

Federated International Network Research Testbeds

Joe Mambretti, Director, (j-mambretti@northwestern.edu)

International Center for Advanced Internet Research (www.icaair.org)

Northwestern University

Director, Metropolitan Research and Education Network (www.mren.org)

Co-Director, StarLight (www.startap.net/starlight), Director, StarLight

International/National Communications Exchange Facility

(www.startap.net/starlight),

PI IRNC: RXP: StarLight SDX, Co-PI Chameleon, PI-iGENI, PI-OMNINet



Global Research Platform Workshop

University of California San Diego

La Jolla, California

September 16-17, 2019



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, and Networks (e.g., StarLight, Specialized Science Networks, the Metropolitan Research and Education Network et al)



Selected iCAIR Research Topics

SDX Development

- ❖ 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 , Including Network Appliances such as data Transfer Nodes (DTNs)
- ❖ 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/SDC/SDS/SDI/SDE



Selected Applications



GENI
www.geni.net



Open Storage Network
www.openstorage.network.org



OSIRIS
www.osris.org



XSEDE
www.xsede.org



Blue Waters
bluewaters.ncsa.illinois.edu



PRAGMA
www.pragma-grid.net



CENTRA
www.globalcentra.org



OSG
www.openscience.grid.org



GRP
theglobalresearchplatform.net/



PRP
pacificresearchplatform.org



CHASE-CI
www.calit2.net/newsroom/article.php?id=2910



SAGE2
sage2.sagecommons.org



Polar Geospatial Center
www.pgc.umn.edu



IceCube
icecube.wisc.edu



Chameleon
www.chameleoncloud.org



Jetstream
www.jetstream-cloud.org



Genomic Science Program
genomicscience.energy.gov



LSST
www.lsst.org



Pierre Auger Observatory
www.auger.org



Belle II
www.belle2.org



LBNF/DUNE/ProtoDUNE
lbnf.fnal.gov



ISS
www.nasa.gov/station



SKA
www.skatelescope.org



XENON
xenon.astro.columbia.edu



NOVA
novaexperiment.fnal.gov



Virgo
www.virgo-gw.eu



LIGO
www.ligo.caltech.edu



SDSS
www.sdss.org



ALMA
www.almaobservatory.org



LHC
home.cern/science/accelerators/large-hadron-collider



LHCONE
twiki.cern.ch/twiki/bin/view/LHCONE/WebHome



LHCOPN
twiki.cern.ch/twiki/bin/view/LHCOPN/WebHome



IVOA
www.ivoa.net

International Federated Testbeds As Instruments for Computer Science/Network Science

- **The StarLight Communications Exchange Facility Supports ~ 25-30 Network Research Testbeds (Instruments For Computer Science/Networking Research)**
- **Software Defined Networking Techniques Are Used To Segment (“Slice”) Distributed Virtual Environments**
- **StarLight Supports Two Software Defined Exchanges (SDXs), An NSF IRNC SDX & a Network Research GENI SDX (Global Environment for Network Innovations)**
- **The GENI SDX Supports National and International Federated Testbeds**



StarLight SDX Overview

- This IRNC StarLight SDX Initiative Is Designing, Implementing, and Operating New Services For Global Data Intensive Sciences, Based On Emerging Next Generation Architecture and Technologies, Including Virtualization, Orchestration, Segmentation (Slicing), Software Defined Resources, Programmability and Customization.
- These Macro Trends Enable Exchanges To Be Agile Platforms For Dynamic Services Provisioning, Real-Time Responsiveness, and Distributed Control Over Core Resources, Including By Applications, Edge Processes and Devices.

This Project Is Transitioning Network Exchanges
To
Open Innovation Platforms

StarLight – “By Researchers For Researchers”

StarLight: Experimental Optical Infrastructure/Proving Ground For Next Gen Network Services
Optimized for High Performance Data Intensive Science
Multiple 100 Gbps (60+ Paths)
StarWave
100 G Exchange
World’s Most Advanced Exchange
Multiple First of a Kind
Services and Capabilities



View from StarLight



Abbott Hall, Northwestern University's Chicago Campus



Global Research Platform: Global Lambda Integrated Facility Available Advanced Network Resources



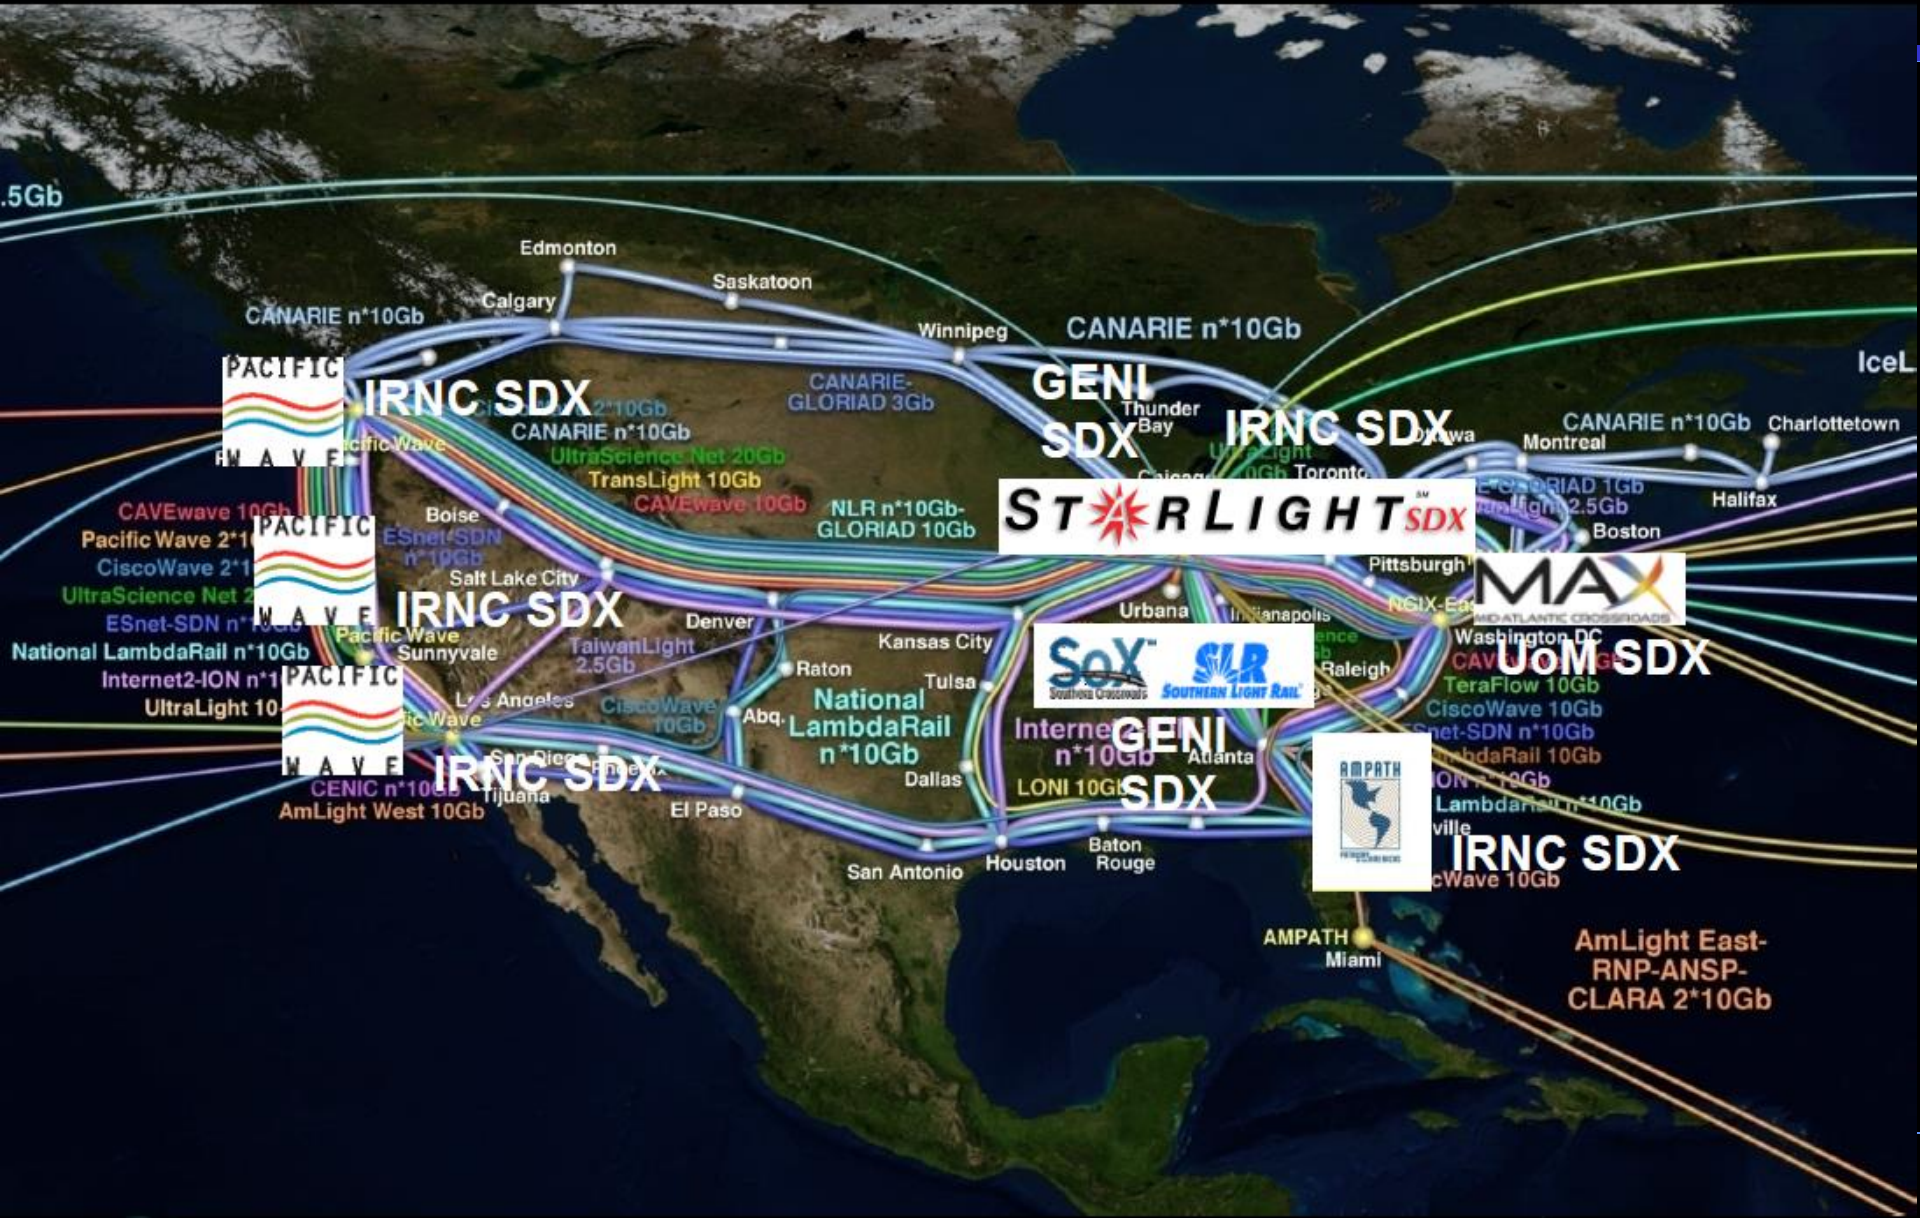
Visualization courtesy of Bob Patterson, NCSA; data compilation by Maxine Brown, UIC.



