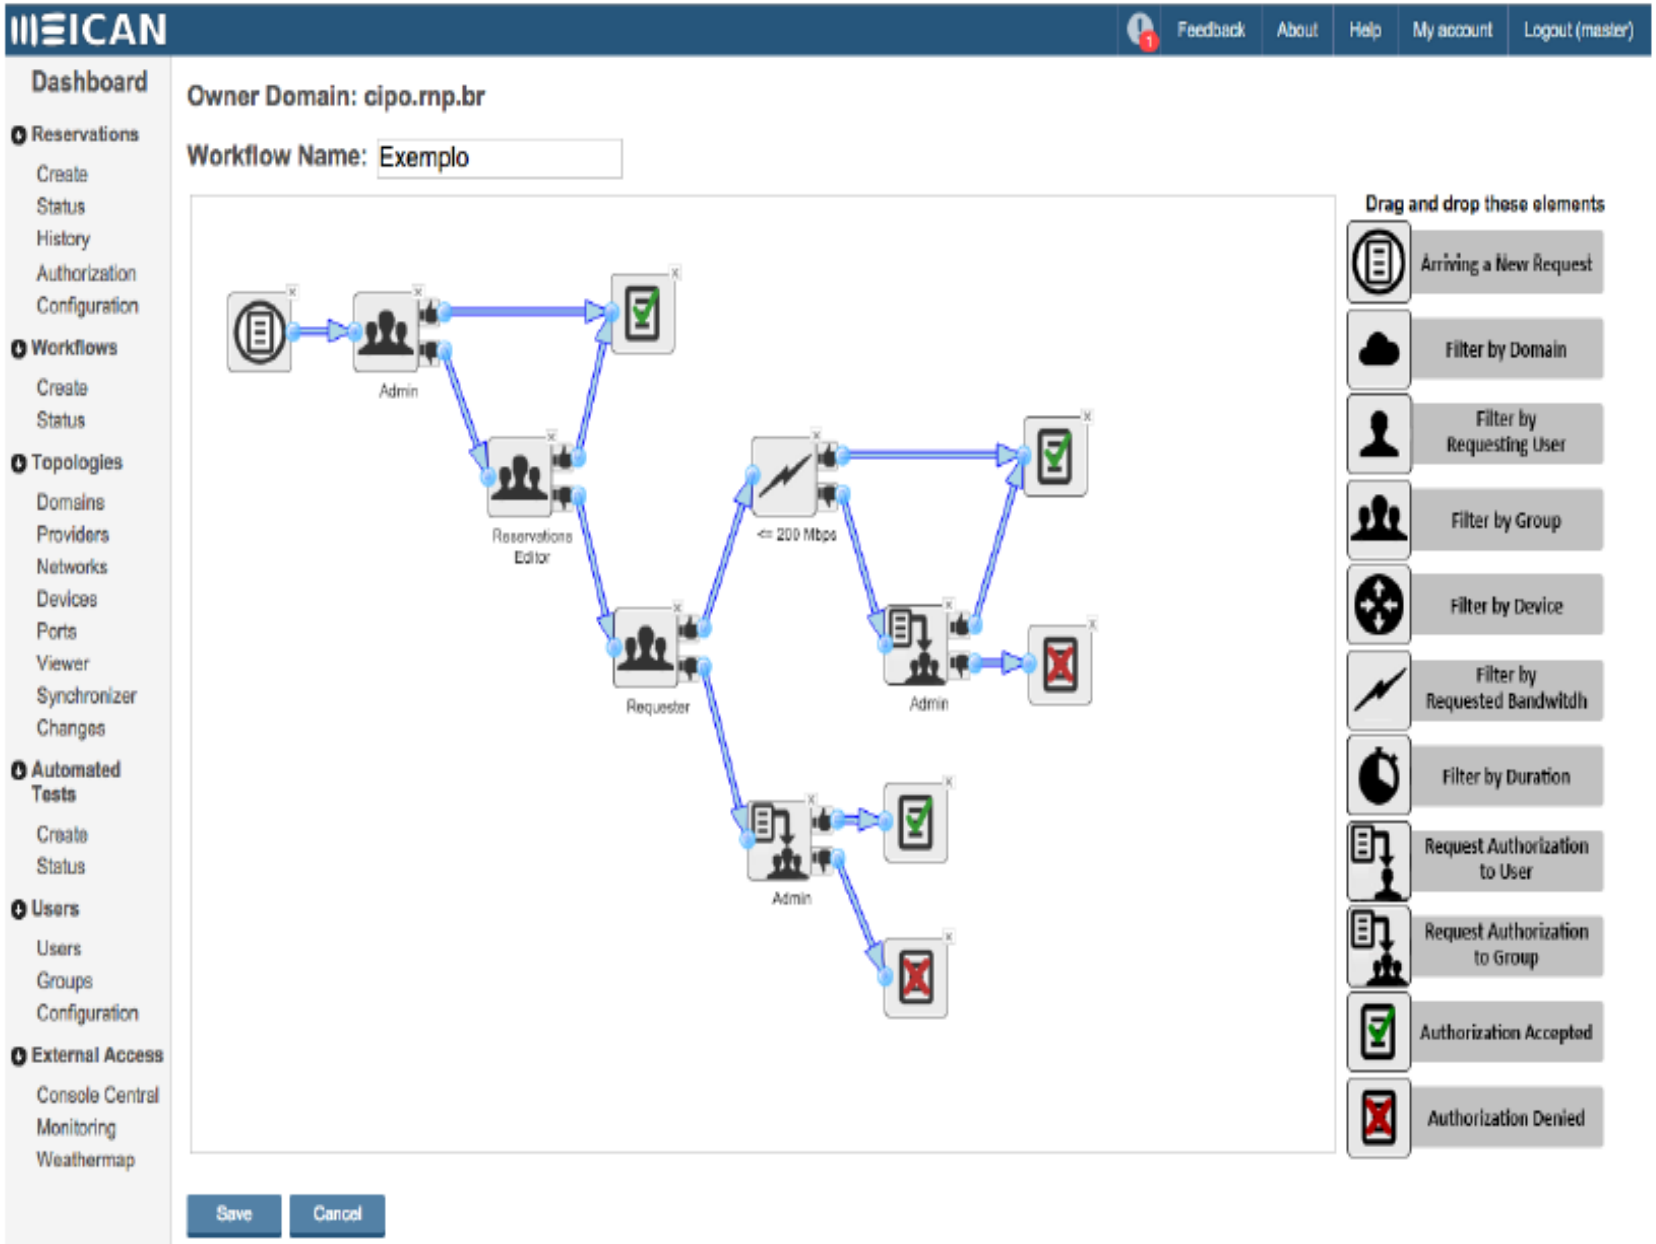
MEICAN
Feedback
About
Help
My account
Logout (master)
Dashboard
Reservations
Create
Status
History
Authorization
Configuration
Workflows
Create
Status
Topologies
Domains
Providers
Networks
Devices
Ports
Viewer
Synchronizer
Changes
Automated Tests
Create
Status
Users
Users
Groups
Configuration
External Access
Console Central
Monitoring
Weathermap
Owner Domain: cipo.rnp.br
Workflow Name: Exemplo
Admin
Reservations Editor
Requester
<= 200 Mbps
Admin
Admin
Drag and drop these elements
Arriving a New Request
Filter by Domain
Filter by Requesting User
Filter by Group
Filter by Device
Filter by Requested Bandwitdh
Filter by Duration
Request Authorization to User
Request Authorization to Group
Authorization Accepted
Authorization Denied
Save
Cancel

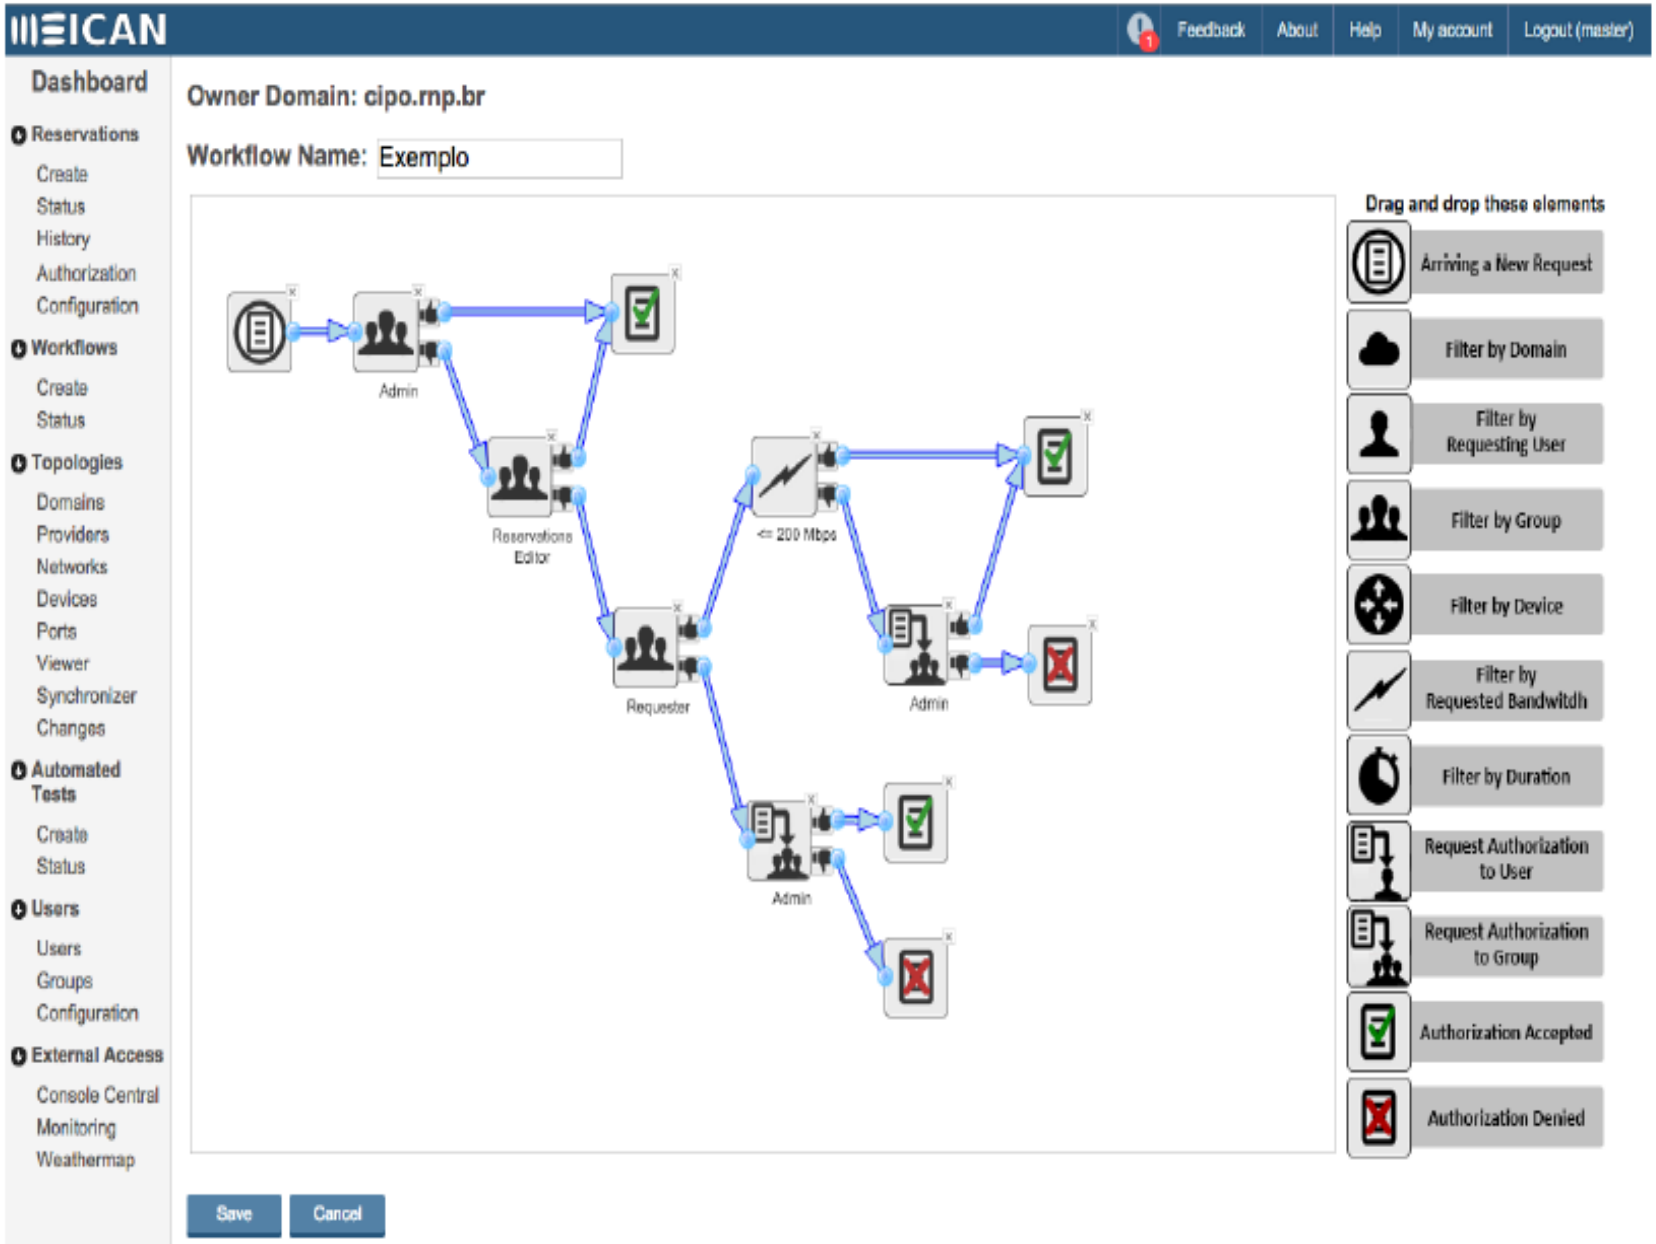
MEICAN
Feedback
About
Help
My account
Logout (master)
Dashboard
Reservations
Create
Status
History
Authorization
Configuration
Workflows
Create
Status
Topologies
Domains
Providers
Networks
Devices
Ports
Viewer
Synchronizer
Changes
Automated Tests
Create
Status
Users
Users
Groups
Configuration
External Access
Console Central
Monitoring
Weathermap
Owner Domain: cipo.rnp.br
Workflow Name: Exemplo
Admin
Reservations Editor
Requester
<= 200 Mbps
Admin
Admin
Drag and drop these elements
Arriving a New Request
Filter by Domain
Filter by Requesting User
Filter by Group
Filter by Device
Filter by Requested Bandwitdh
Filter by Duration
Request Authorization to User
Request Authorization to Group
Authorization Accepted
Authorization Denied
Save
Cancel

AARNet - Pacific Wave - Starlight
Inter-domain SDX Topology v0.4

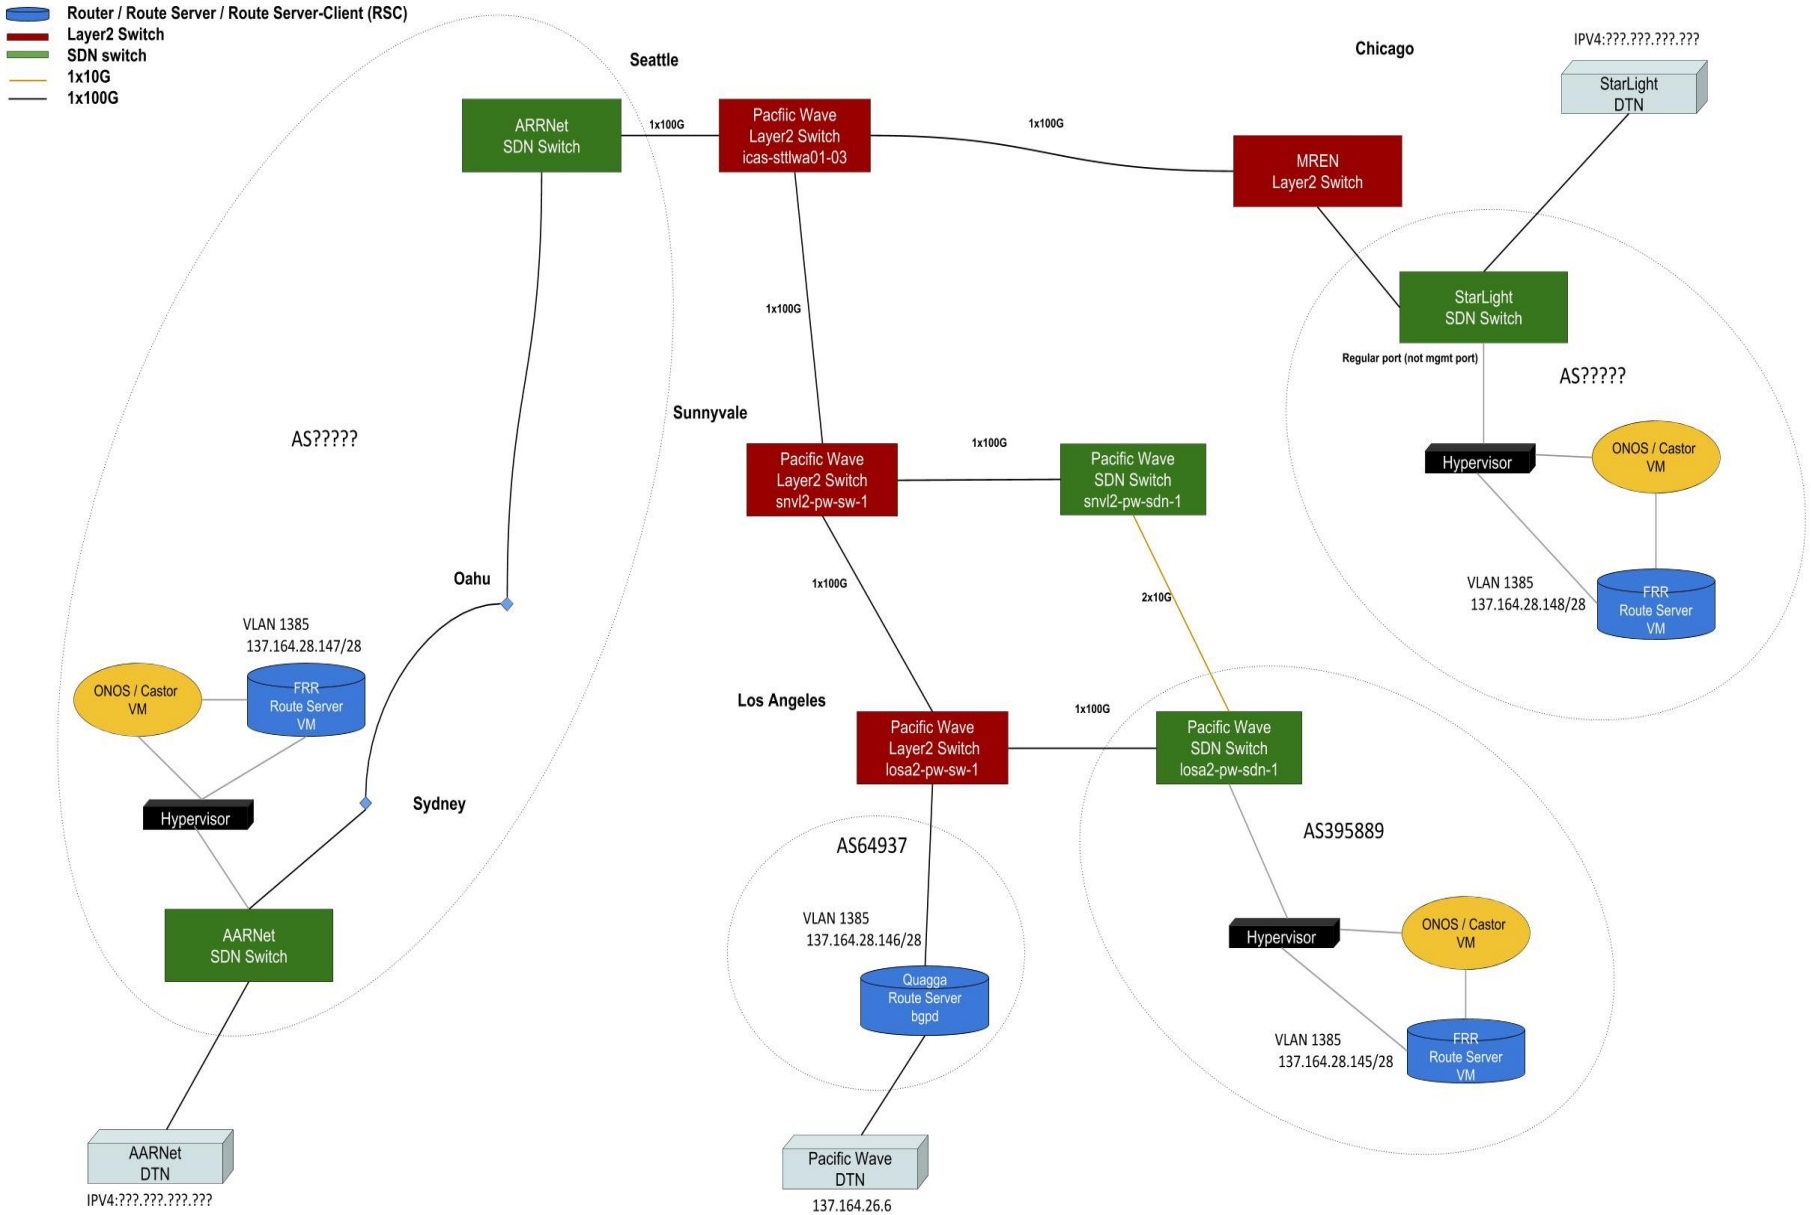

NOTE: this diagram represents a subset of sites, devices, and connections

https://docs.google.com/drawings/d/1NIcvdKg8Zy4mFH-uIYmYsJdz5HKpmgjfk1QyBKpntel/edit?usp=sharing
V0.04 20170801

