DevOps and Software Defined Networking (SDNs) and Software Defined Exchanges (SDXs)

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)

MAGIC Meeting
Washington DC
October 3, 2018

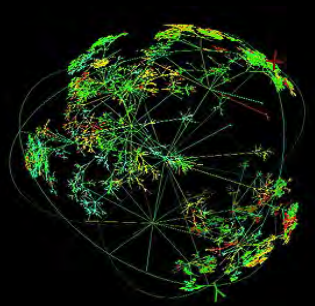

iCAIR

STARLIGHT SM

Introduction to iCAIR:

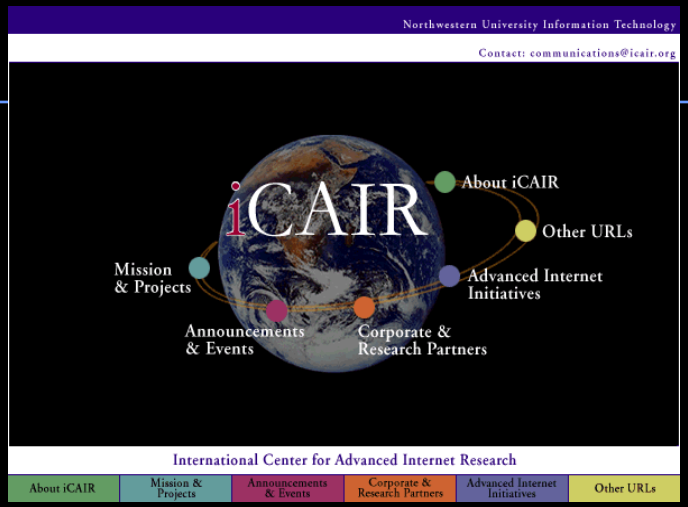

Accelerating Leading Edge Innovation and Enhanced Global Communications through Advanced Internet Technologies, in Partnership with the Global Community

- **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/Science Research Platforms)**

ALMA: Atacama Large Millimeter Array

ANDRILL: Antarctic Geological Drilling
www.andrill.org

BIRN: Biomedical Informatics Research Network
www.nbirn.net

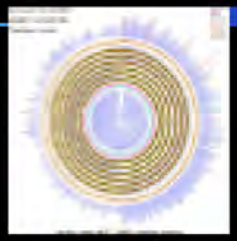

CAMERA metagenomics
camera.calit2.net

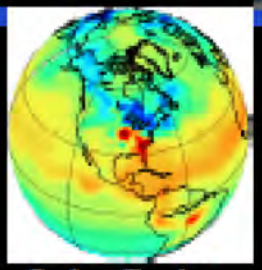

Carbon Tracker
www.esrl.noaa.gov/gmd/ccgg/carbontracker

CineGrid
www.cinegrid.org

LHCONE
www.lhcone.net

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

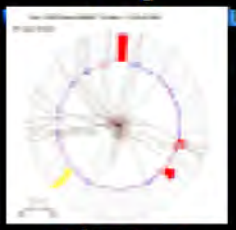

DØ (DZero)
www-d0.fnal.gov

LIGO
www.ligo.org

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

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

TeraGrid
www.teragrid.org

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

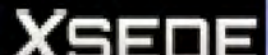

XSEDE
www.xsede.org

Compilation By Maxine Brown

STARLIGHT SM

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℠

Global Research Platform: Based On The Global Lambda Integrated Facility

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

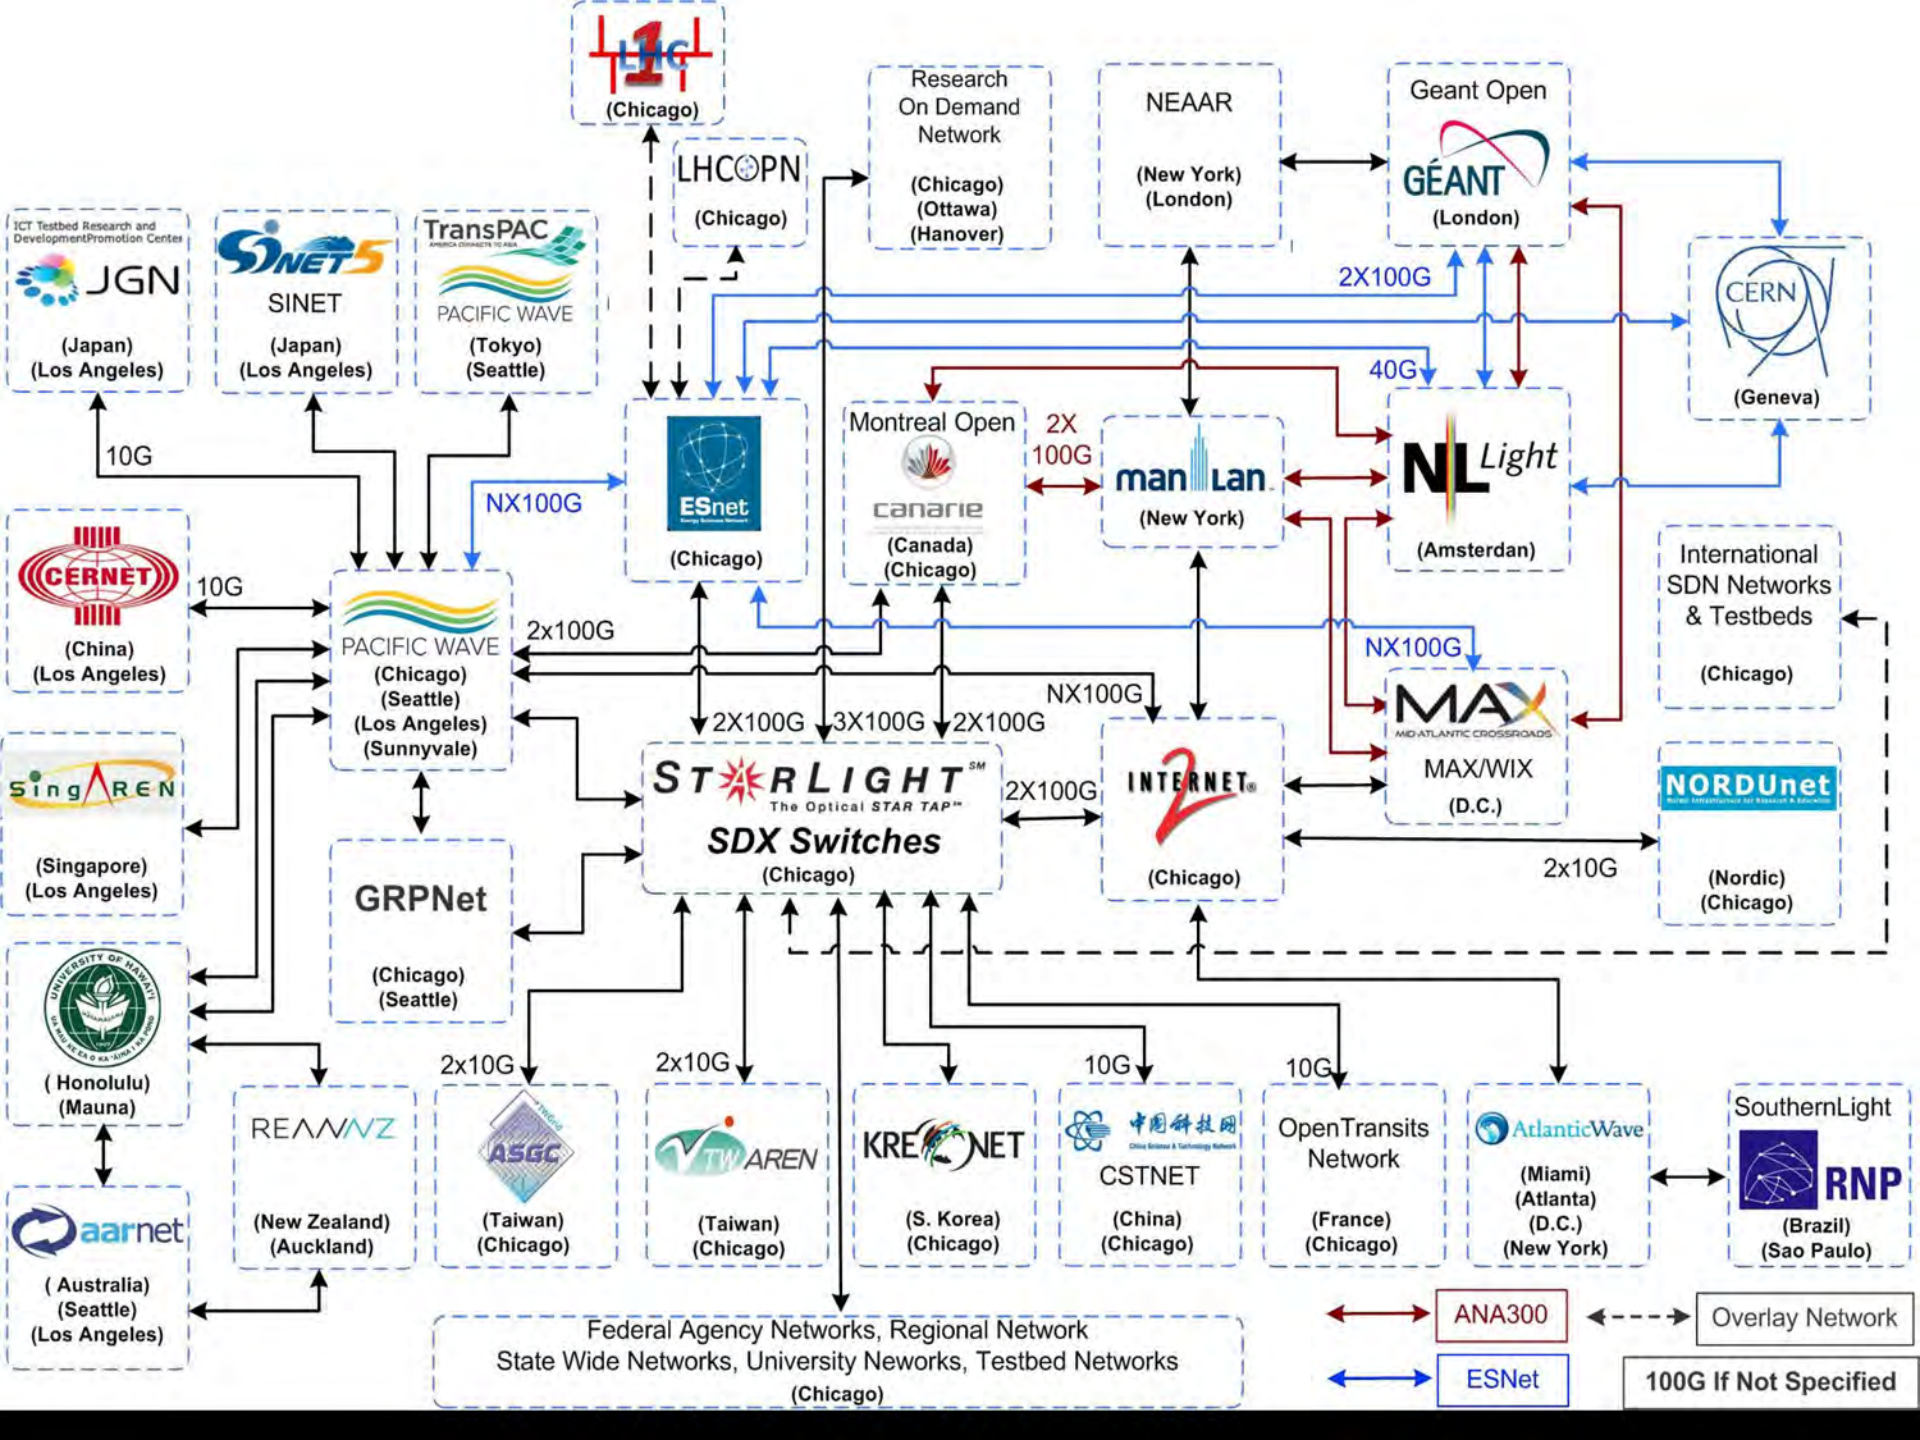
LHC
(Chicago)
LHCOPN
(Chicago)
Research On Demand Network
(Chicago)
(Ottawa)
(Hanover)
NEAAR
(New York)
(London)
Geant Open
GÉANT
(London)
CERN
(Geneva)
2X100G
40G
ICT Testbed Research and DevelopmentPromotion Center
JGN
(Japan)
(Los Angeles)
SINET
(Japan)
(Los Angeles)
TransPAC
PACIFIC WAVE
(Tokyo)
(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)
NORDUnet
(Nordic)
(Chicago)
2x10G
GRPNet
(Chicago)
(Seattle)
(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
ESNet
Overlay Network
100G If Not Specified