www.glif.is

STARLIGHTSM

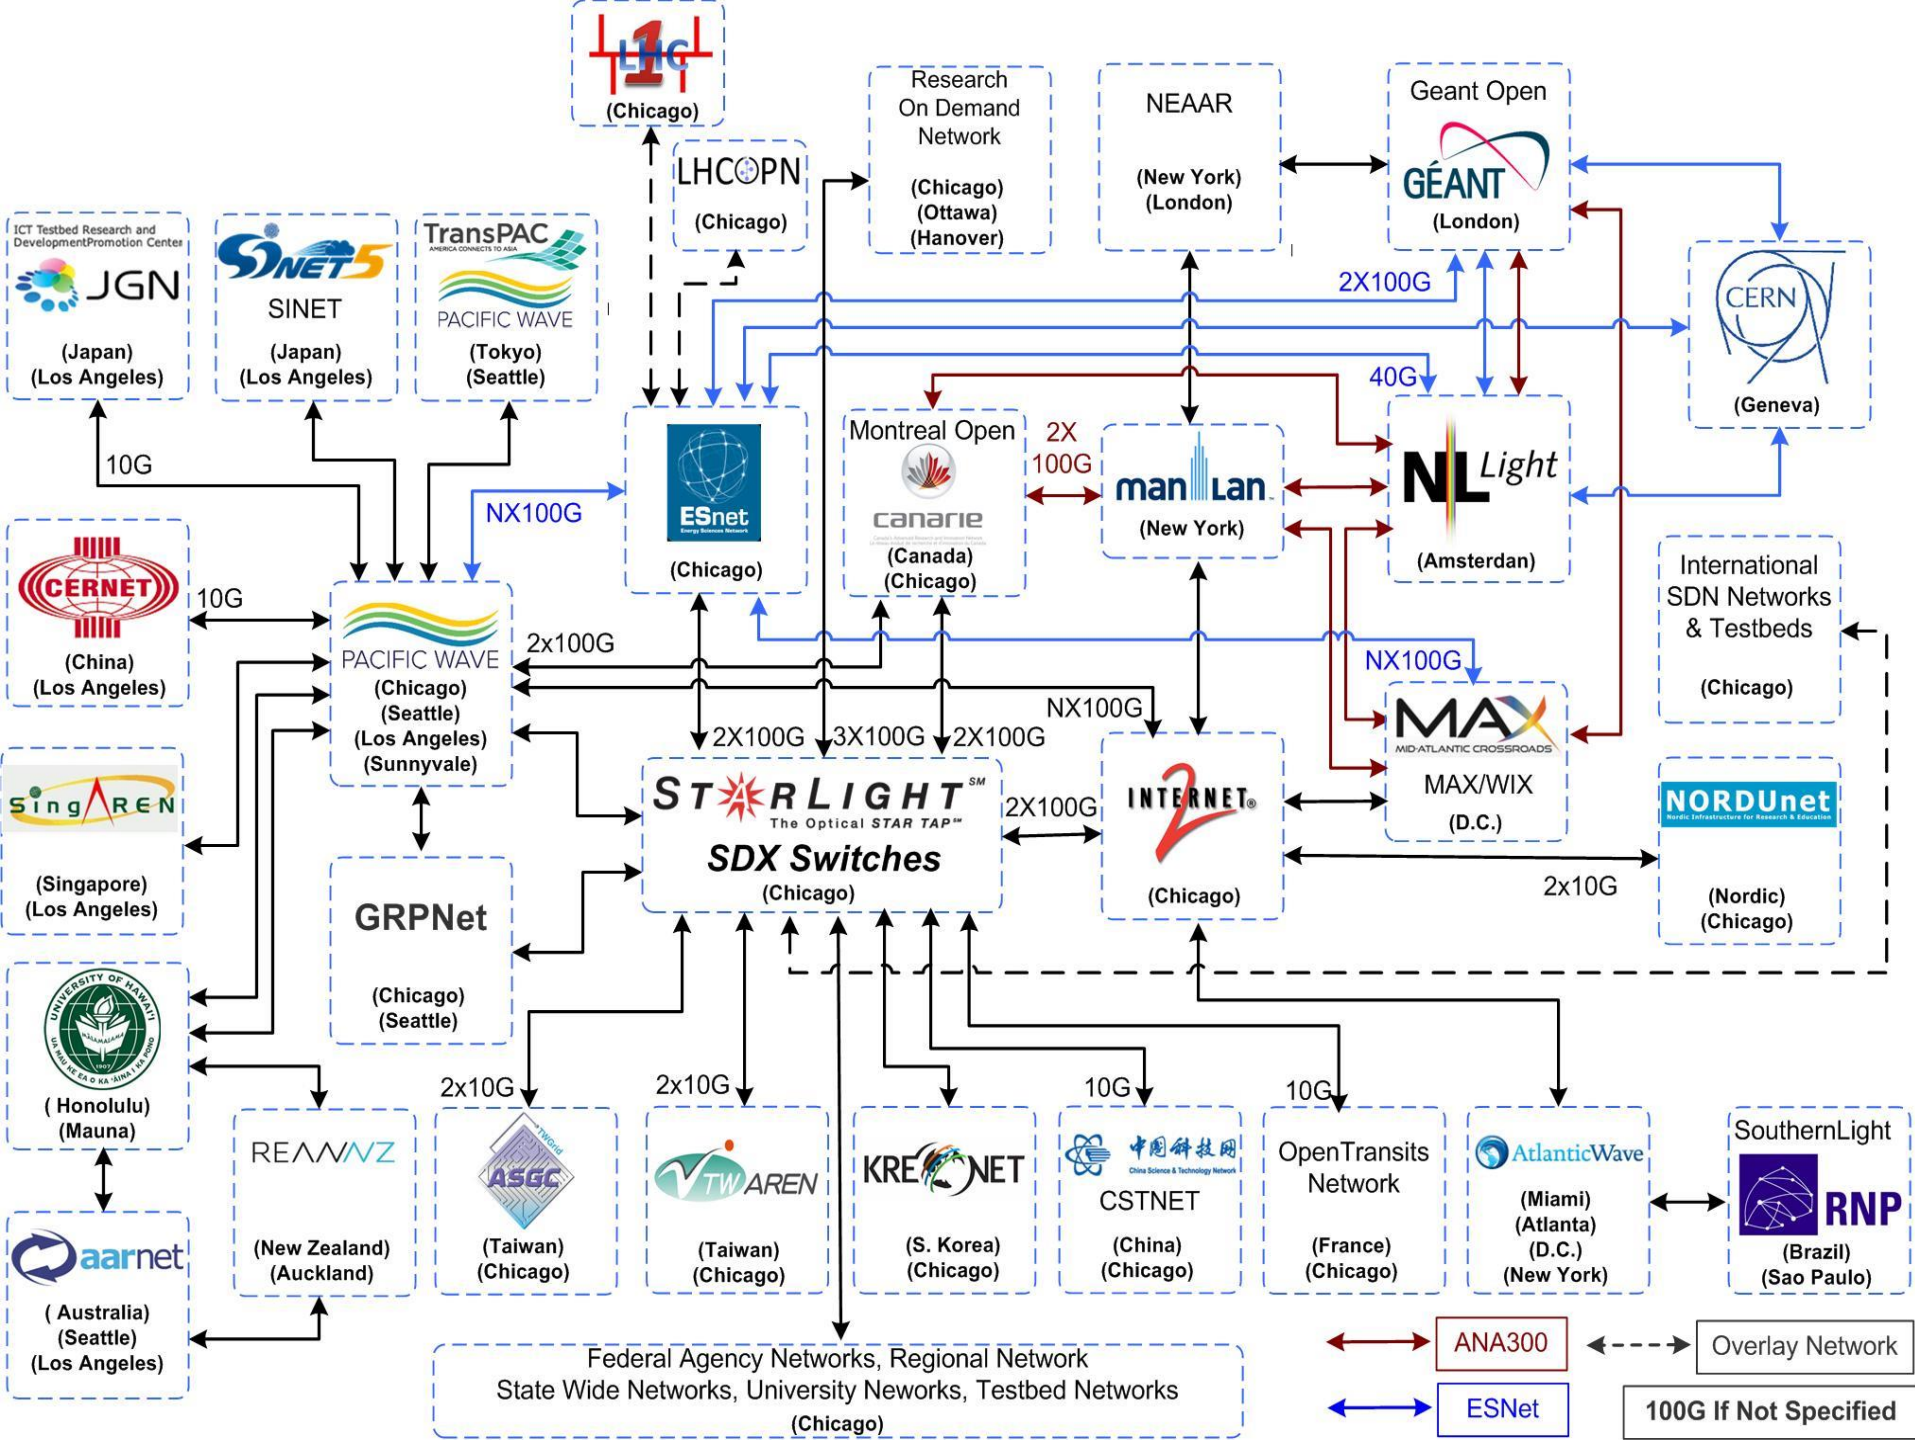
Emerging US SDX Interoperable Fabric



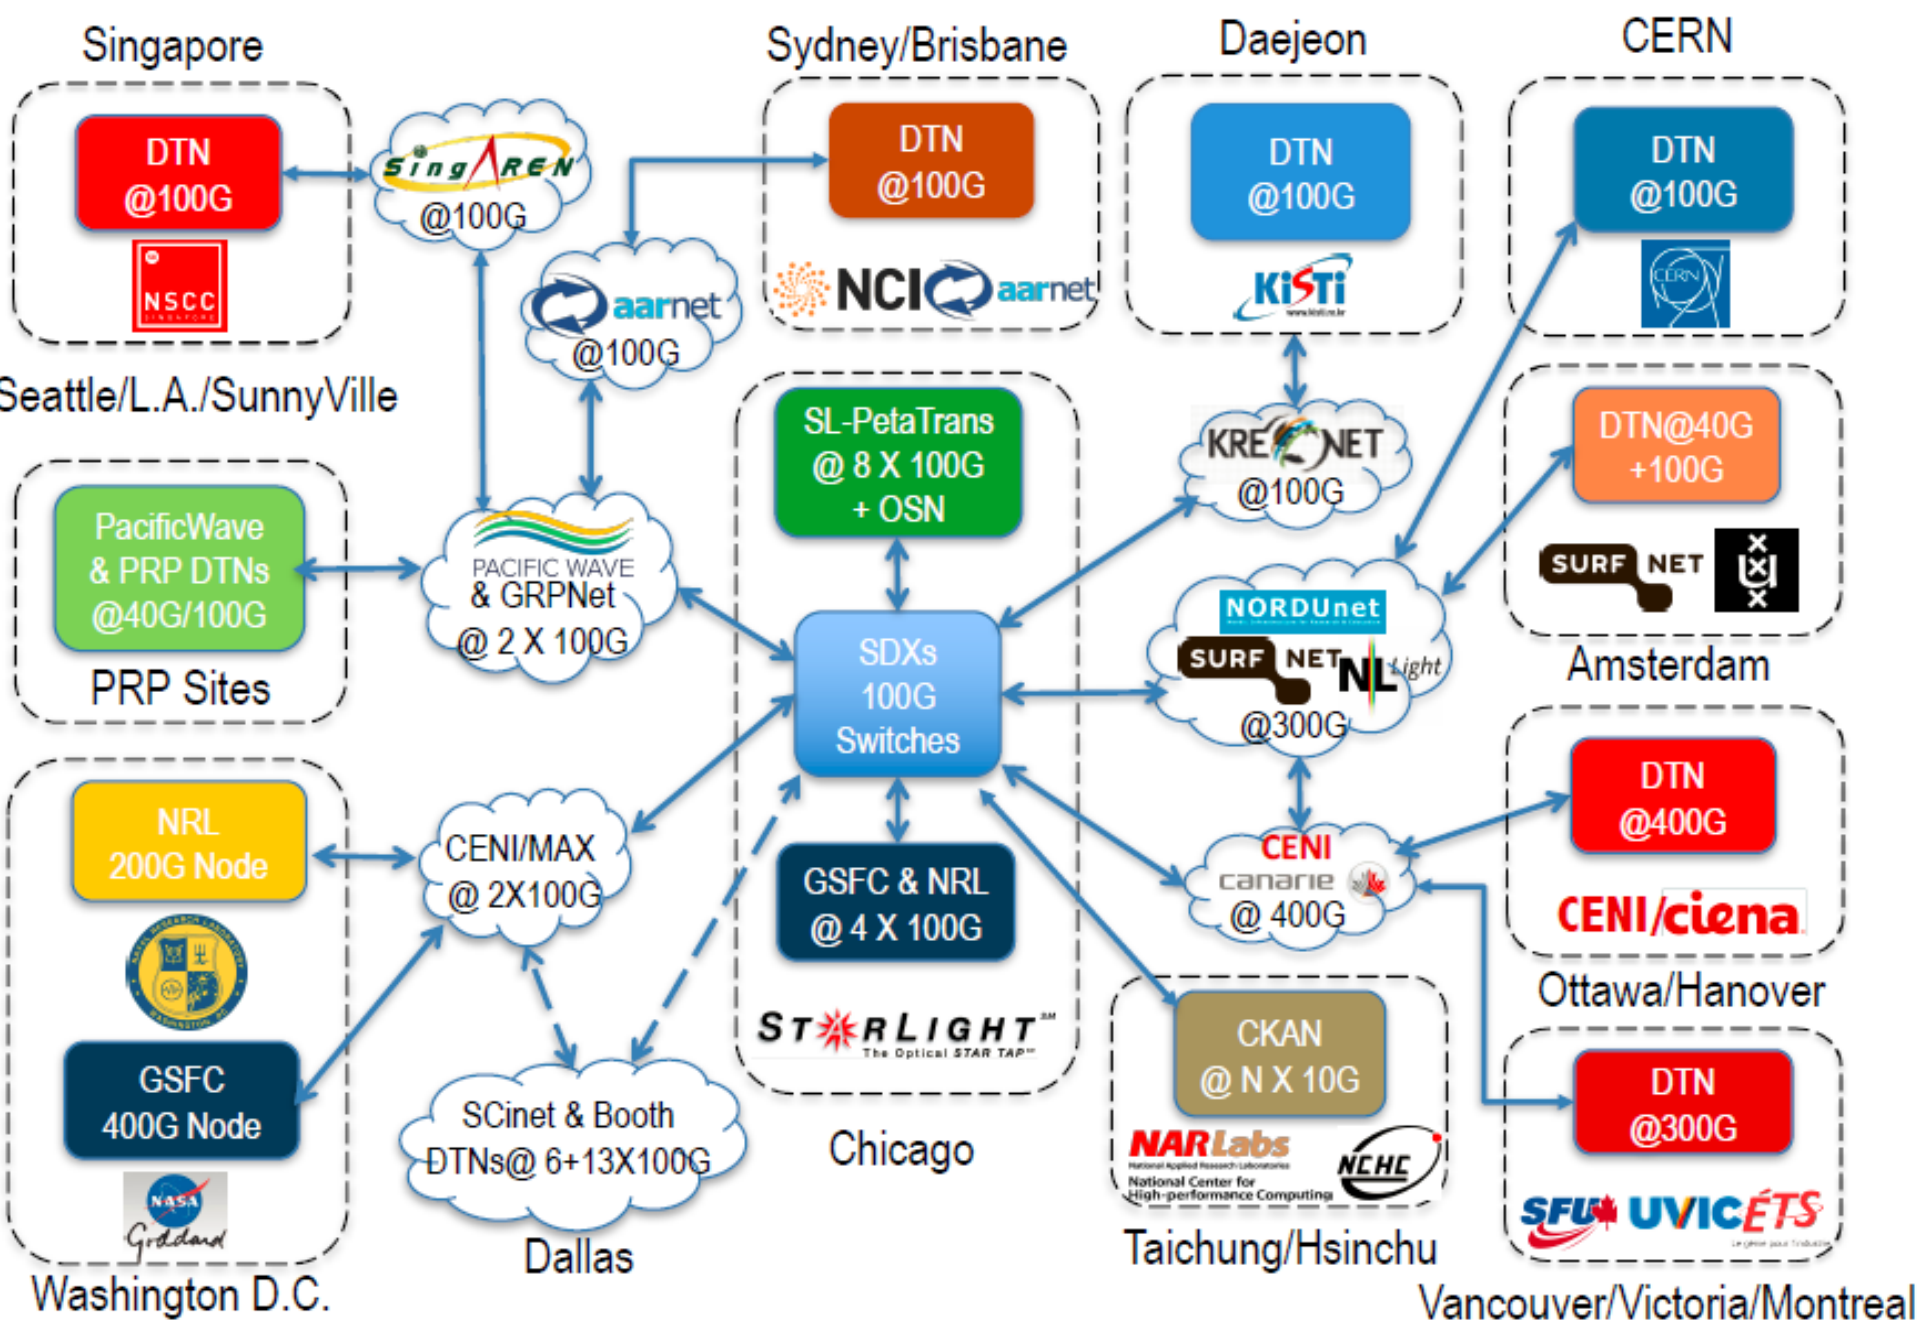
GLIF Automated GOLE Fabric 2018




SURF




PetaTrans: Petascale Sciences Data Transfer





-  Campus Network
-  WIMAX/LTE
-  Advanced Layer2 Service POP
-  InstaGENI Rack
-  ExoGENI Rack
-  OpenGENI Rack
-  CiscoGENI Rack
-  Regional Network

EU FED4FIRE International Federation



Also Note: GEANT Testbed-as-a-Service



www.chameleoncloud.org