AutoGOLE Fabric: Another View

AutoGOLE Dashboard

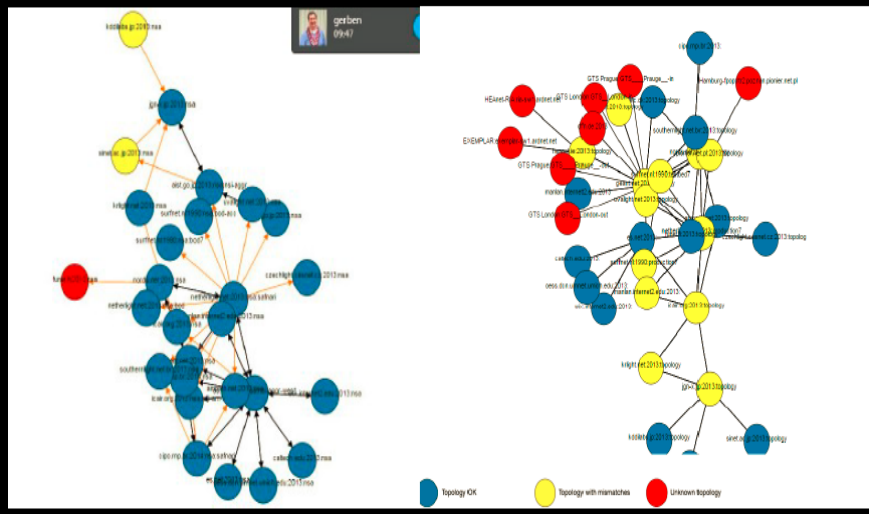

Control Plane

Data Plane

GLIF AutoGOLE Initiative

Source: John MacAuley

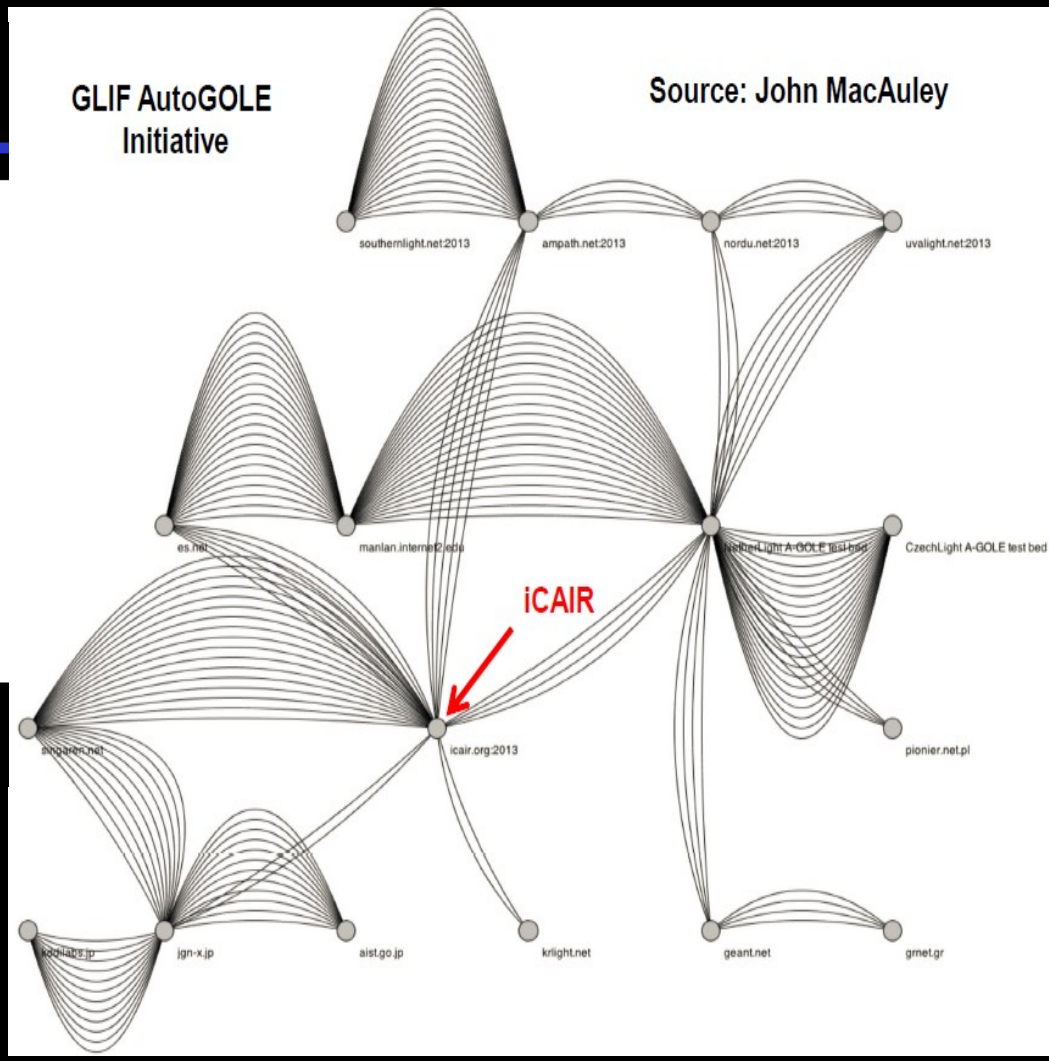

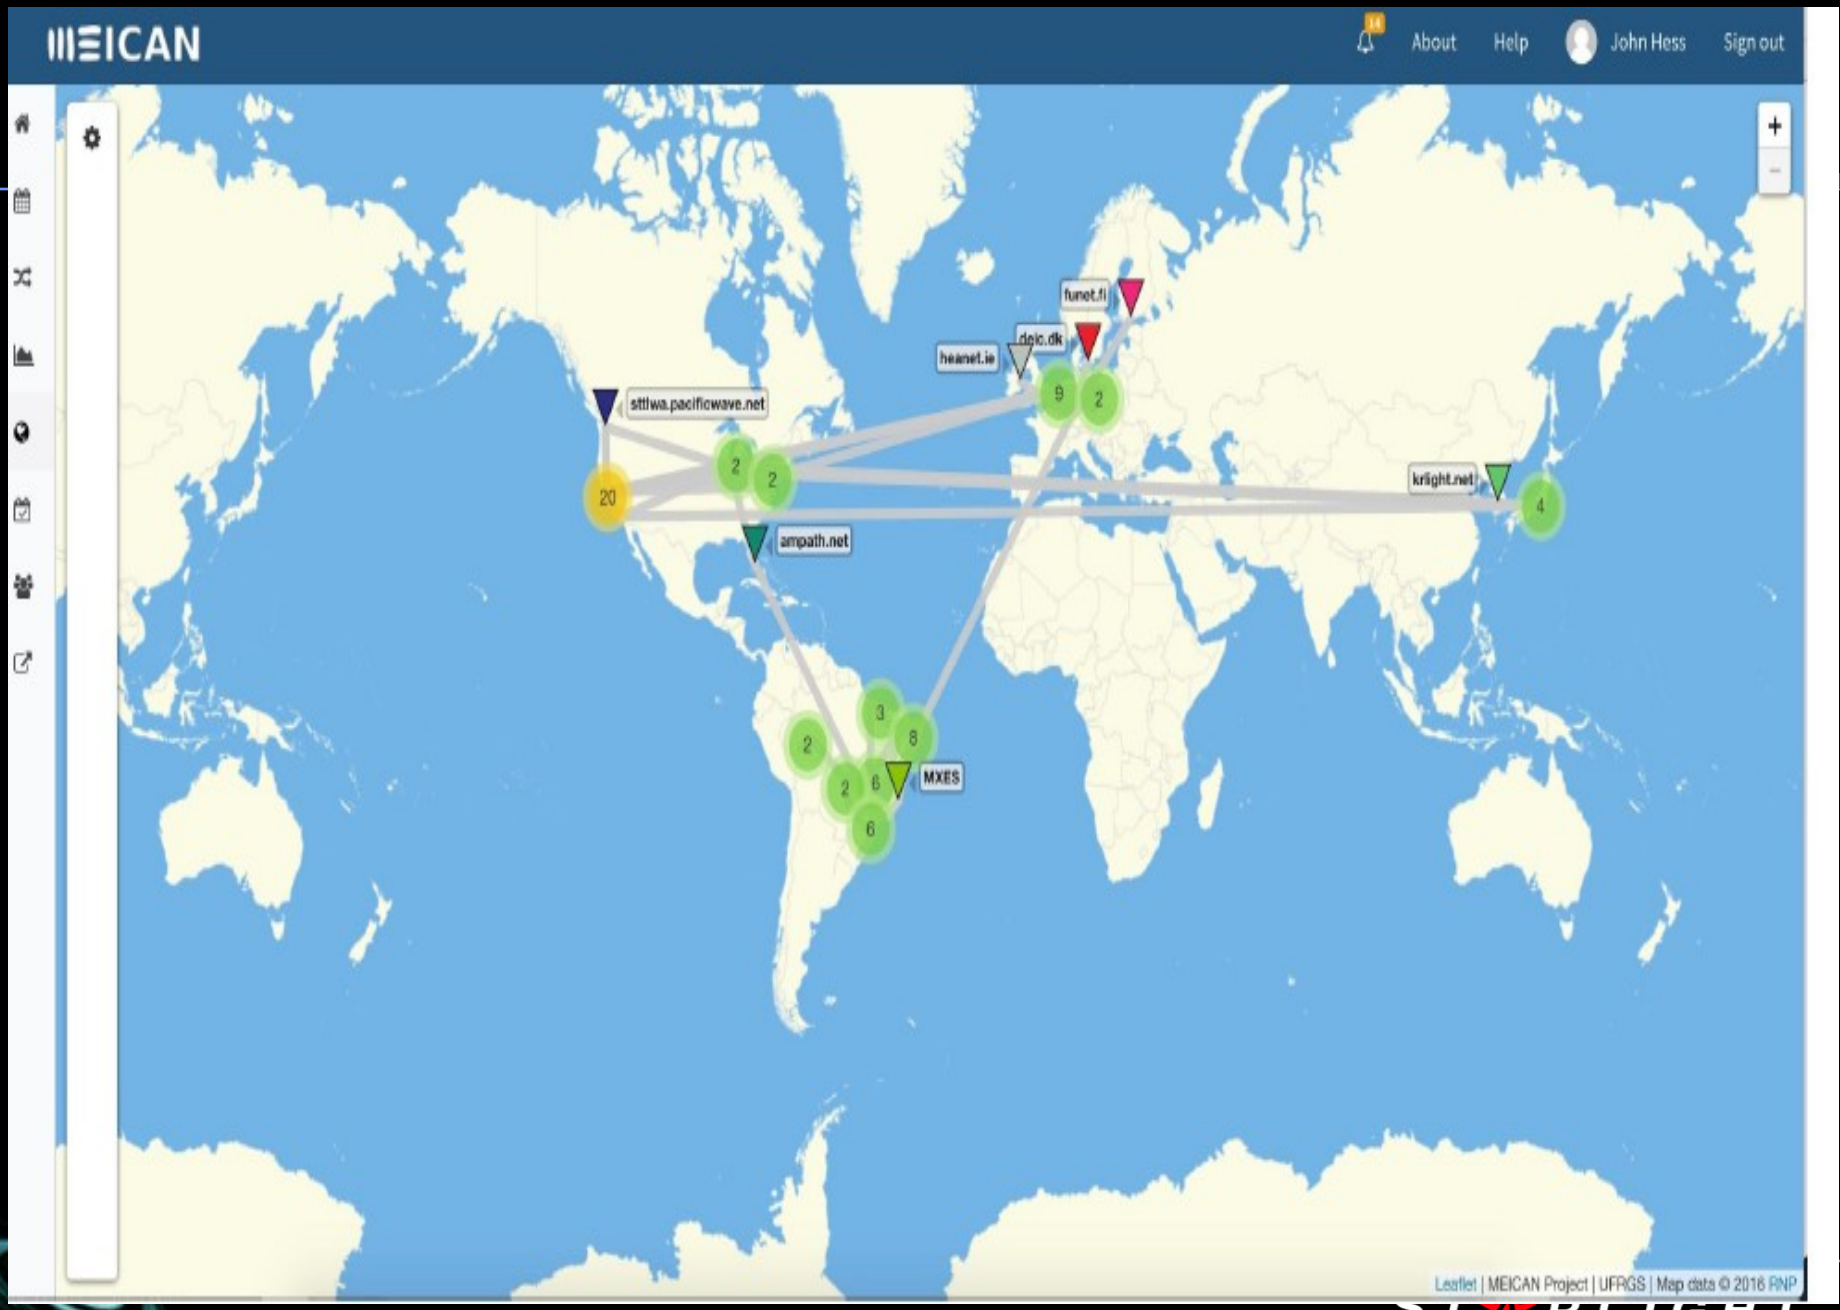
MEICAN
About
Help
John Hess
Sign out
funet.fi
deic.dk
heanet.ie
sttlwa.pacificwave.net
ampath.net
krlight.net
MXES
Leaflet | MEICAN Project | UFRGS | Map data © 2016 RNP

Designing and Deploying A

Bioinformatics Software-Defined Network Exchange (SDX): Architecture, Services, Capabilities, and Foundation Technologies

Joe Mambretti, Jim Chen, Fei Yeh
International Center for Advanced Internet Research
Northwestern University
Robert Grossman, Piers Nash, Alison Heath, Renuka Arya, Stuti Agrawal, Zhenyu Zhang
Center for Data Intensive Science
University of Chicago
Chicago, Illinois, USA

iCAIR

STARLIGHT SM

2016 Bioinformatics SDXs Network

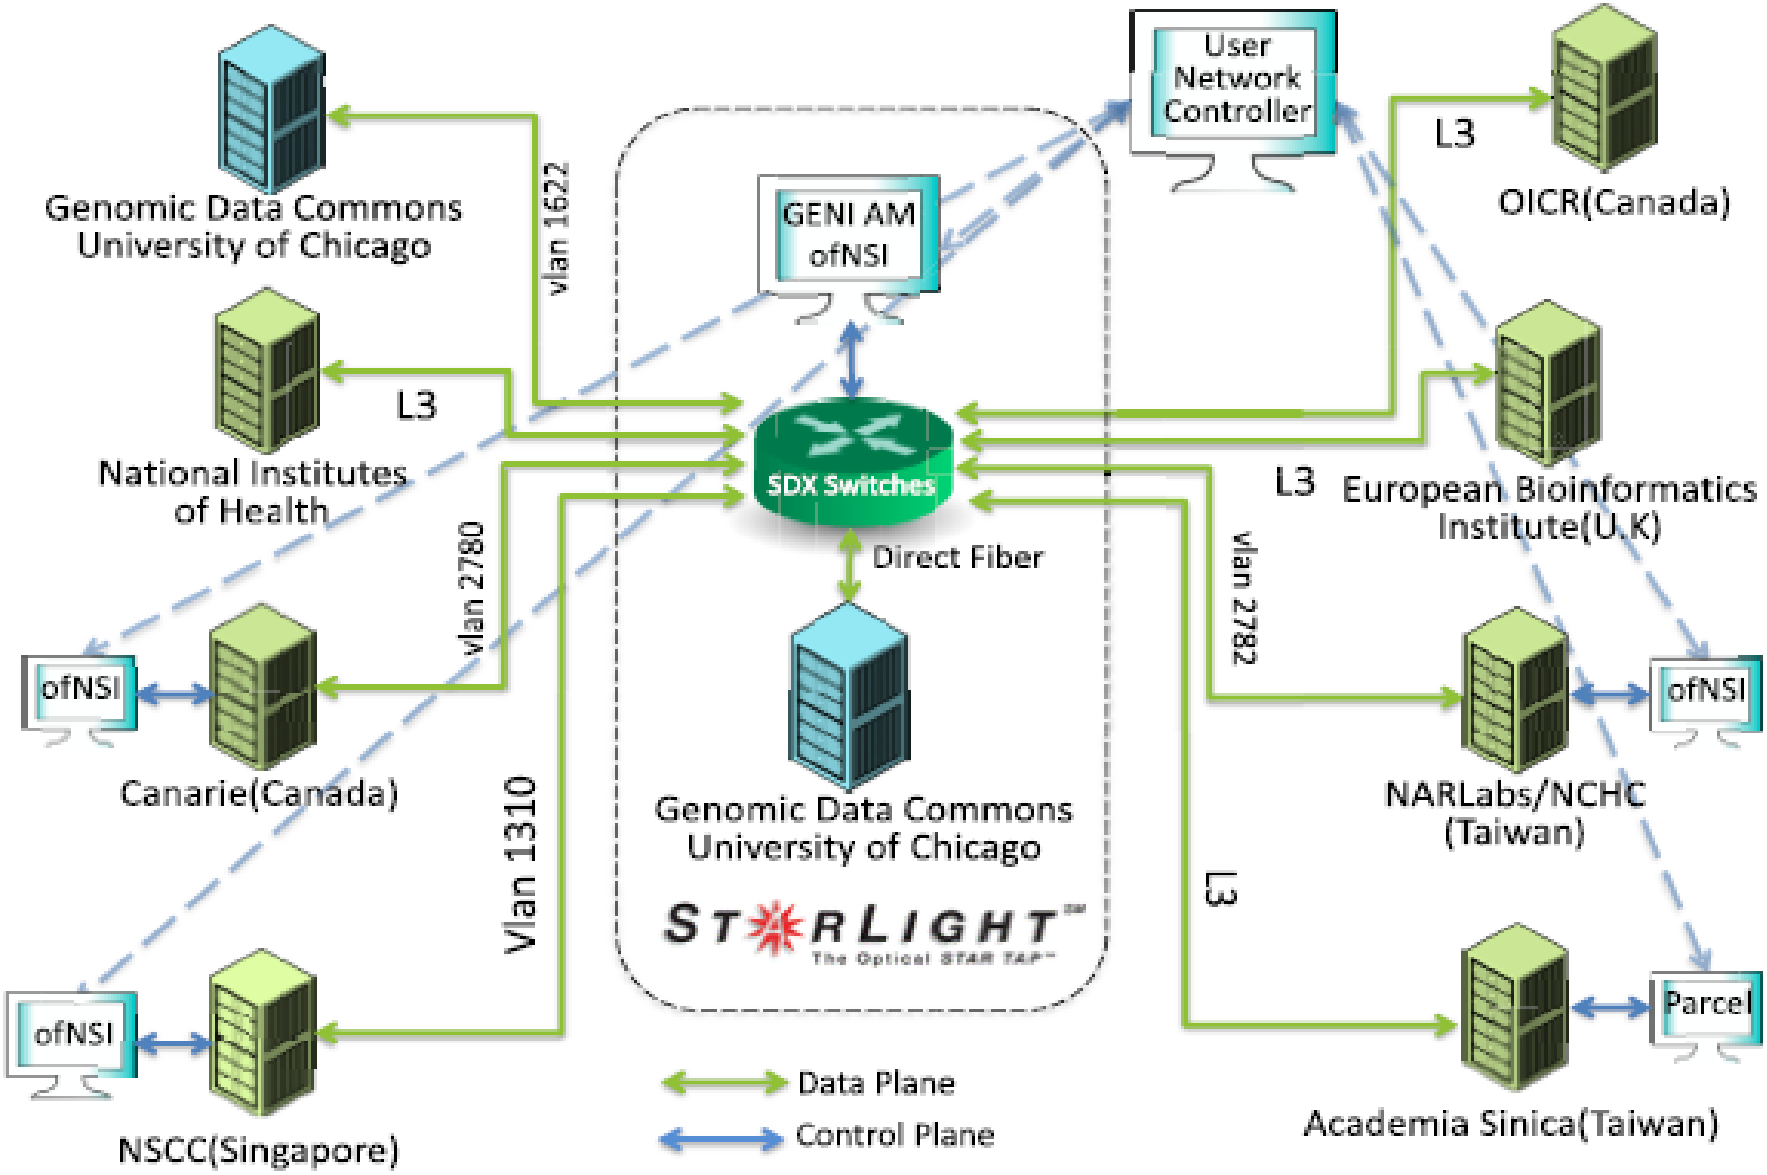

Global LambdaGrid Workshop 2017 Demonstrations, Sydney Australia

International Multi-Domain Provisioning Using AutoGOLE Based Network Service Interface (NSI 2.0)

