Next Generation Data Transfer Nodes (DTNs) For Global Science: Architecture, Technology, Enabling Capabilities

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)

TNC 2018
Trondheim, Norway
June 10-14, 2018

Introduction to iCAIR:

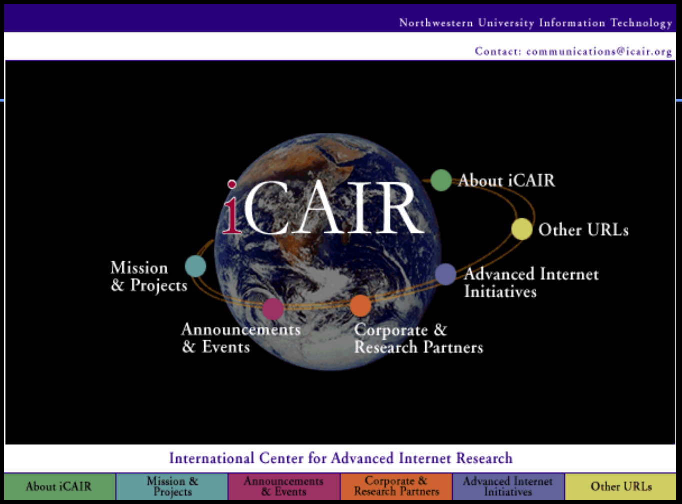

Accelerating Leading Edge Innovation and Enhanced Global Communications through Advanced Internet Technologies, in Partnership with the Global Community

- **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 STARLIGHT℠

Topics

- Motivation For Data Transfer Nodes (DTNs)
- Science DMZ 101
- Regional/National/Global DMZs (i.e., Regional, National and Global Research Platforms)
- Programmable Network Techniques and Devices (DTNs In Context)
- DTN 101
- Current DTNs
- Emerging Next Generation DTNs
- Results of Recent DTN Research Experiments
- Future DTNs
- Conclusions

Motivations

- Data Intensive Science - 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
- Also, However,
- AI Training/DL
- Distributed Storage Systems
- NFV
- Low Latency Services and Applications (Increasing Robotics)
- Low Latency <u>And</u> High Capacity/Big Data Applications

GriPhyN
Data Intensive Science

NEPTUNE
A Fiber-Optic Telescope to Inner Space

NEES

GENOMES to LIFE
BIOLOGICAL SOLUTIONS FOR ENERGY CHALLENGES
U.S. DEPARTMENT OF ENERGY

NEON
National Ecological Observatory Network

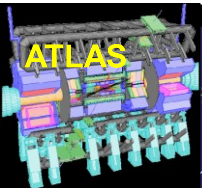

ALMA

ALMA: Atacama Large Millimeter Array

ANDRILL: Antarctic Geological Drilling
www.andrill.org

BIRN

BIRN: Biomedical Informatics Research Network
www.nbirn.net

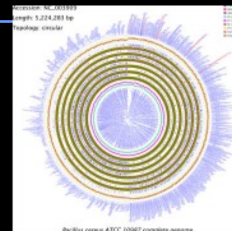

CAMERA metagenomics
camera.calit2.net

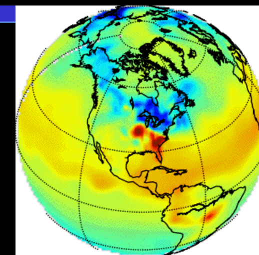

Carbon Tracker
www.esrl.noaa.gov/gmd/ccgg/carbontrack

CineGrid

CineGrid
www.cinegrid.org

LHC

LHCONE
www.lhcone.net

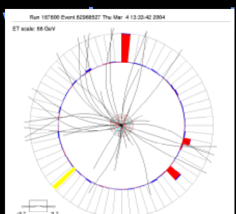

DØ (DZero)
www-d0.fnal.gov

GEON

GEON: Geosciences Network
www.geongrid.org

GLEON

GLEON: Global Lake Ecological Observatory Network

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

OOI-CI
ci.oceanobservatories.org

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

CLASS

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

PRAGMA

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

IVOA

IVOA: International Virtual Observatory
www.ivoa.net

LIGO
www.ligo.org

WLCG
Worldwide LHC Computing Grid

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

TeraGrid

TeraGrid
www.teragrid.org

Open Science Grid

OSG
www.opensciencegrid.org

the globus alliance

Globus Alliance
www.globus.org

SKA
SQUARE KILOMETRE ARRAY

SKA
www.skatelescope.org

Sloan Digital Sky Survey
www.sdss.org

XSEDE

XSEDE
www.xsede.org

Compilation By Maxine Brown

STARLIGHT℠

Science DMZ 101– Segmentation Of Data Intensive Flows (Ref:fasterdata.es.net)