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

Kate Keahey

Joe Mambretti, Pierre Riteau, Paul Ruth, Dan Stanzione

SEPTEMBER 28, 2017


1



TACC


renci



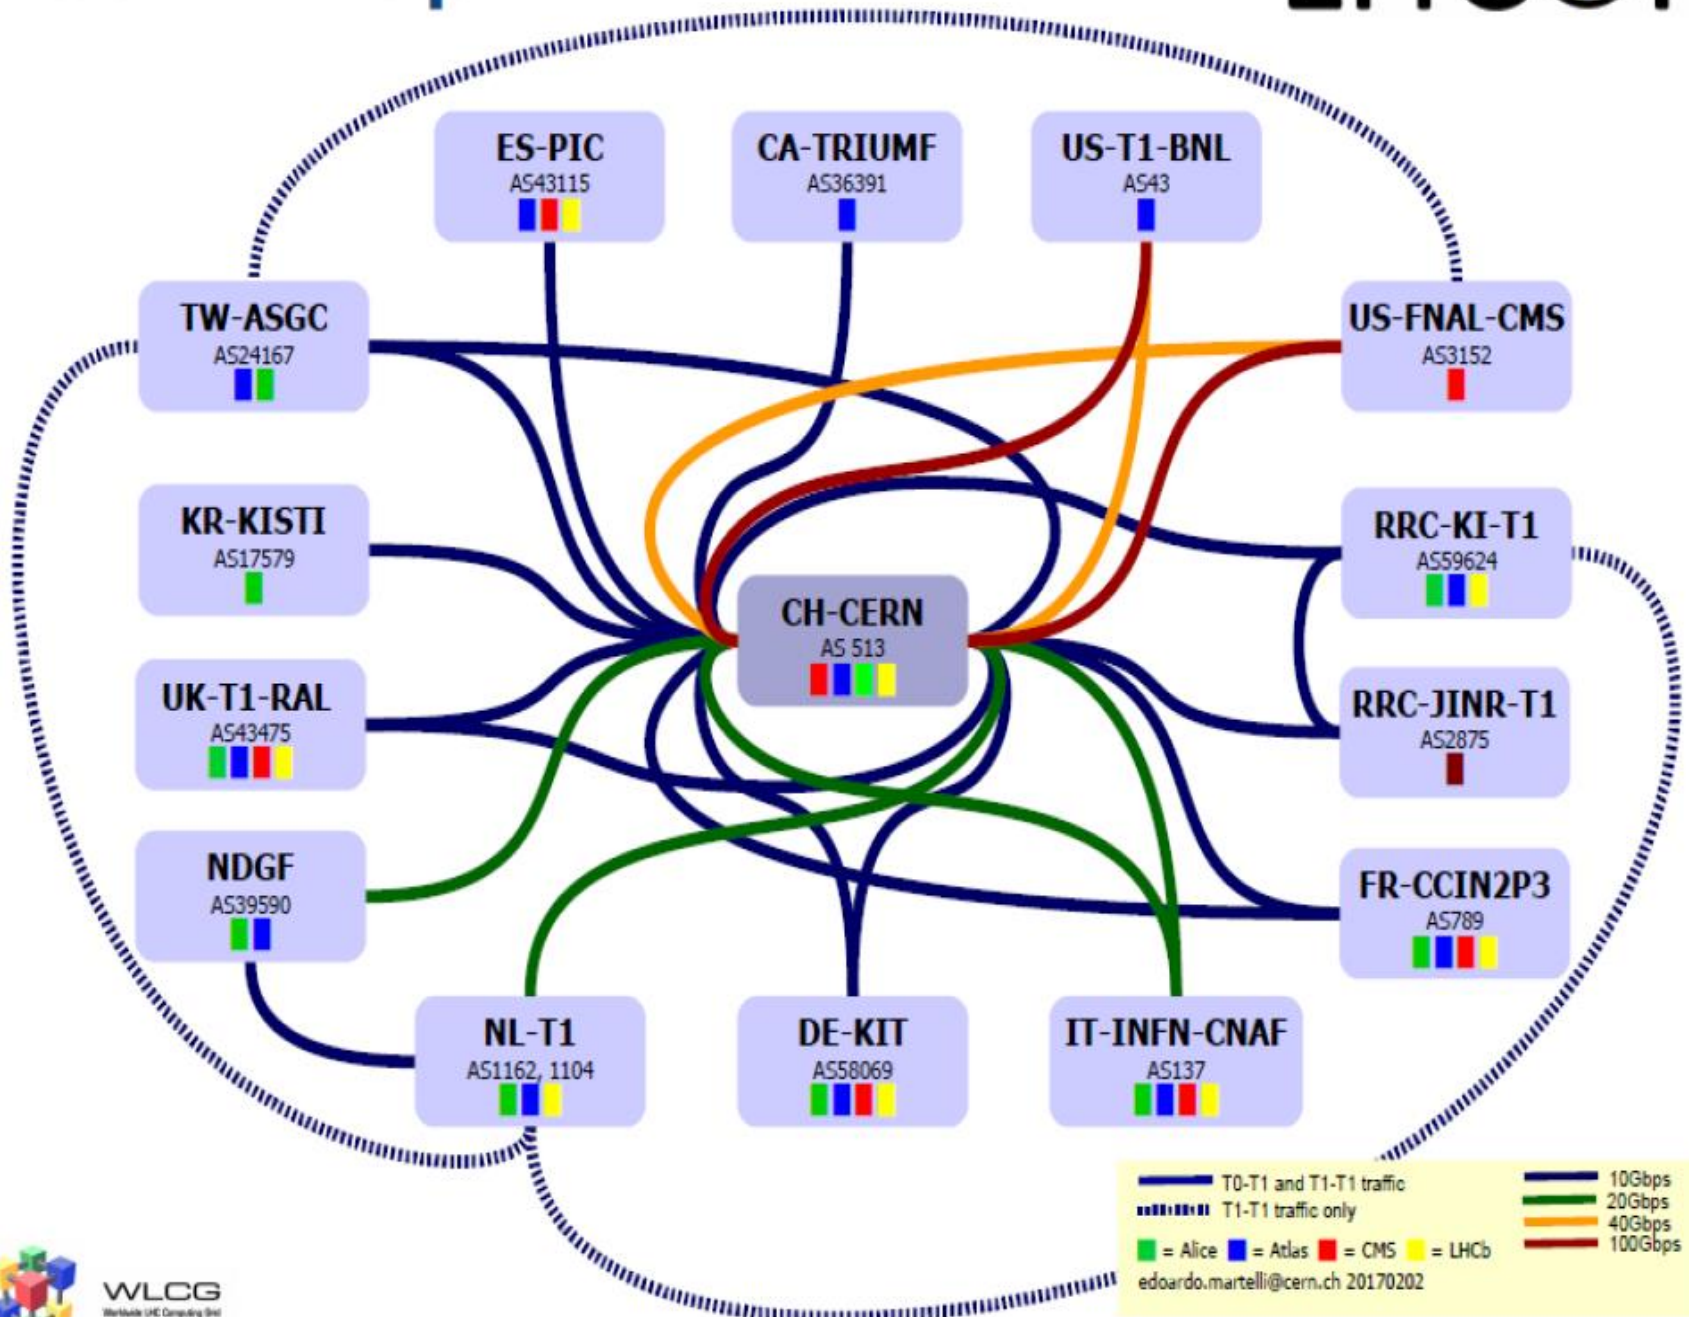





A Cross-Pacific SDN Testbed



LHCOPN map



LHCONE L3VPN: A global infrastructure for High Energy Physics data analysis (LHC, Belle II, Pierre Auger Observatory, NoVA, XENON)



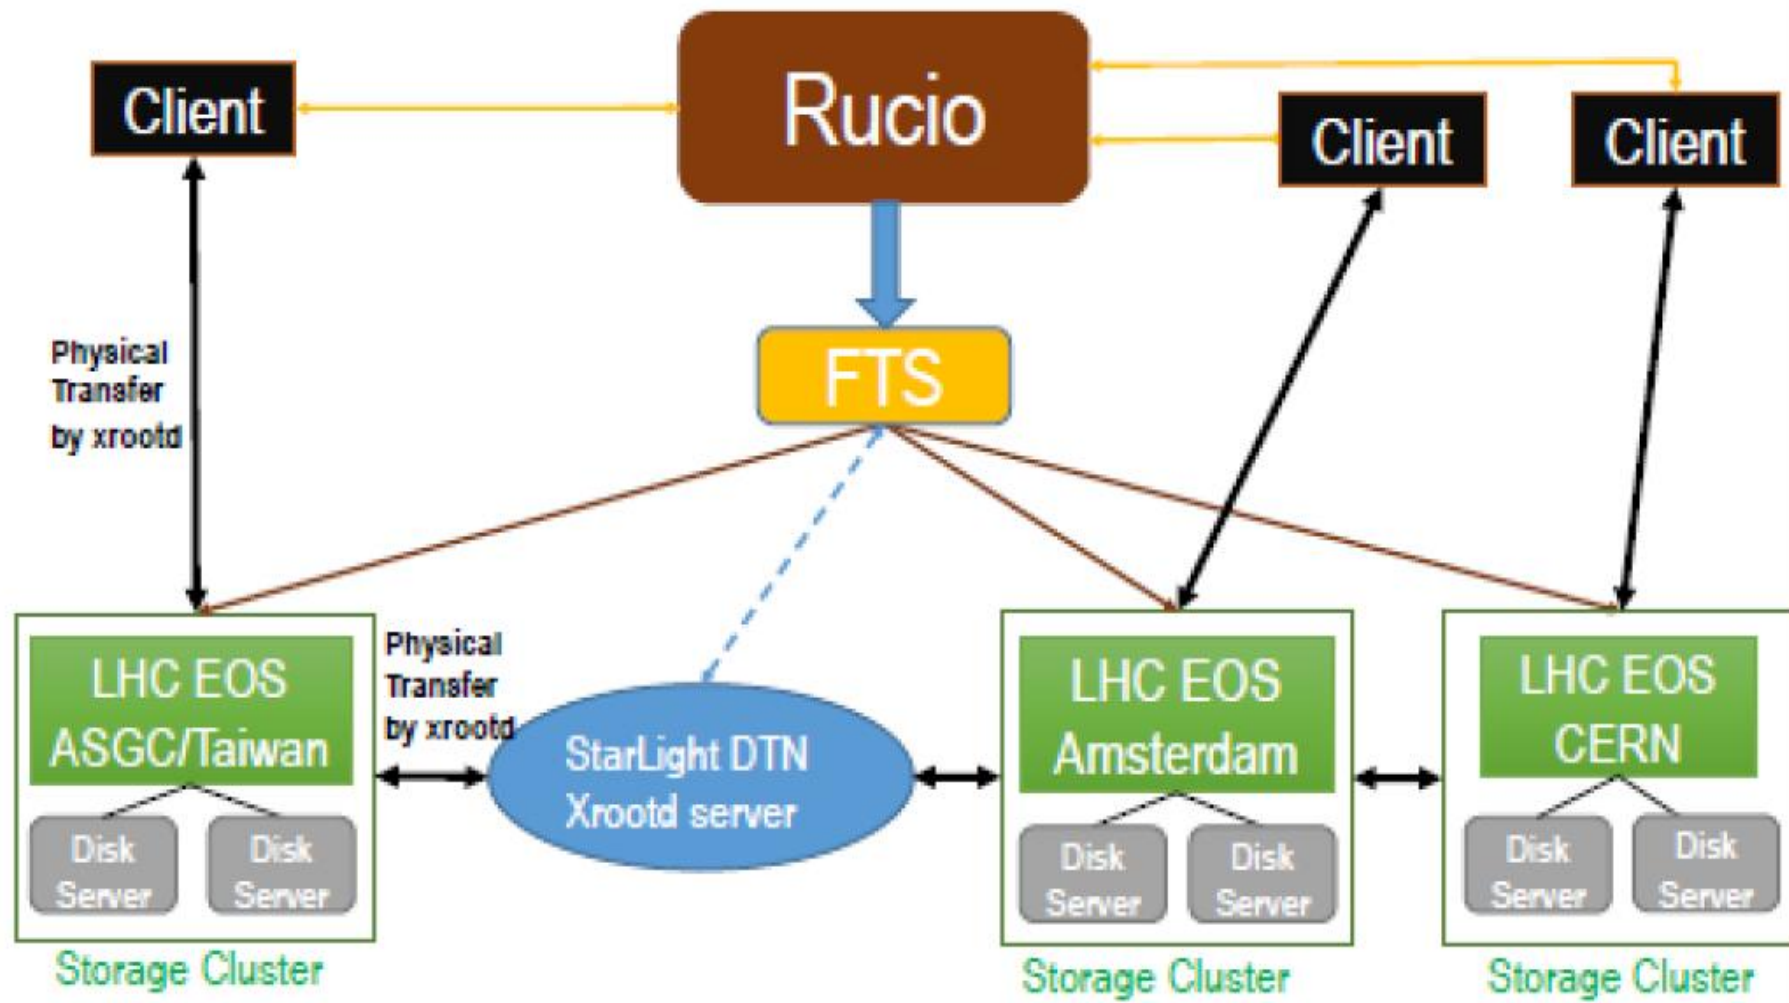
New Science Communities Using LHCONE

- ❖ 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.