Using RNP MEICAN Tools for NSI Provisioning

Large Scale Airline Data Transport Over SD-WANs Using NSI and DTNs

Large Scale Science Data Transport Over SD-WANs Using NSI and DTNs

SDX Interdomain Interoperability At L3

Transferring Large Files E2E Across WANs Enabled By SD-WANs and SDXs

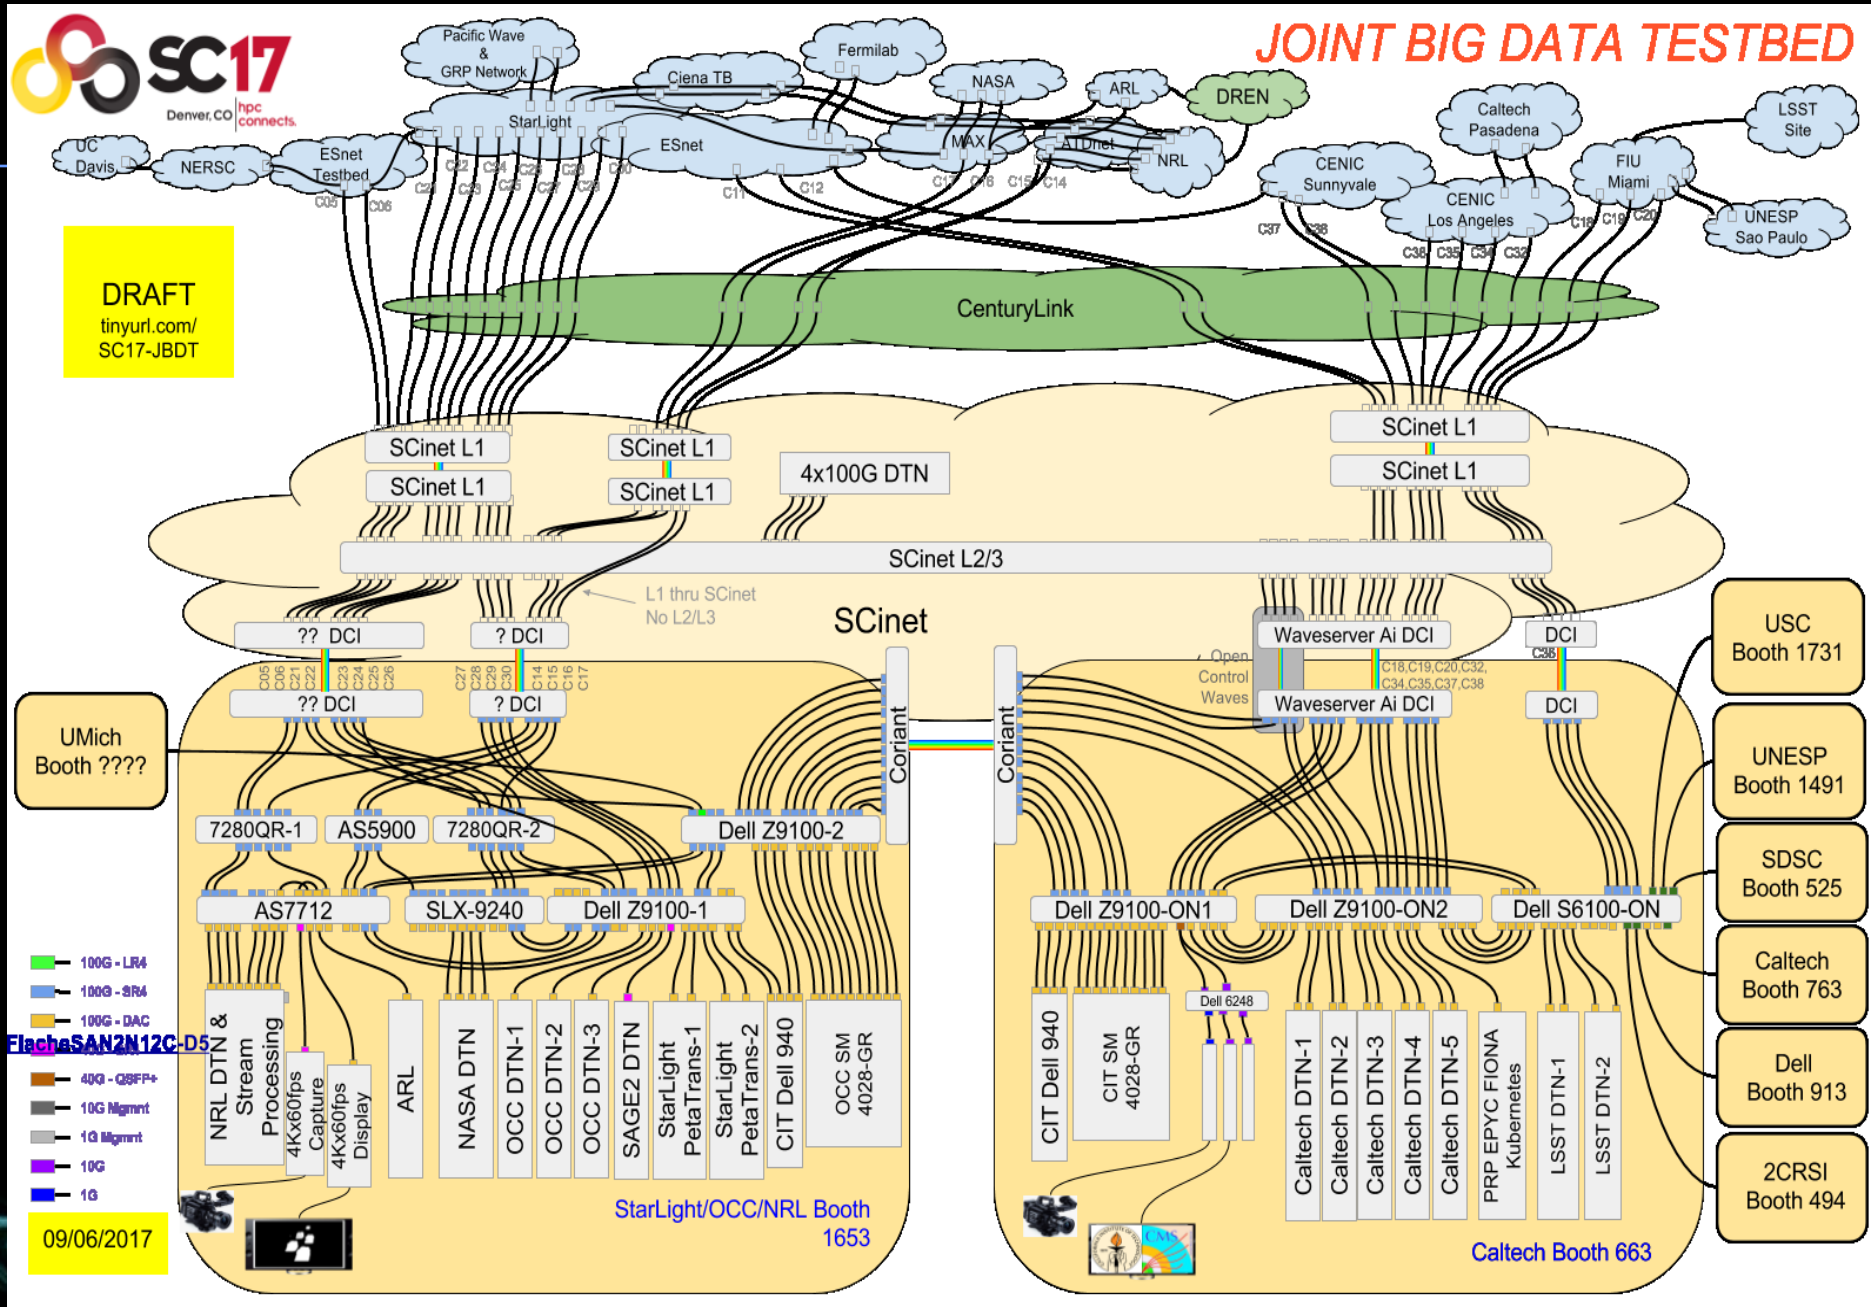
SC17
Denver, CO hpc connects.
JOINT BIG DATA TESTBED
DRAFT
tinyurl.com/
SC17-JBDT
Pacific Wave & GRP Network
Ciena TB
Fermilab
NASA
ARL
DREN
StarLight
ESnet
MAX
ATDnet
NRL
UC Davis
NERSC
ESnet Testbed
Caltech Pasadena
LSST Site
CENIC Sunnyvale
CENIC Los Angeles
FIU Miami
UNESP Sao Paulo
CenturyLink
SCinet L1
SCinet L1
SCinet L1
SCinet L1
4x100G DTN
SCinet L1
SCinet L1
SCinet L2/3
L1 thru SCinet
No L2/L3
SCinet
?? DCI
? DCI
?? DCI
? DCI
Waveserver Ai DCI
Waveserver Ai DCI
DCI
DCI
Open Control Waves
C18,C19,C20,C32, C34,C35,C37,C38
Coriant
Coriant
UMich Booth ????
7280QR-1
AS5900
7280QR-2
Dell Z9100-2
AS7712
SLX-9240
Dell Z9100-1
Dell Z9100-ON1
Dell Z9100-ON2
Dell S6100-ON
Dell 6248
NRL DTN & Stream Processing
4Kx60fps Capture
4Kx60fps Display
ARL
NASA DTN
OCC DTN-1
OCC DTN-2
OCC DTN-3
SAGE2 DTN
StarLight PetaTrans-1
StarLight PetaTrans-2
CIT Dell 940
OCC SM 4028-GR
CIT Dell 940
CIT SM 4028-GR
Caltech DTN-1
Caltech DTN-2
Caltech DTN-3
Caltech DTN-4
Caltech DTN-5
PRP EPYC FIONA Kubernetes
LSST DTN-1
LSST DTN-2
100G - LR4
100G - SR4
100G - DAC
40G - QSFP+
10G Mgmnt
1G Mgmnt
10G
1G
09/06/2017
StarLight/OCC/NRL Booth 1653
Caltech Booth 663
USC Booth 1731
UNESP Booth 1491
SDSC Booth 525
Caltech Booth 763
Dell Booth 913
2CRSI Booth 494

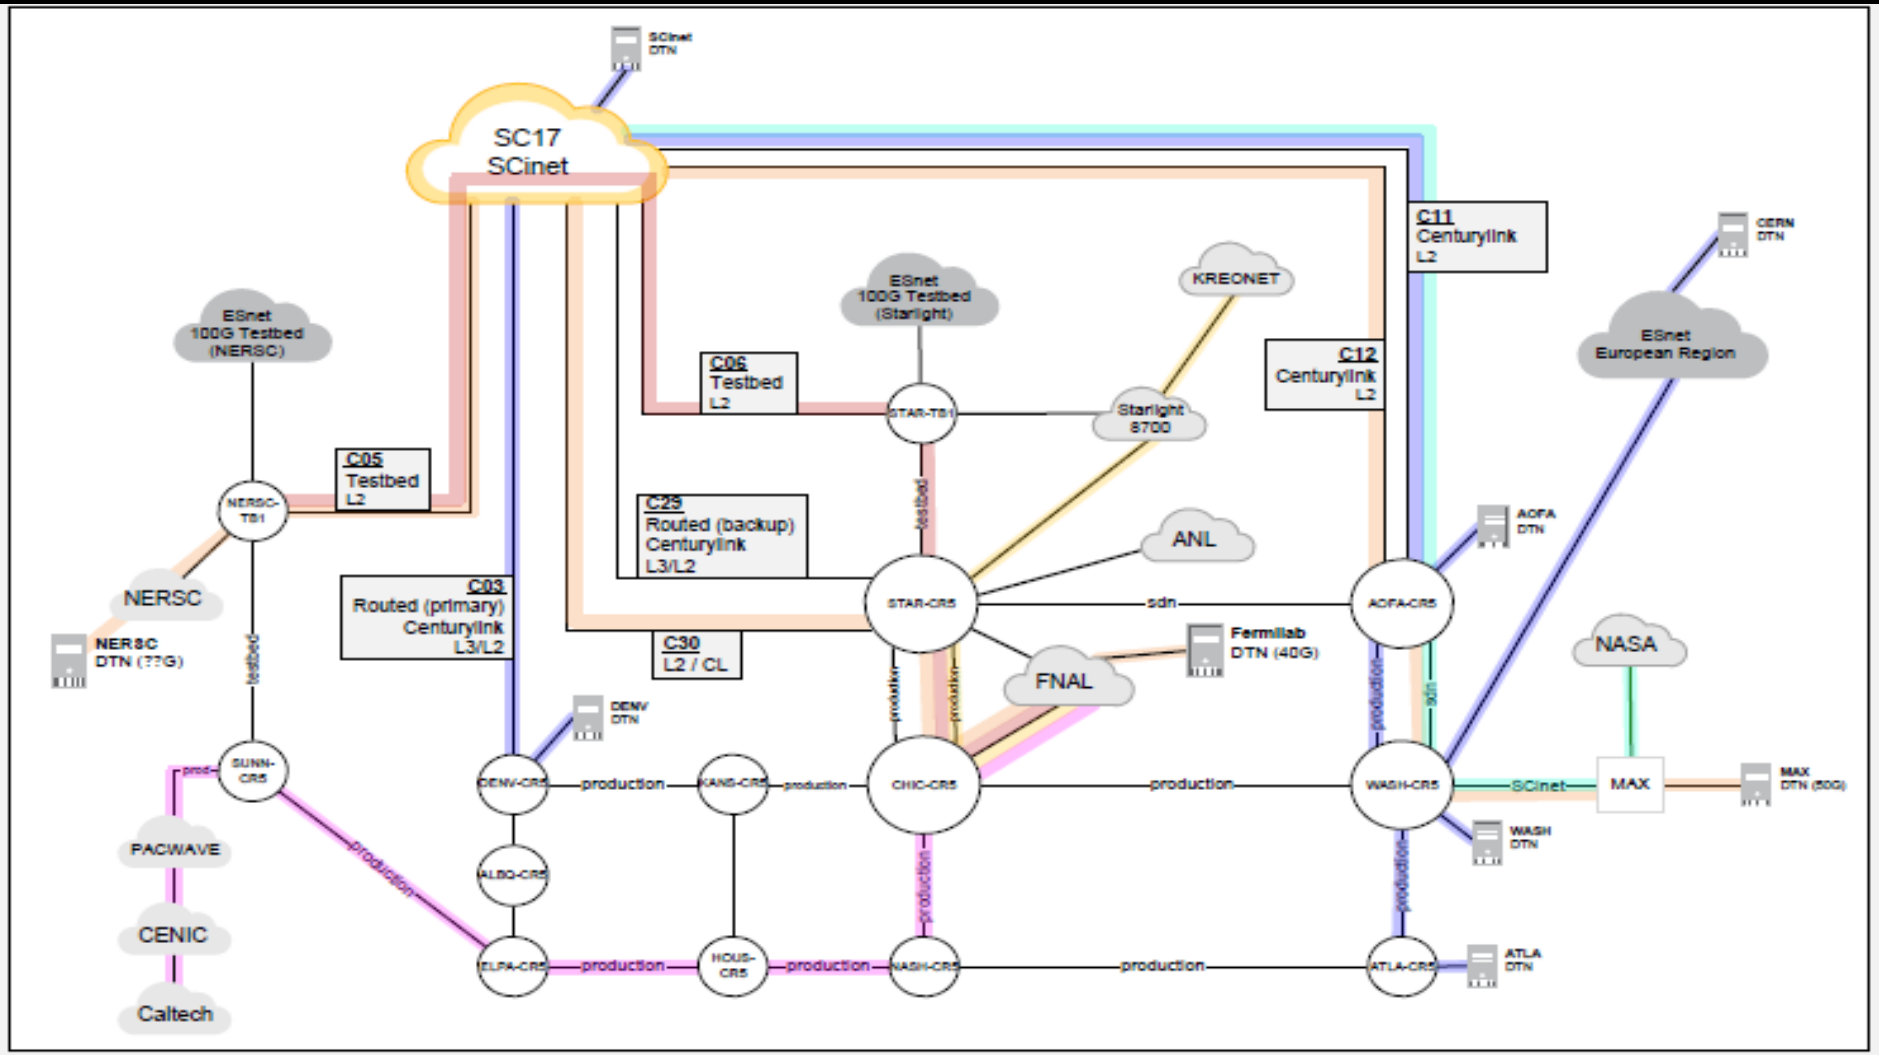

- NASA GODDARD 400G Disk Transfer (100G ESnet path, not active during show)
- SENSE (40G FNAL, 50G MAX)
- FNAL BigData Experess AmoebaNet (2x10G OSCARS)
- BDE FNAL KISTI SDN (2x1G vlans 1662, 1663)
- HEPcloud (10G)(vlan 3602)
- Calibers (1G vlan TBD)