Emerging US SDX Interoperable Fabric

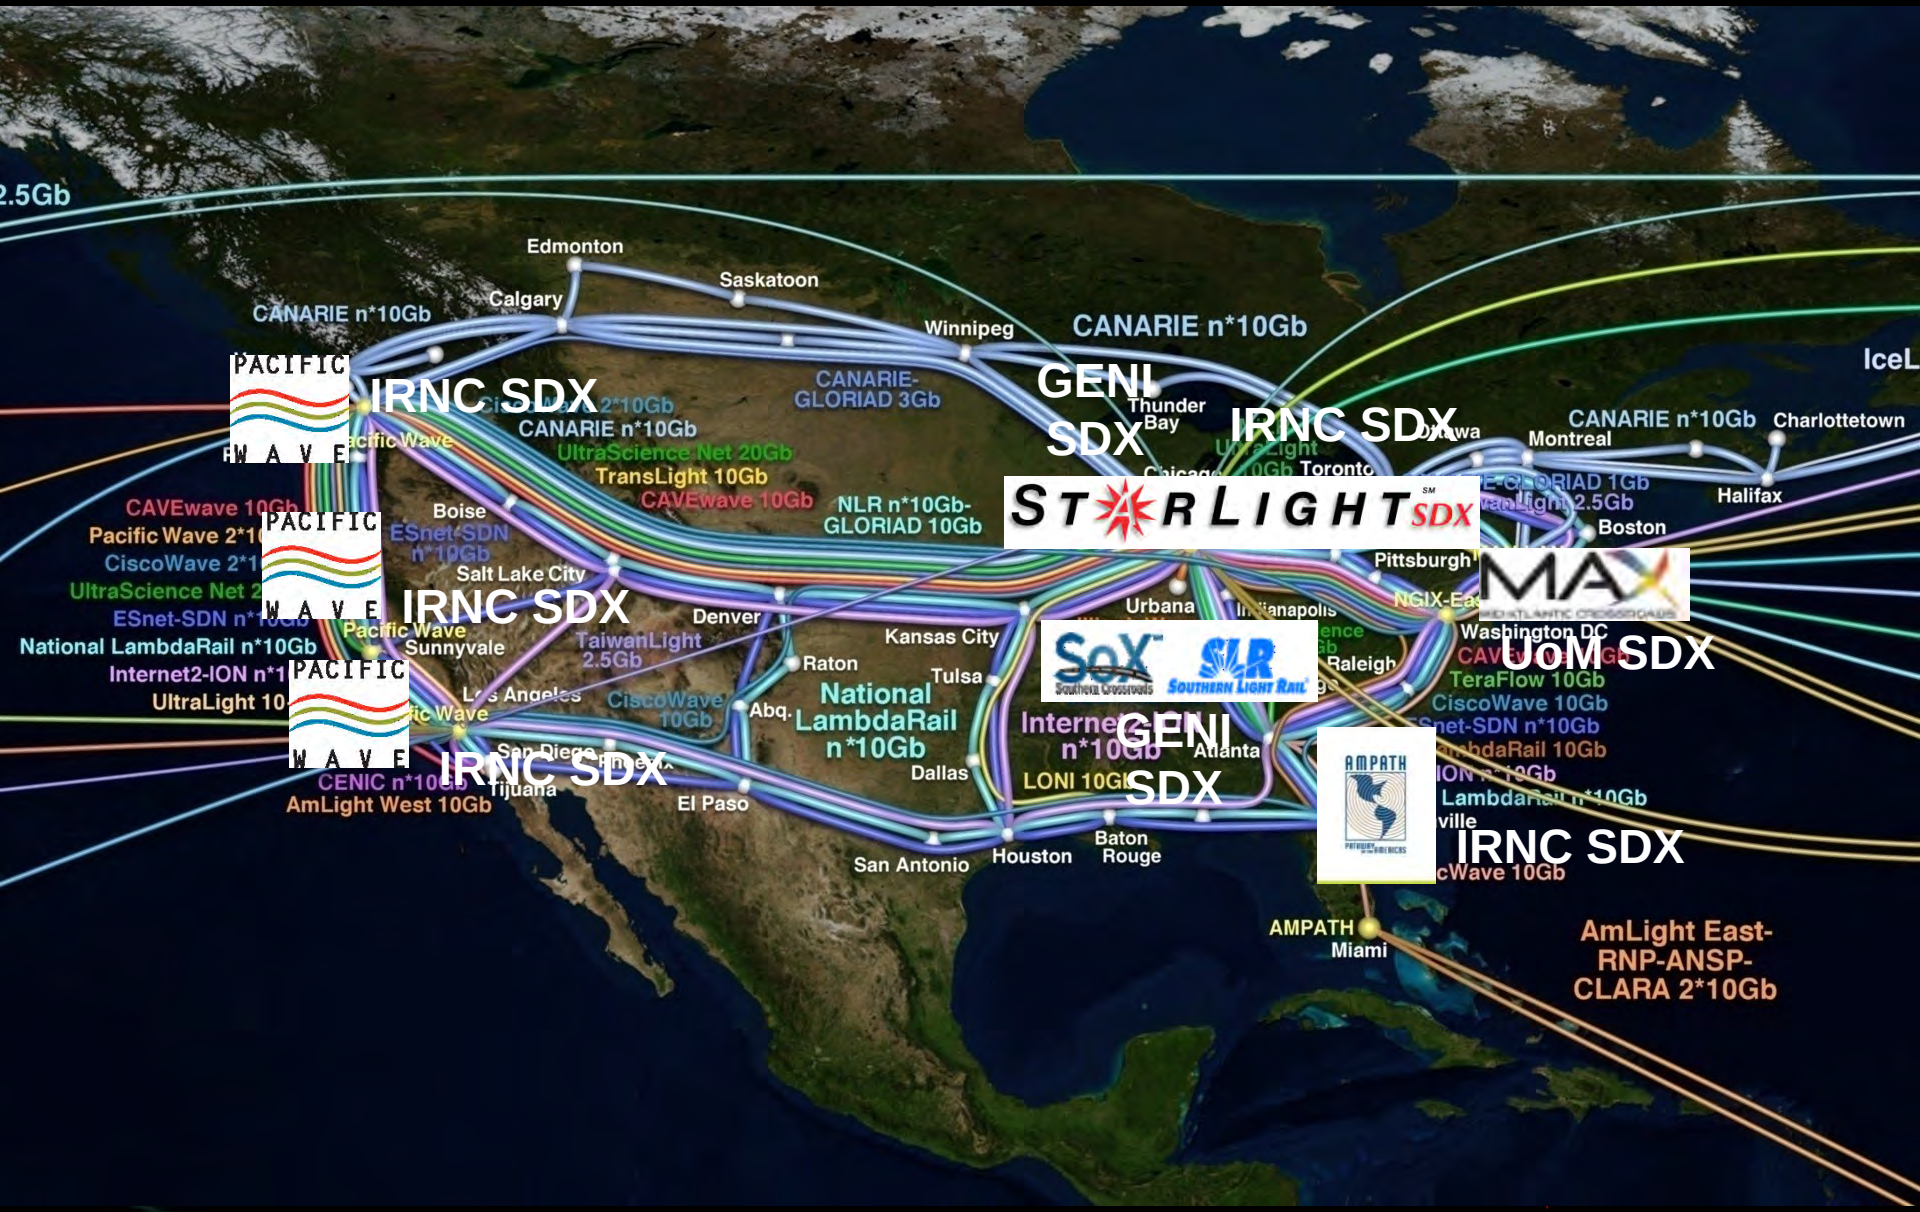

LHCONE: A global infrastructure for the High Energy Physics (LHC and Belle II) data management