New=> DUNE/ProtoDUNE – Deep Underground Neutrino Experiment



High Energy Physics SDX DTN Prototype Service



StarLight SDX Geoscience Research Workflow



20th Innovations in Clouds, Internet and Networks

PARIS

March 7 - 9, 2017



Designing and Deploying



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

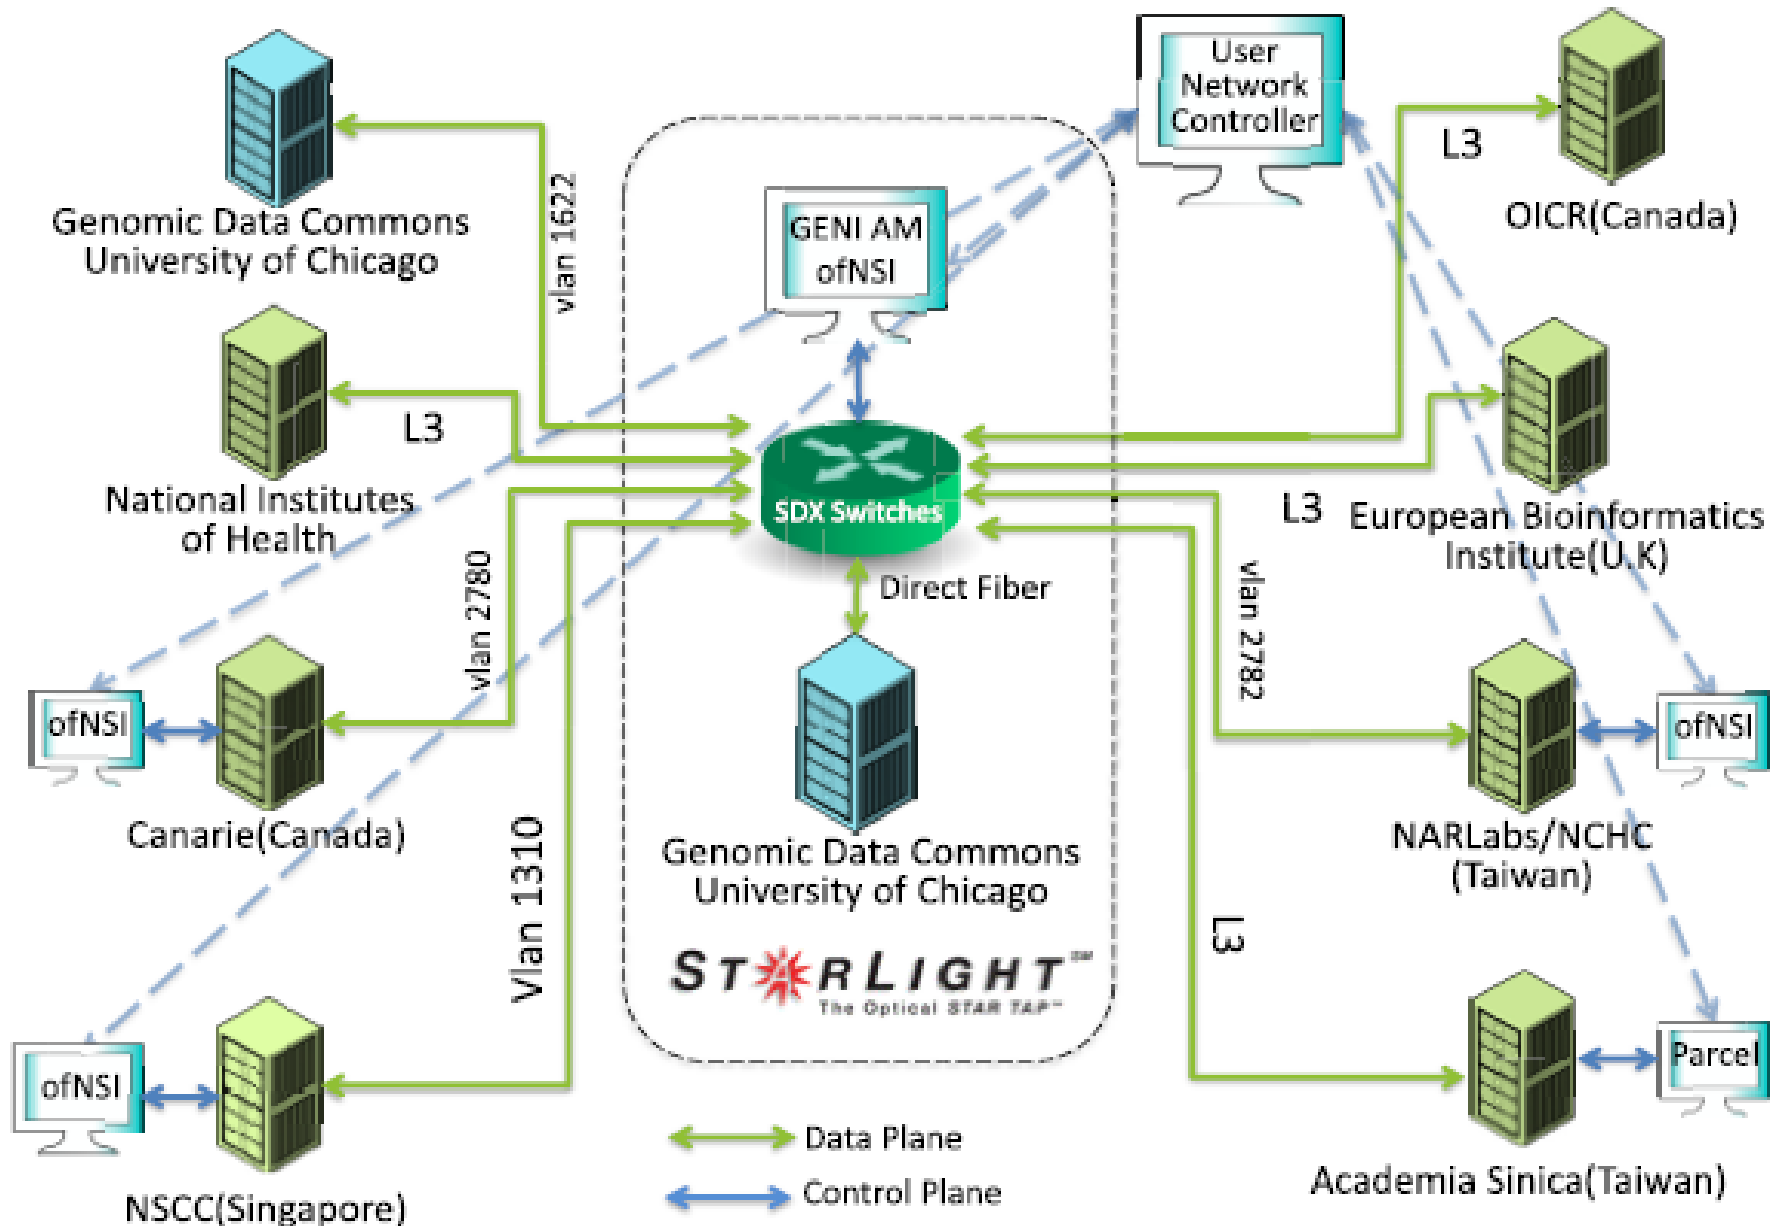
Robert Grossman, Piers Nash, Alison Heath, Renuka Arya, Stuti Agrawal,
Zhenyu Zhang