ESnet

SC17 demos – ESnet

Zaoh Harlan, ESnet 8/26/2017

FILENAME SC17-ESNET-DEMOS-V1.0.1.VSD

PetaTrans: Petascale Sciences Data Transfer

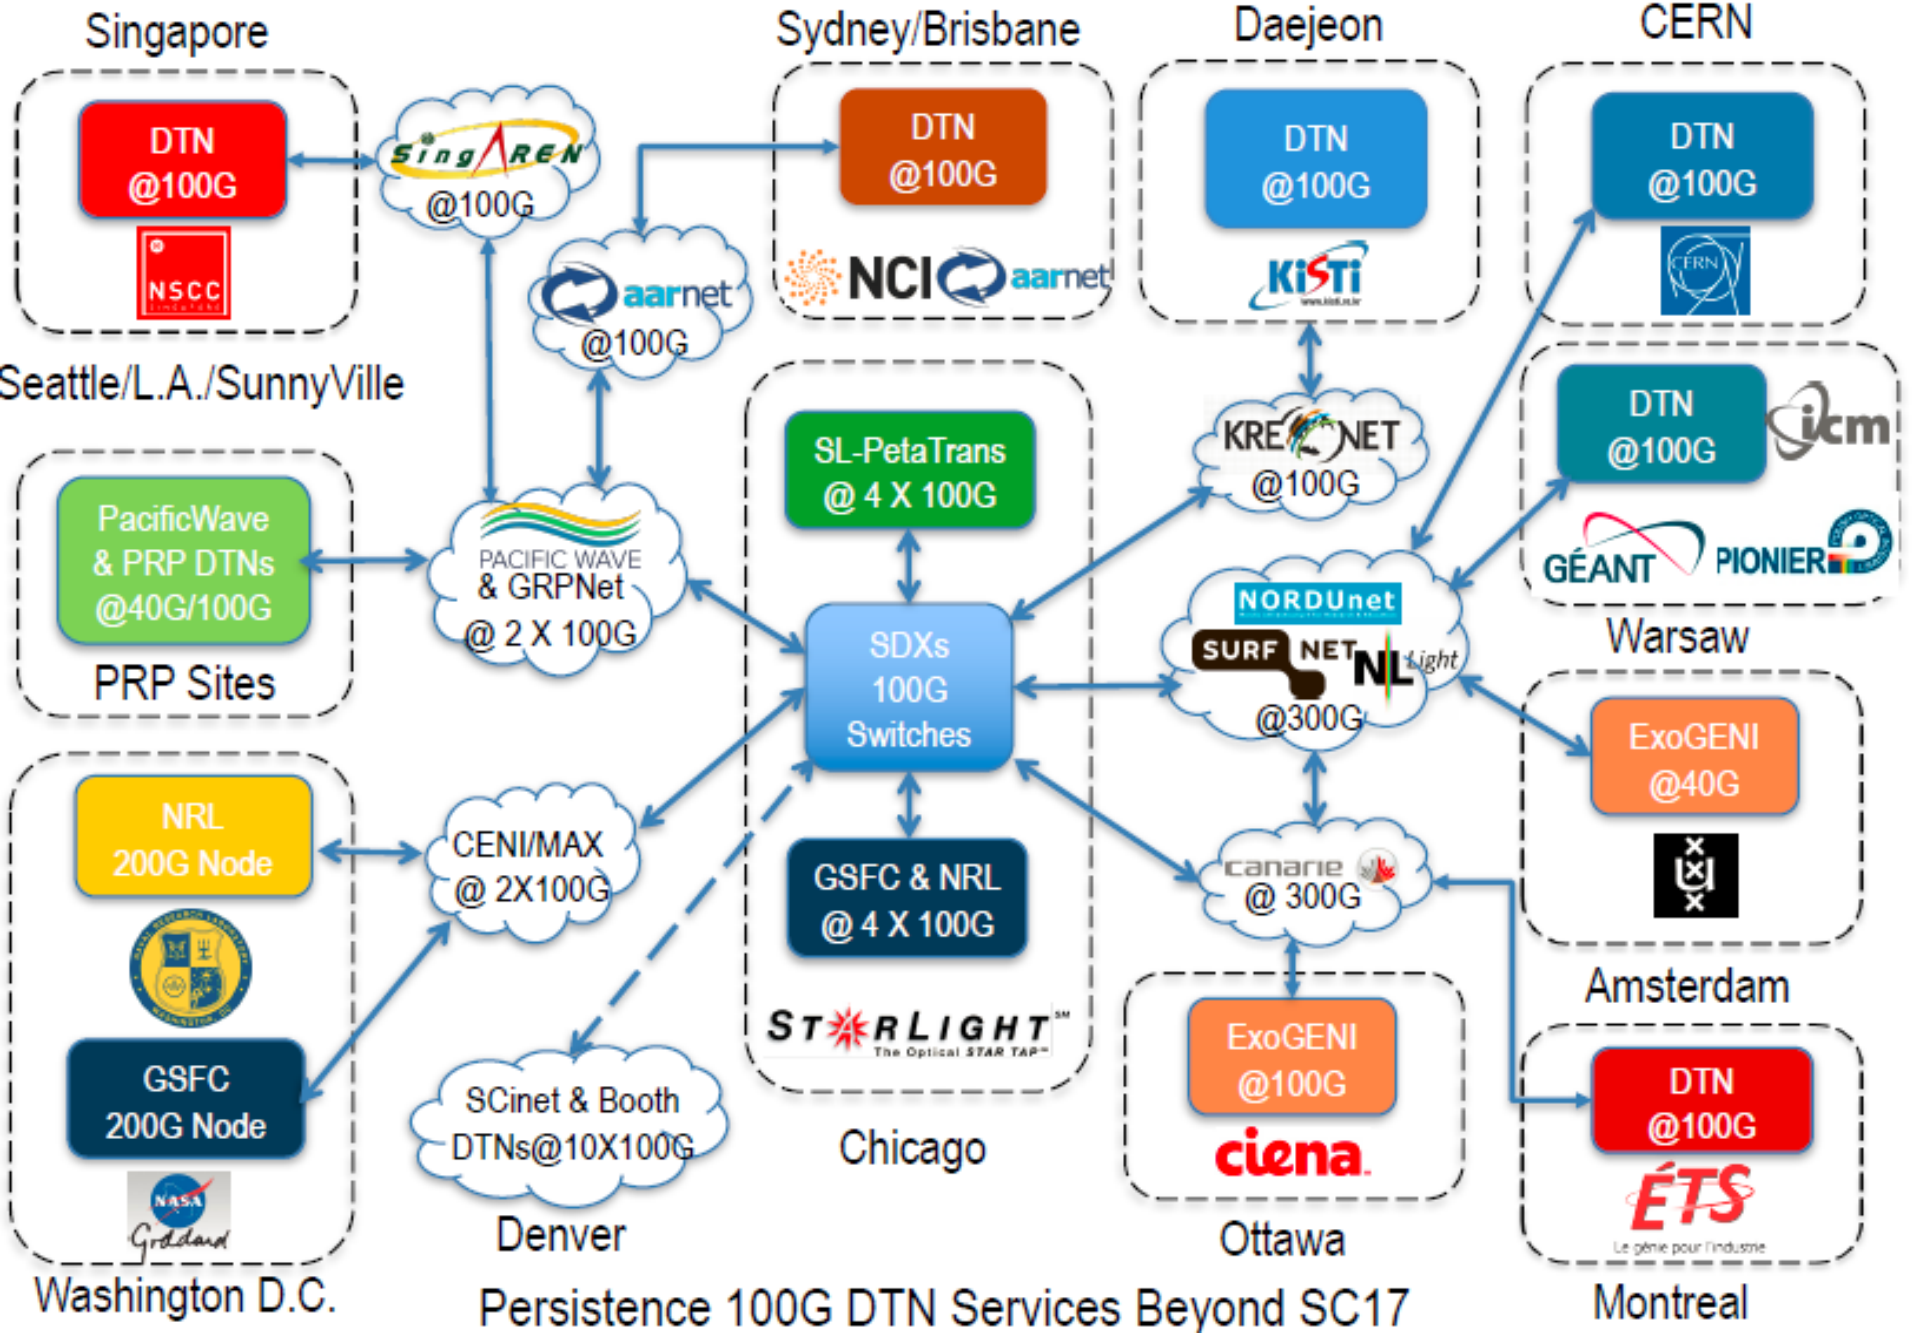

A Disaggregated SCinet Optical Layer

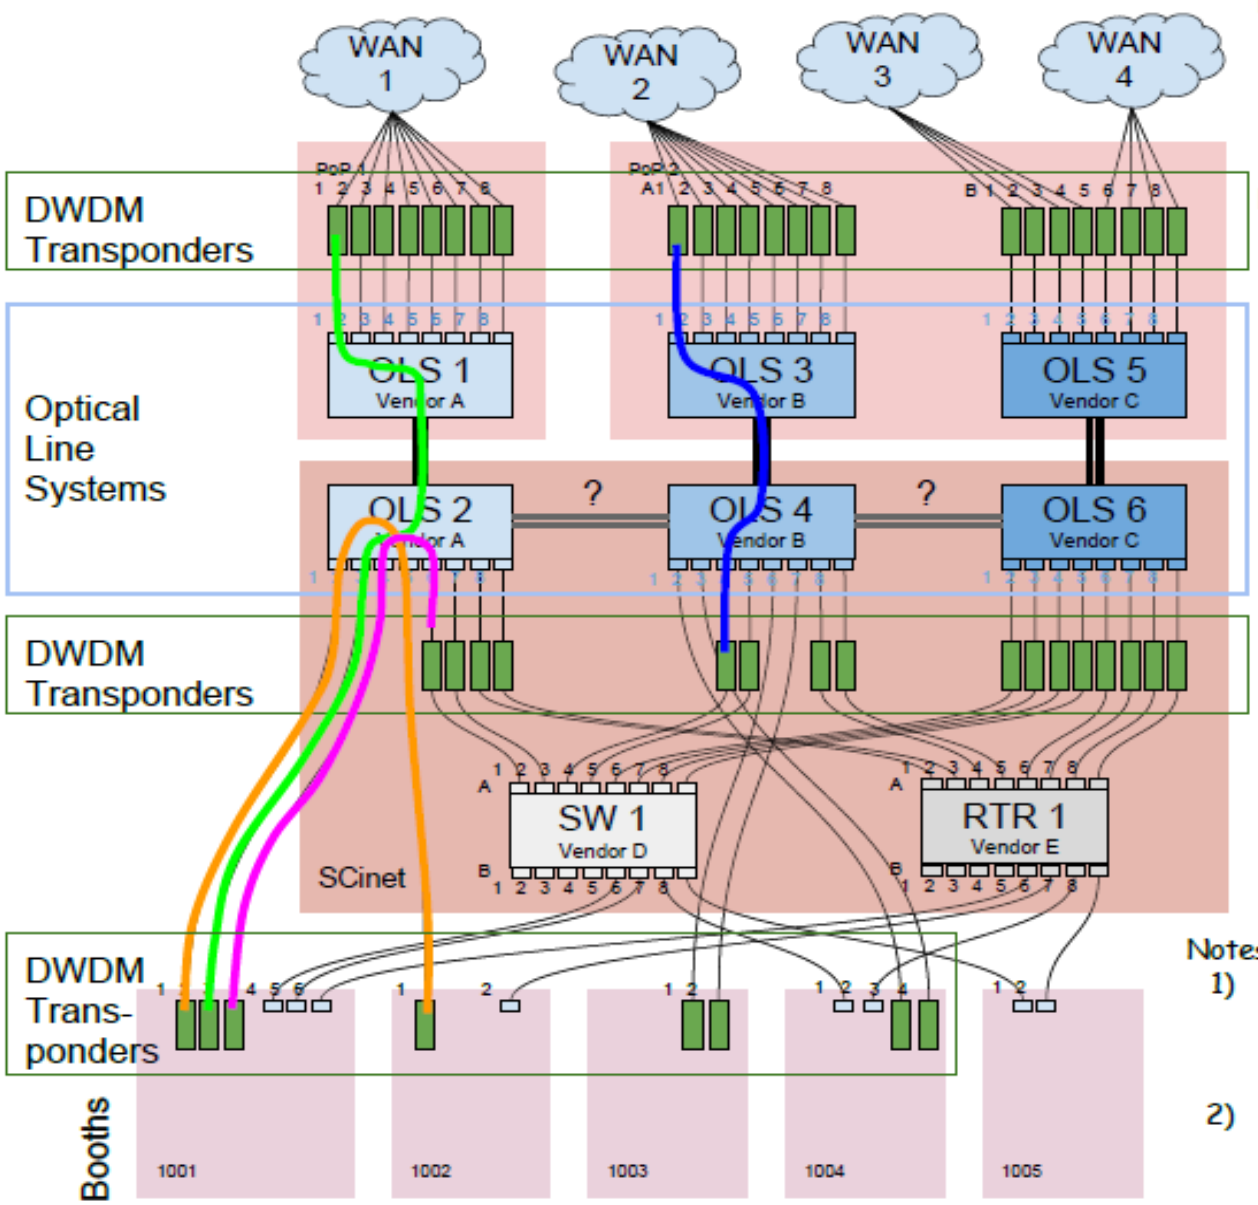

Reconfiguration options

A. Booth to booth connections
B. Booth to WAN connections
C. Booth to switch or router connections
D. WAN to switch or router connections

Examples

A. B-B
   a. Booth 1001-1 to 1002-1 via optical layer
   b. Booth 1001-1 to 1004-3 via optical layer (assumes OLS2 to OLS4 path)
B. Booth to WAN
   a. Booth 1001-2 to PoP1-1 via OLS2-2 and OLS1-1
   b. Booth 1001-2 to PoP2-B1 via OLS2-2, OLS4, OLS6 and OLS5-1
C. Booth to switch/router
   a. Booth 1001-3 to SW1-A1
   b. Booth 1003-1 to RTR1-A5 (assumes OLS4 to OLS6 path)
D. WAN to switch/router
   a. PoP2-A1 (WAN2) to SW1-3 via OLS3-1 and OLS4-3
   b. PoP2-A2 (WAN2) to RTR1-3 via OLS3-2 and OLS4-7

Notes

1) Transponders could be from multiple vendors but for near term the links would need to be built with matching transponders.
2) **Controllers and orchestration systems are not shown** but all Tpndr/OLS systems must be connected

STARLIGHT SM

SC17

Demonstrations of 400 Gbps Disk-to-Disk WAN File Transfers using iWARP and NVMe Drives

An SC17 Collaborative Initiative Among NASA and Several Partners

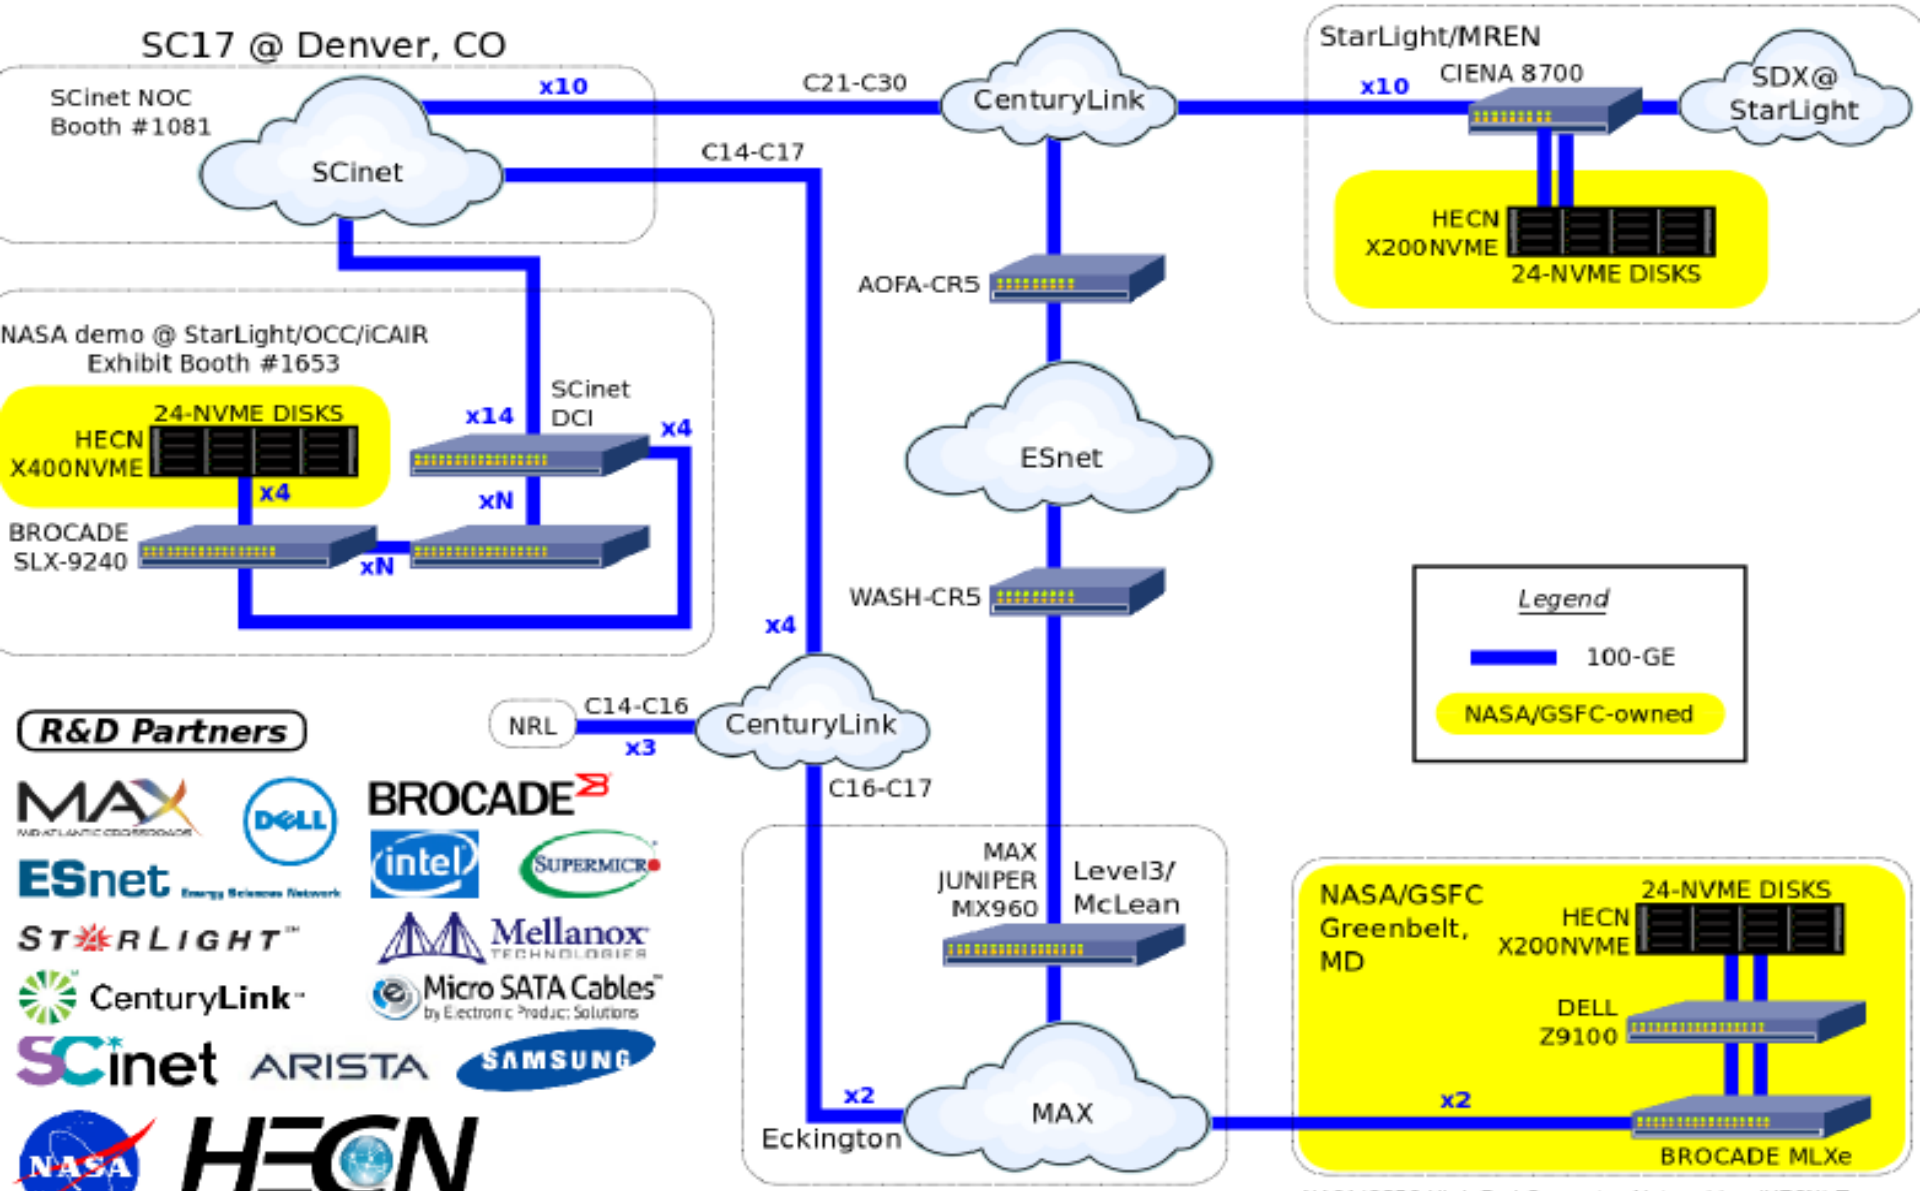

NASA/GSFC High End Computer Networking (HECN) Team
Diagram by Bill Fink / Paul Lang - 8/14/2017