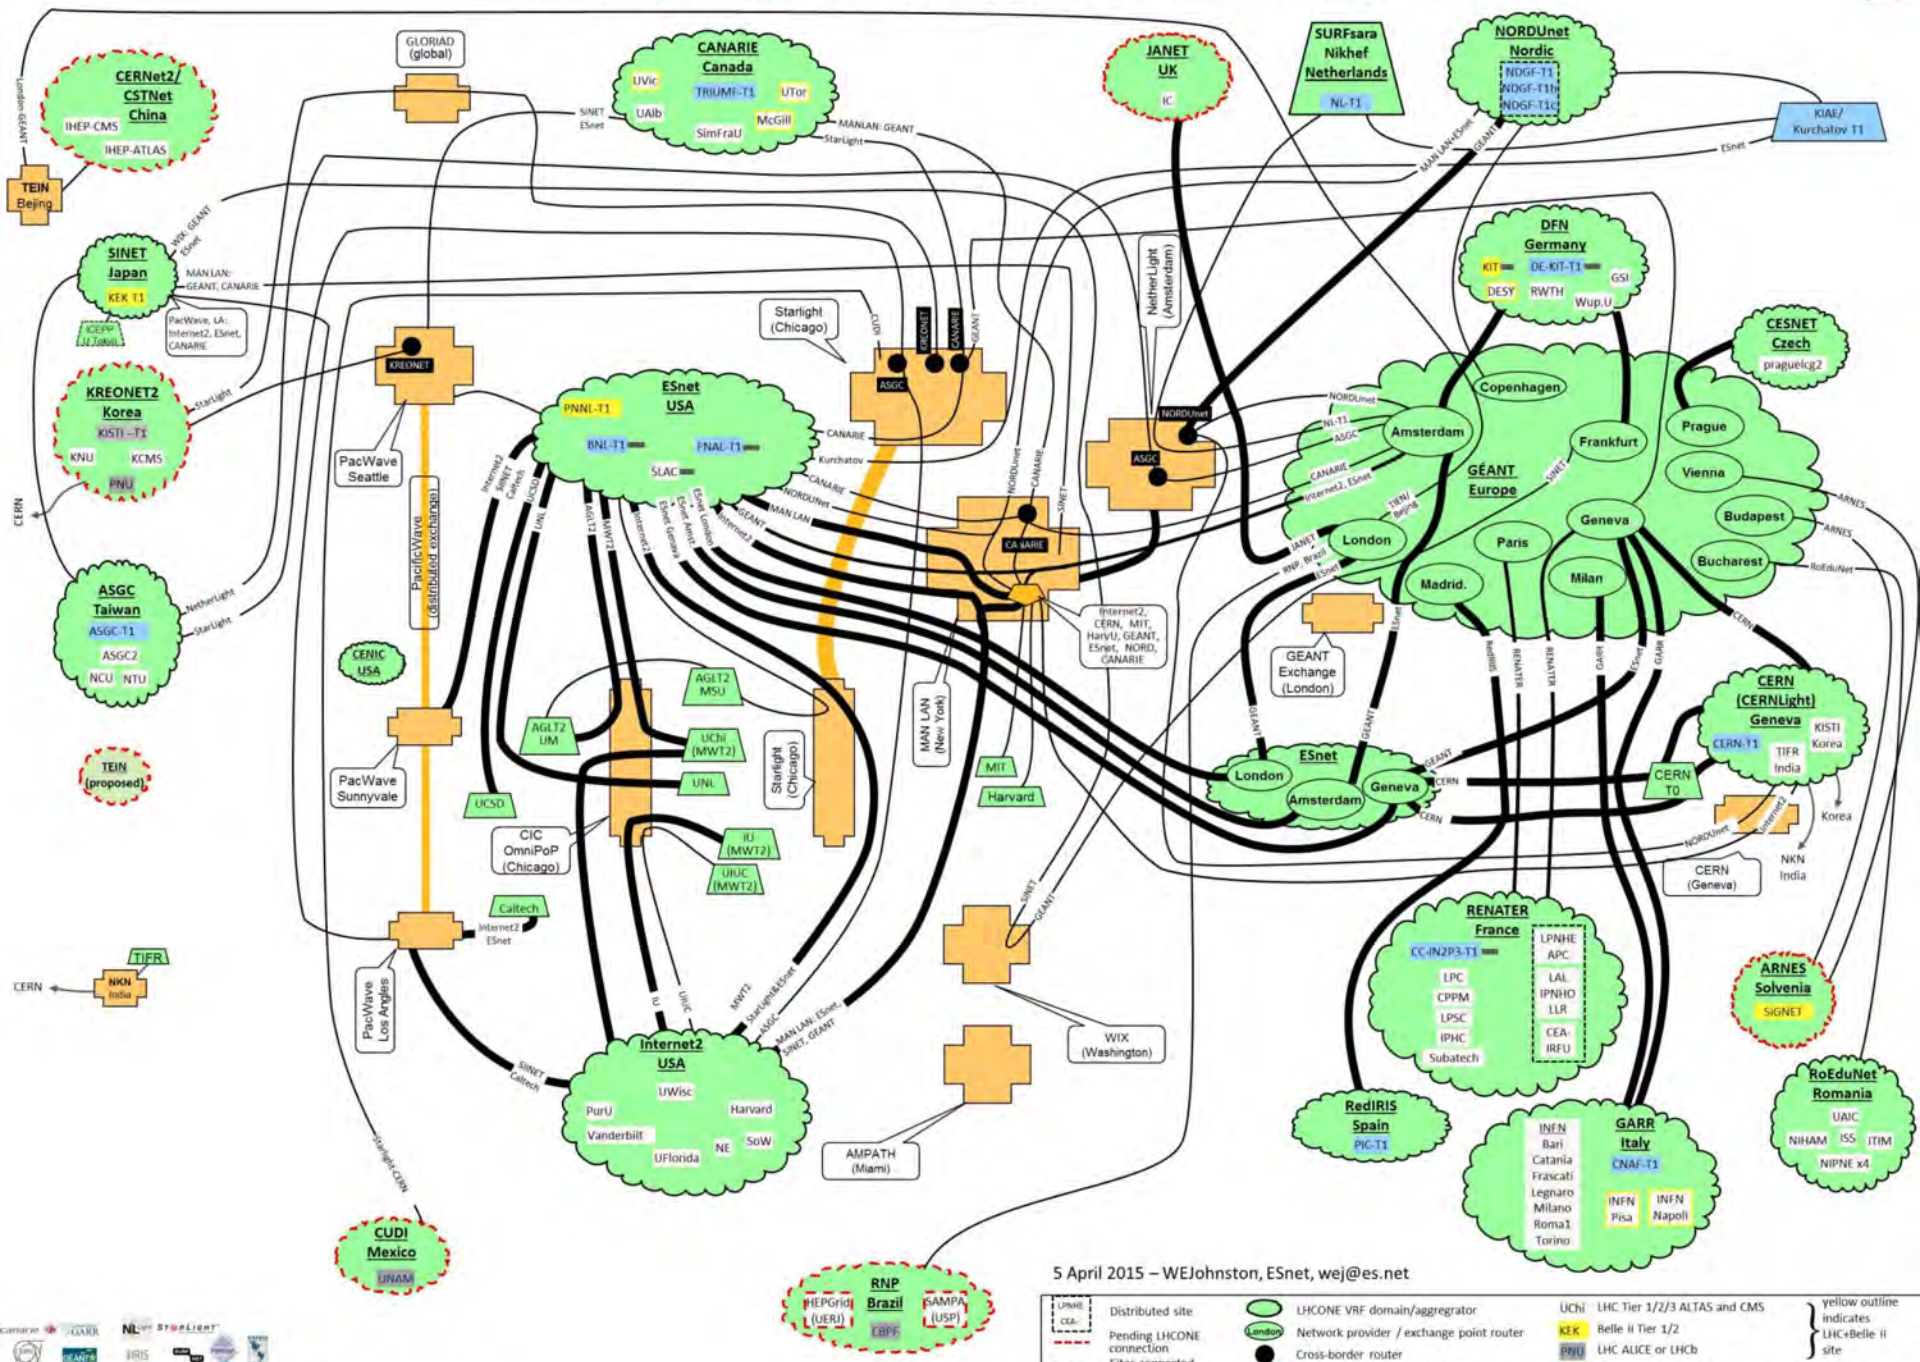

PetaTrans: Petascale Sciences Data Transfer

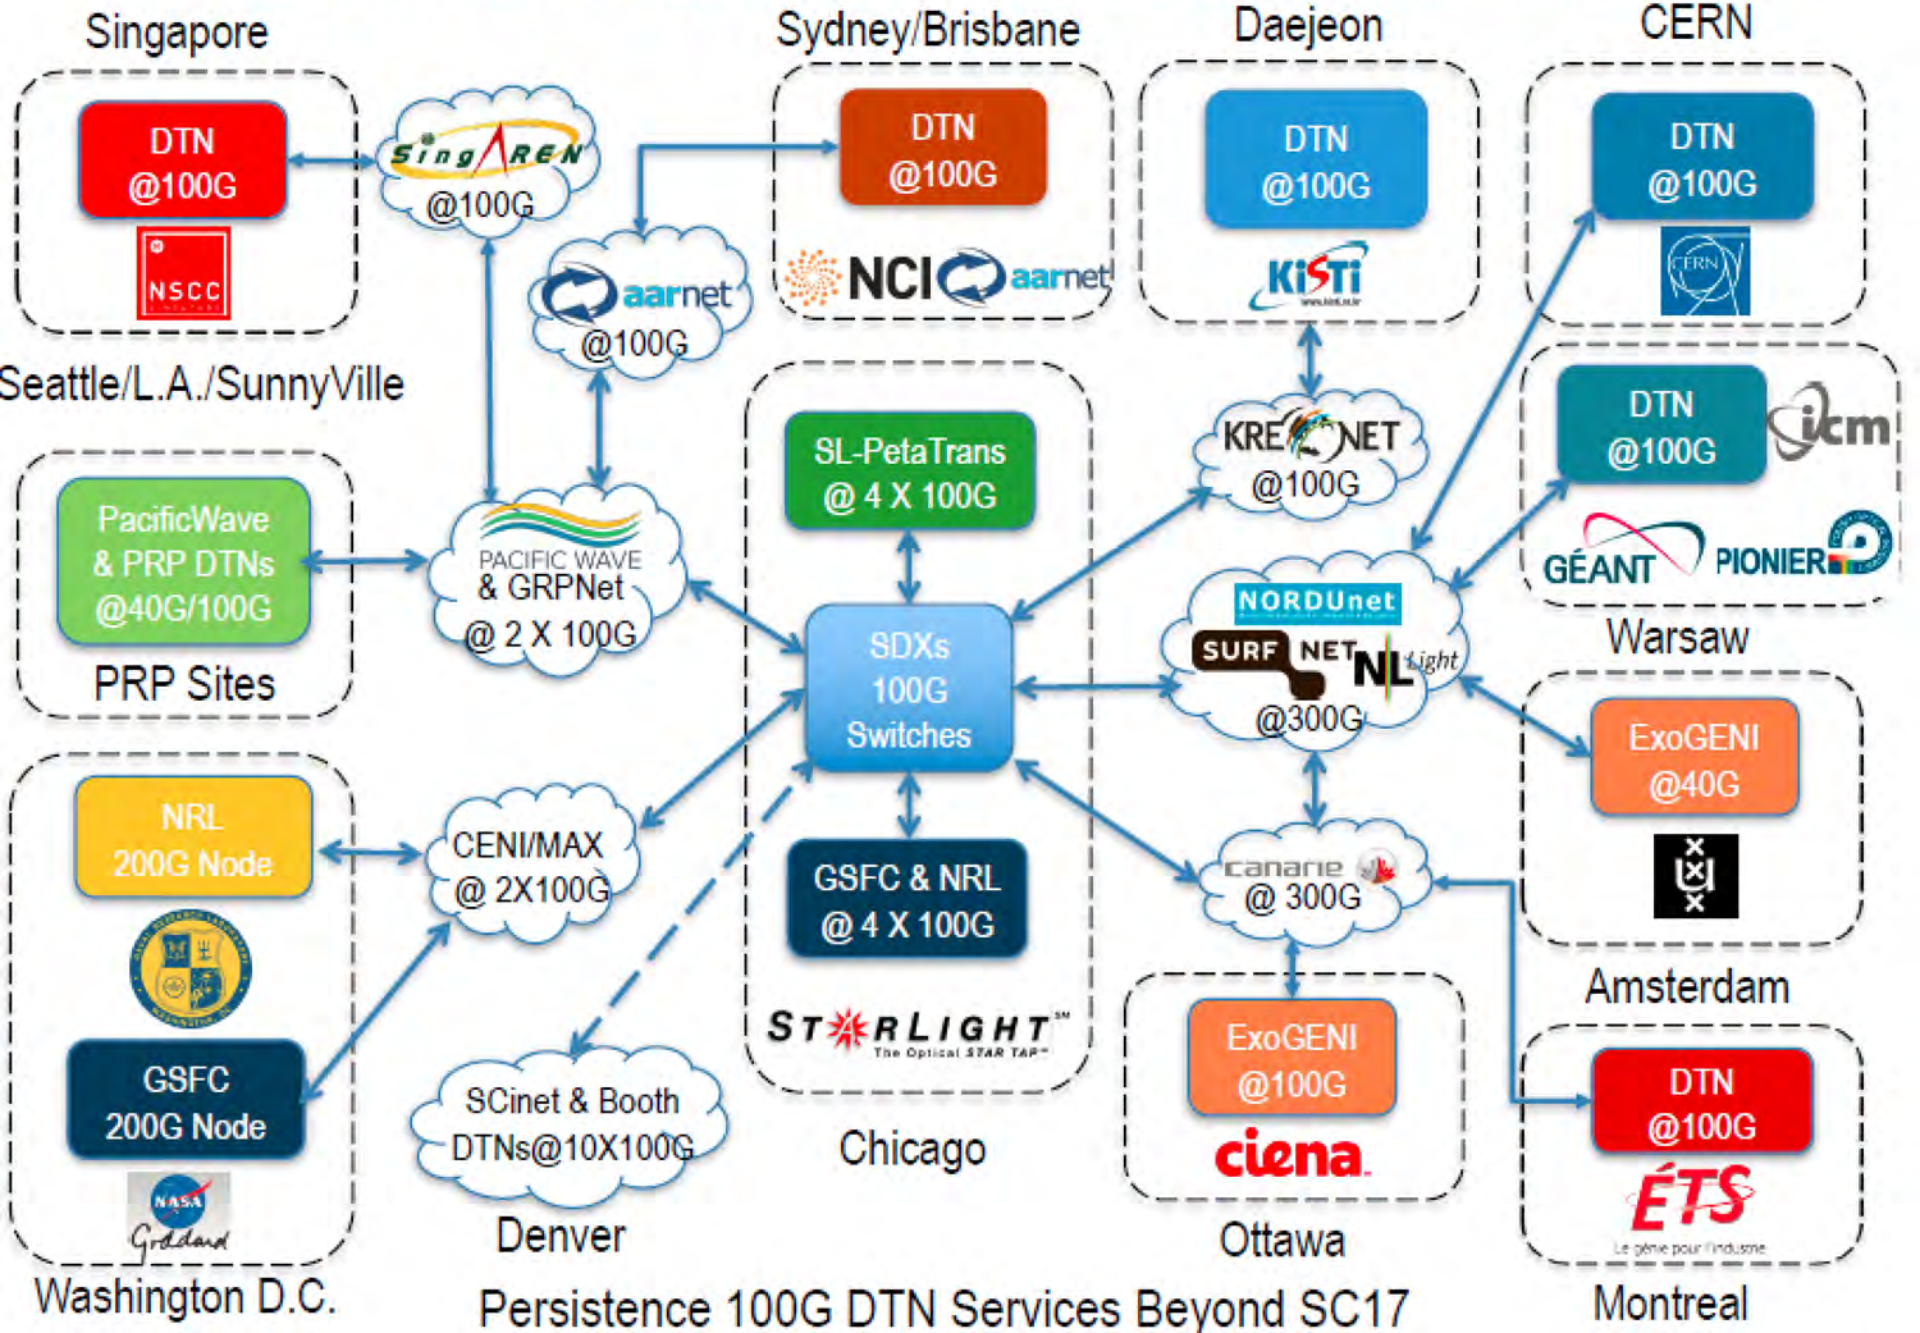

Emerging Topics In Advanced Networking

- 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

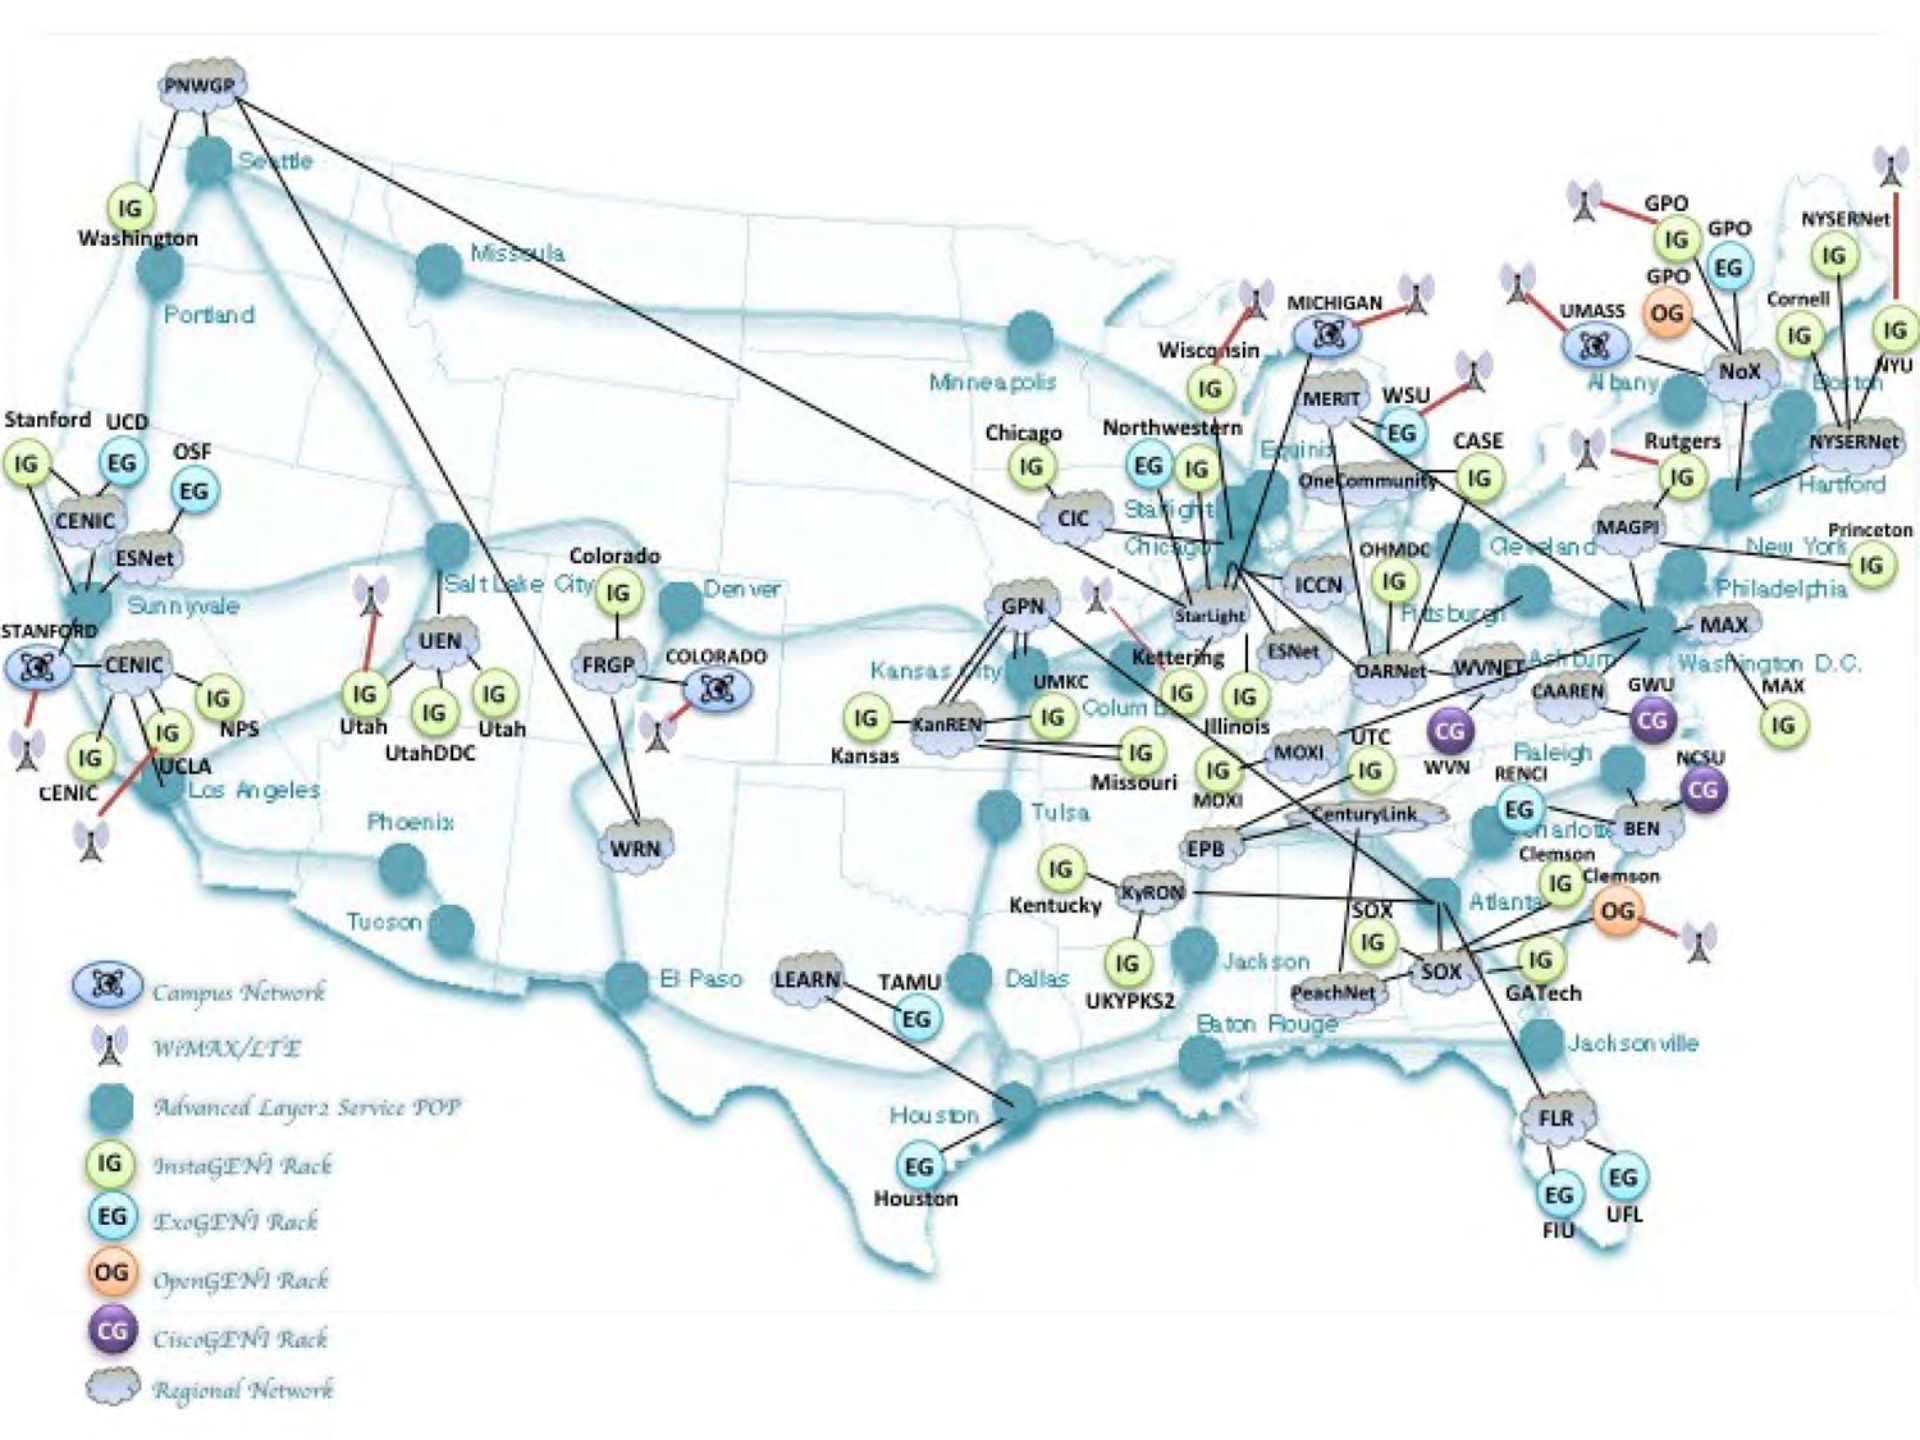
PNWGP
Seattle
IG
Washington
Portland
Missoula
Minneapolis
Stanford
UCD
OSF
IG
EG
EG
CENIC
ESNet
Sunnyvale
STANFORD
CENIC
IG
NPS
IG
UCLA
CENIC
Los Angeles
Salt Lake City
UEN
IG
Utah
IG
UtahDDC
IG
Utah
Colorado
IG
FRGP
COLORADO
Denver
WRN
Phoenix
Tucson
El Paso
LEARN
TAMU
EG
Dallas
Houston
EG
Houston
Chicago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