Center for Data Intensive Science
University of Chicago
Chicago, Illinois, USA

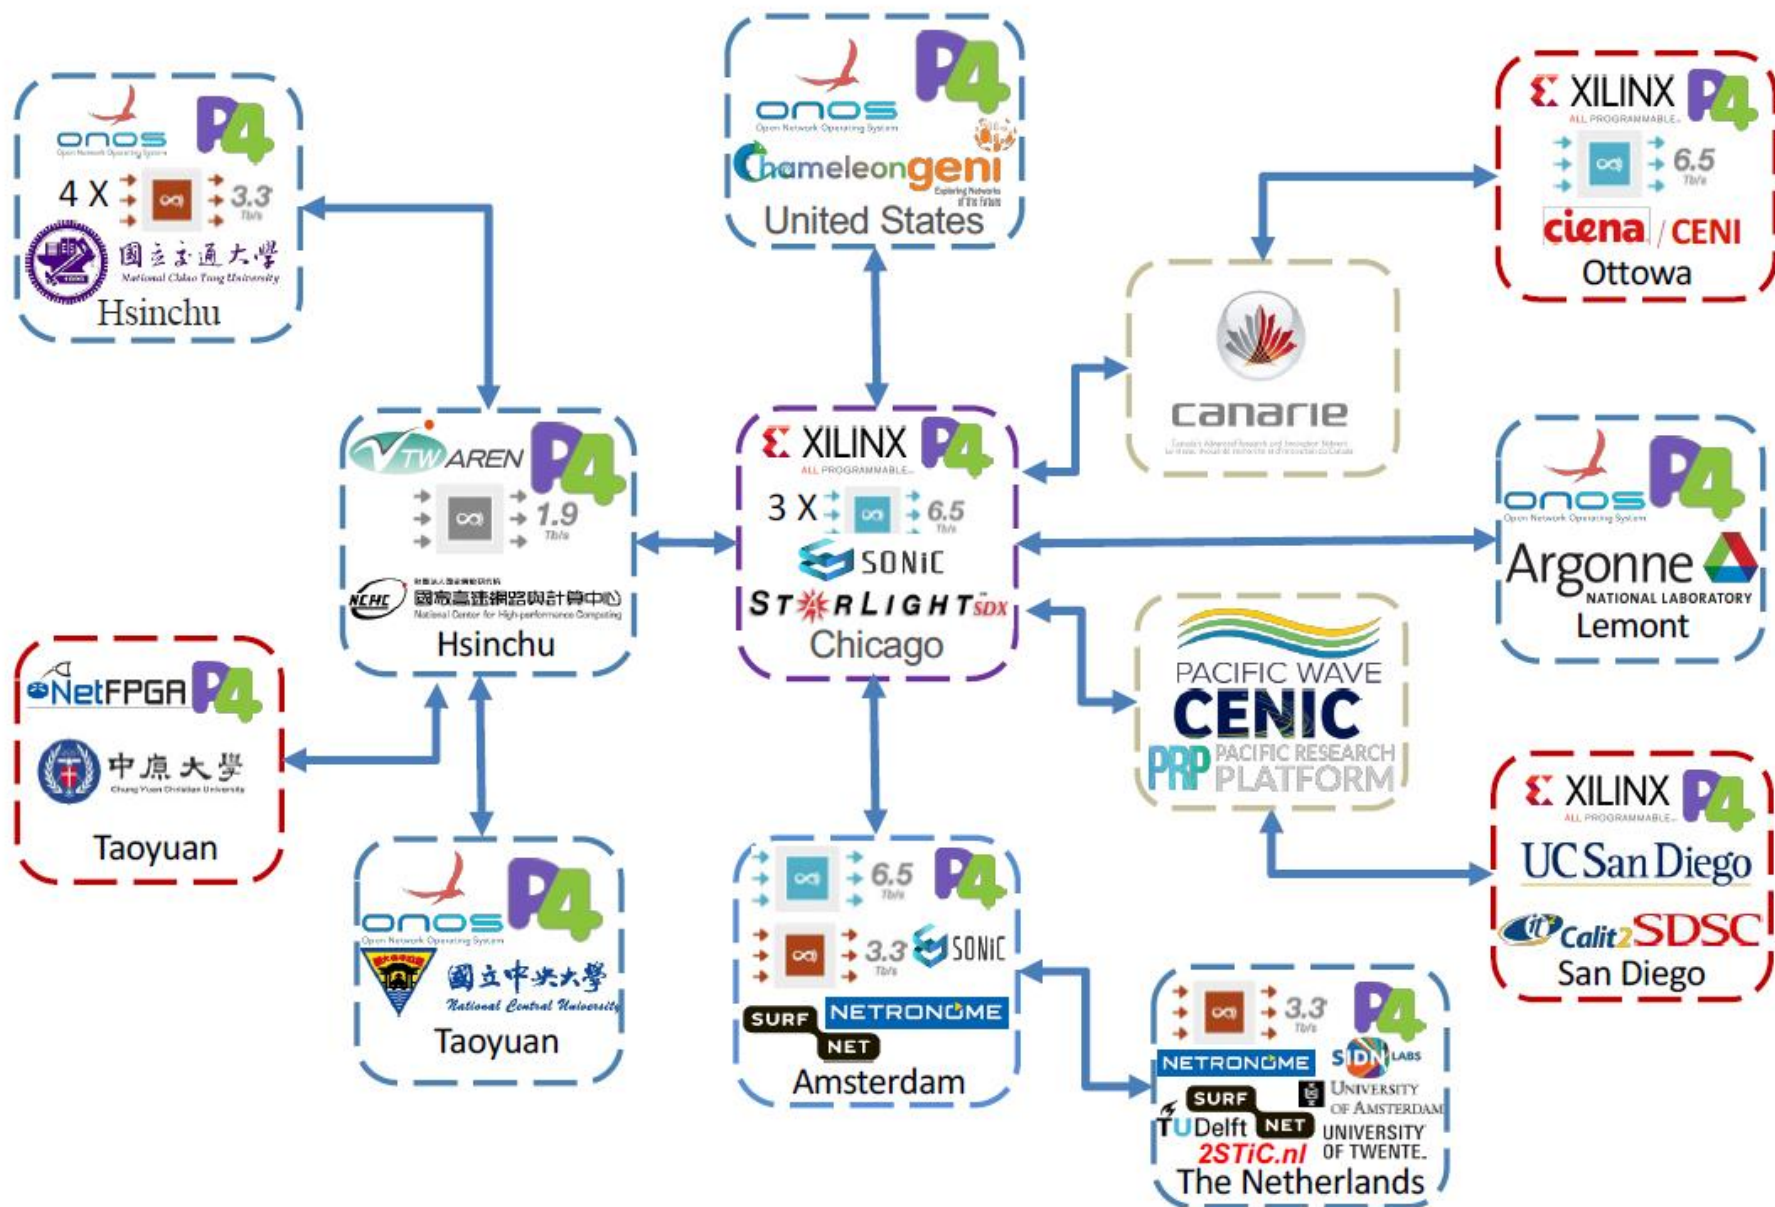


STARLIGHTSM

2016 Bioinformatics SDXs Network




International P4 Experimental Networks (i-P4EN)