Dynamic Distributed Data Processing

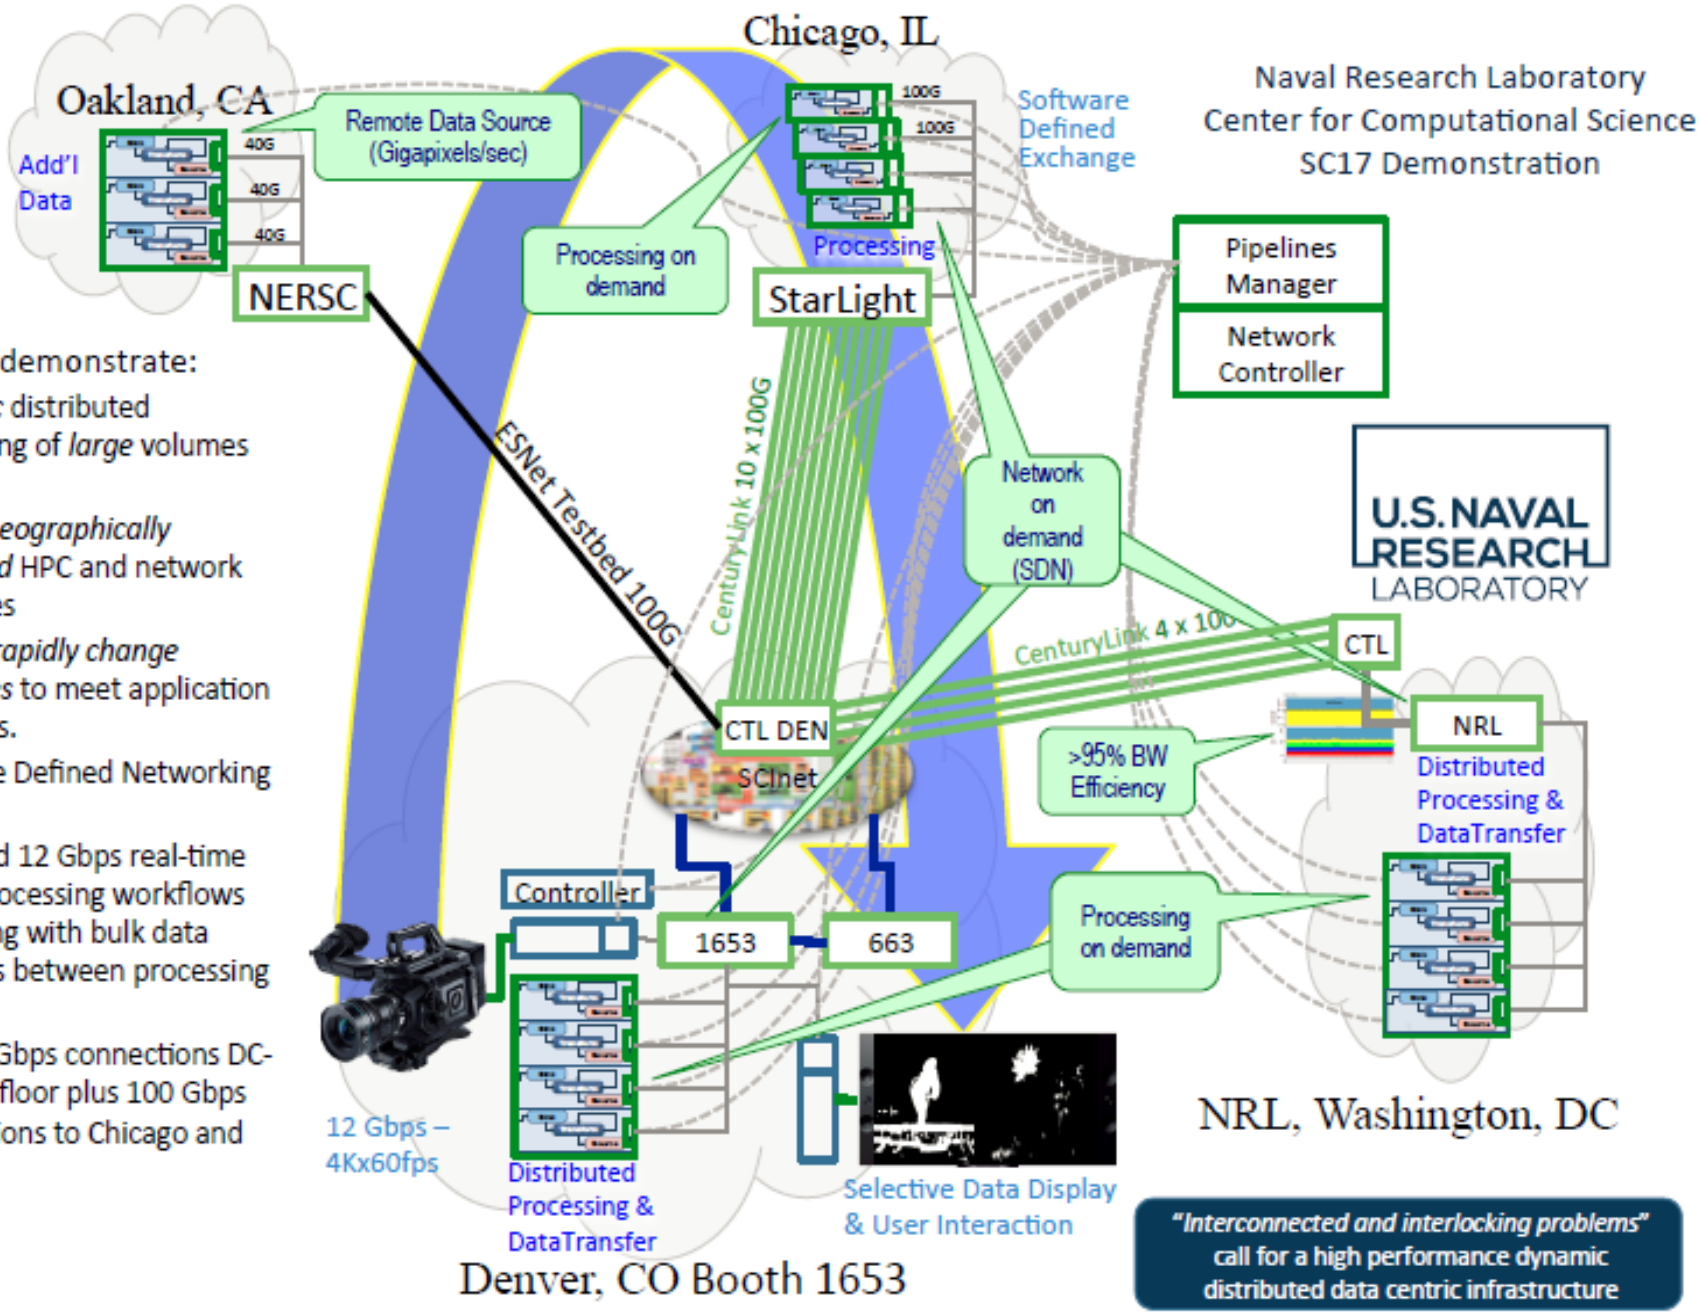

NRL will demonstrate:

- *Dynamic* distributed processing of *large* volumes of data
- Across *geographically dispersed* HPC and network resources
- Able to *rapidly change resources* to meet application demands.
- Software Defined Networking (SDN)
- Pipelined 12 Gbps real-time video processing workflows coexisting with bulk data transfers between processing nodes.
- 3 – 100 Gbps connections DC-to-SC17 floor plus 100 Gbps connections to Chicago and Oakland

"Interconnected and interlocking problems" call for a high performance dynamic distributed data centric infrastructure

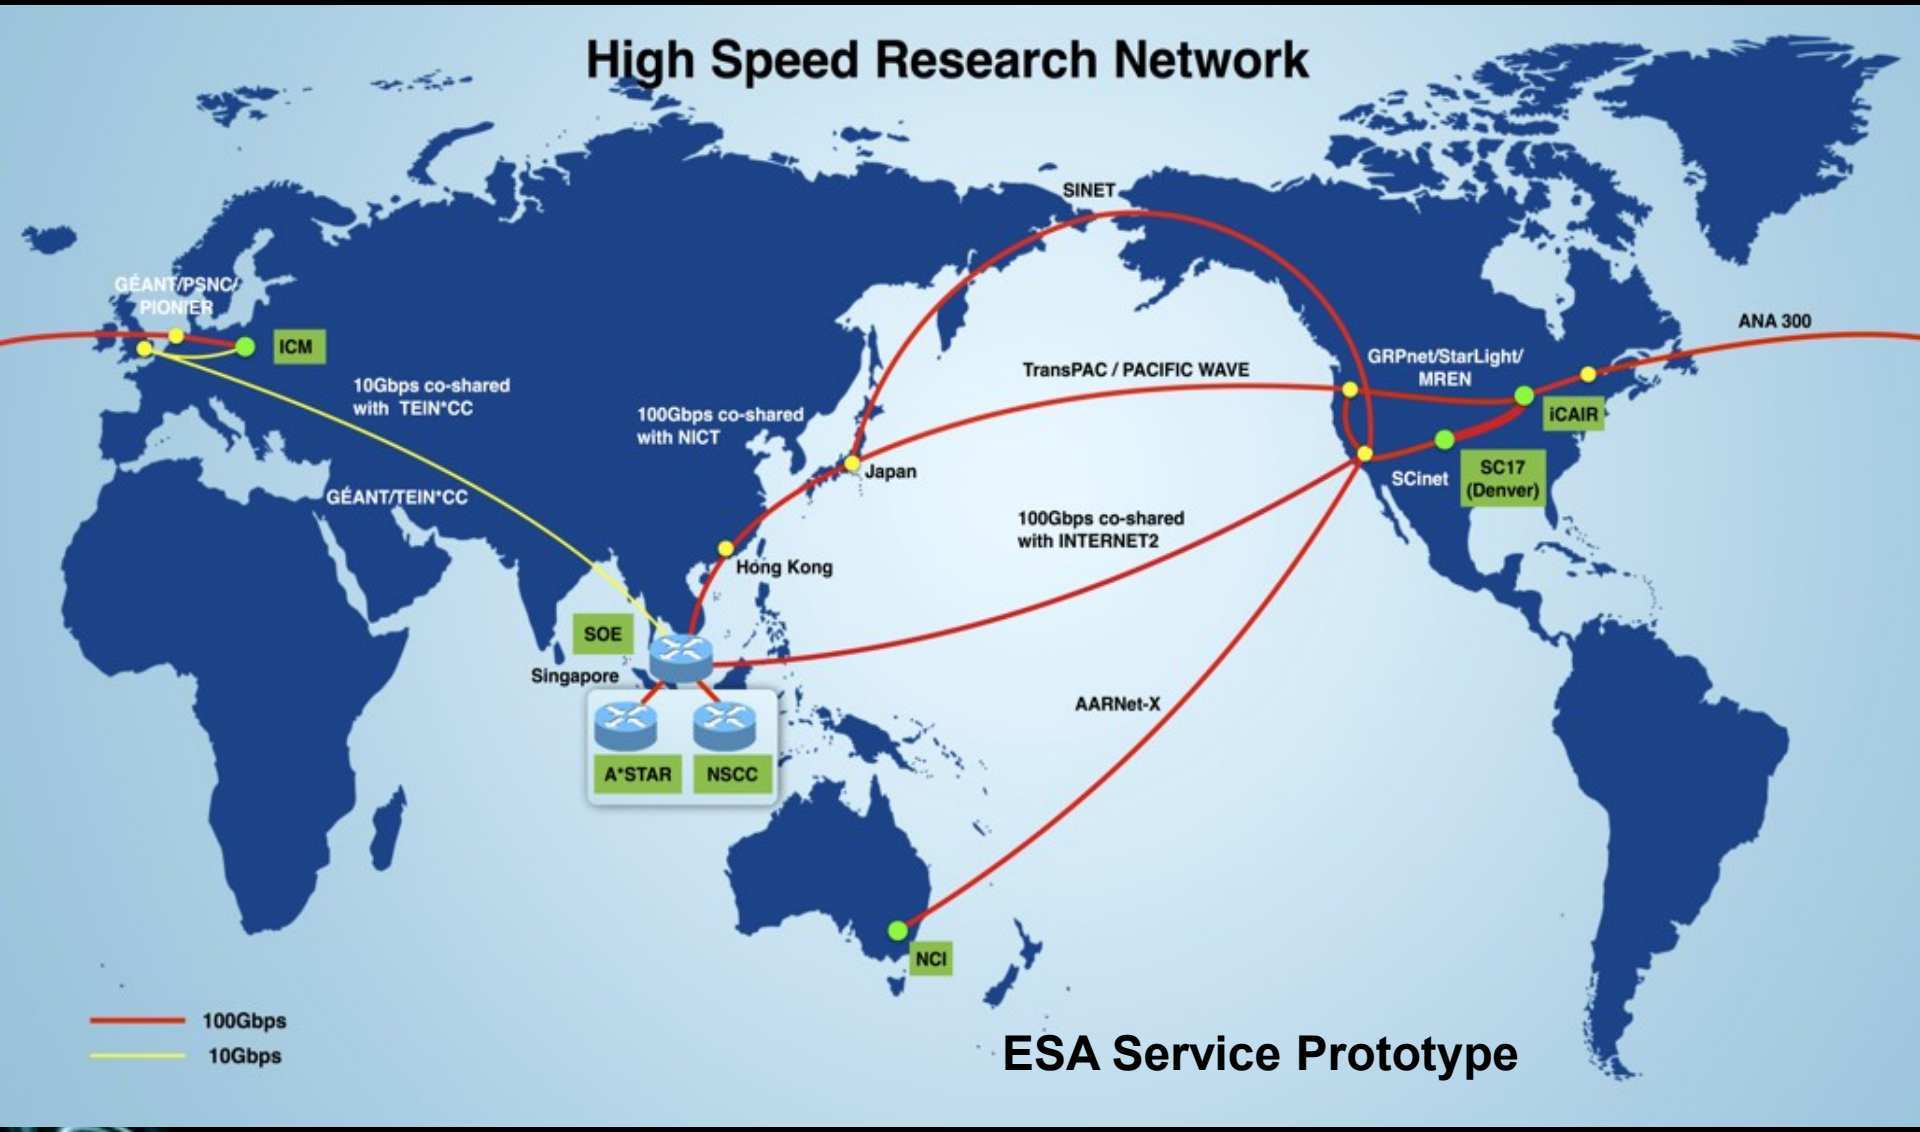
High Speed Research Network
SINET
GÉANT/PSNC/
PIONIER
ICM
10Gbps co-shared
with TEIN*CC
GÉANT/TEIN*CC
100Gbps co-shared
with NICT
Japan
Hong Kong
SOE
Singapore
A*STAR
NSCC
TransPAC / PACIFIC WAVE
GRPnet/StarLight/
MREN
iCAIR
ANA 300
SCinet
SC17
(Denver)
100Gbps co-shared
with INTERNET2
AARNet-X
NCI
100Gbps
10Gbps
ESA Service Prototype

StarLight SDX Geoscience Research Workflow

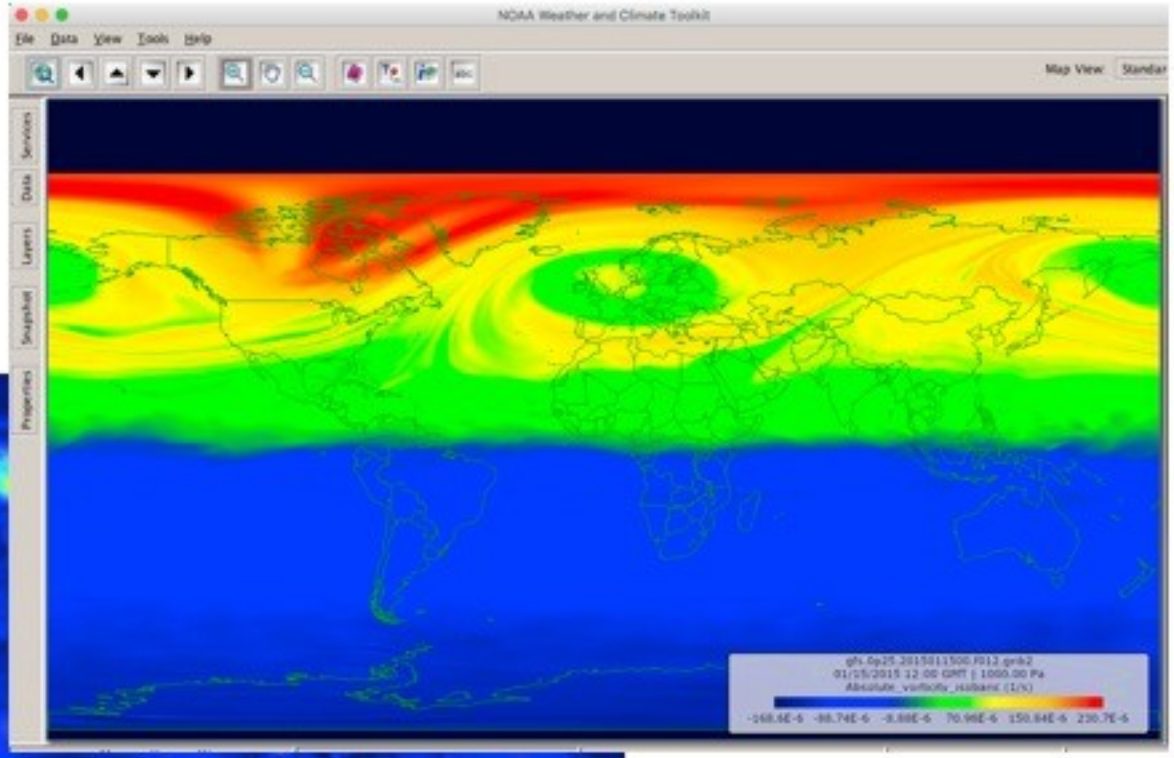

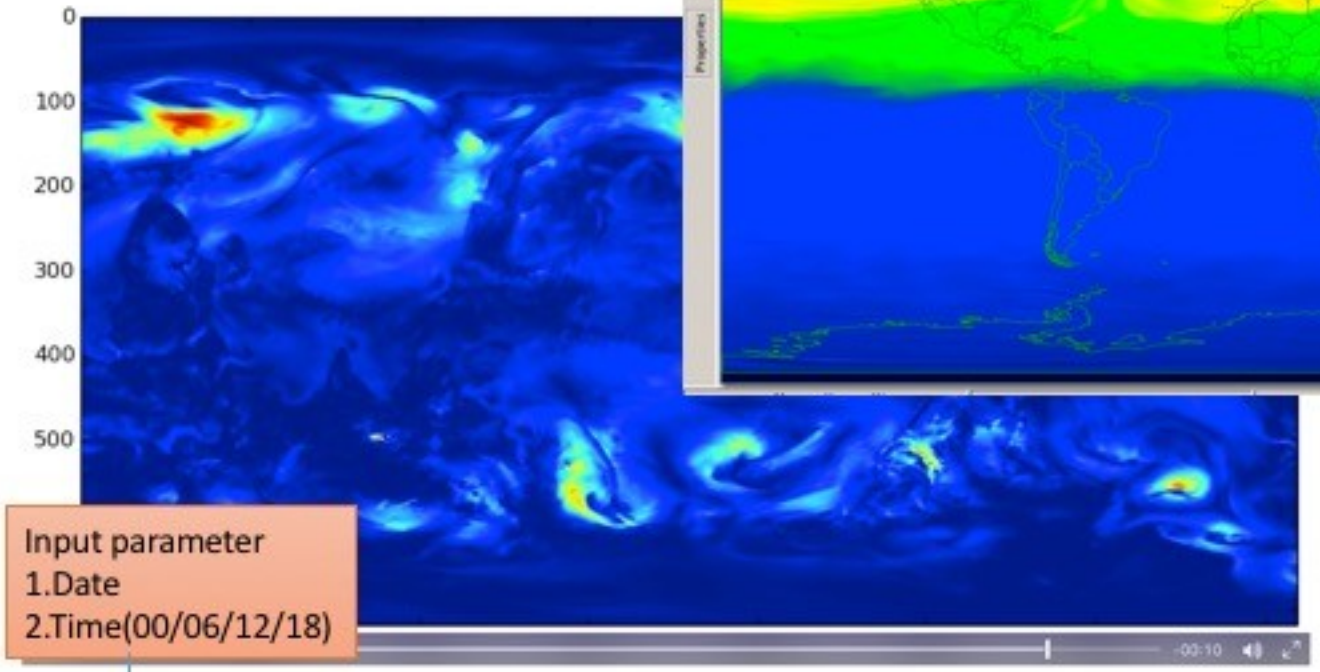

Input parameter
1.Date
2.Time(00/06/12/18)

Download To SDX DTN → grib2 → Data convert1 → netCDF → Data convert2 → png → Convert to video → mp4 → Display → Select Data → Data Transfer To Other DTN