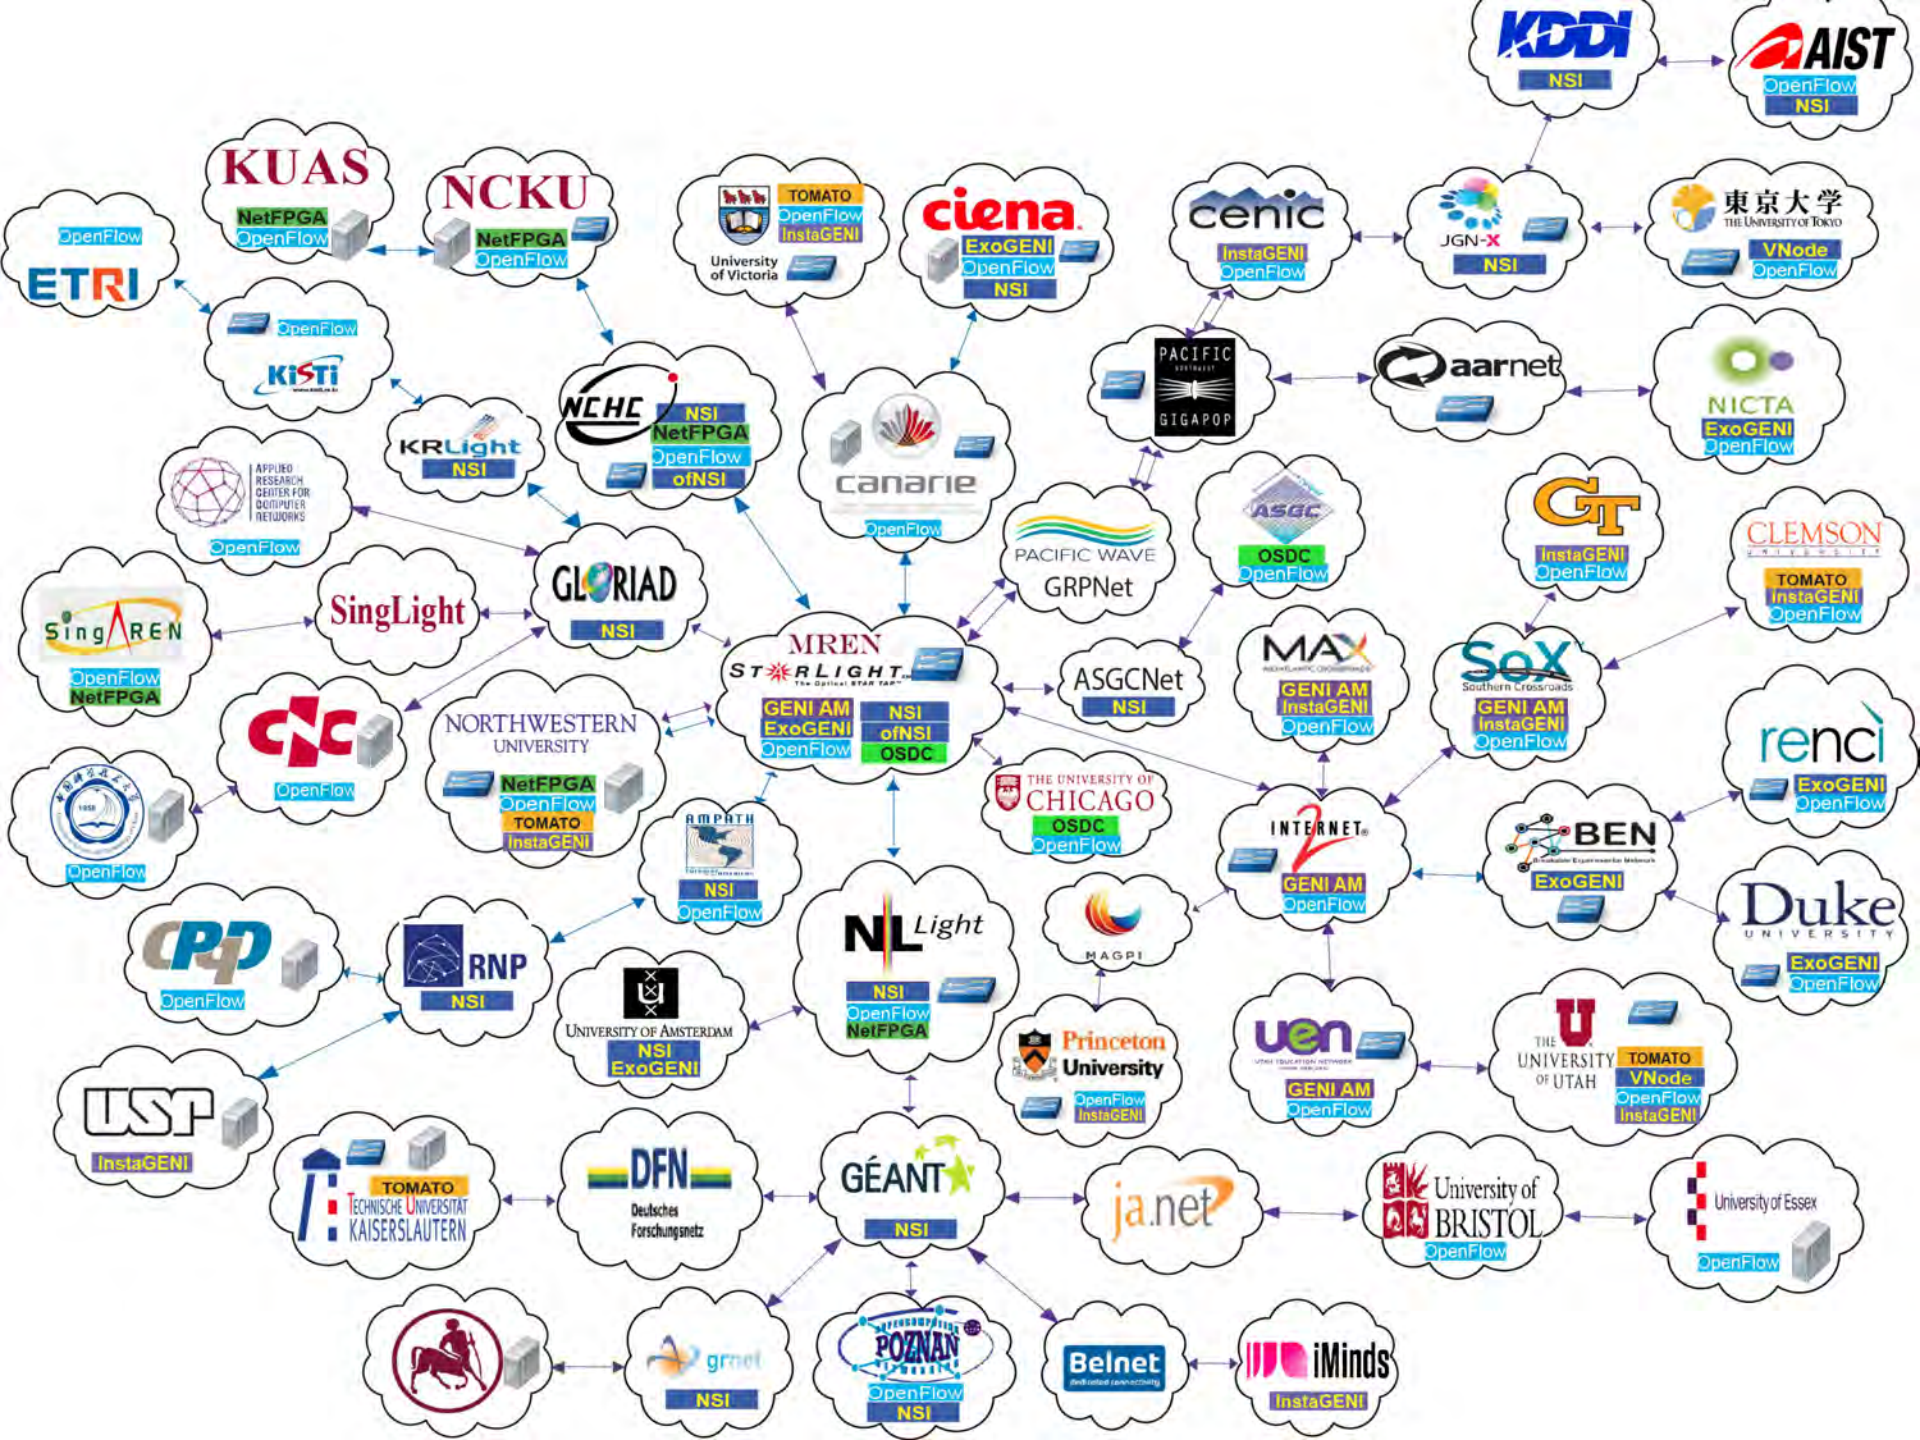
KDDI
NSI
AIST
OpenFlow
NSI
KUAS
NetFPGA
OpenFlow
NCKU
NetFPGA
OpenFlow
TOMATO
OpenFlow
InstaGENI
University of Victoria
ciena
ExoGENI
OpenFlow
NSI
cenic
InstaGENI
OpenFlow
JGN-X
NSI
東京大学
THE UNIVERSITY OF TOKYO
VNode
OpenFlow
OpenFlow
ETRI
OpenFlow
KISTI
NCHC
NSI
NetFPGA
OpenFlow
ofNSI
PACIFIC GIGAPOP
aarnet
NICTA
ExoGENI
OpenFlow
KRLight
NSI
canarie
OpenFlow
APPLIED RESEARCH CENTER FOR COMPUTER NETWORKS
OpenFlow
GLORIAD
NSI
PACIFIC WAVE
GRPNet
ASGC
OSDC
OpenFlow
GT
InstaGENI
OpenFlow
CLEMSON
TOMATO
InstaGENI
OpenFlow
SingAREN
OpenFlow
NetFPGA
SingLight
MREN
STARLIGHT
GENI AM
ExoGENI
OpenFlow
NSI
ofNSI
OSDC
ASGCNet
NSI
MAX
GENI AM
InstaGENI
OpenFlow
SoX
Southern Crossroads
GENI AM
InstaGENI
OpenFlow
CNC
OpenFlow
NORTHWESTERN UNIVERSITY
NetFPGA
OpenFlow
TOMATO
InstaGENI
THE UNIVERSITY OF CHICAGO
OSDC
OpenFlow
renci
ExoGENI
OpenFlow
OpenFlow
AMPATH
NSI
OpenFlow
INTERNET2
GENI AM
OpenFlow
BEN
ExoGENI
CPqD
OpenFlow
RNP
NSI
NL Light
NSI
OpenFlow
NetFPGA
MAGPI
Duke
UNIVERSITY
ExoGENI
OpenFlow
UNIVERSITY OF AMSTERDAM
NSI
ExoGENI
Princeton University
OpenFlow
InstaGENI
uen
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
iMinds
InstaGENI