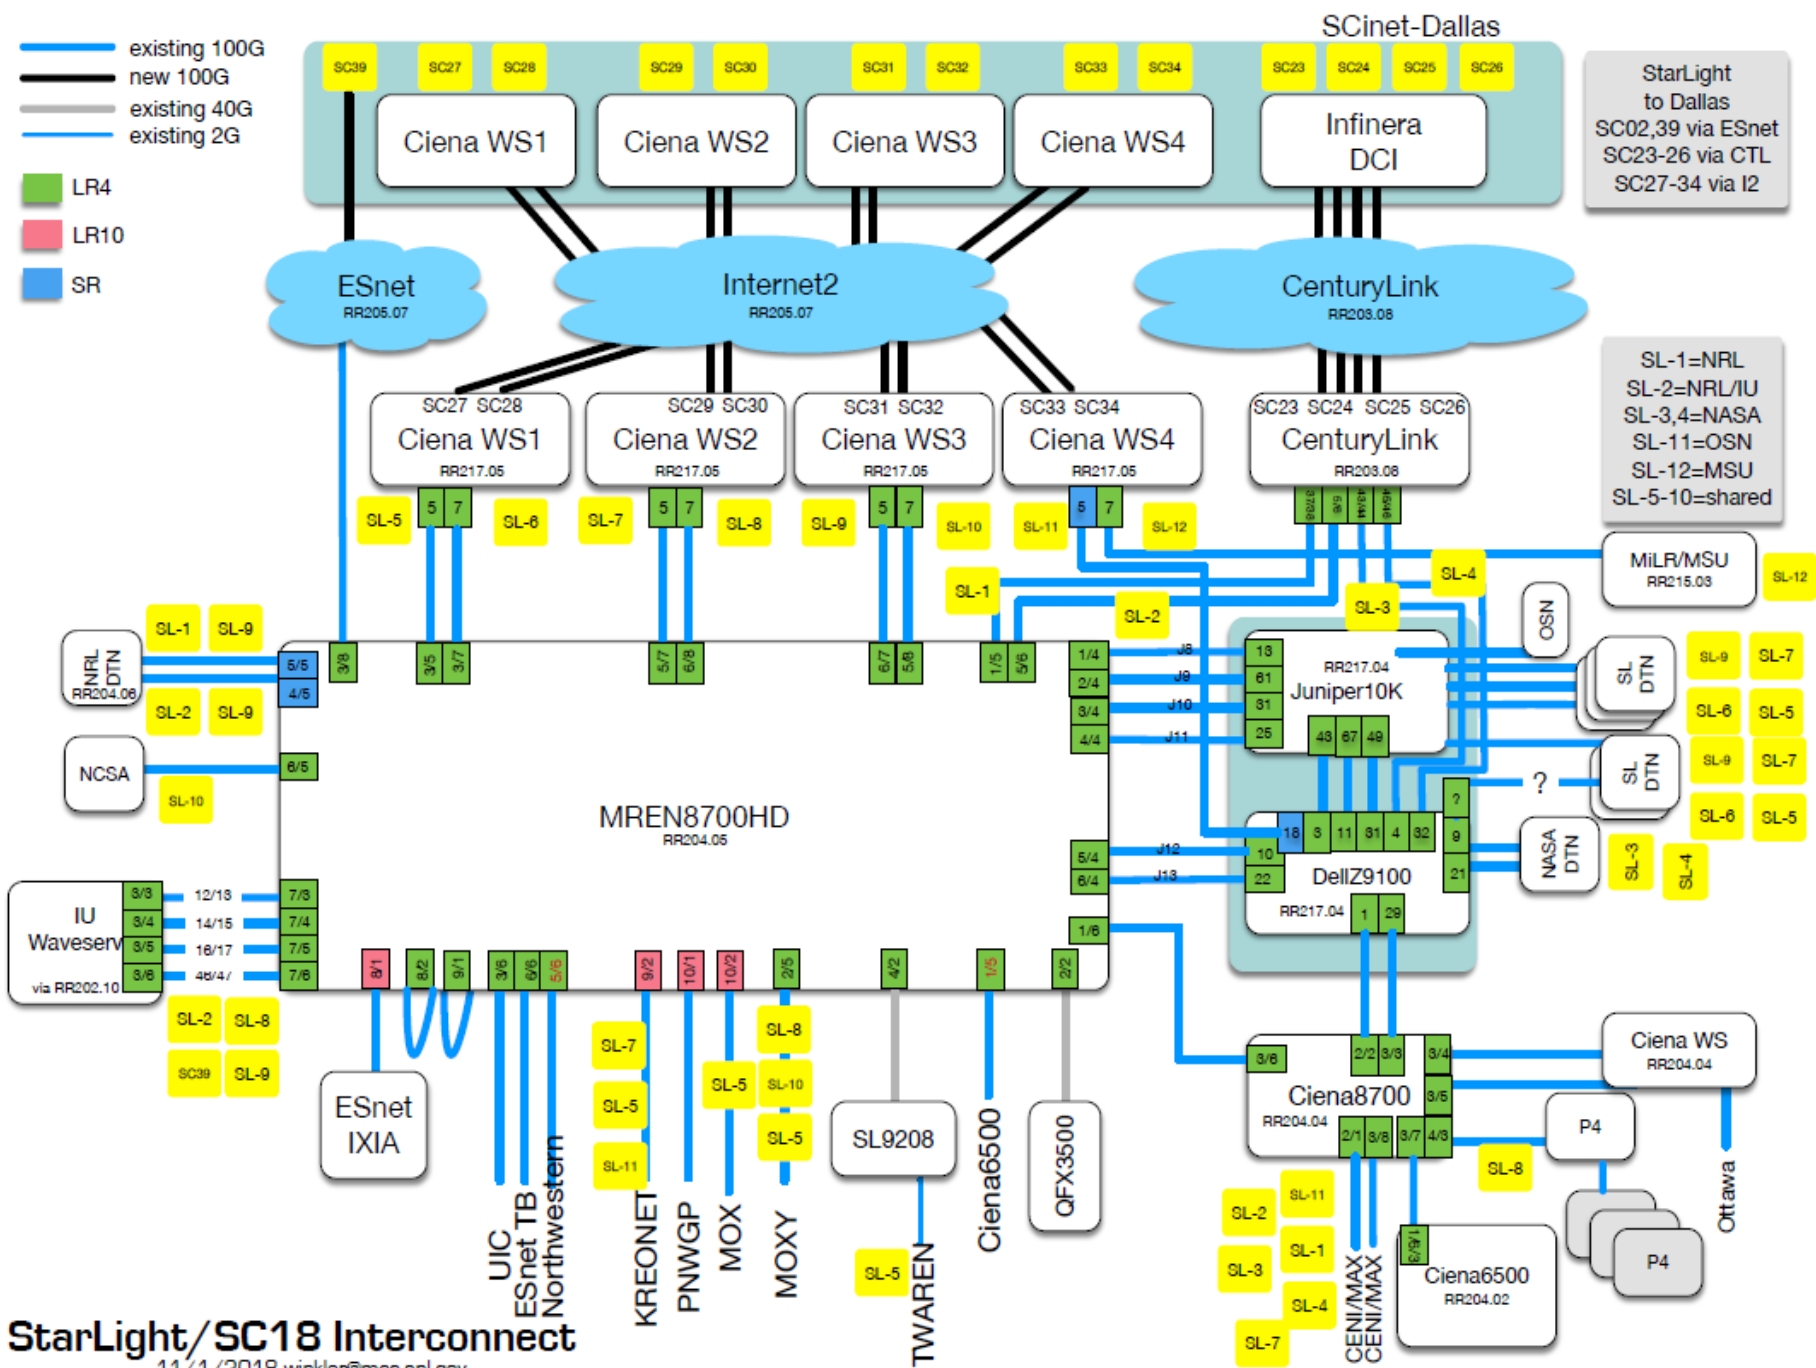


100 Gbps DTN Optical Testbed

Ciena's OPⁿ research network testbed



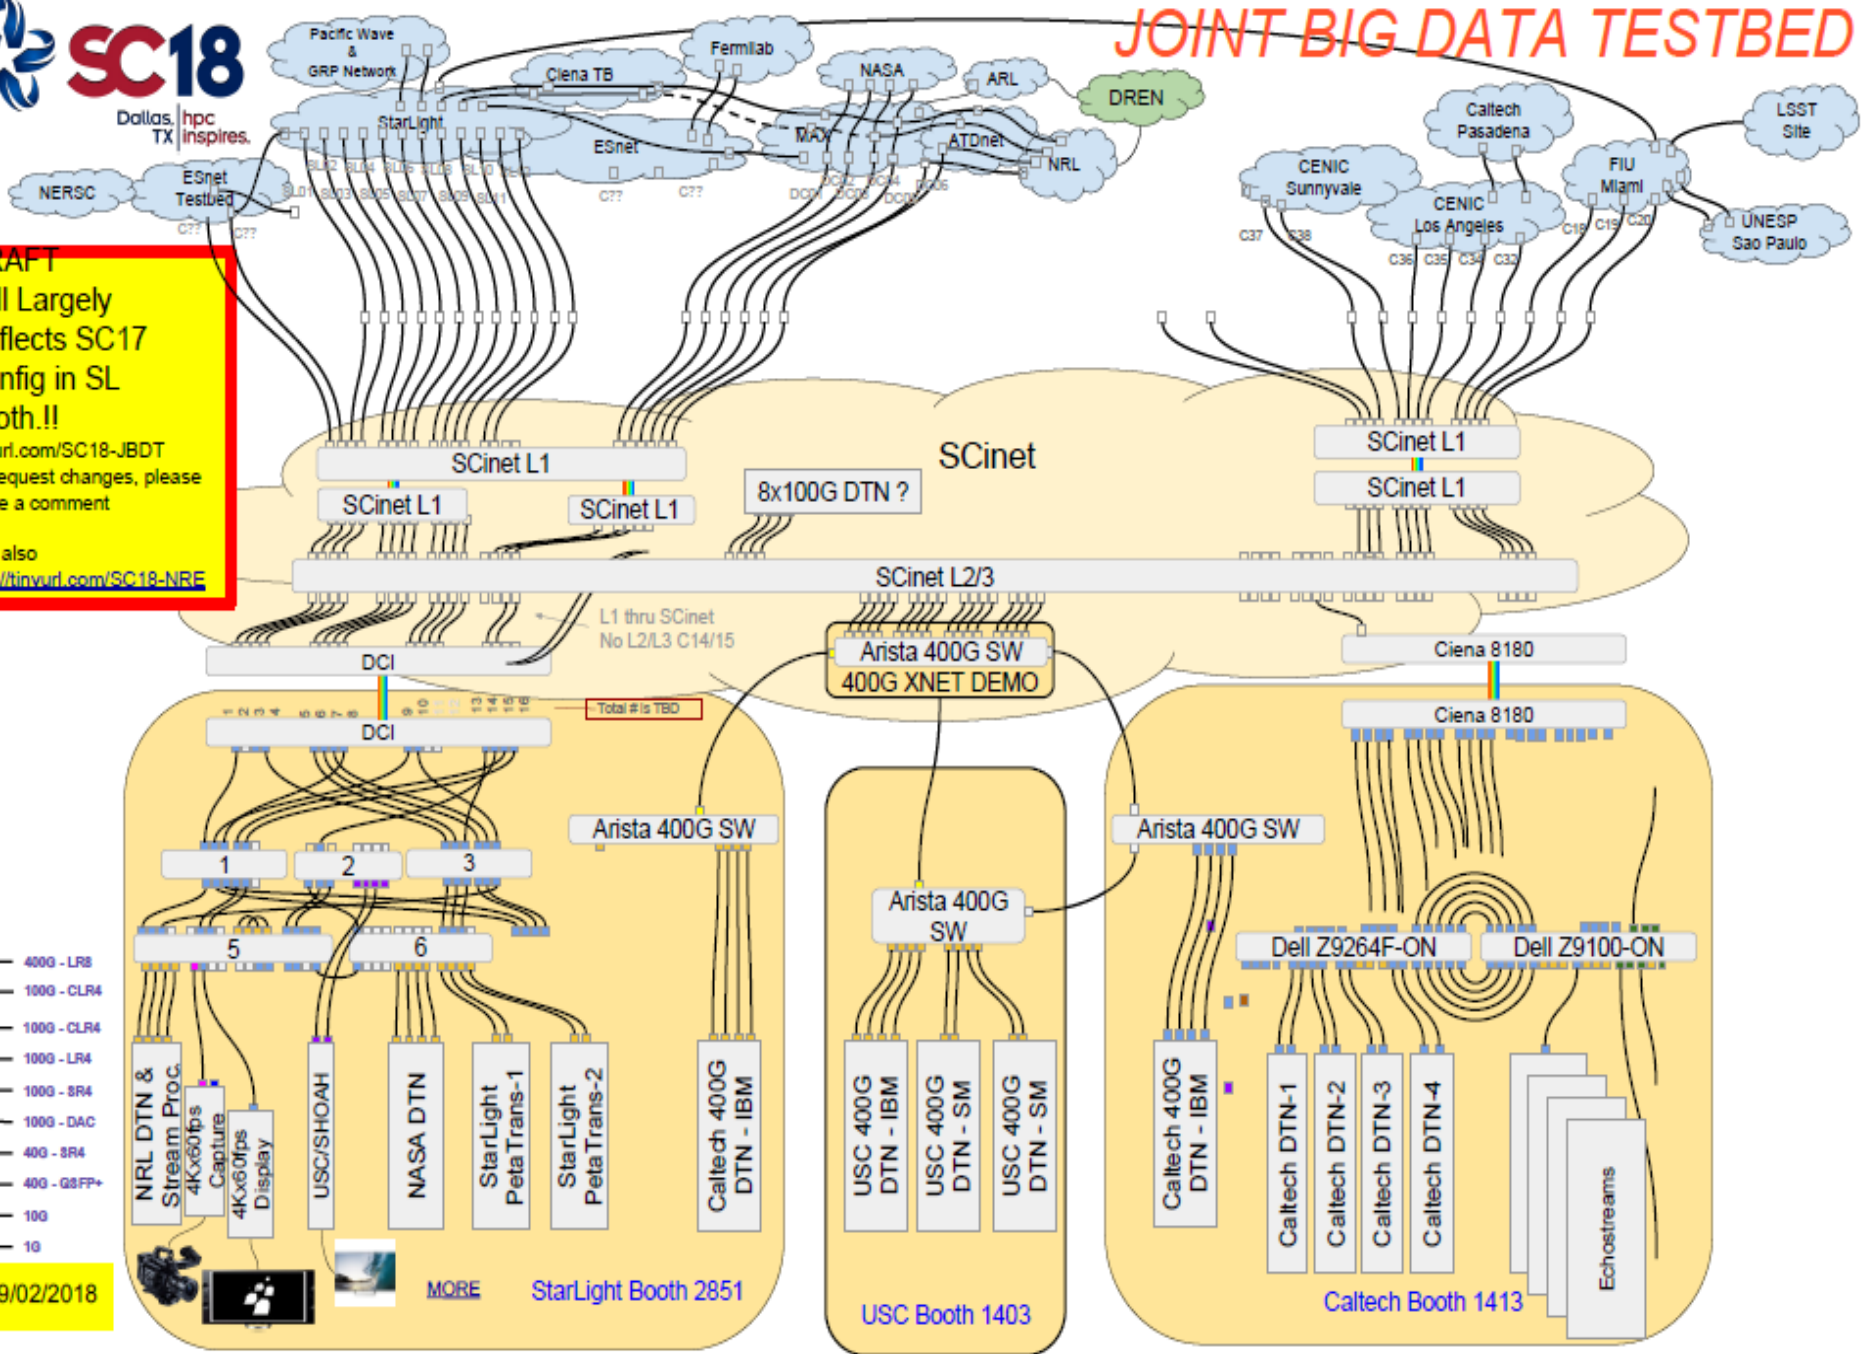
- existing 100G
- new 100G
- existing 40G
- existing 2G
- LR4
- LR10
- SR



StarLight to Dallas
 SC02,39 via ESnet
 SC23-26 via CTL
 SC27-34 via I2

SL-1=NRL
 SL-2=NRL/IU
 SL-3,4=NASA
 SL-11=OSN
 SL-12=MSU
 SL-5-10=shared

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>



09/02/2018

StarLight Booth 2851

USC Booth 1403

Caltech Booth 1413

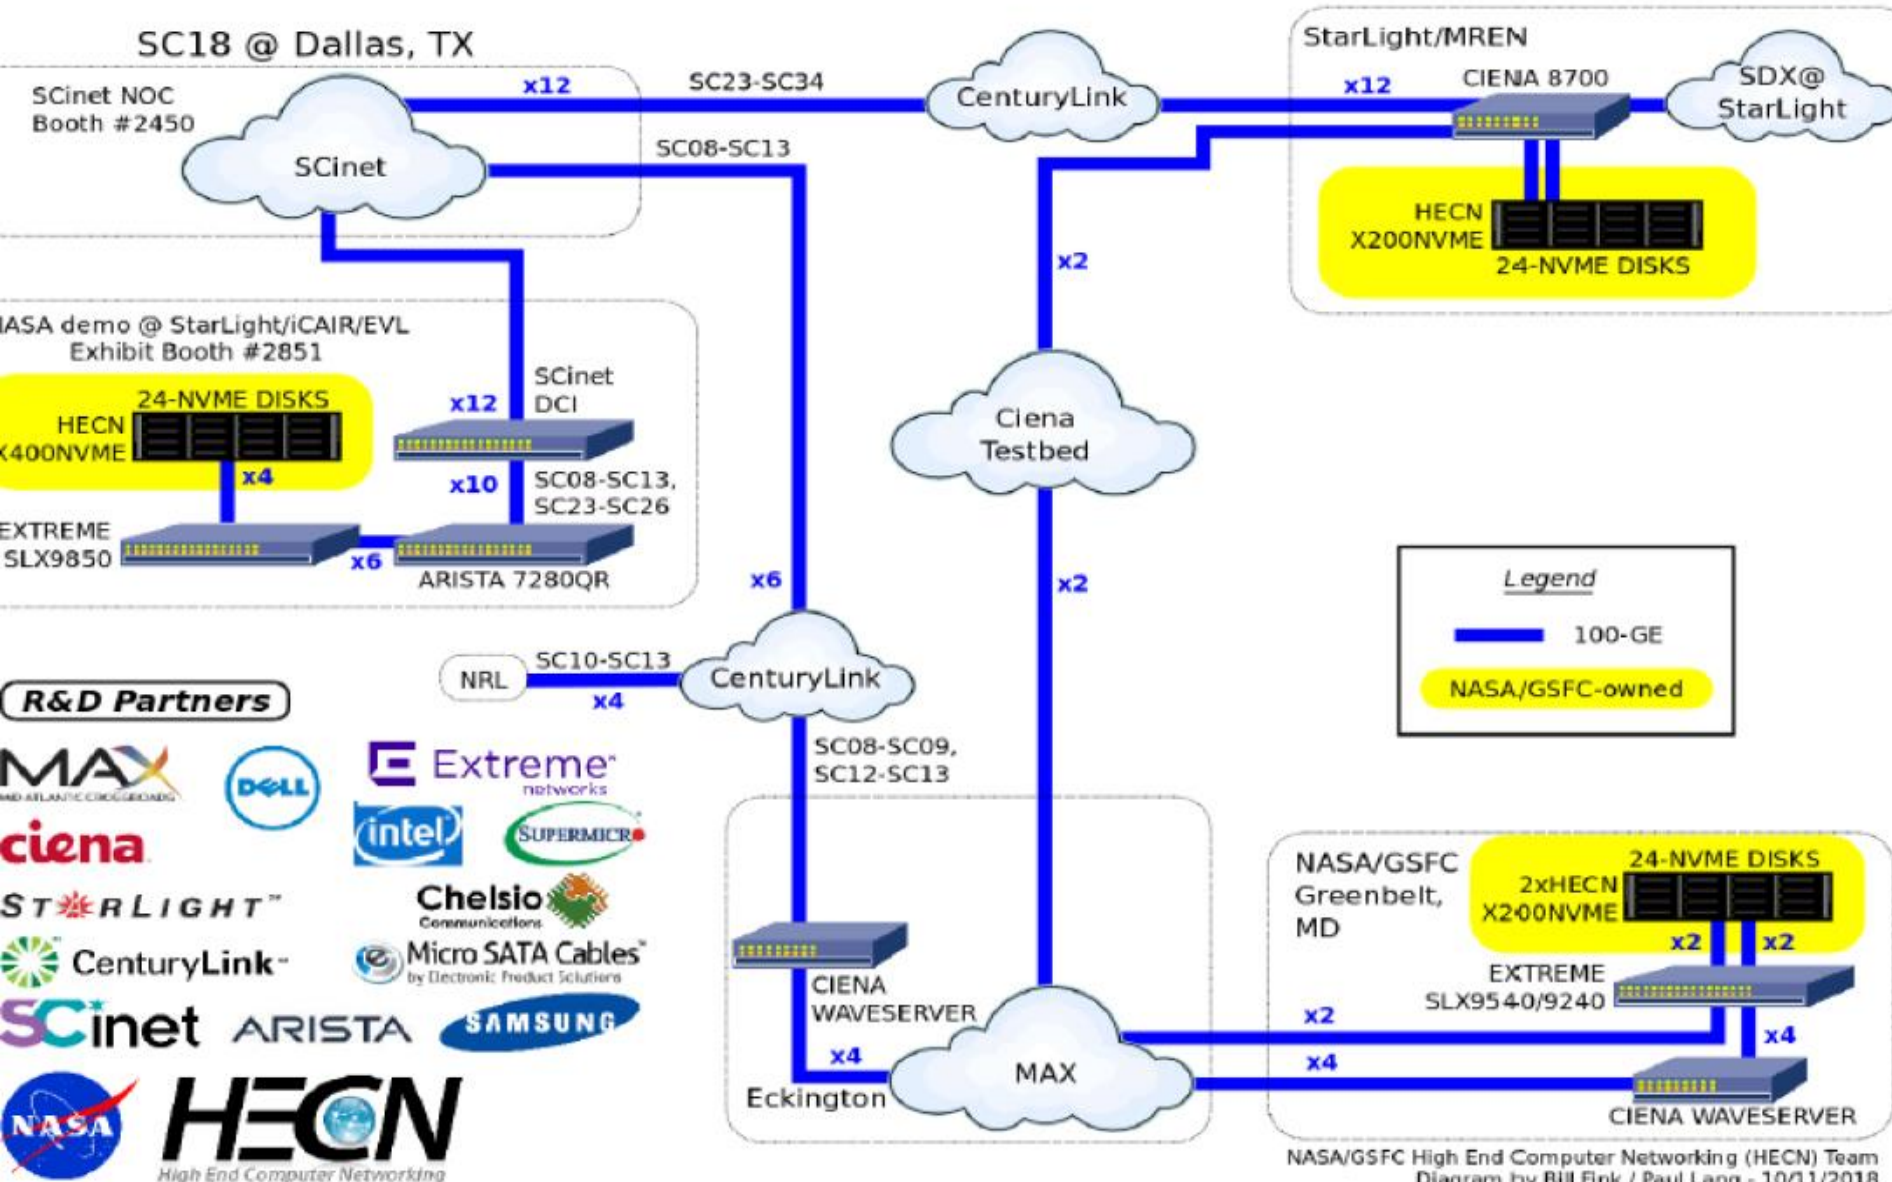
Echostreams

- 400G - LRS
- 100G - CLR4
- 100G - CLR4
- 100G - LR4
- 100G - SR4
- 100G - DAC
- 40G - SR4
- 40G - QSFP+
- 10G
- 10G

SC18

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

An SC18 Collaborative Initiative Among NASA and Several Partners




Implementing a SCinet DTN (DTN-as-a-Service, DaaS)



SC18

X-NET:

SCinet Data Transfer Node(DTN) Service




SCinet

Dallas, TX | hpc inspires.