GeoScience SDX DTN Service Prototype

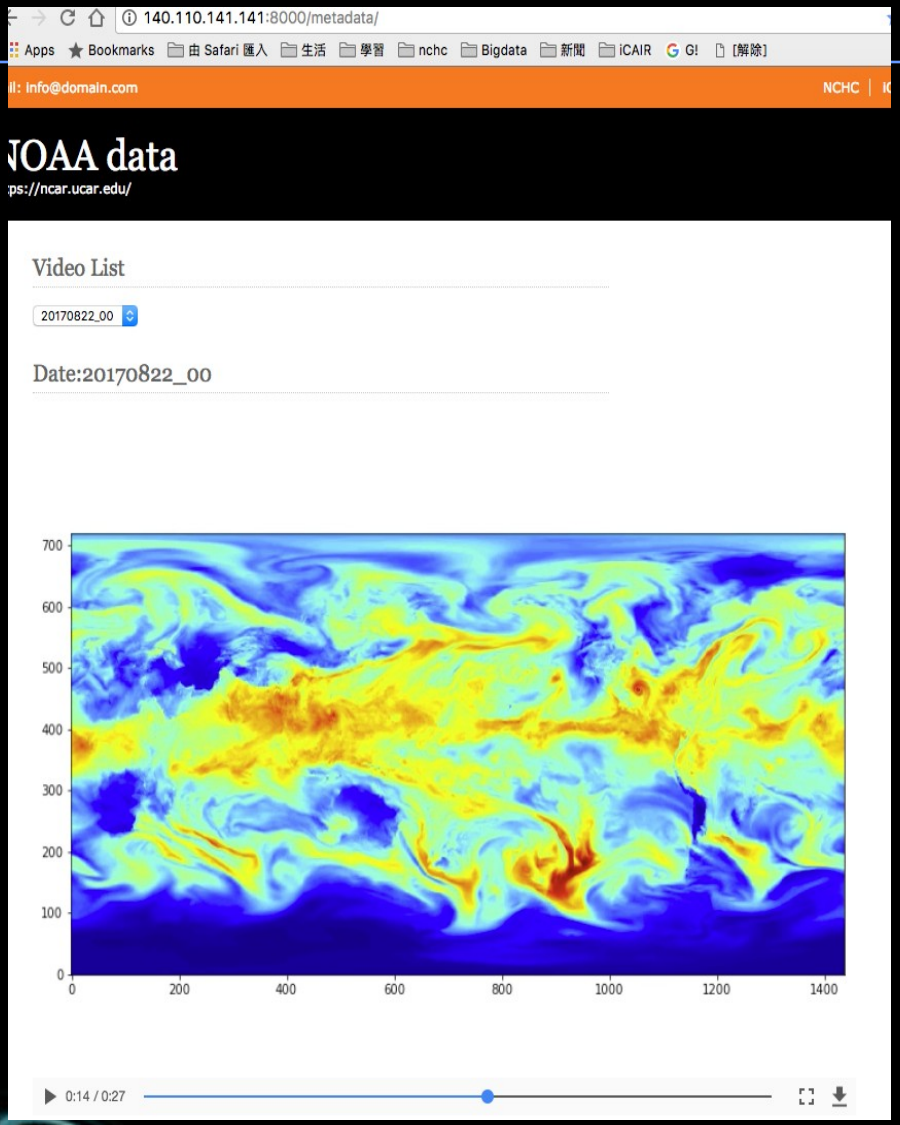

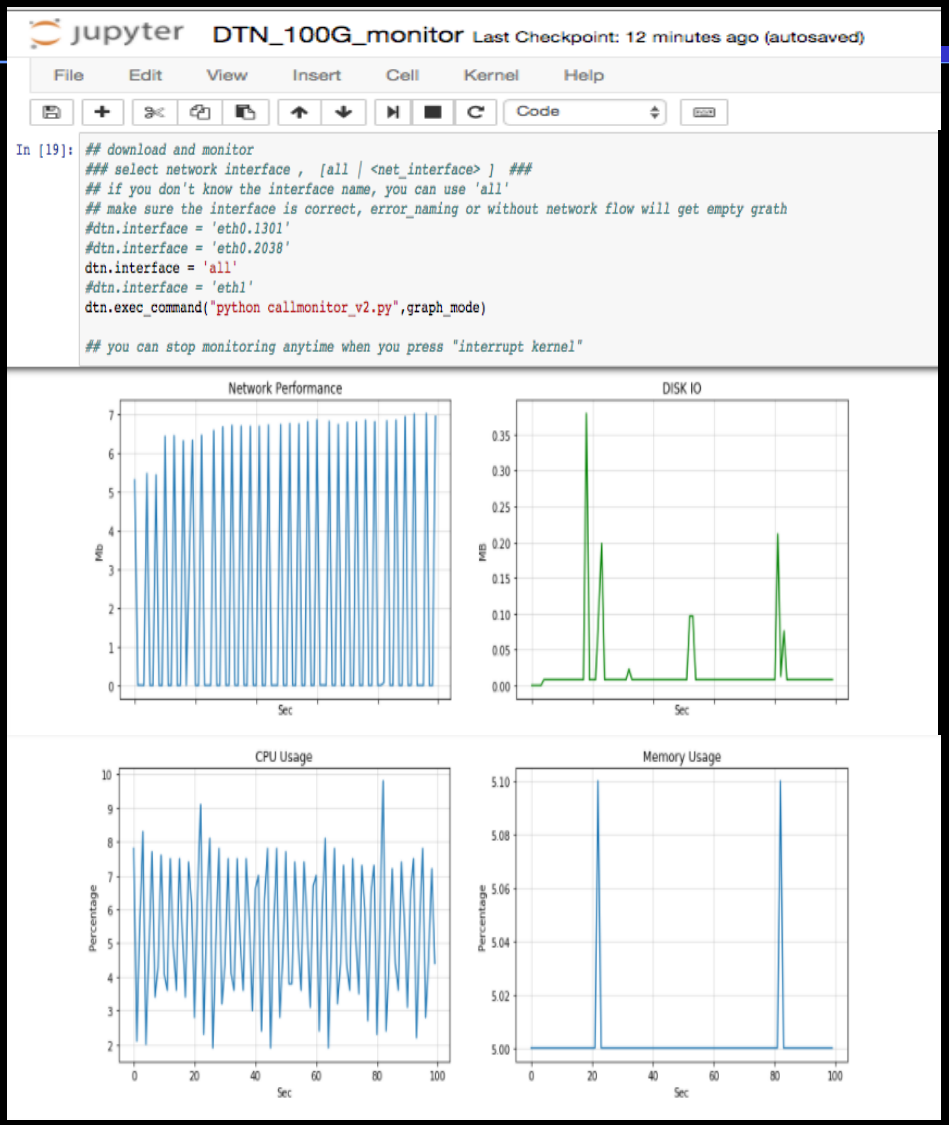

File Screening workflow

File Transfer workflow

NOAA data

2017.6.5_06

Band 10 Ozone Mixing Ratio [kg/kg]
Isobaric surface

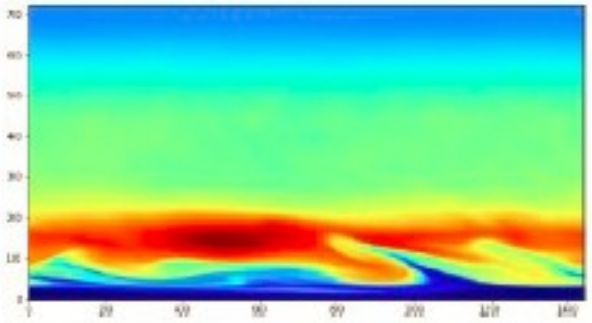

Band 47 Absolute vorticity [1/s]
Isobaric surface

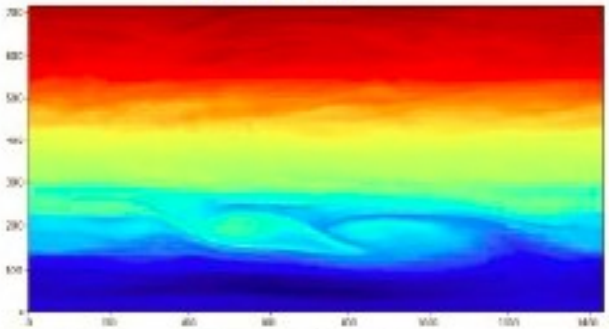

Band 20 u-component of wind [m/s]
Isobaric surface

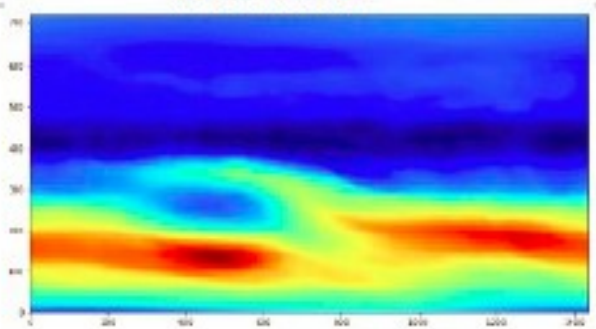

Band 315 Temperature [C]
Level at specified pressure difference

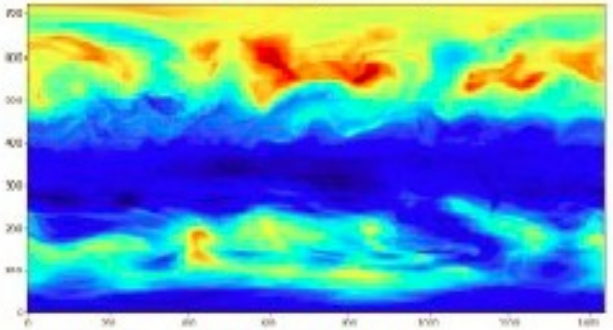

Band 49 Geopotential height [gpm]
Isobaric surface

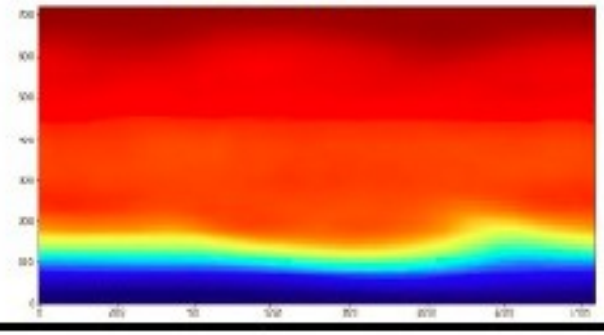

Band 328 Temperature [C]
Sigma level

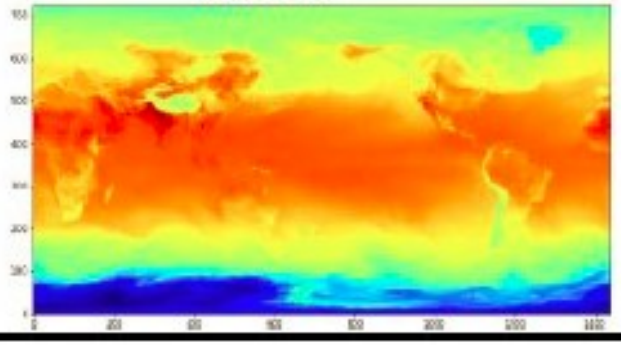

BigData Express – Toward Schedulable, Predictable, and High-performance Data Transfer

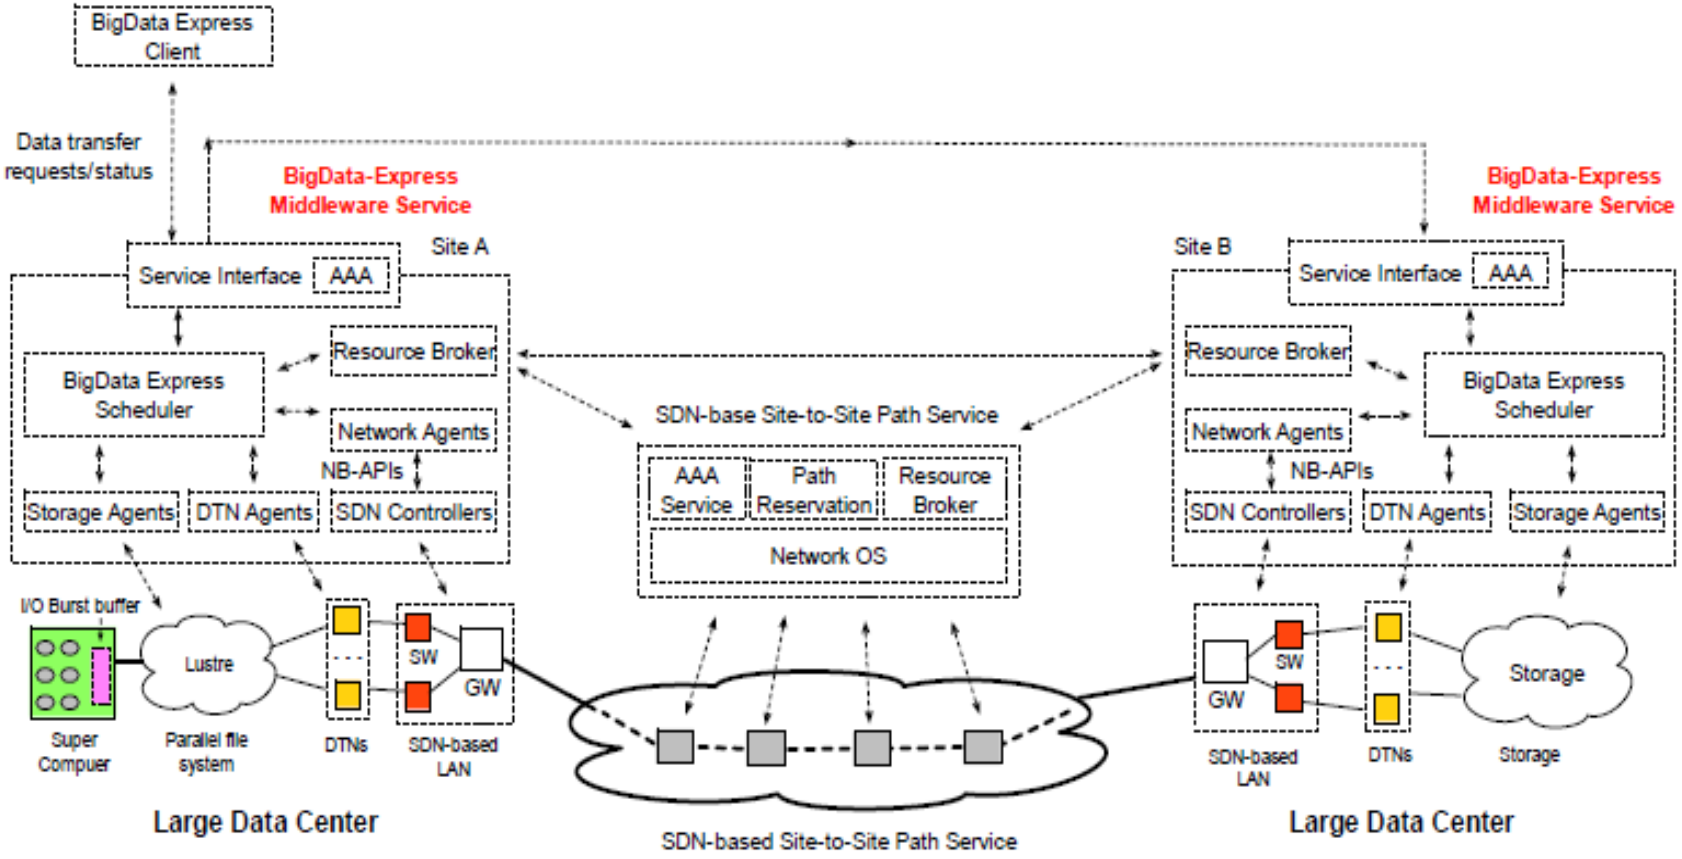

A Cross-Pacific SDN Testbed

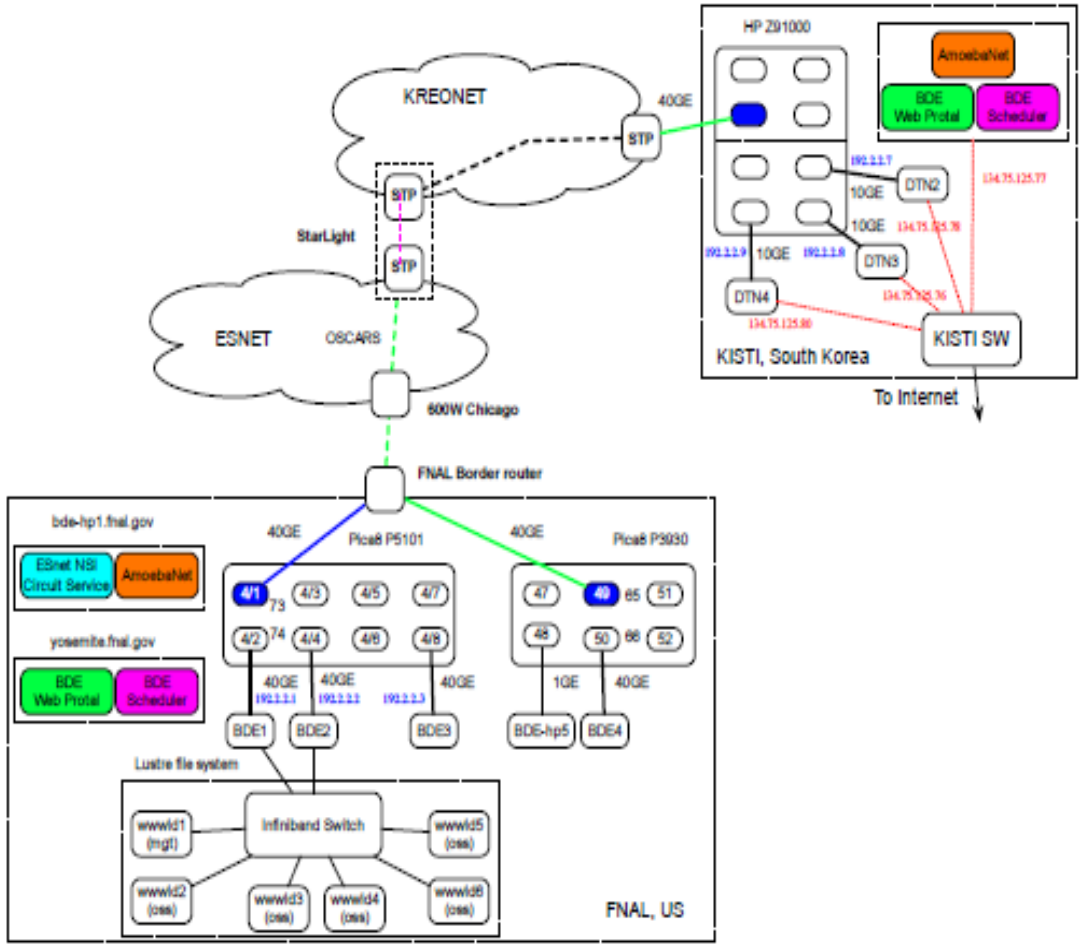

BigData Express SC'17 DEMO

- BigData Express: a schedulable, predictable, and high-performance data transfer service
  - QoS-guaranteed data transfer
  - DTN as a service
  - Network as a service
  - Distributed resource brokering/matching

Fermilab · Oak Ridge National Laboratory · KISTI (www.kisti.re.kr) · KREONET · ESnet · StarLight The Optical STAR TAP℠

A DOE/SC/ASCR-sponsored research project

Software is available at: http://bigdataexpress.fnal.gov

Caltech, StarLightOCC and Partners at SC17

- ~3 Tbps each at the Caltech and OCC booths
- *Connection to the Dell booth*
- 1+ Tbps between the Booths and ~3Tbps to the WAN
- Caltech Booth:
  200G dedicated to Caltech campus;
  300G to PRP (UCSD, Stanford, UCSC, et al);
  300G to Brazil+Chile via Miami;
  200G to ESnet
- Waveserver Ai + other DCIs in the booths: *N X 100GE to 400G, 200G waves*
- Wavelength Sensitive Switching in the Ciena 6500 platform