Chameleon

www.chameleoncloud.org

CHAMELEON: A LARGE SCALE, RECONFIGURABLE EXPERIMENTAL INSTRUMENT FOR COMPUTER SCIENCE

Kate Keahey

Joe Mambretti, Pierre Riteau, Paul Ruth, Dan Stanzione

2016 Bioinformatics SDXs Network

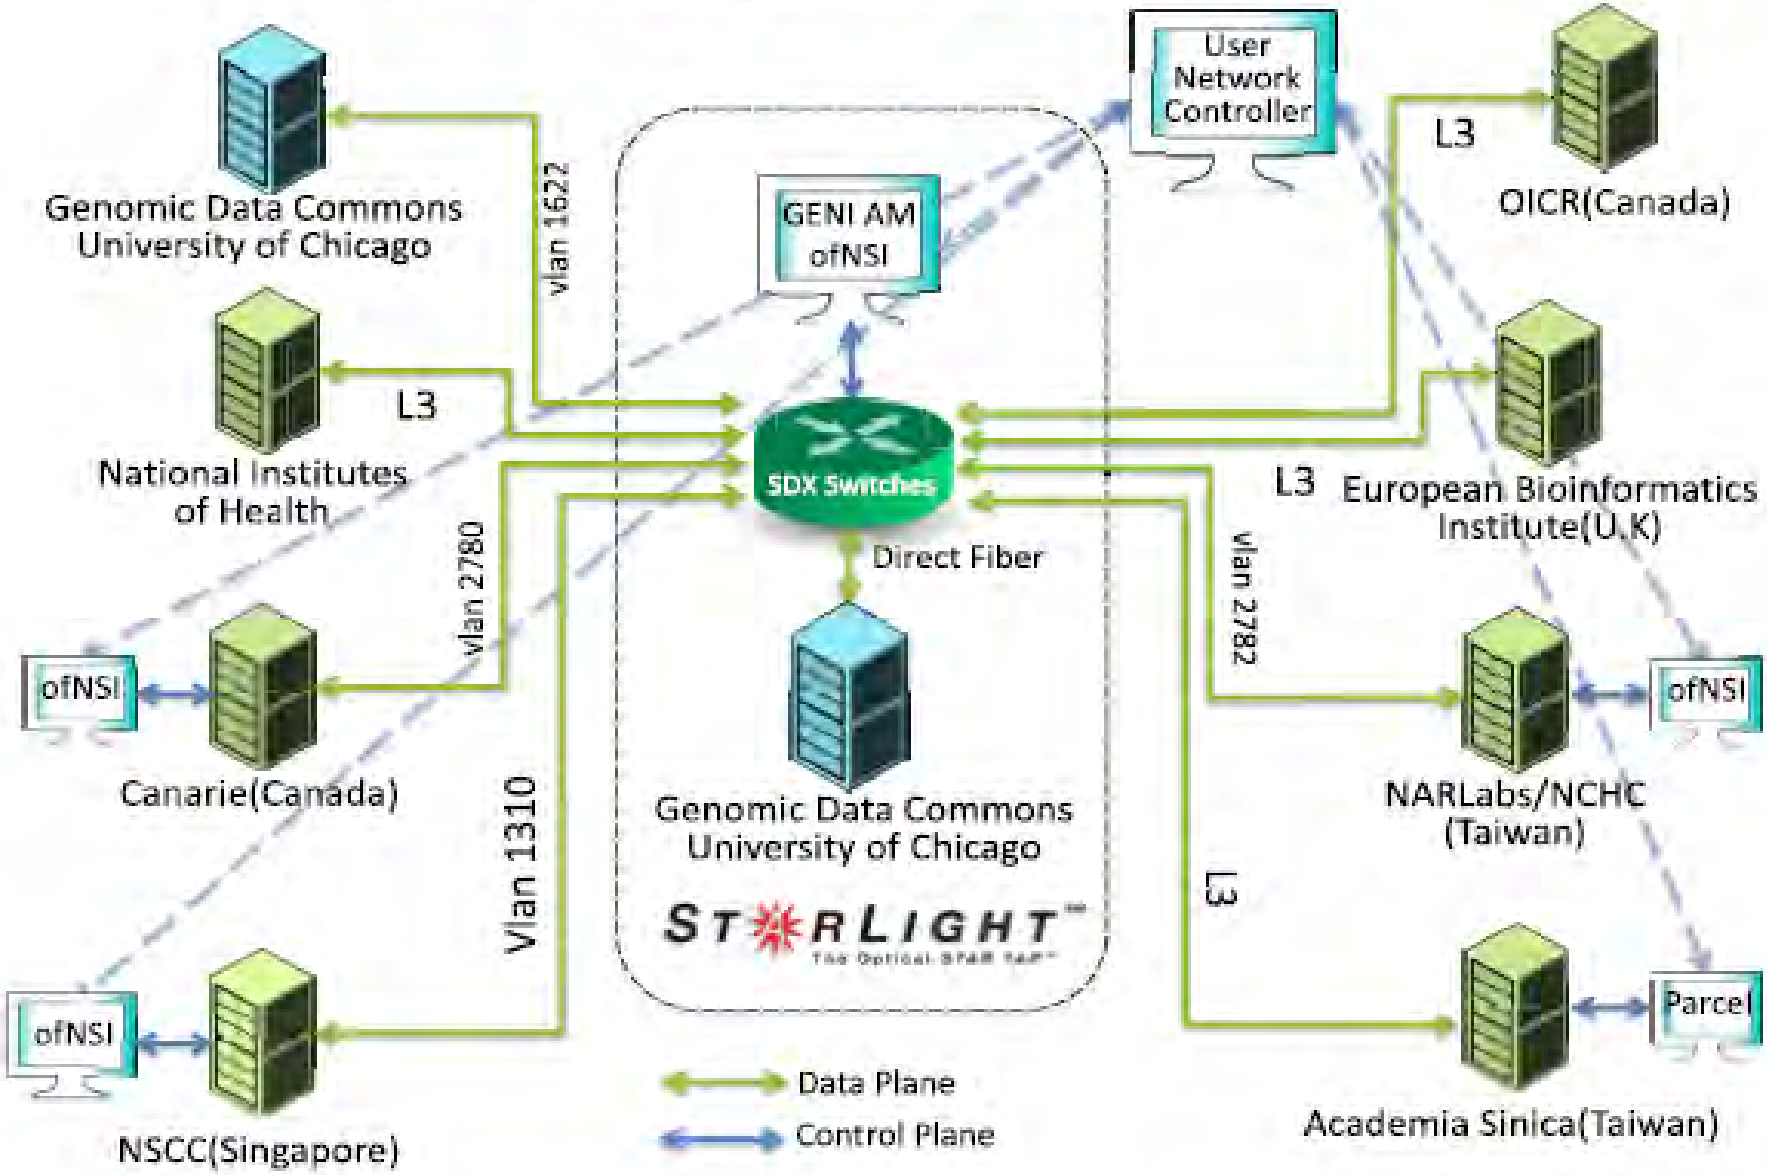

A Disaggregated SCinet Optical Layer

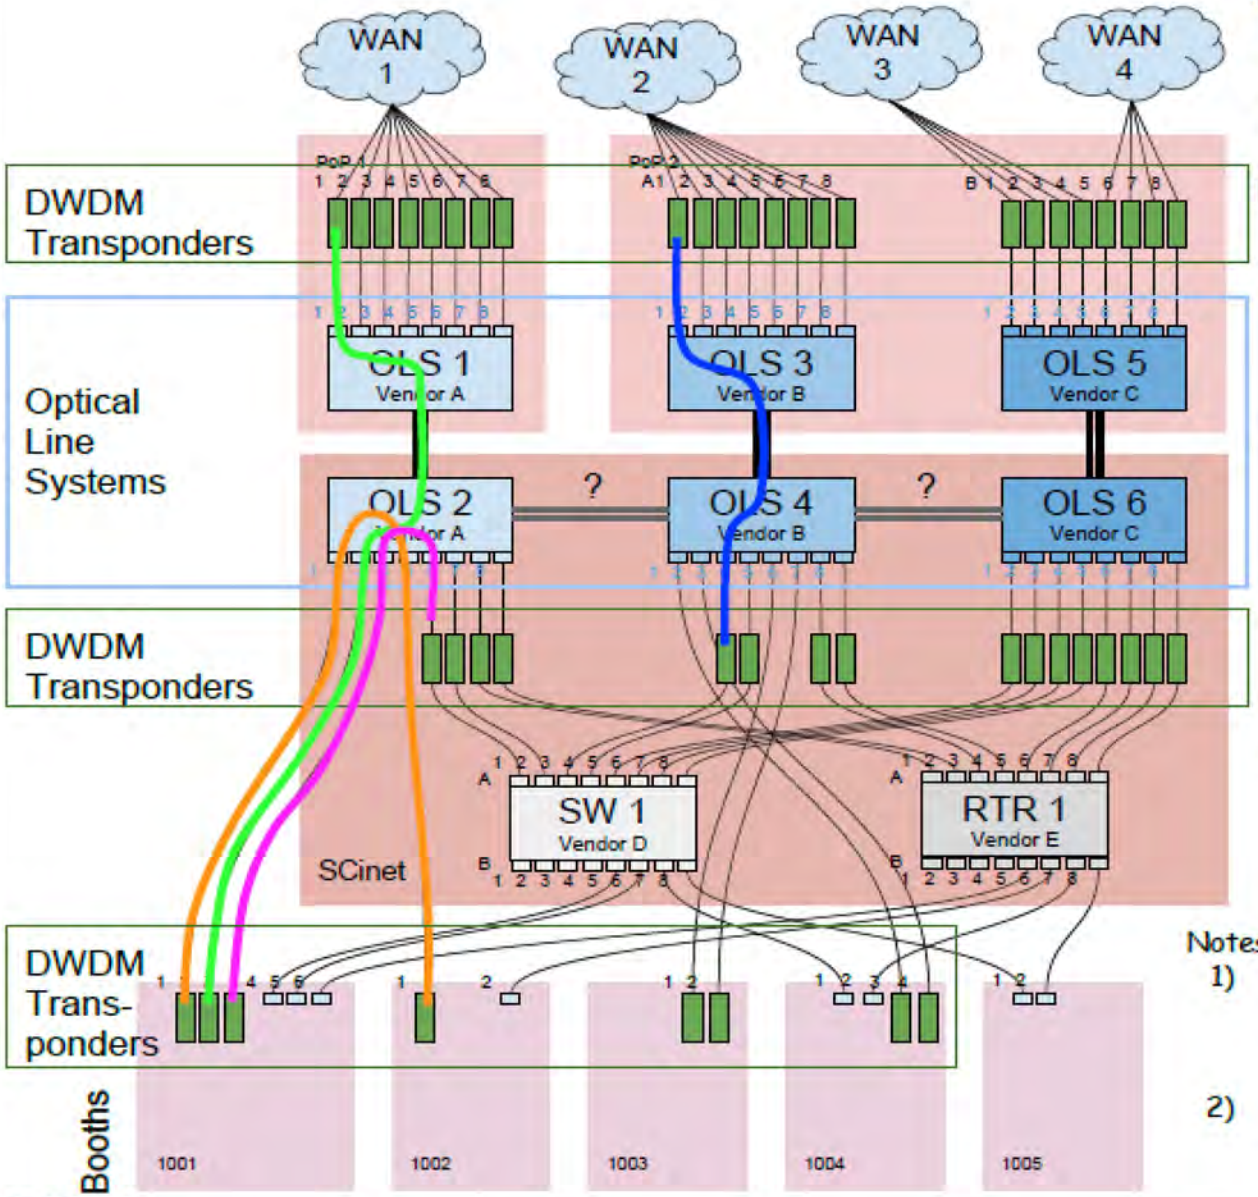

Reconfiguration options

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 SDX Geoscience Research Workflow

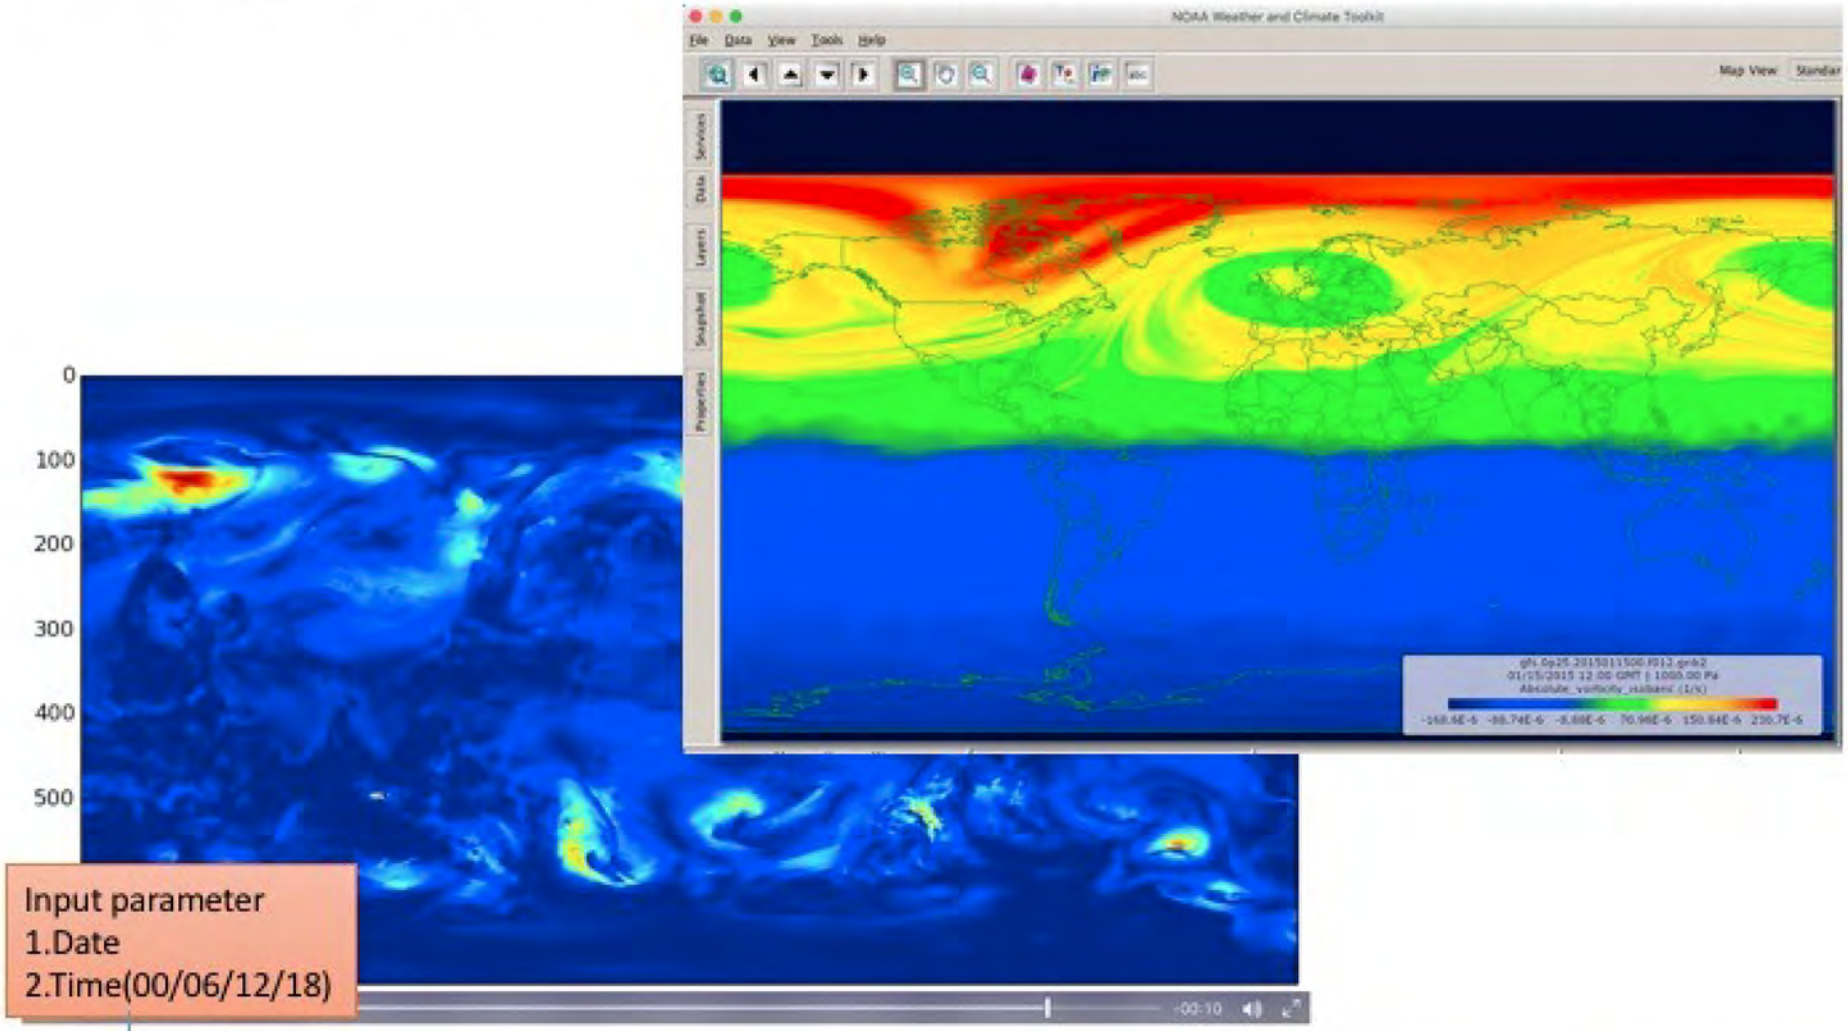

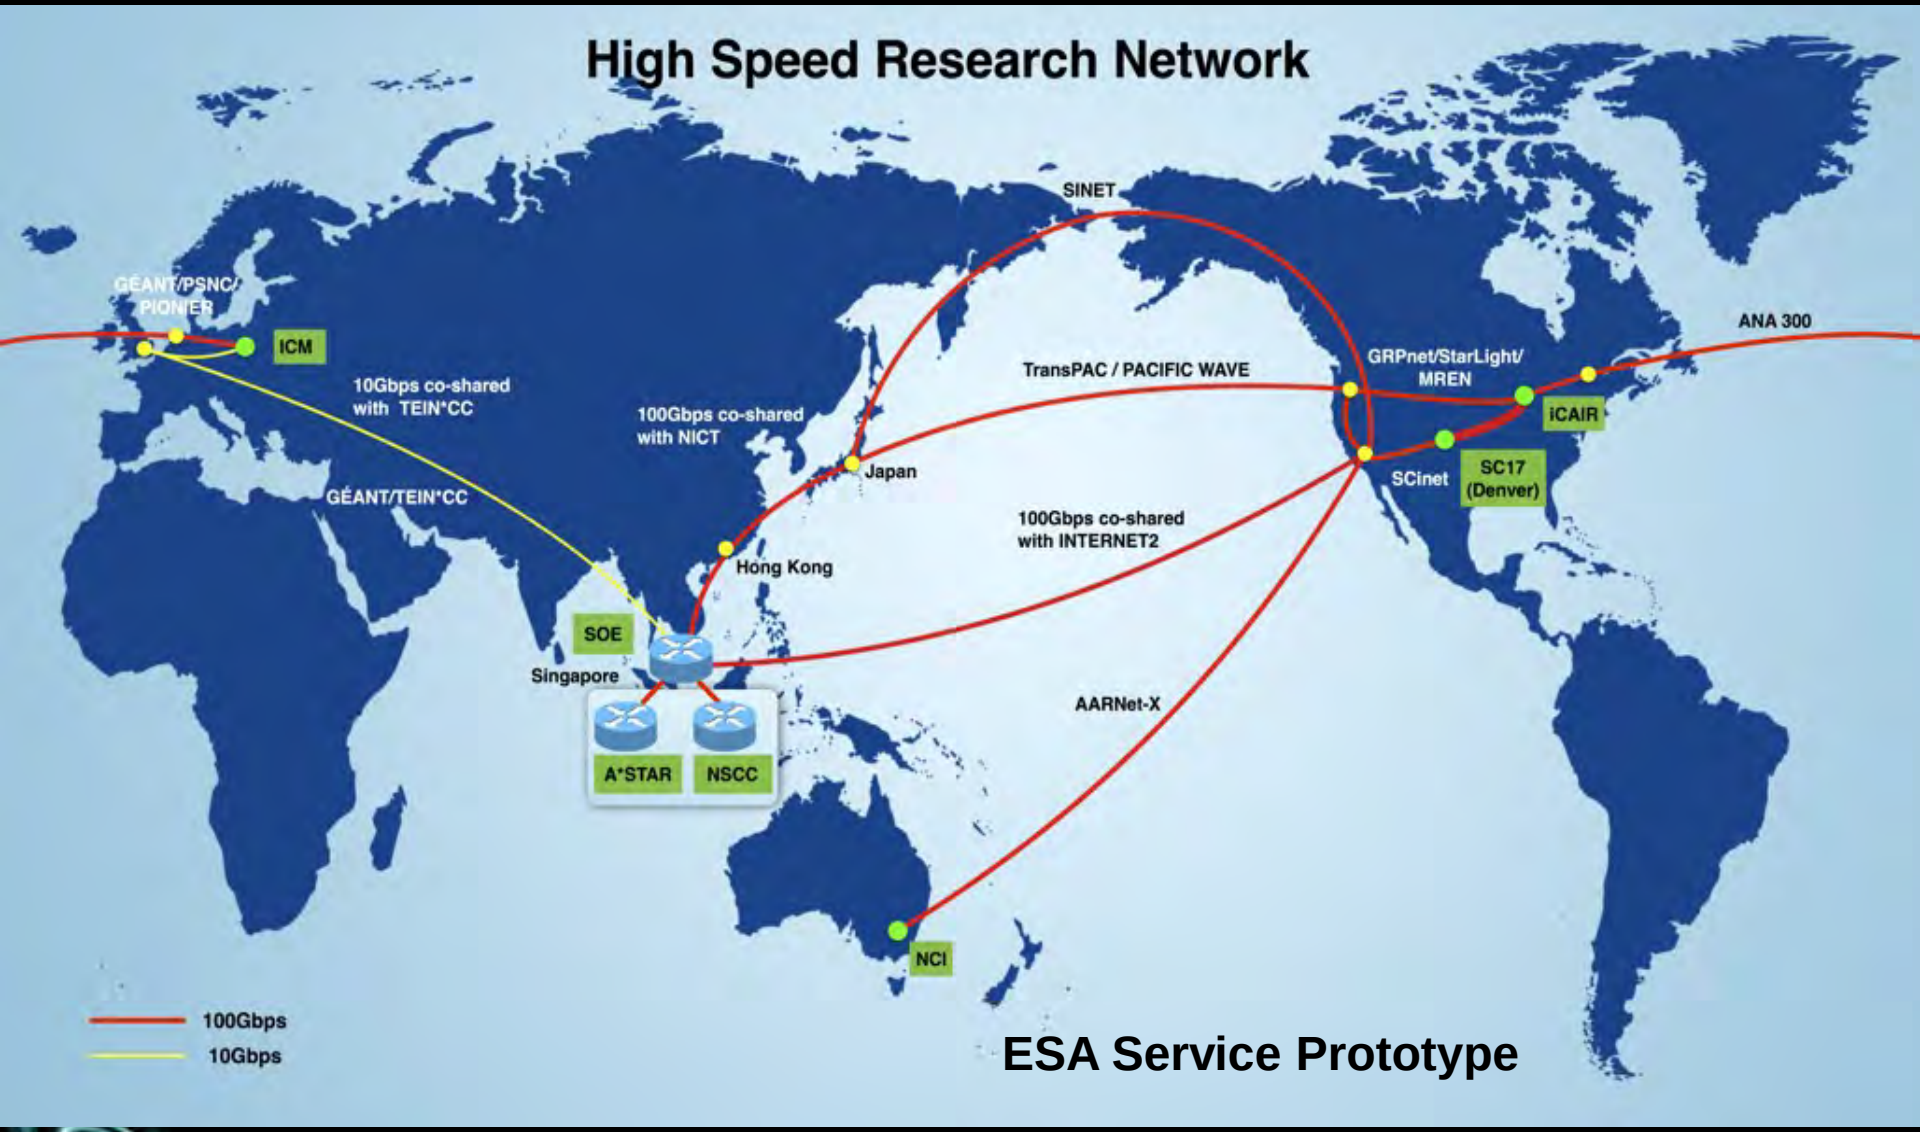
High Speed Research Network
SINET
GÉANT/PSNC/
PIONIER
ICM
10Gbps co-shared
with TEIN*CC
GÉANT/TEIN*CC
100Gbps co-shared
with NICT
Japan
TransPAC / PACIFIC WAVE
GRPnet/StarLight/
MREN
ANA 300
iCAIR
SCinet
SC17
(Denver)
100Gbps co-shared
with INTERNET2
Hong Kong
SOE
Singapore
A*STAR
NSCC
AARNet-X
NCI
100Gbps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