Source; Jim Chen, iCAIR

STARLIGHTSM

Supercomputing Asia Data Mover Challenge 2020




Composable Platform as a Service

Instrument for Deep Learning & Visualization (COMPaaS DLV)


Using the *Liquid* Composable Infrastructure System to Demonstrate...



...interoperable Science DMZs and Data Transfer Nodes (DTNs) at a national/international scale, using best strategies from the Pacific Research Platform (PRP) over the GLIF infrastructure.



...SAGE2™, the *de facto* operating system for managing Big Data on tiled-display walls, enabling local and remote collaborators to access, share and interact with a variety of digital media.



...on-demand post-processing and visualization of a UIC Mechanical Engineering high-fidelity CFD simulation that, on a supercomputer, generates ~11TB worth of data for four simulations.


Electronic Visualization Laboratory
University of Illinois at Chicago
<https://www.evl.uic.edu>



NSF award CNS-1828265 to UIC/EVL for COMPaaS DLV
<https://www.evl.uic.edu/compaas-dlv>

SAGE2 is trademark of the University of Illinois Board of Trustees

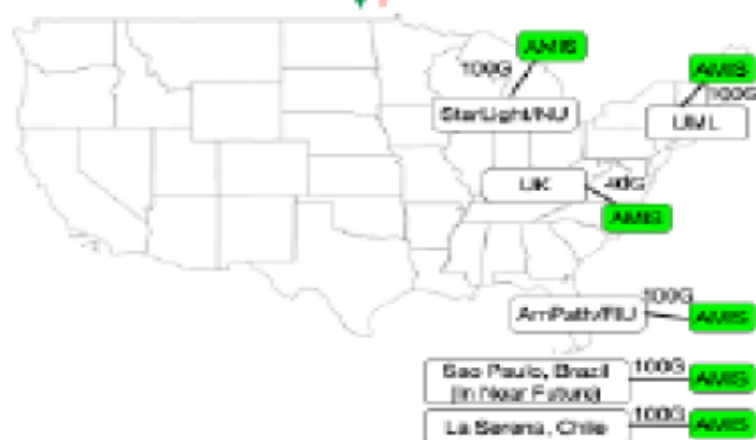

Liquid: <https://www.liquid.com/>
PRP: <https://prp.ucsd.edu>
GLIF: <https://www.glif.is>
SAGE2: <http://sage2.sagecommons.org/>
CFD simulation: <https://cmtl.uic.edu>

Global P4 Experiment Network(G-P4EN)

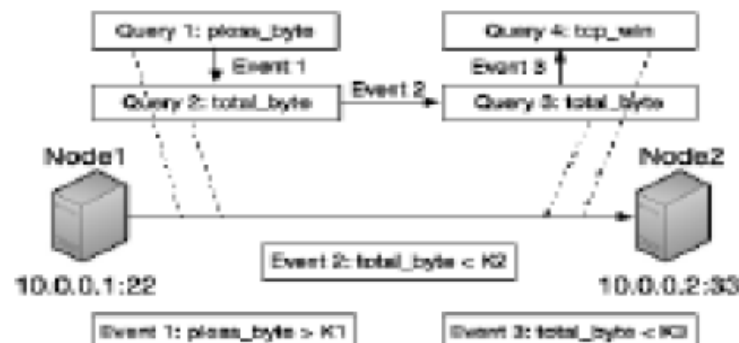
AMIS: Advanced Measurement Instrument and Services for Programmable Network Measurement of Data Intensive Flows

Yan Luo, U. of Massachusetts Lowell; Gabriel Ghinita, U. of Massachusetts Boston;
Cody Bumgardner, U. of Kentucky; Michael McGarry, U. of Texas El Paso

AMIS framework for programmable measurement Equery: event driven measu



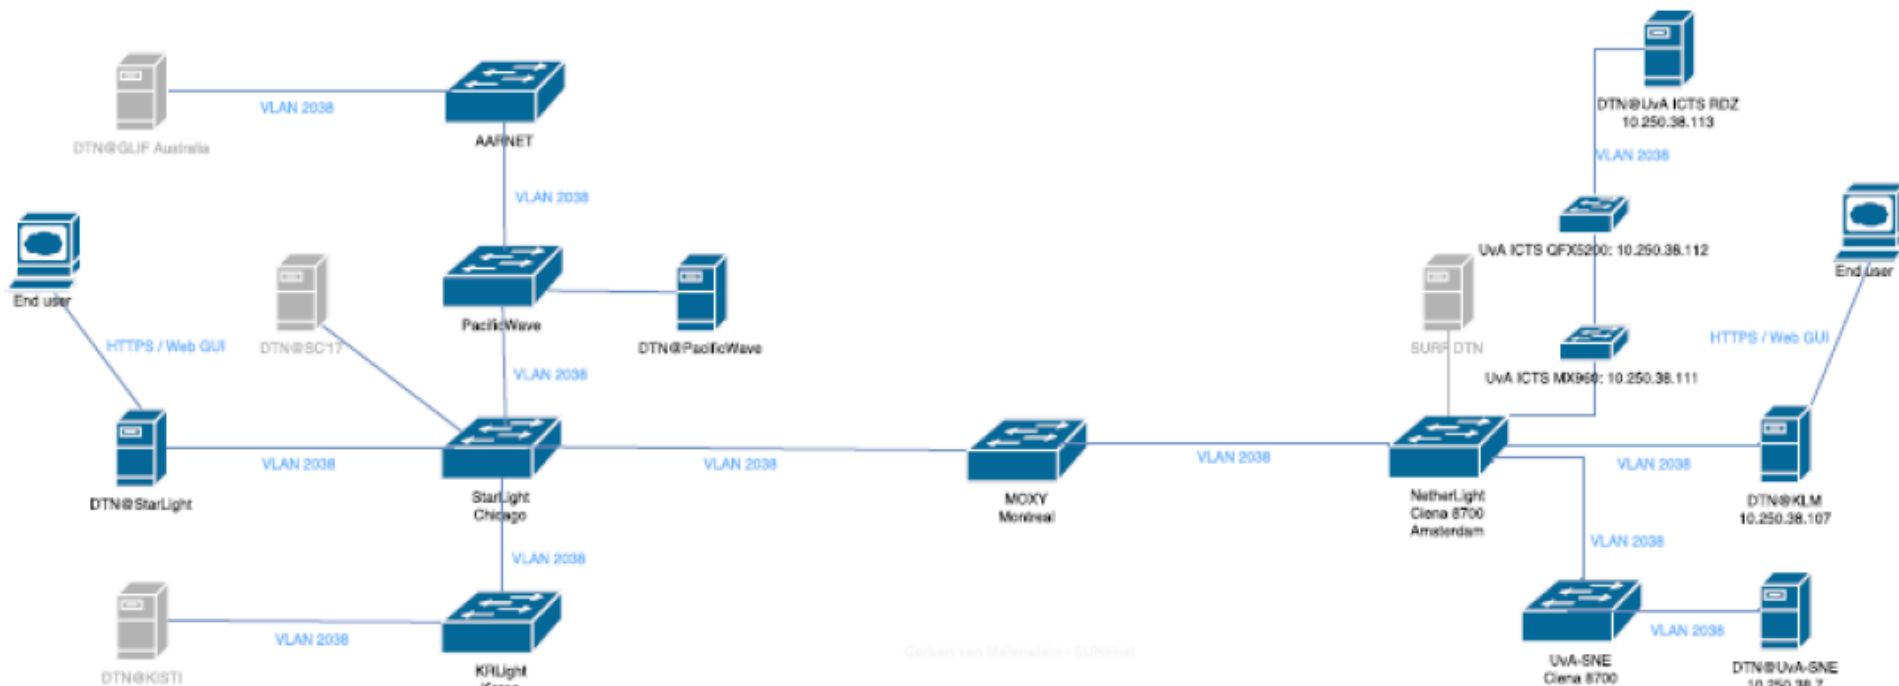
- q1: select ploss_byte where src_addr=10.0.0.1, dst_addr=10.0.0.2, src_port=22, dst_port=33, protocol=TCP, node_id=Node1;
- q2: select total_byte where src_addr=10.0.0.1, dst_addr=10.0.0.2, src_port=22, dst_port=33, protocol=TCP, node_id=Node1 when q1.ploss_byte > K1;
- q3: select total_byte where src_addr=10.0.0.1, dst_addr=10.0.0.2, src_port=22, dst_port=33, protocol=TCP, node_id=Node2 when q2.total_byte < K2;
- q4: select tcp_win where src_addr=10.0.0.1, dst_addr=10.0.0.2, src_port=22, dst_port=33, protocol=TCP, node_id=Node2 when q3.total_byte < K3;



This project is supported by US National Science Foundation (No. 1450937,1450975,1450996,1450997)
Collaborators: Joe Mambretti, Jim Chen and Fei Yeh, StarLight/CAIR/Northwestern University,
Jeronimo Bezerra, Julio Ibarra, AMPATH/Florida International University

Transferring LargeScale Airline Data E2E Across WANs Using DTNs

v5, 21 SEP 2017



Garben van Malen - SURFnet

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/extended
- 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