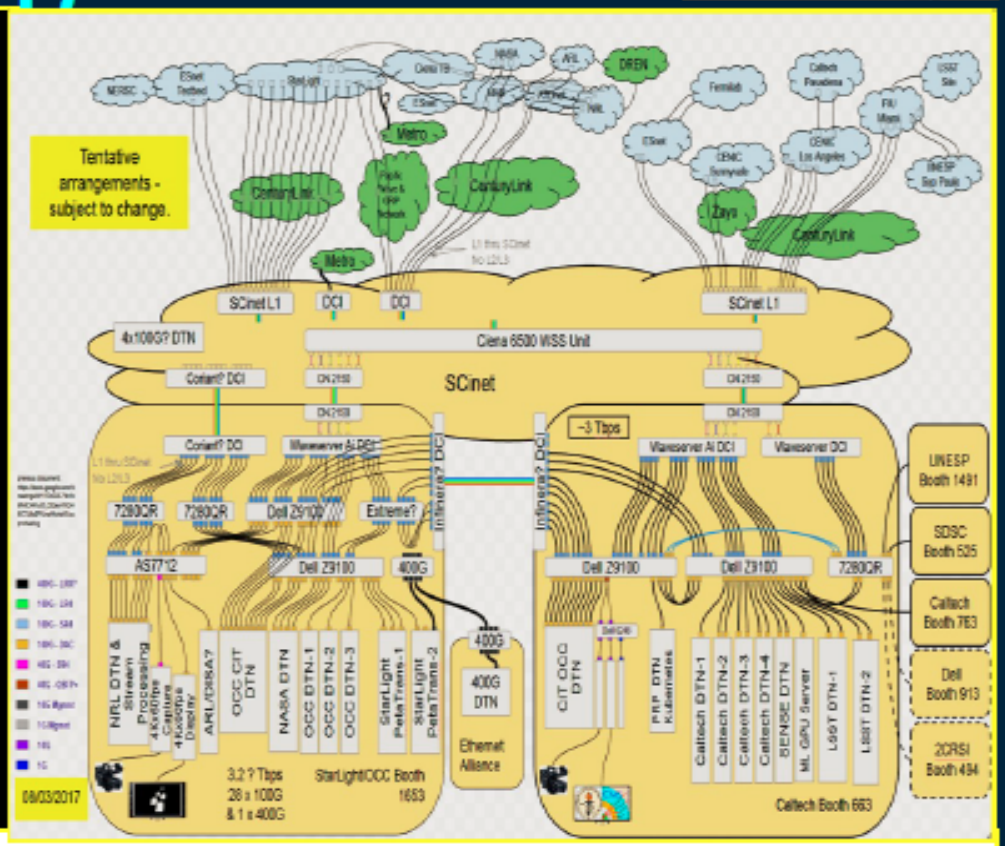

★ Microcosm: Creating the Future SCinet and the Future of Networks for Science

STARLIGHT℠

Transferring LargeScale Airline Data E2E Across WANs Using DTNs

v5, 21 SEP 2017

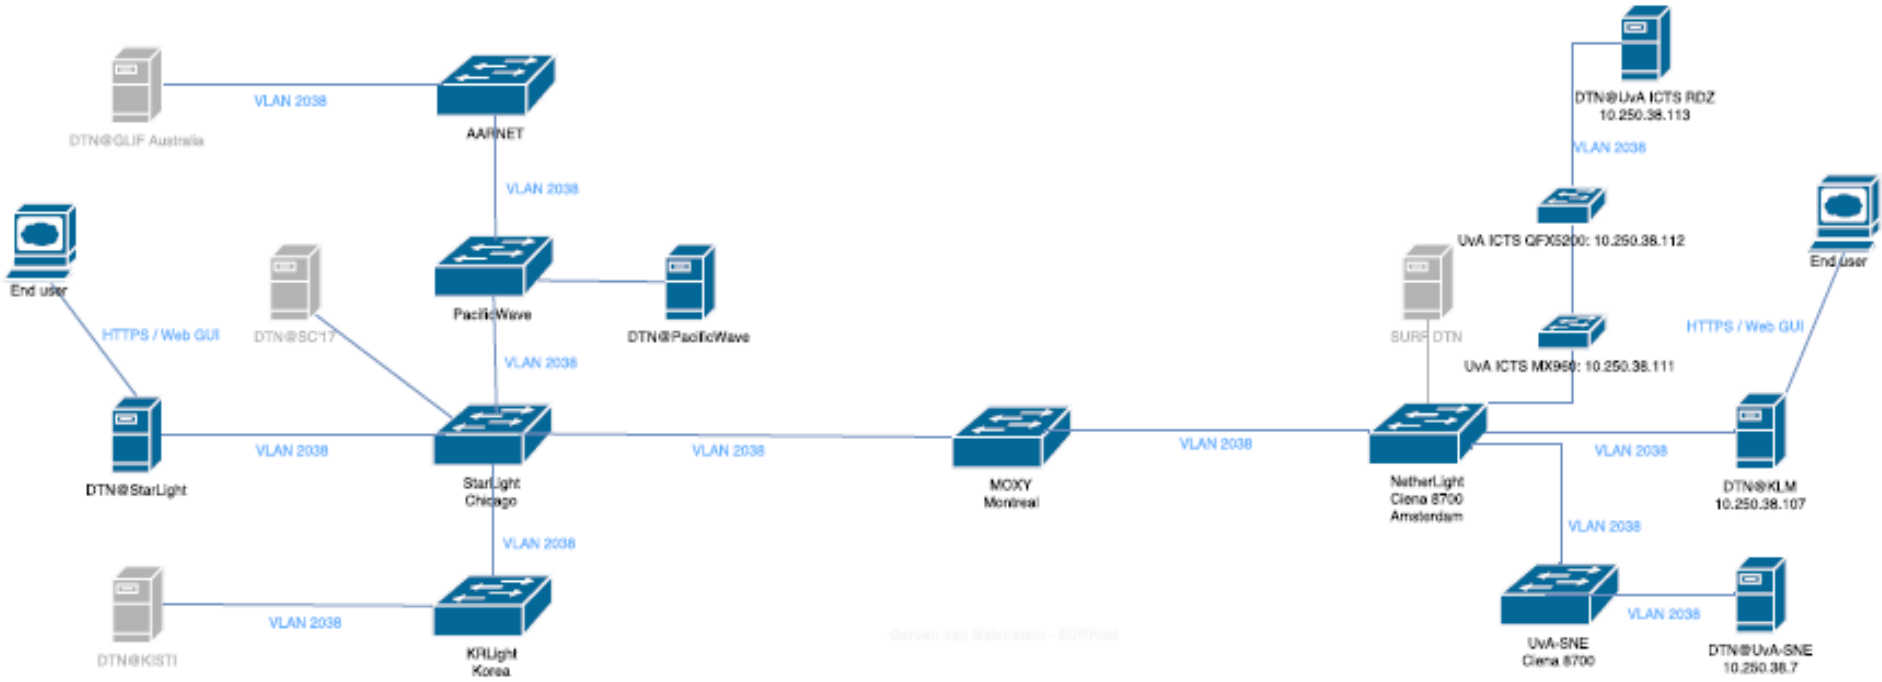

Ingredients
- Using Globus Toolkit (NOT Globus Online)
- Has GridFTP under the hood
- Under Globus license (must be evaluated)
- 40Gbit/s data transfer expected
- VLAN 2038, multipoint/extending
- Including authentication/authorization framework, e.g. SURFconext

Minimal setup
- Data transfer between DTN@UvA to DTN@StarLight at 40G
- Compare this to IPv4 performance Chicago-Amsterdam

Additional features
- Single Sign-On
- Comparison to IPv6
- Auto-deletion of file when transfer completed
- >40Gbps data transfer
- Expanding sites for GLIF and/or SC