Source Corby Schmidt

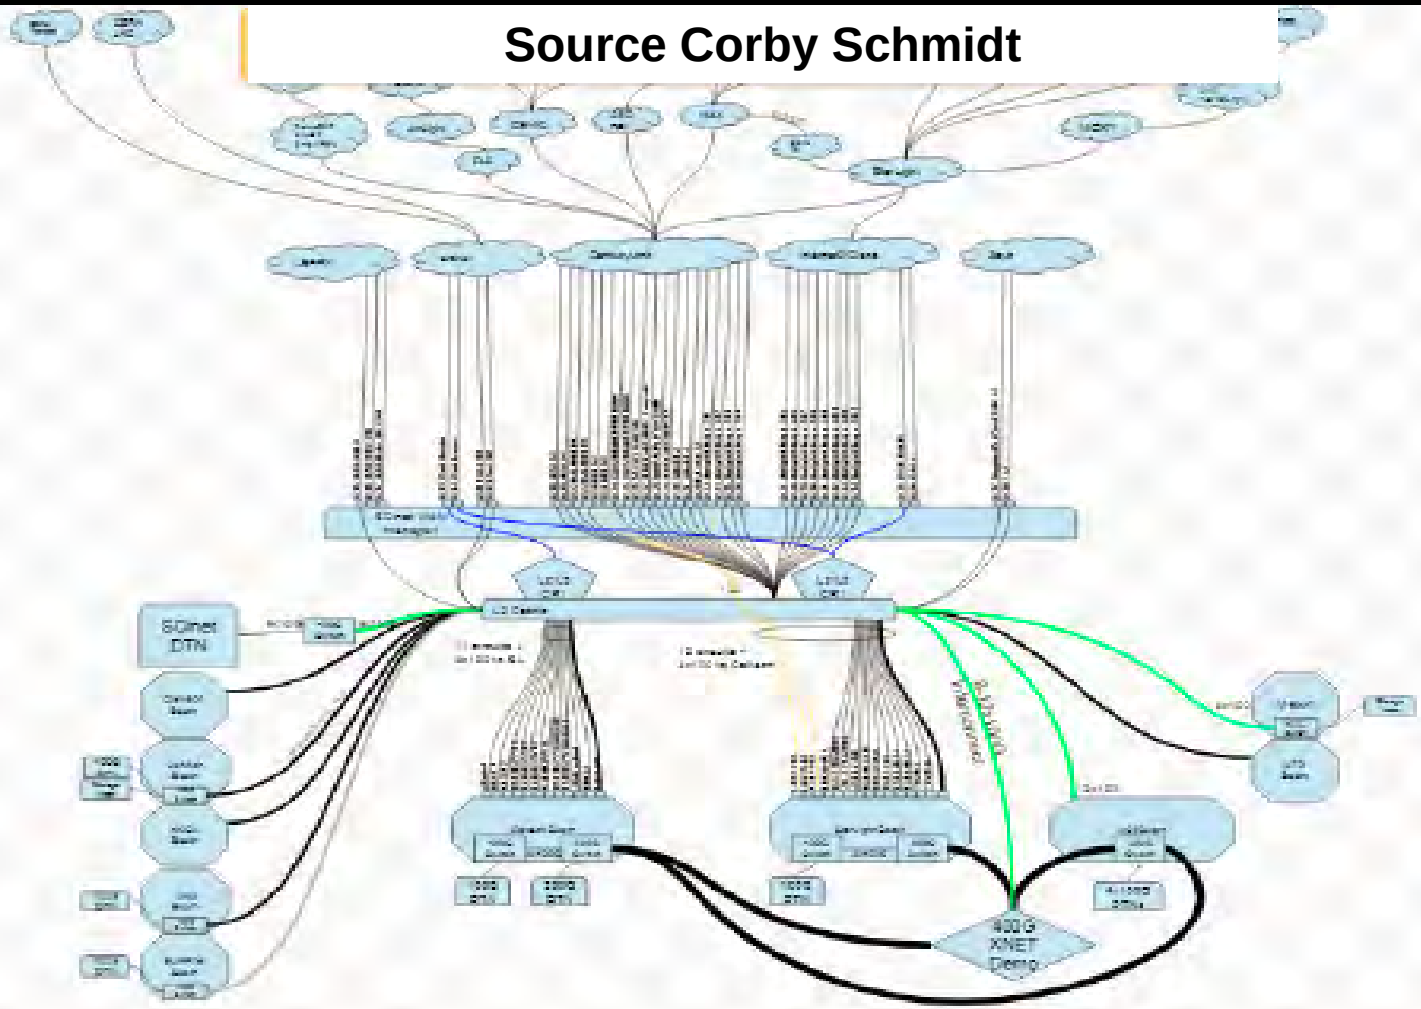

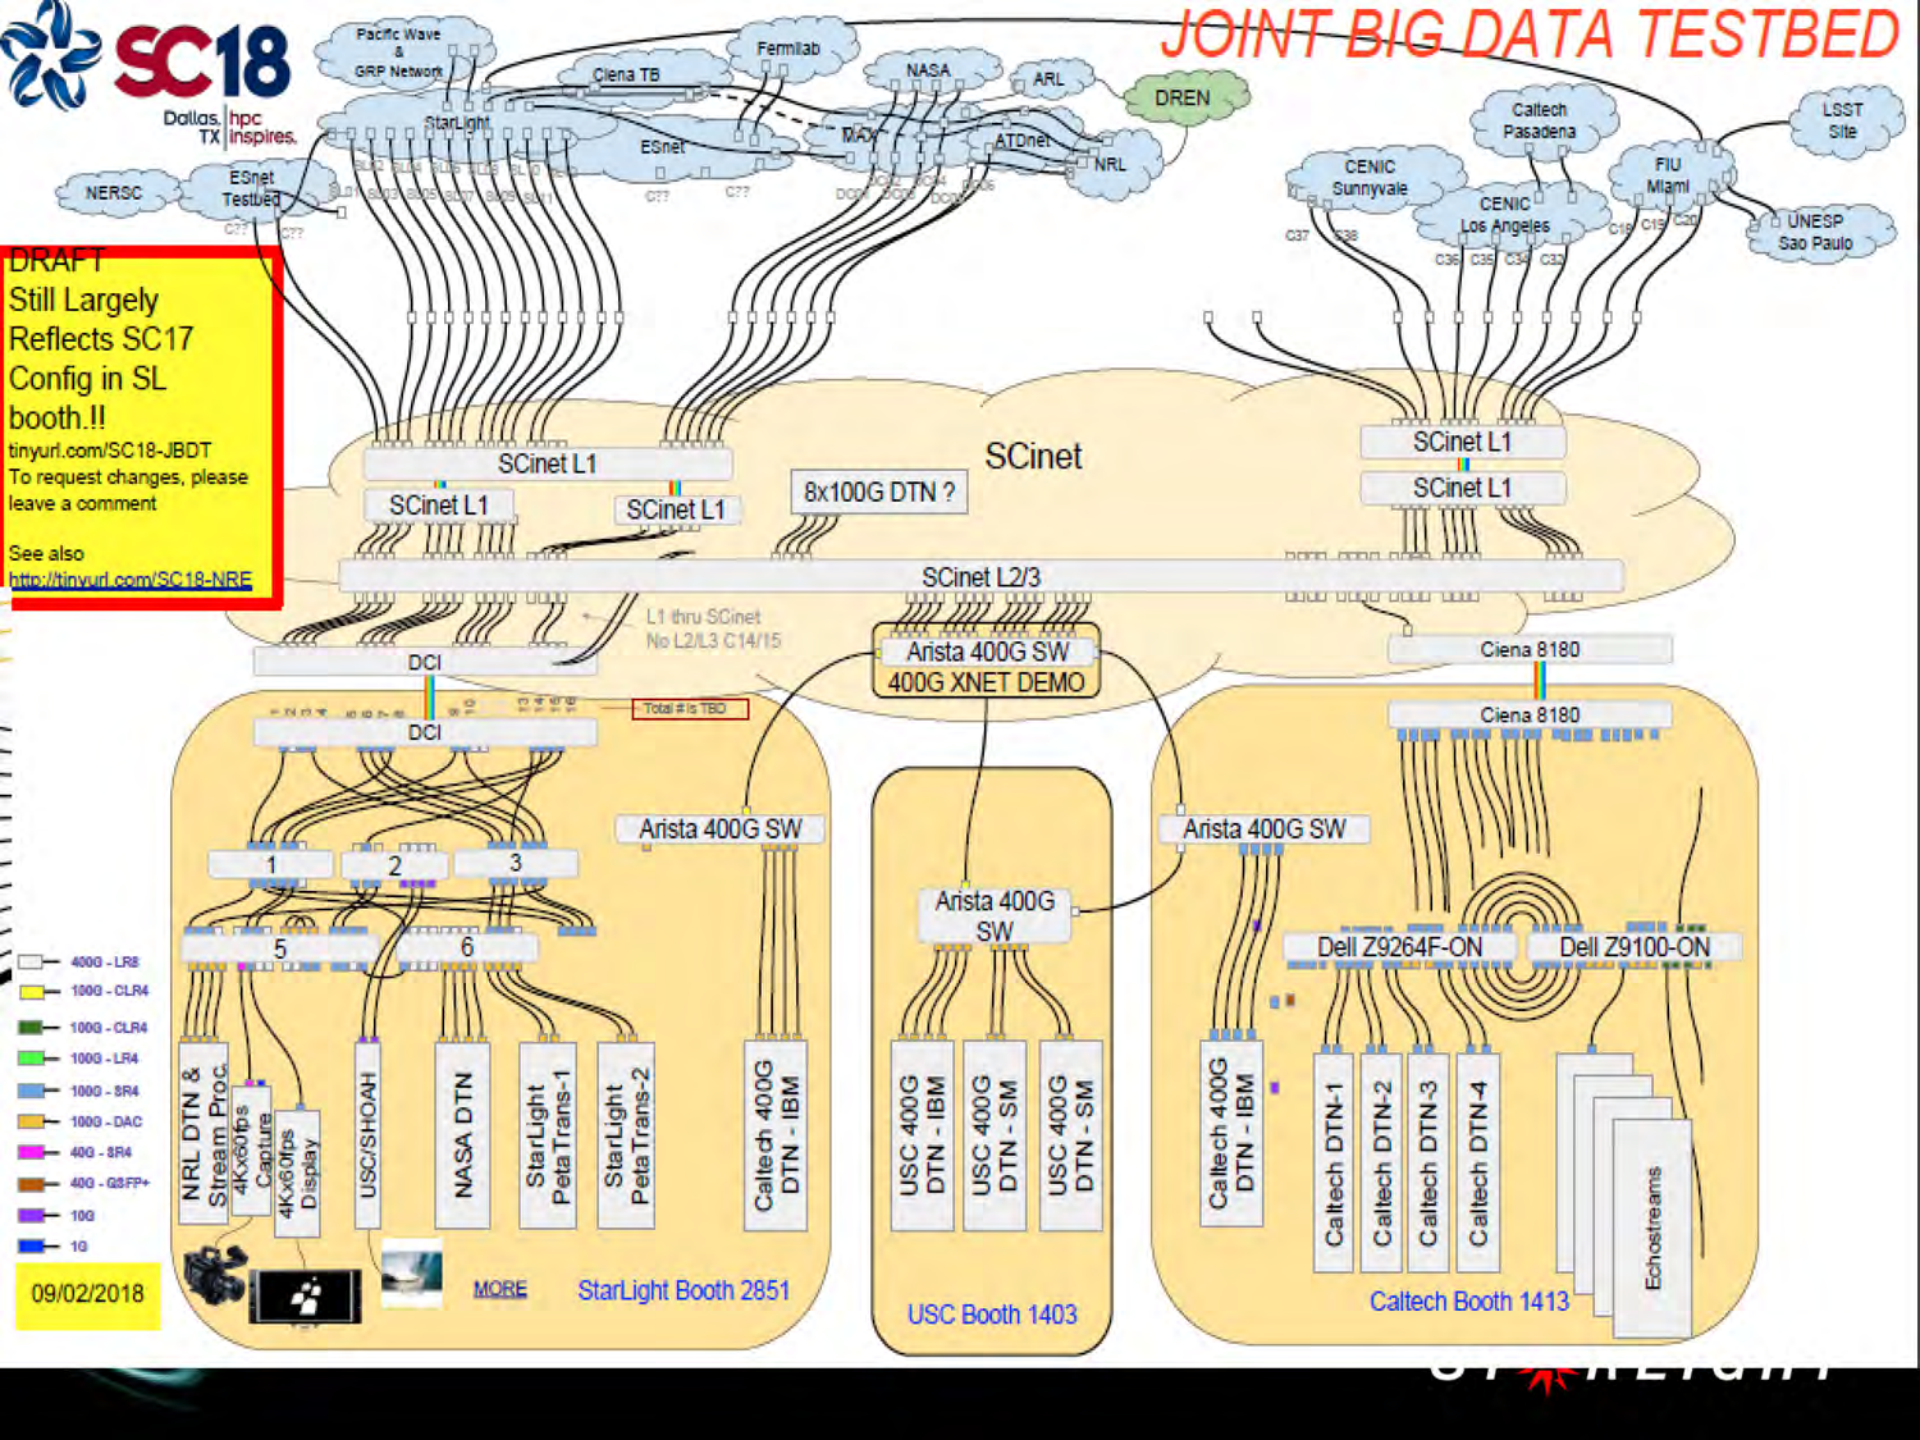
SC18
Dallas, TX | hpc inspires.
JOINT BIG DATA TESTBED
DRAFT
Still Largely Reflects SC17 Config in SL booth.!!
tinyurl.com/SC18-JBDT
To request changes, please leave a comment
See also http://tinyurl.com/SC18-NRE
Pacific Wave & GRP Network
Fermilab
Ciena TB
NASA
ARL
DREN
StarLight
ESnet
MAX
ATDnet
NRL
NERSC
ESnet Testbed
Caltech Pasadena
LSST Site
CENIC Sunnyvale
FIU Miami
CENIC Los Angeles
UNESP Sao Paulo
SCinet
SCinet L1
SCinet L1
SCinet L1
8x100G DTN ?
SCinet L1
SCinet L1
SCinet L2/3
L1 thru SCinet
No L2/L3 C14/15
DCI
DCI
Total # is TBD
Arista 400G SW
400G XNET DEMO
Ciena 8180
Ciena 8180
Arista 400G SW
Arista 400G SW
Arista 400G SW
1
2
3
5
6
Dell Z9264F-ON
Dell Z9100-ON
NRL DTN & Stream Proc.
4Kx60fps Capture
4Kx60fps Display
USC/SHOAH
NASA DTN
StarLight PetaTrans-1
StarLight PetaTrans-2
Caltech 400G DTN - IBM
USC 400G DTN - IBM
USC 400G DTN - SM
USC 400G DTN - SM
Caltech 400G DTN - IBM
Caltech DTN-1
Caltech DTN-2
Caltech DTN-3
Caltech DTN-4
Echostreams
MORE
StarLight Booth 2851
USC Booth 1403
Caltech Booth 1413
400G - LR8
100G - CLR4
100G - CLR4
100G - LR4
100G - SR4
100G - DAC
40G - SR4
40G - QSFP+
10G
1G
09/02/2018