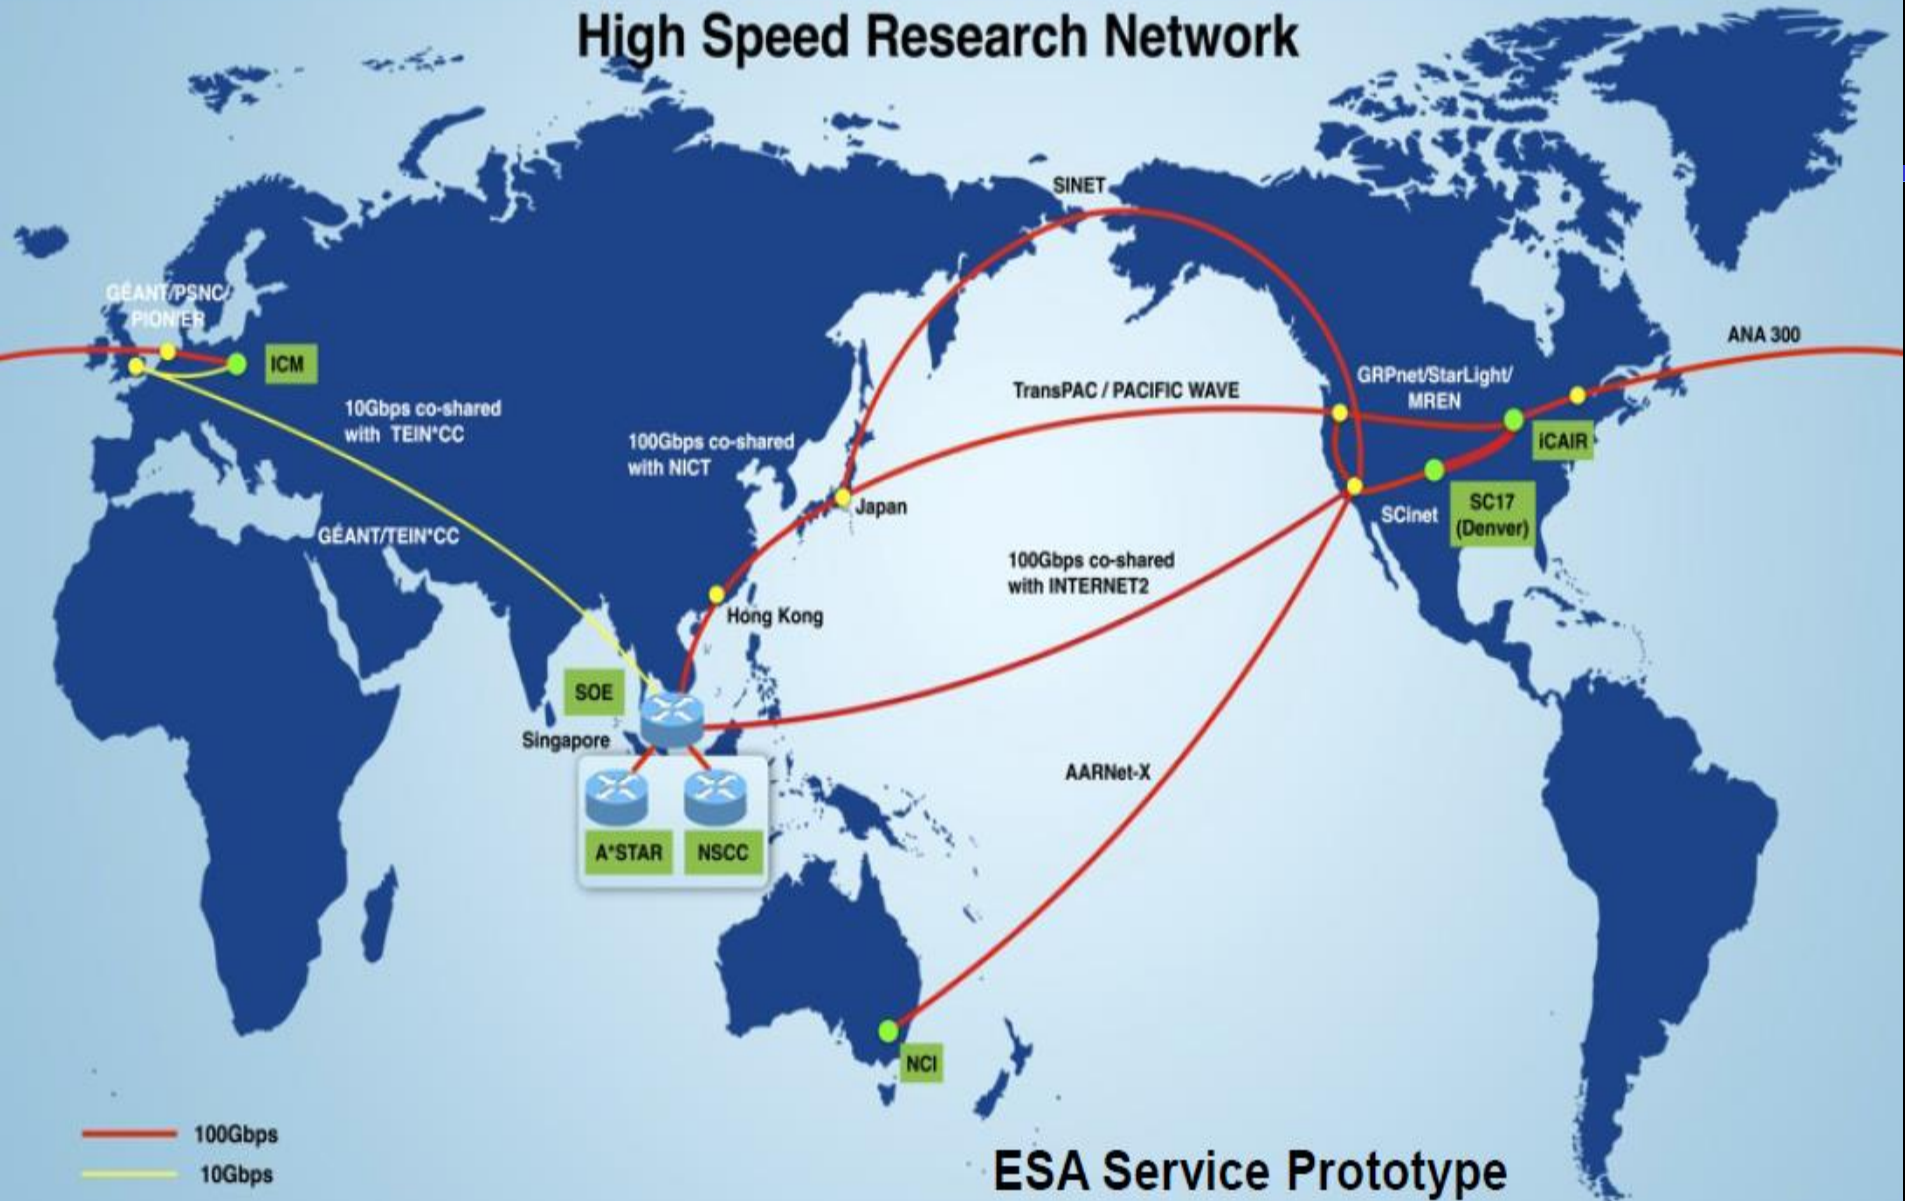
Ideas

- Dutch Research LAN Project

SENSE Orchestration Testbed



High Speed Research Network

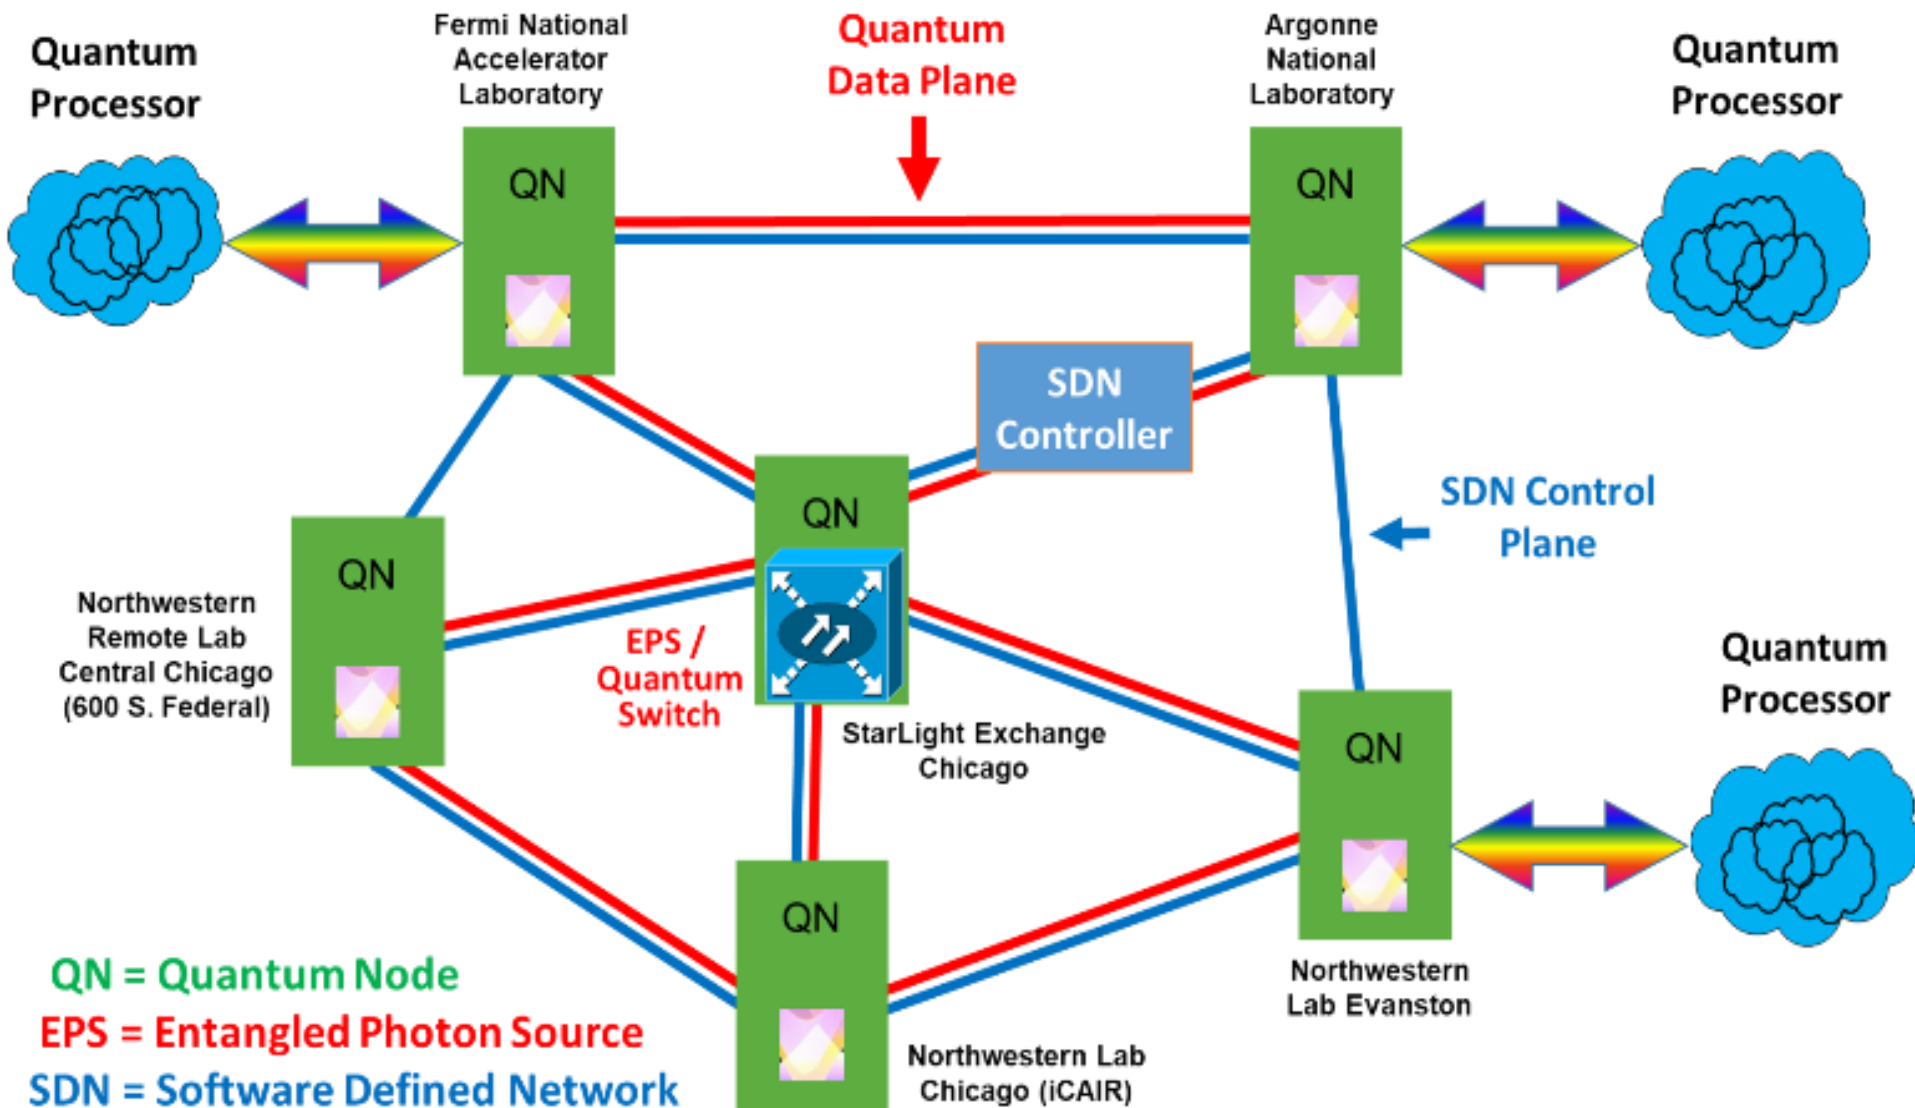


National Quantum Science Initiative

- **Multiple Federal Agency Initiatives**
- **Multiple Consortium Projects**
- **Multiple Topics**
 - **Physics**
 - **Quantum Science**
 - **Information Theory**
 - **Information Science**
 - **Optimization Theory**
 - **Materials Science**
 - **Quantum Components**
 - **Quantum Computing**
 - **Quantum Communications**
 - **Quantum Networking**
- **Multiple Projects: Chicago Quantum Exchange, Illinois Quantum Express Et Al**



Planned Chicago Quantum Exchange Testbed



Ilya Baldine PI, RENCI: FABRIC



Next Generation Distributed Environment For Global Science



GLOBAL RESEARCH PLATFORM





INTERNATIONAL PEERING EXCHANGE



Global Research Platform Network
100 Gbps

StarLight

Global Research Platform

LEGEND


- Pacific Wave POP
- Pacific Research Network (PRN) POP to Asia/US/UK
- 100 Gbps
- 40 Gbps
- 10 Gbps
- Software Defined Network
- Commercial Peering Point

WESTERN REGIONAL NETWORK

- Meridian/Polystar/Pharos
- Aerial/Oregon/Starlink
- RBCP/Colombia and Whyming
- ABOC/Albuquerque GigPop
- LUN Hawaii
- CERN - California

KREONet2 SD-WAN GLORIAD-KR

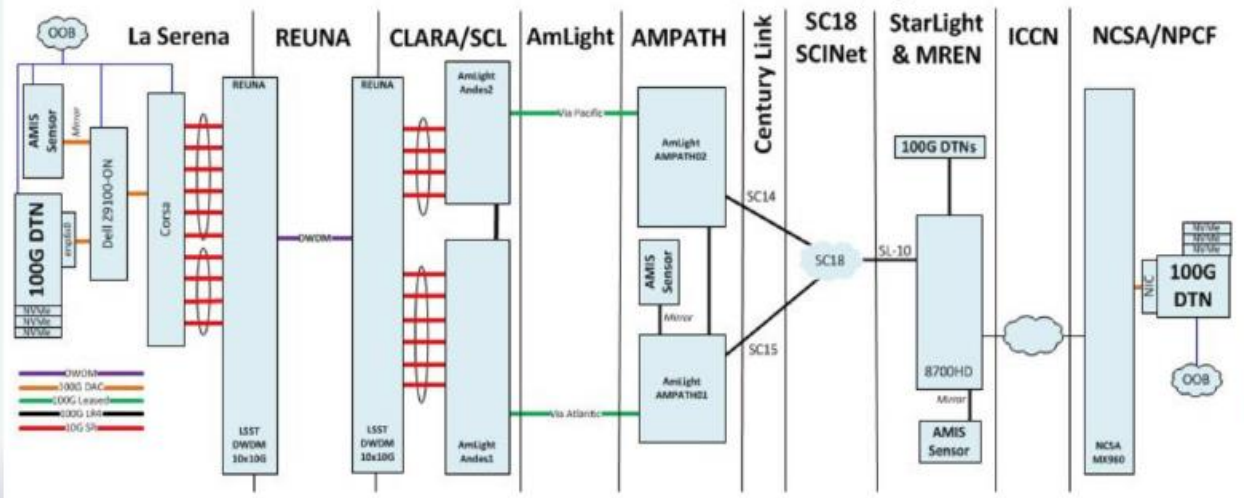
KISTI Daejeon ↔ 100 G ↔ StarLight



Soon: Daejeon SDX ↔ StarLight SDX



LSST/SC18 Demo Topology



www.startup.net/starlight

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

STARLIGHTSM