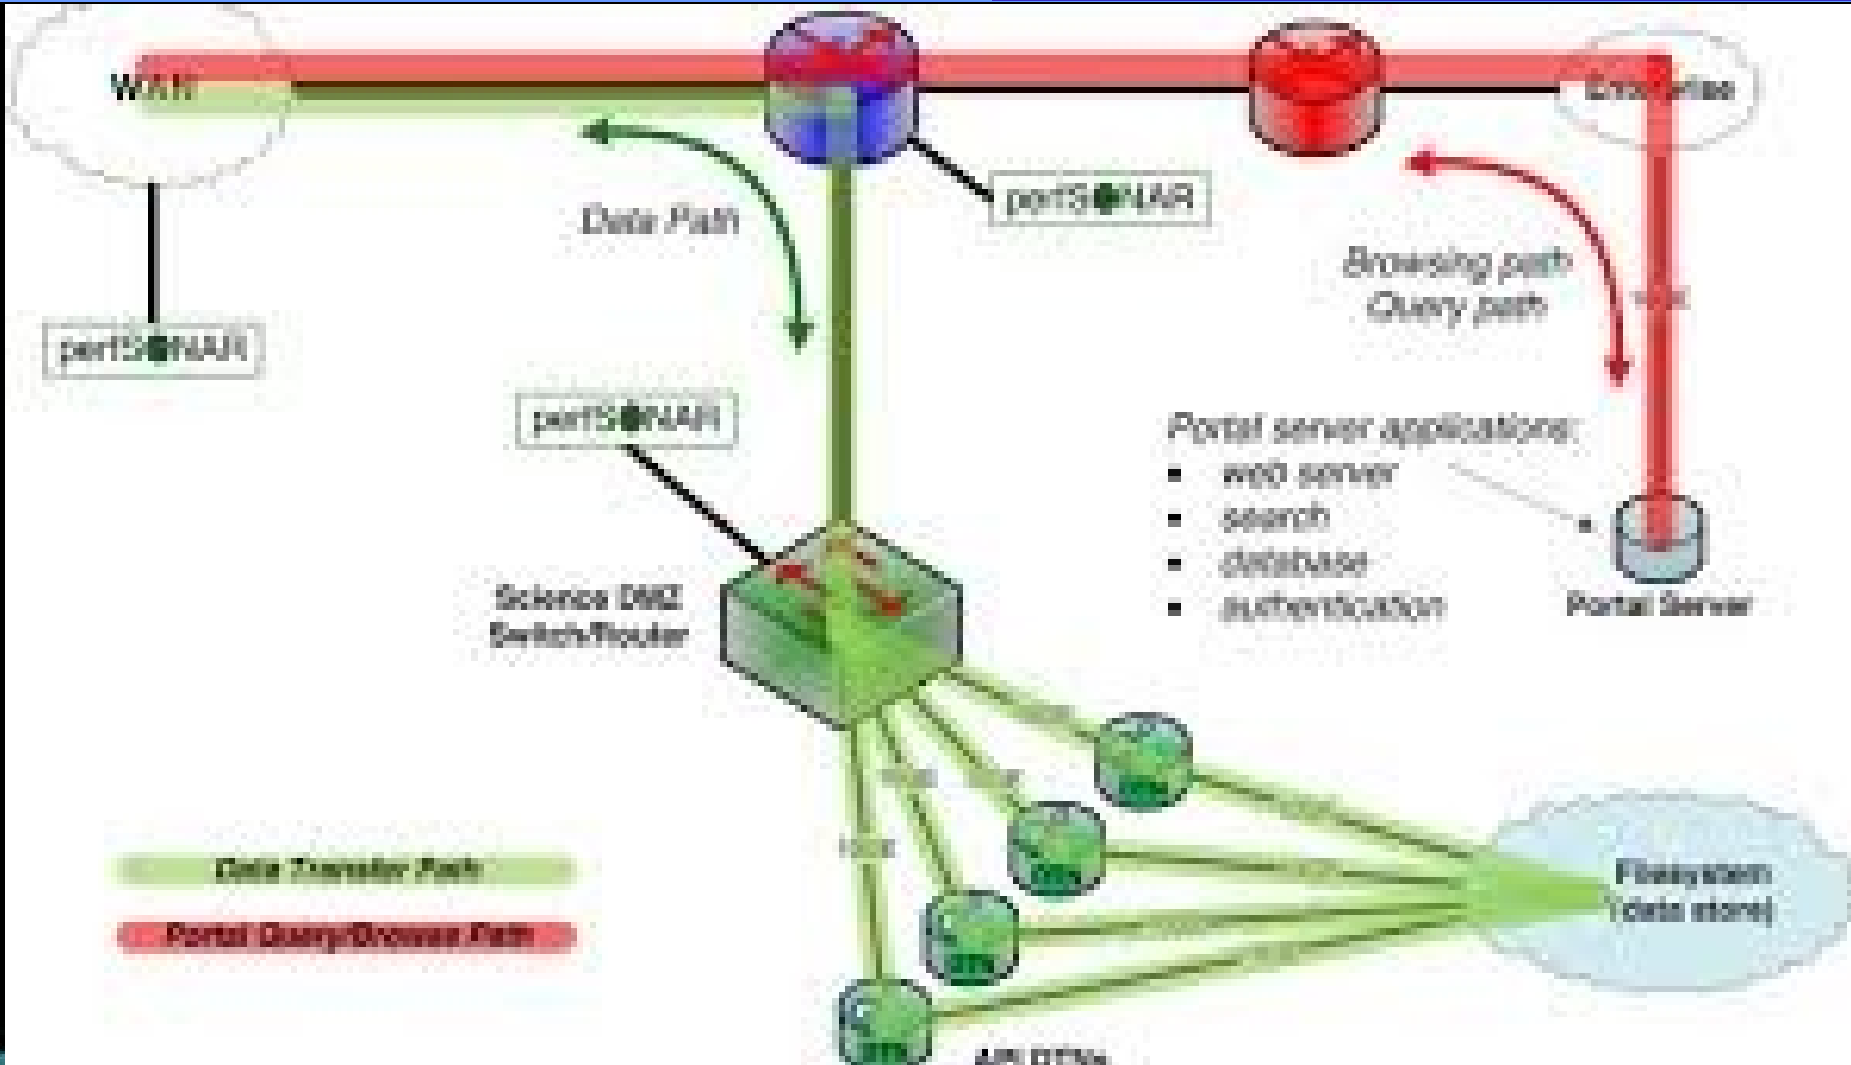

Energy Science Network (ESnet Innovation)

STARLIGHT℠

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

Current International GRP Partners

Programmable Network Techniques and Devices (DTNs In Context)

- Kernel By Pass
- SDN/SDX & Network Programming Languages
- Programmable Switch ASICs
- Programmable Network Processors
- FPGAs
- Smart/Programmable NICs
- Ref: Barefoot Tofino, Intel FlexPipe, Cavium XPliant, Netronome Agilio.
- P4 Based In-Network Telemetry
- AI/ML/DL Integrated With Network Programming
- RDMA (Remote Direct Memory Access)
- DPDK (Data Plane Development Kit)
- New Backplane/Switch Fabrics
- RoCE (RDMA Over Converged Ethernet)

DTN 101

- DTNs Are Purpose Built Servers Optimized For Data Transfer
- No "Standard" DTN Exists: Many Designs/Configs/ Architectures Have Been Implemented (e.g. Memory Intensive, Storage Intensive, Memory and Storage Intensive, Neither Memory or Storage Intensive, 1 G, 10 G 40 G, 100 G, Edge Device, Intermediate Device, WAN Device)
- Multiple OS Stacks (Custom Tuned)
- Multiple DTN Middleware Stacks
- Multiple Transport Protocols
- Multiple Environmental Contexts (Ref Previous Slide)

Flash I/O Network Appliance (FIONA)

- FIONAs PCs [ESnet DTNs]:
  - ~$8,000 Big Data PC with:
    - 10/40 Gbps Network Interface Cards
    - 3 TB SSDs
  - Higher Performance at higher cost:
    - +NVMe SSDs & 100Gbps NICs Disk-to-Disk
    - +Up to 8 GPUs [4M GPU Core Hours/Week]
    - +Up to 196 TB of Disks used as Data *Capacitors*