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

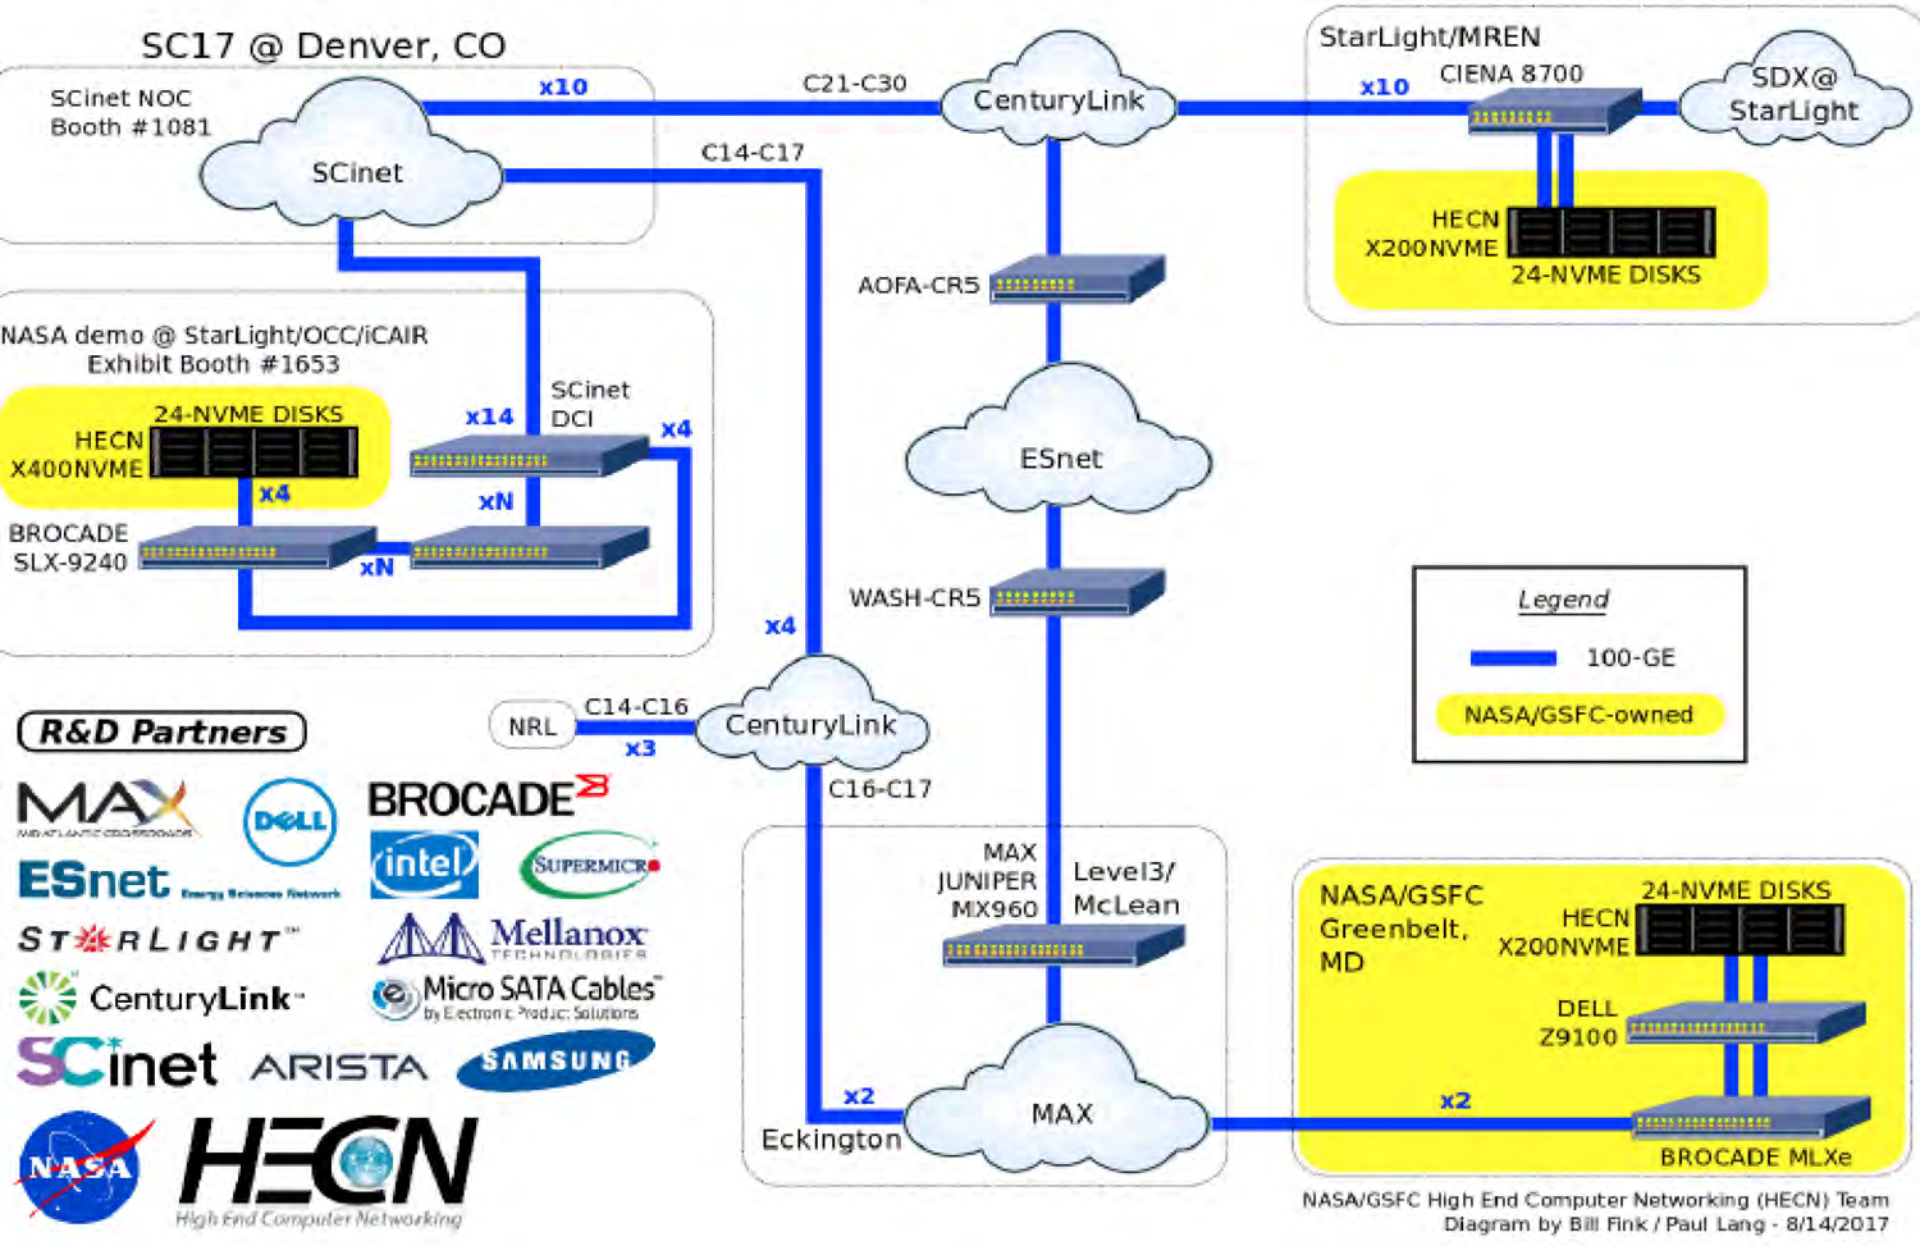

NASA/GSFC High End Computer Networking (HECN) Team
Diagram by Bill Fink / Paul Lang - 8/14/2017

Dynamic Distributed Data Processing

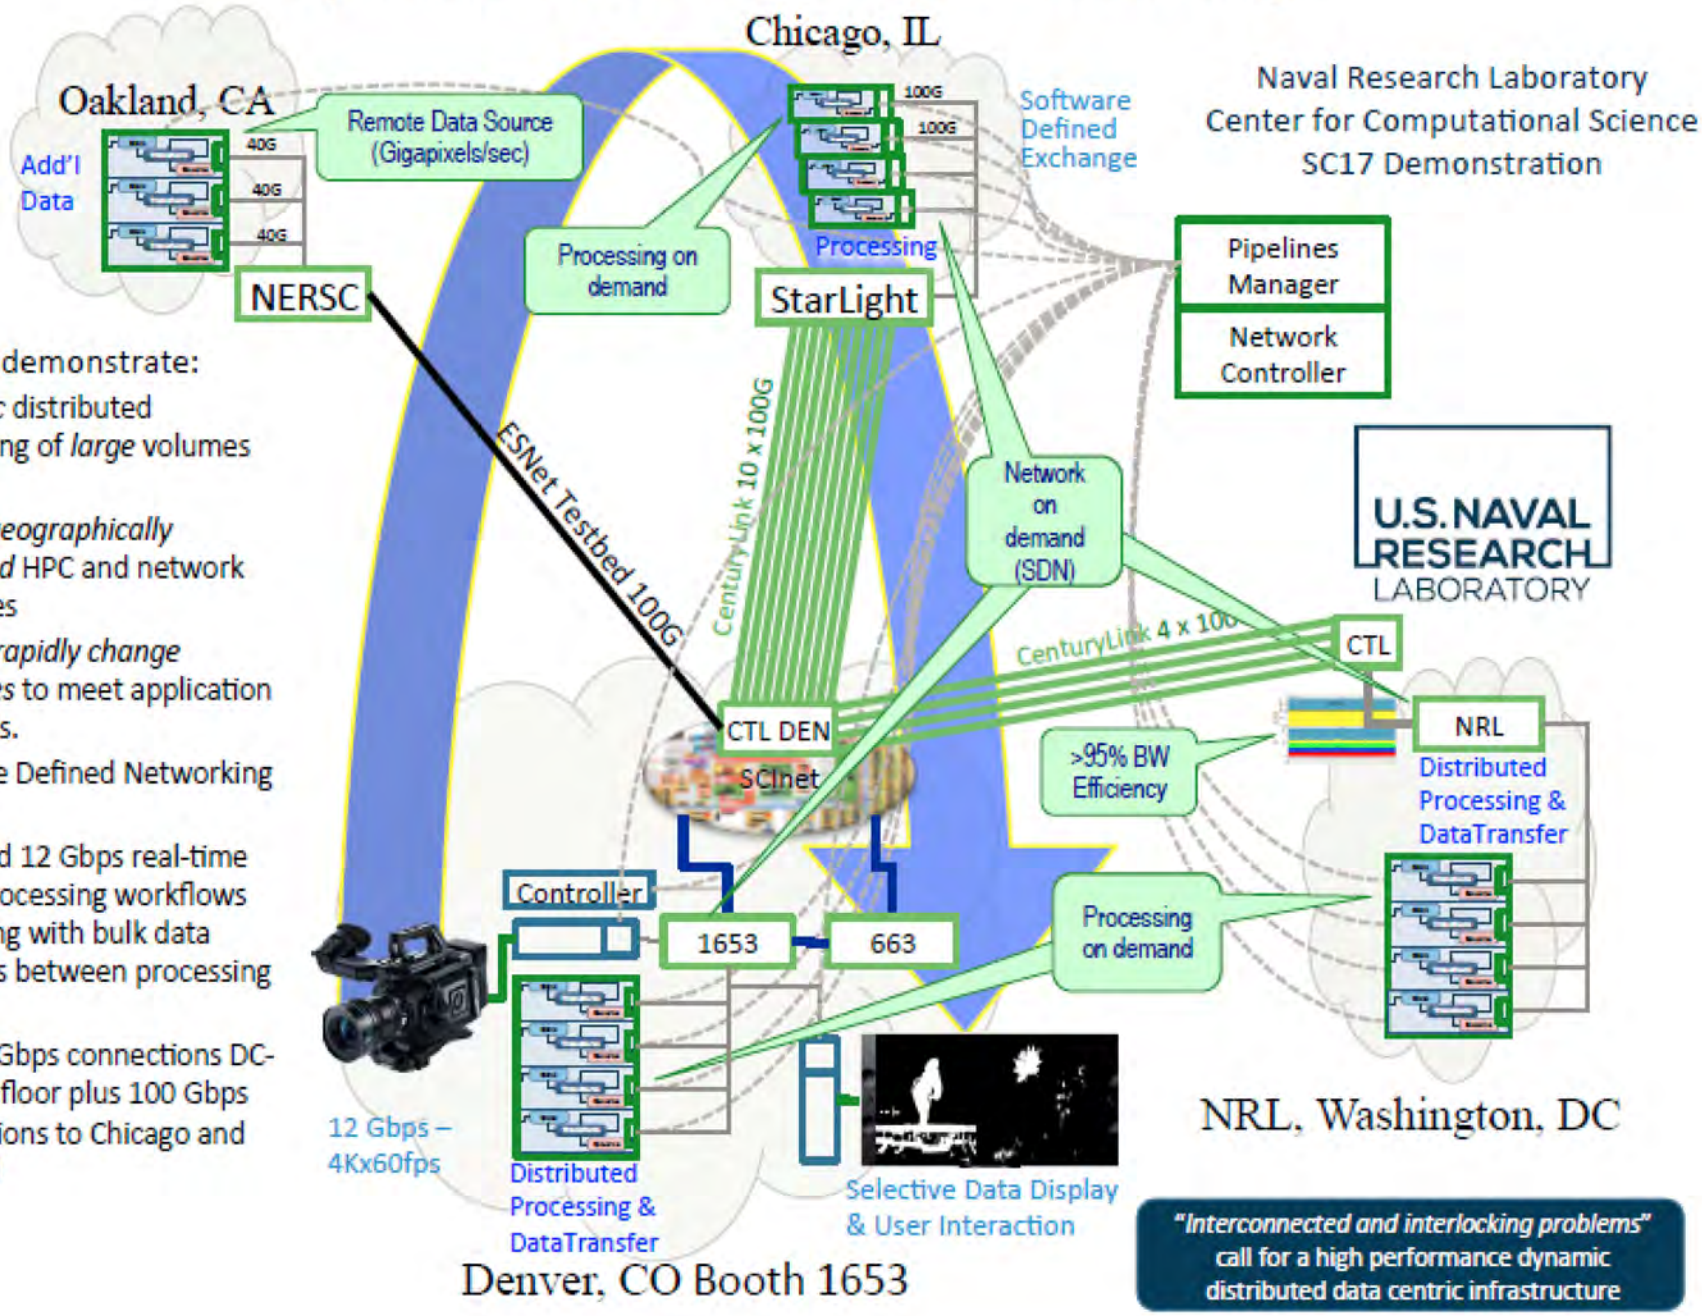

NRL will demonstrate:

- *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

A Cross-Pacific SDN Testbed

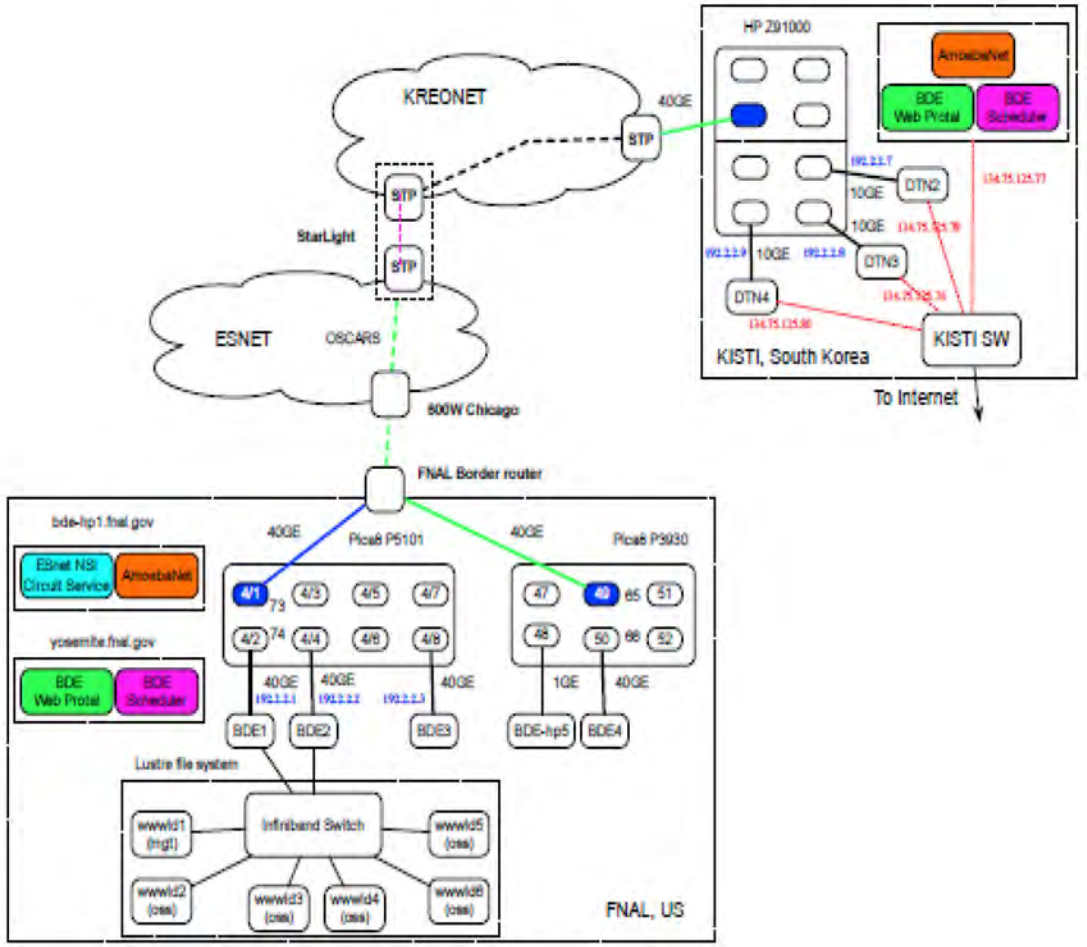

Transferring LargeScale Airline Data E2E Across WANs Using DTNs

v5, 21 SEP 2017

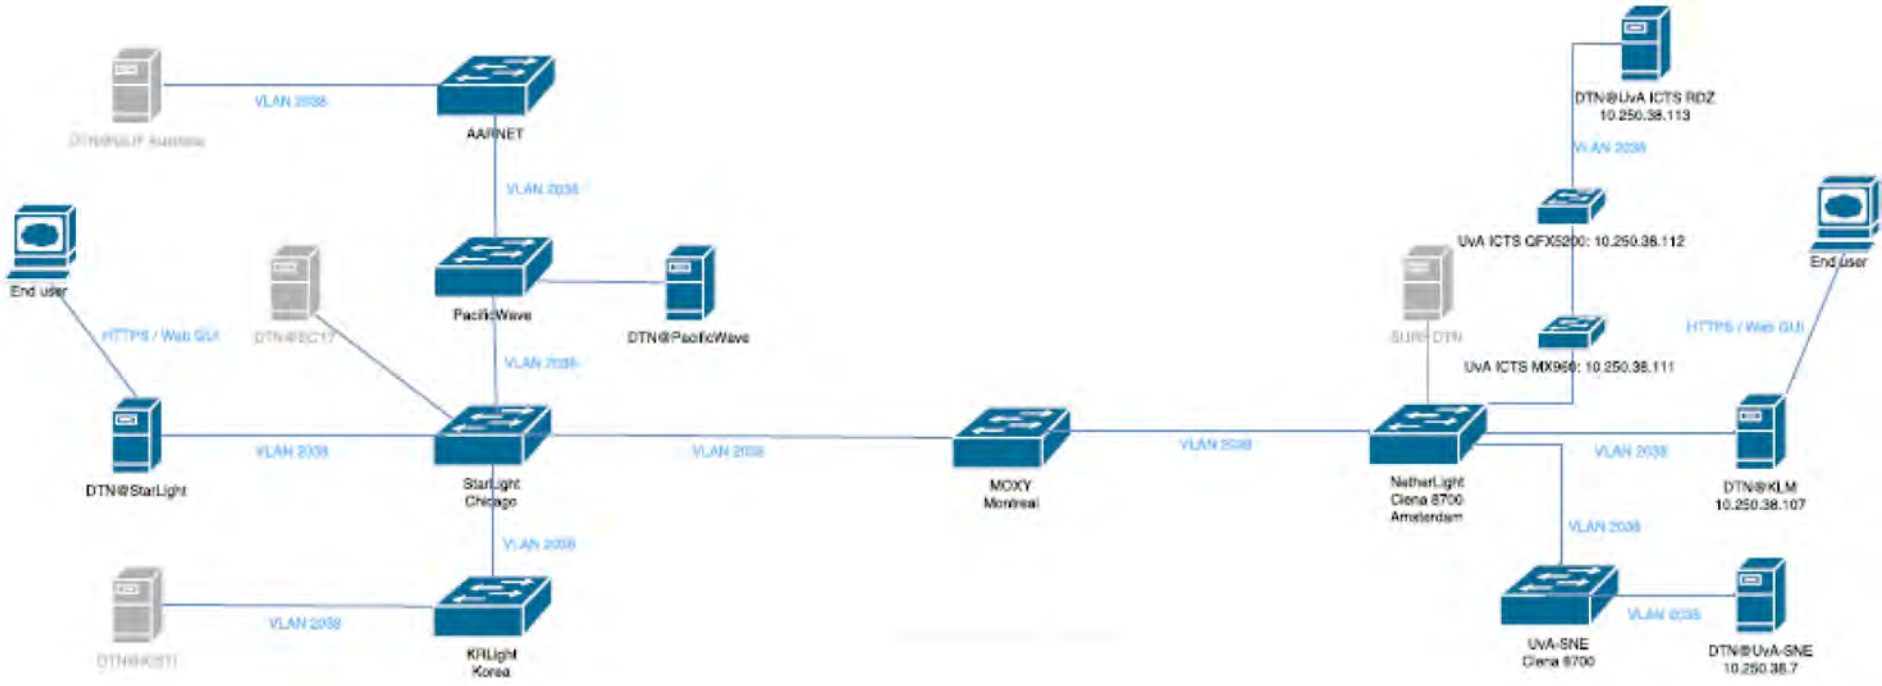

Ingredients
- Using Globus Toolkit (NOT Globus Online)
- Has GridFTP under the hood
- Under Globus license (must be evaluated)
- 40Gbit/s data transfer expected
- VLAN 2038, multipoint/extending
- Including authentication/authorization framework, e.g. SURFconext

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

100 Gbps DTN Optical Testbed

Ciena's OPⁿ research network testbed

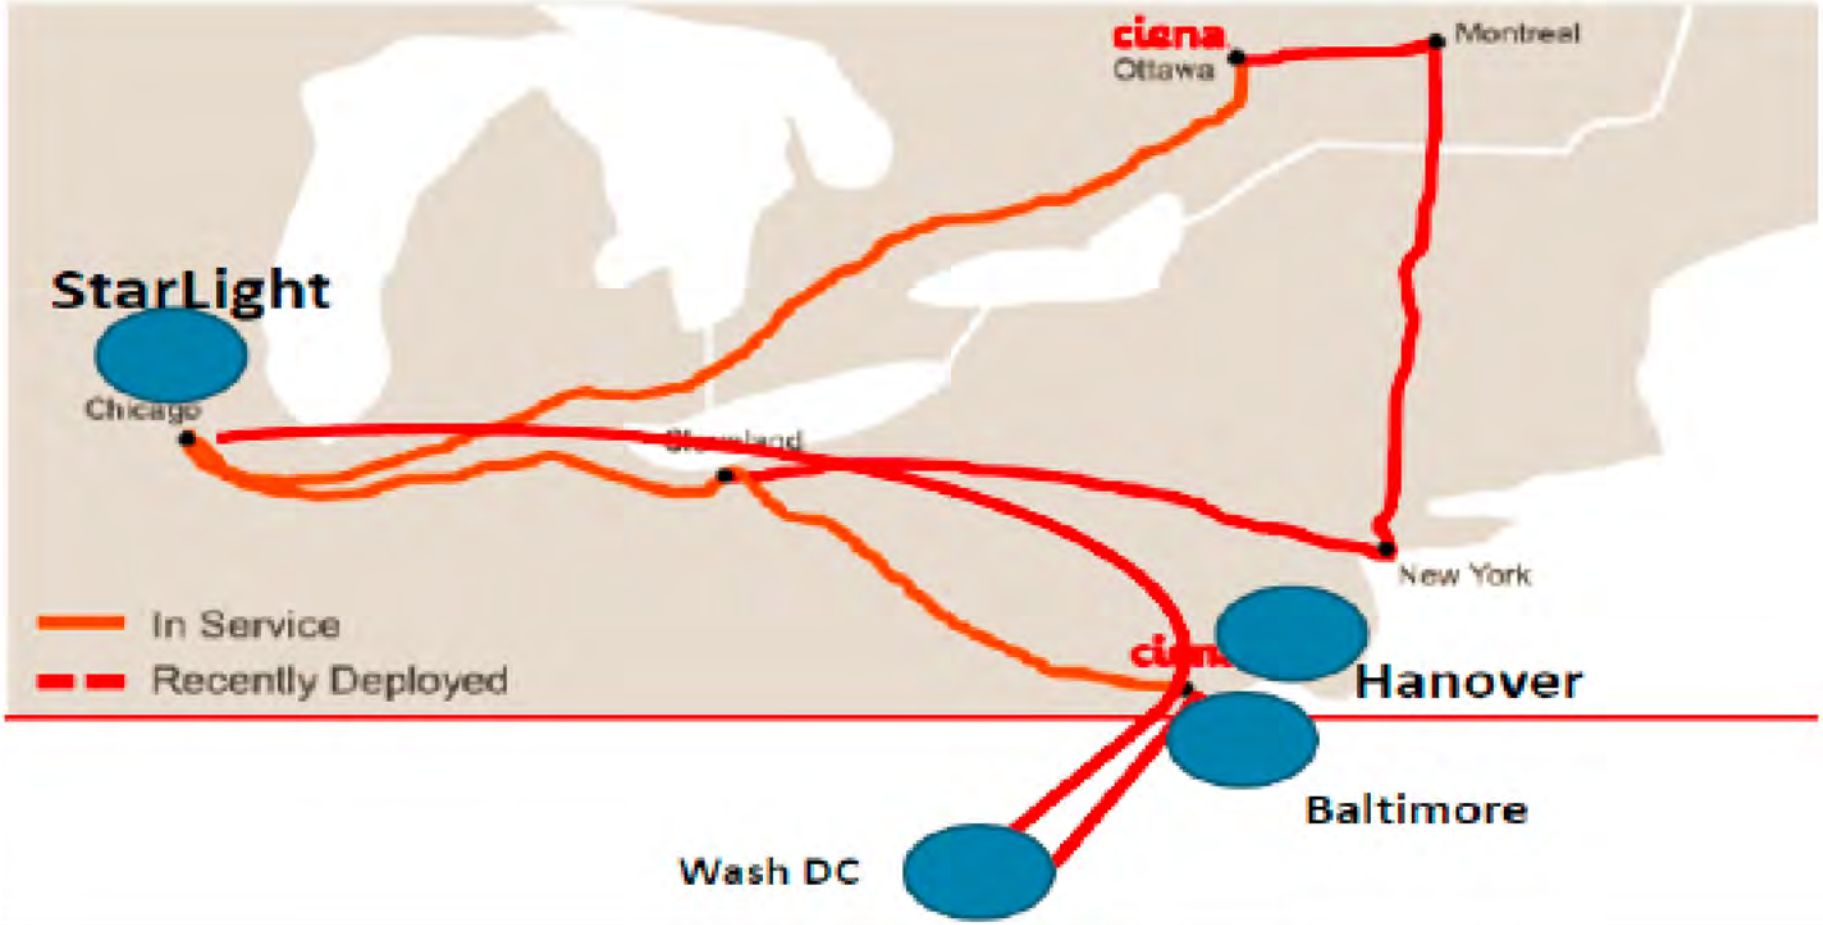

GENI – iCAIR P4 Testbed – Integrated With GENI StarLight SDX

- In Partnership With The GENI Initiative, iCAIR Is Developing a P4 Testbed for Computer Science Research.
- The Testbed Will Be Integrated With the GENI SDX At StarLight
- P4 (Programming Protocol-Independent Packet Processors).
- An Emerging Networking Programming Language,
- A Domain Specific Language for Network Protocols.
- Highly Flexible In Contrast To OpenFlow
- Testbed Based on Tofino (Barefoot Networks) Switches
- Compiler (V16) Enables Rules To Be Dynamically Implemented In Chip

Other Notable iCAIR Suported Testbeds

- International AI Testbed (Currently Being Designed)
- International DTN Testbed
- LHC P2P Service
- High Performance Digital Media Network (HPDMnet)
- Geophysical Sciences Testbed
- Content Routing Network
- AutoGOLE
- MEICAN/Network Service Interface (NSI)
- AmoebaNet
- Cisco Information Centric Networking Testbed (ICN)
- SD-WAN Testbed
- ToMaTo
- Et Al

Building the Open Storage Network

Alex Szalay
The Johns Hopkins University

Institute for Data Intensive Engineering and Science
idies

STARLIGHT SM

www.startap.net/starlight

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

"Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Networking and Information Technology Research and Development Program."

The Networking and Information Technology Research and Development (NITRD) Program

Mailing Address: NCO/NITRD, 2415 Eisenhower Avenue, Alexandria, VA 22314

Physical Address: 490 L'Enfant Plaza SW, Suite 8001, Washington, DC 20024, USA Tel: 202-459-9674, Fax: 202-459-9673, Email: nco@nitrd.gov, Website: https://www.nitrd.gov