Ideas
- Dutch Research LAN Project

Compute Canada/CANARIE/StarLight SC17 Demonstrations

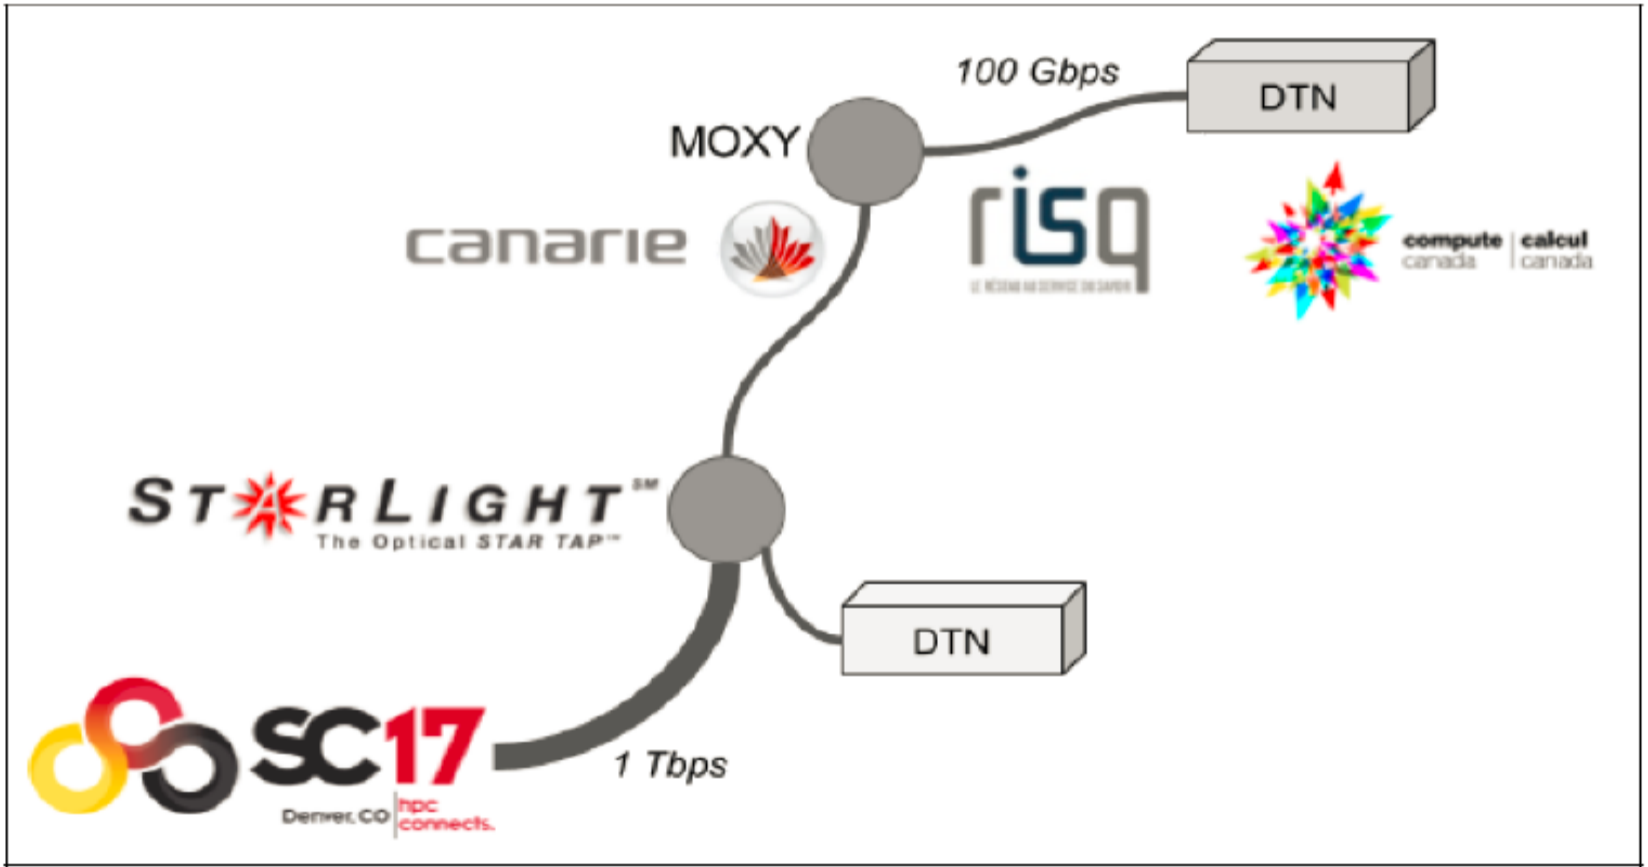

Figure 1: A high-level network map

Compute Canada/CANARIE/StarLight SC17 Demonstrations

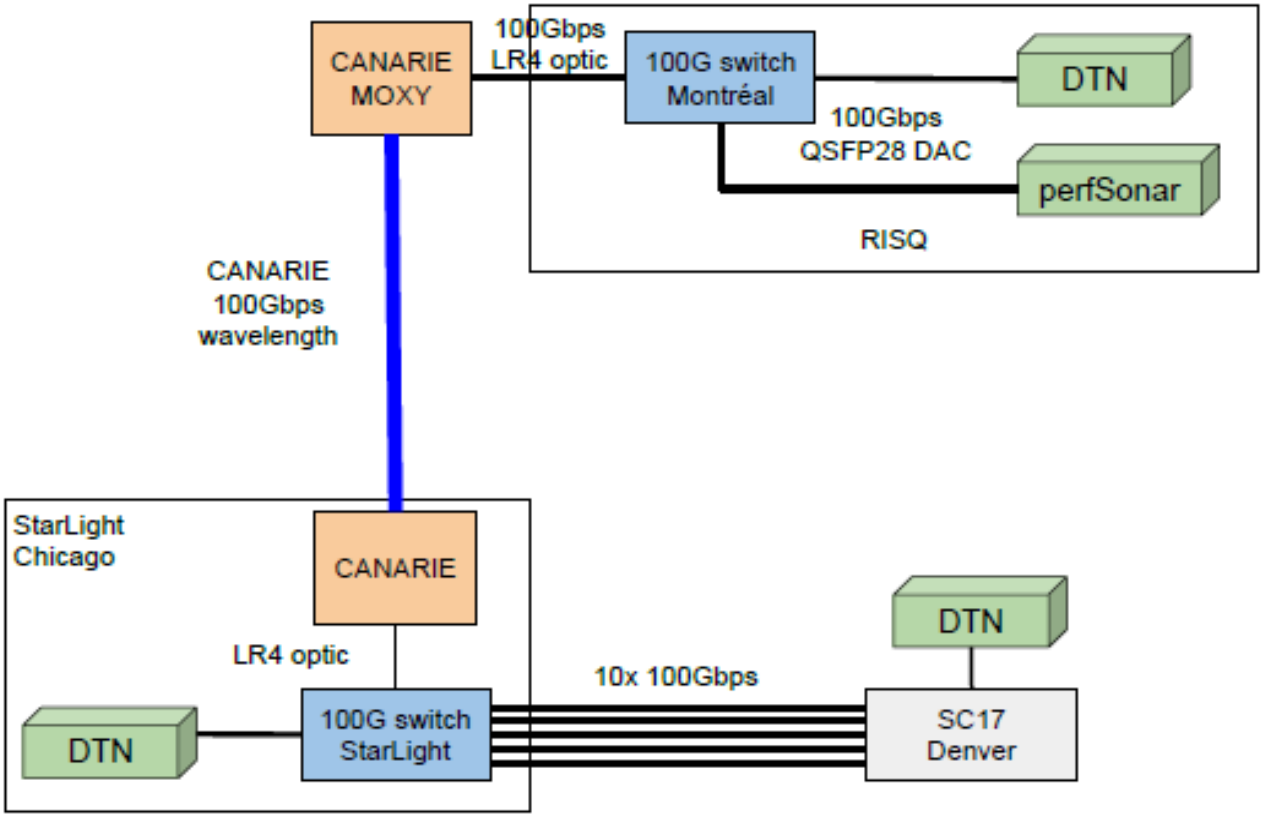

Multi-Institution, Hyper-Converged ScienceDMZ

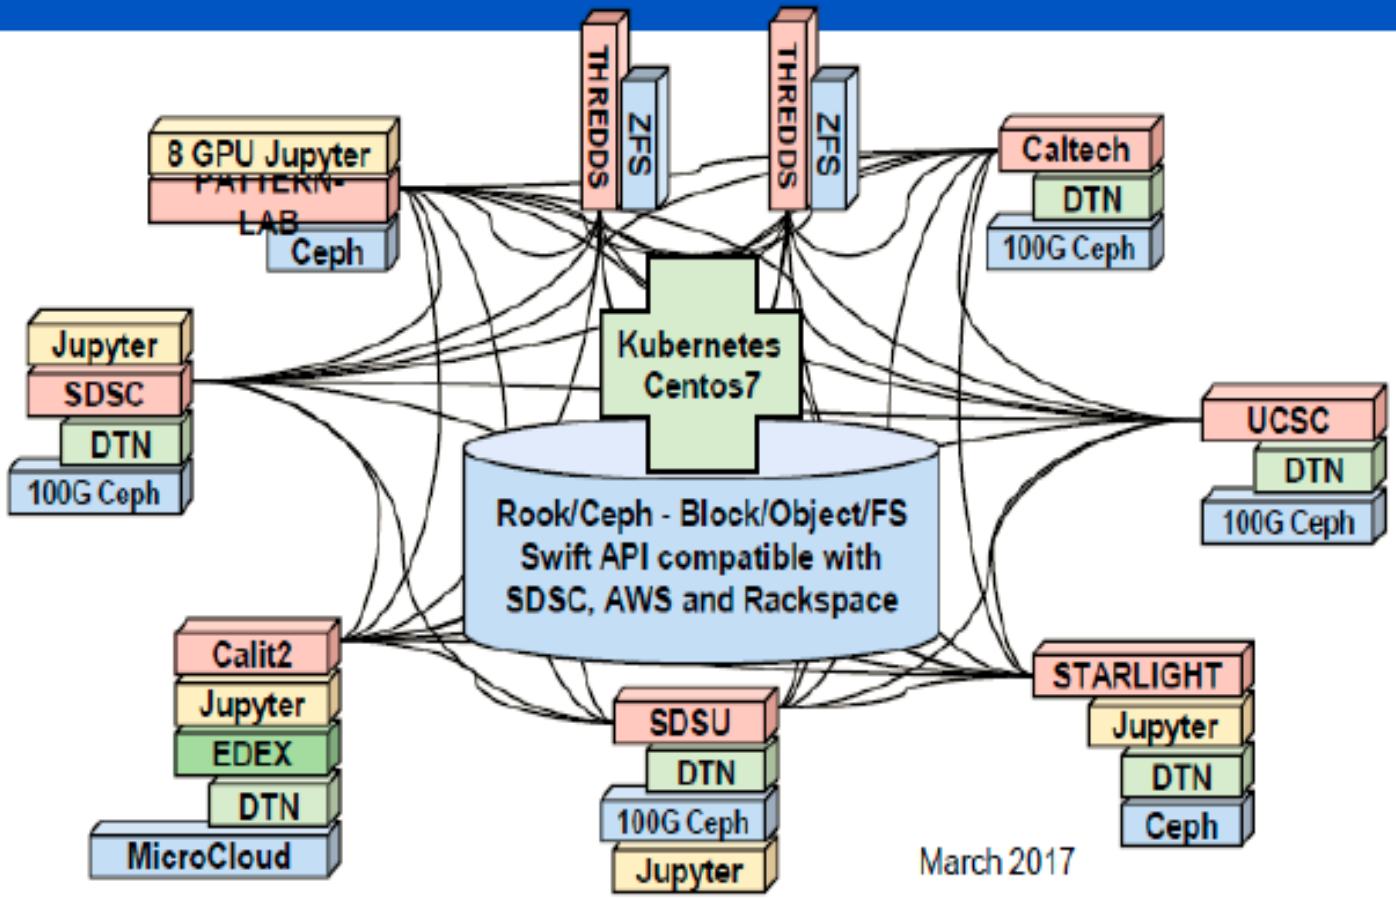

Source; John Graham UCSD

it2 Innovation
STARLIGHT SM

Implementing a SCinet DTN

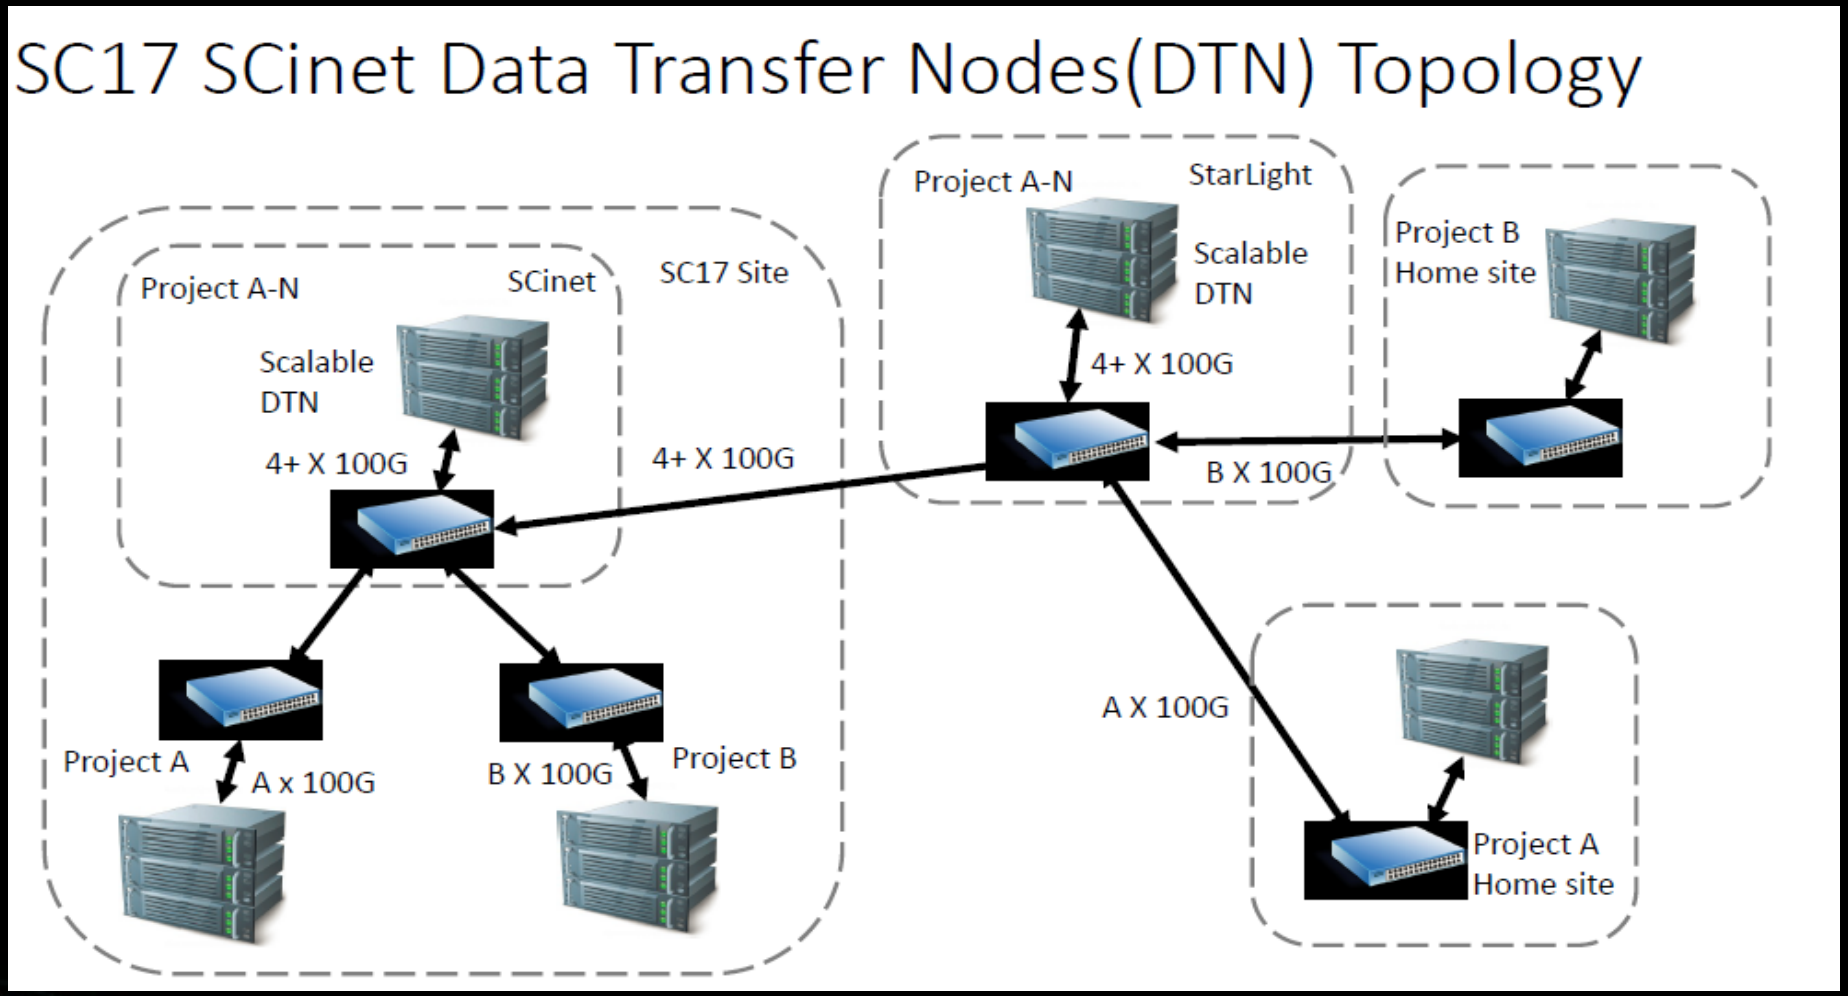

Source; Jim Chen, iCAIR

STARLIGHT SM

The OSiRIS Project: Collaborative Access to Data

CNSECCS Symposium – May 18, 2017

OSiRIS

Open Storage Research Infrastructure

Project Outline

- rationale
- general goals

Project Components

- technical goals
- science users

Future Plans

- current science domains
- science domain roadmap
- technical focus areas

Community Benefits

- technical contributions
- scientific advancement

STARLIGHT℠

Programmable Network Measurement of Data Intensive Flows on 100Gbps Networks

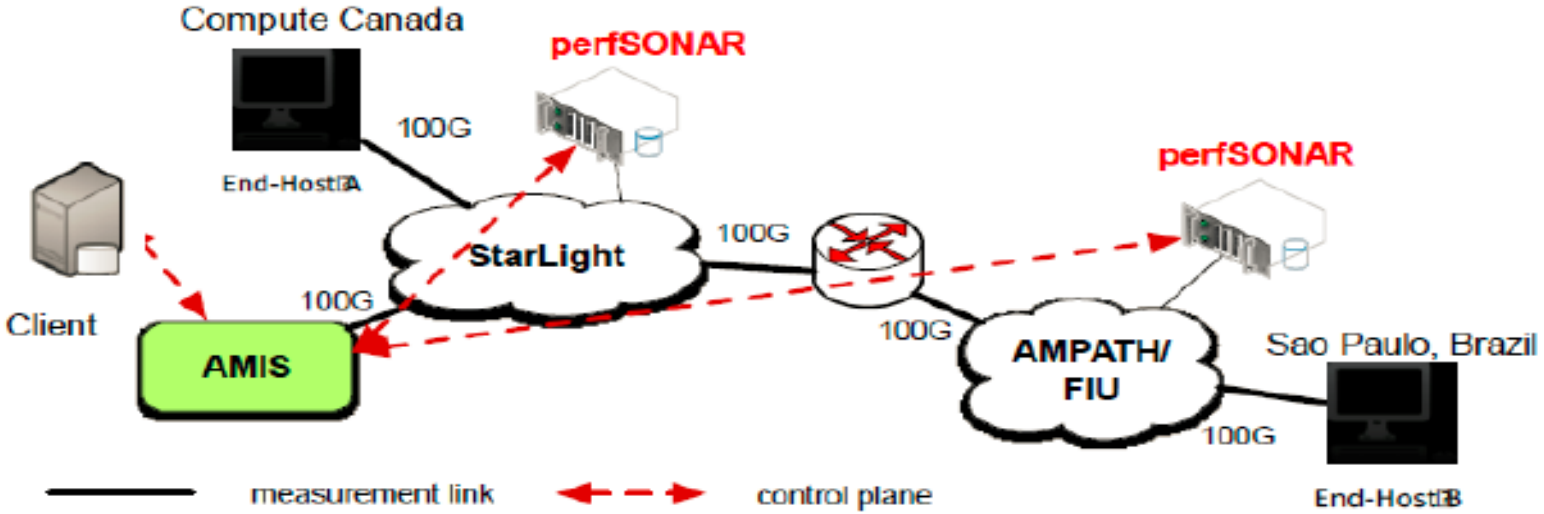

Demo1: Programmable Measurement with RESTful APIs

Demo2: Passive & Active Measurement (TCP window size)

Demo3: Passive & Active Measurement (TCP packet loss)

RNC AMIS Team: Yan Luo, PI, University of Massachusetts Lowell; Gabriel Ghinita, Co-
I, University of Massachusetts Boston; Cody Bumgardner, Co-PI, University of Kentucky;
Michael McGarry, Co-PI, University of Texas El Paso. Contact: **Yan_Luo@uml.edu**
ollaborators: Jeo Mambretti, Jim Chen and Fei Yeh, StarLight/iCAIR/Northwestern
niversity; Jeronimo Bezerra, AMPATH/Florida International University

Programmable Privacy-Preserving Network Measurement for Network Usage Analysis and Troubleshooting

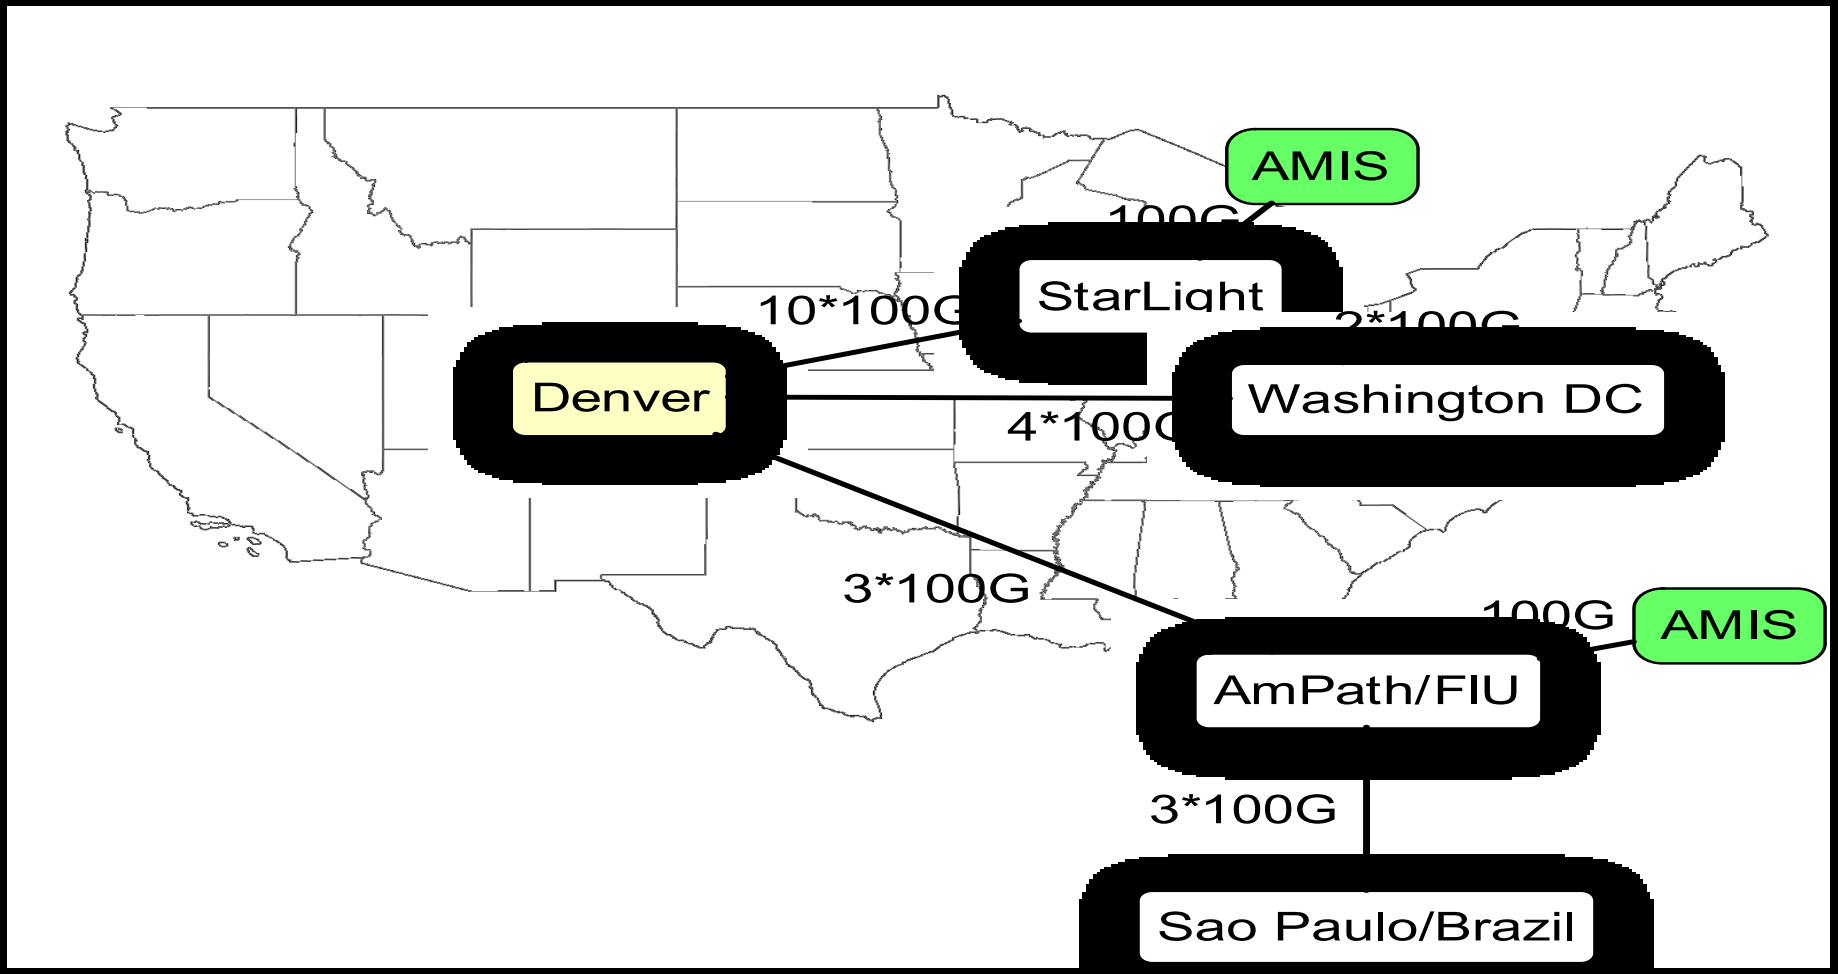


v3.0.0-master-0fbecb9 2017/11/10 [Electron]
SAGE2@SC17
Webview: appear.in – one click video conversations
Invite to start a meeting
Deep Zoom Image: GLIF2017_World_32k_9-21_wrap.dzi 98,304 x 16,384 pixels
Webview: Potree Viewer
Automated GOLE Fabric

Building the Open Storage Network

Alex Szalay
The Johns Hopkins University

Institute for Data Intensive Engineering and Science
idies

Connections

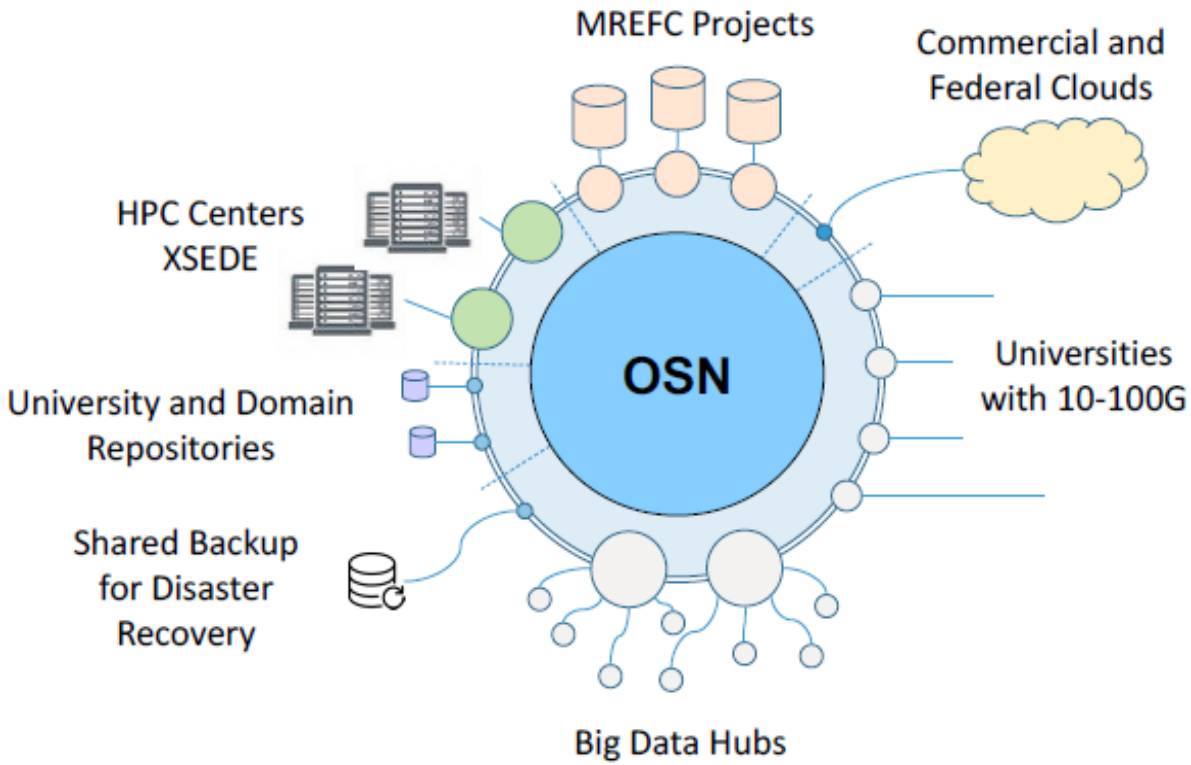

Nationwide 100 Gbps and Minimized Latency

◆ SINET5 will be a nationwide 100-Gbps backbone network using 100-Gigabit Ethernet technology and connect each pair of nodes with a minimized latency.

SINET4

- Star-like topology
- Resource-consuming secondary circuits

SINET5

- Fully-meshed topology with redundancy
- Non-resource-consuming secondary paths

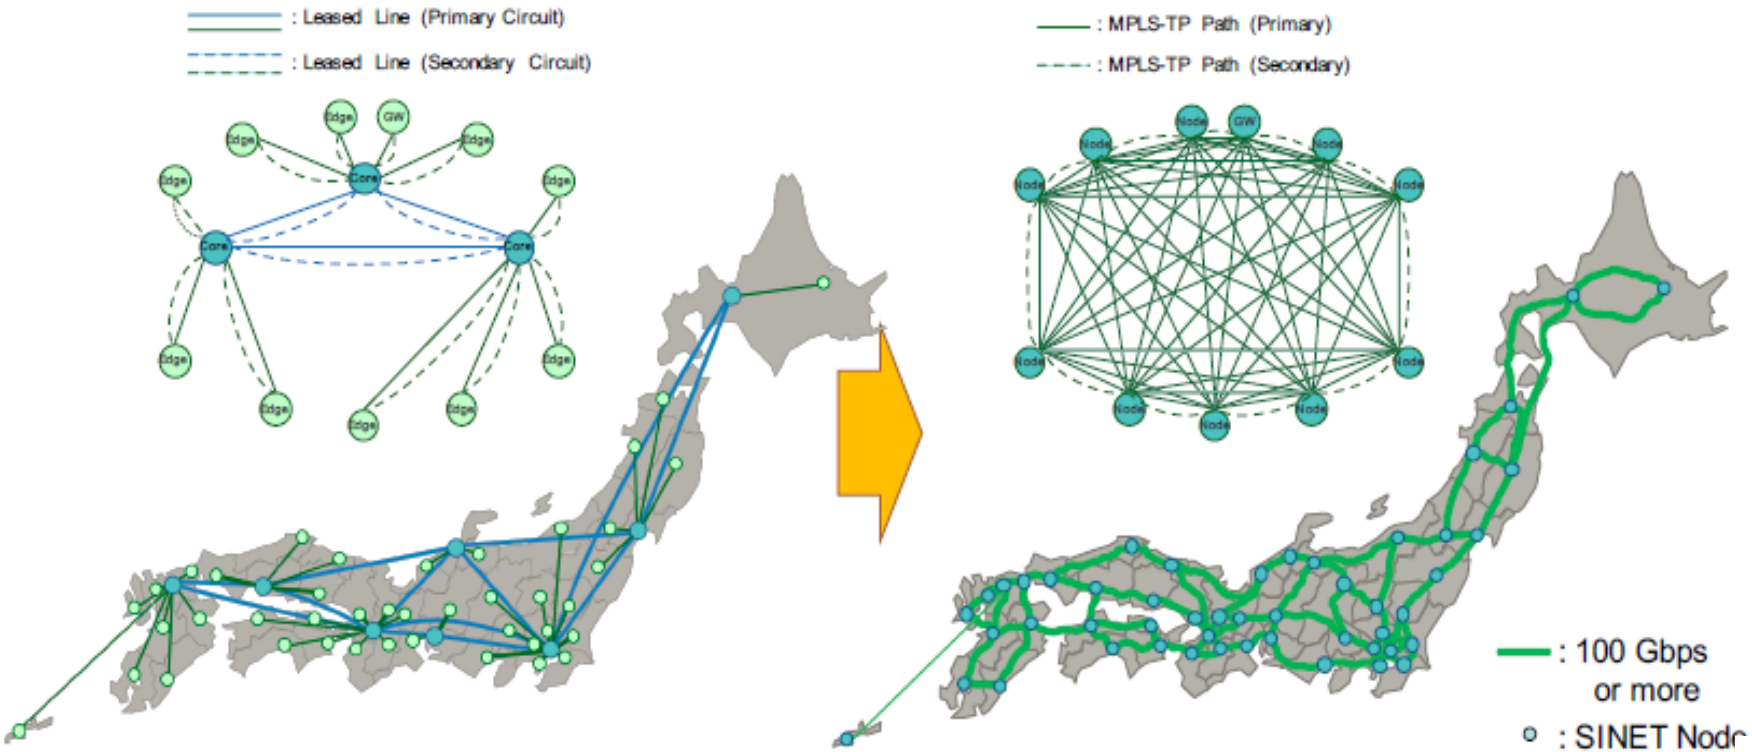

KREONet2 SD-WAN GLORIAD-KR
KISTI Daejeon ⇔ 100 G ⇔ StarLight

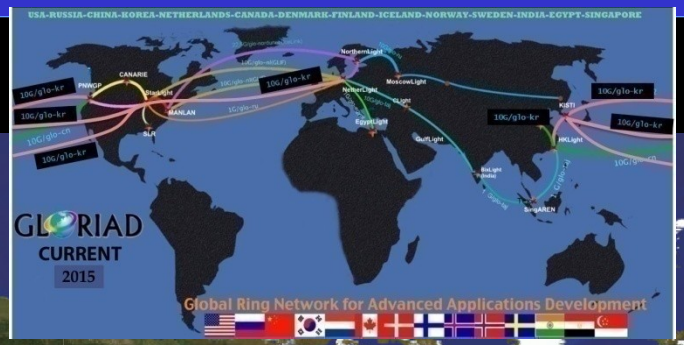

Soon: Daejeon SDX ⇔ StarLight SDX

International KREONET-S Connections to StarLight: SD-WAN Federations

Summary

- Data Intensive Science Can Benefit By Support From Enhanced Services/Techniques/Technologies Provided By A Global Research Platform, Including Services For High Performance WAN Data Transport
- One Approach Relys On L2 WAN Transport Channels
- Another Complementary Enabling Capability Uses DTNs Integrated With Specialized WAN Paths To Optimize E2E Data Flows
- These Core Components Can Be Supplemented By Enhancing Software Stacks, e.g., Jupyter, NSI, MEICAN, P4 Programs, BDE, AI/ML/DL, etc
- Today, All the Components Exist To Create An E2E Transport Service For Data Intensive Science
- It Is Possible To Create This Service and Place It Into Production

www.startap.net/starlight

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