    - +Up to 38 Intel CPU cores or AMD Epyc cores
  - US$1,100 10Gbps FIONA (if 10G is fast enough)
- FIONettes are US$300 EL-30-based FIONAs
  - 1Gbps NIC With USB-3 for Flash Storage or SSD
  - Perfect for Training and smaller campuses

Phil Papadopoulos, SDSC &
Tom DeFanti, Joe Keefe & John Graham, Calit2

PRP PACIFIC RESEARCH PLATFORM

SDSC CITRIS BANATAO INSTITUTE

STARLIGHT℠

Emerging Next Gen DTNs

- New Experimental Research Using DTNs
- iCAIR Innovations In DTNs For Data Intensive Science (Current: Developing 5th Gen)
- Focus Is On 100 Gbps And Multiple 100 Gbps
- Agency Innovations: NRL, NASA GSFC, ESnet
- Advanced Industry Components (e.g., NUMA, NVMe, Path Fabrics)
- Design Innovations (Thread Management, Optimal Affinity Bindings, NUMA Optimization)

Non-Volatile Memory Express (NVMe)

- Non-Volatile Memory Host Controller Interface (NVMHCI) Work Group
- Standard Architecture Specification For PCIe SSDs (Designed Specifically For faster Devices vs Traditional)
  - Register Interface
  - Streamlined Commands
  - Attributes
- Optimizes Host ⇔ Storage
- NVMe Sends I/O Commands/Results To Shared Memory In Server via PCIe Interface
- Parallel I/O Using Multicore Processors
- One Message Queue Supports: 64,000 Commands
- Supports 65,535 I/O Queues

SUPERMICRO 24X NVMe SUPER SERVER

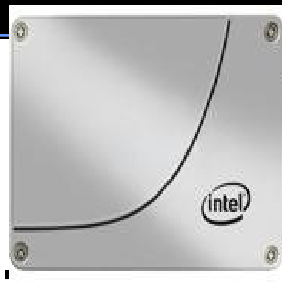

NVMe Type A: 8 X Intel P3700 800G

NVMe Type B: 8 X SamSung 950 512G

+ M.2 to U.2 Adopter

Dell 14G Solution Configuration (Co-Research & Development in Collaboration with Dell)

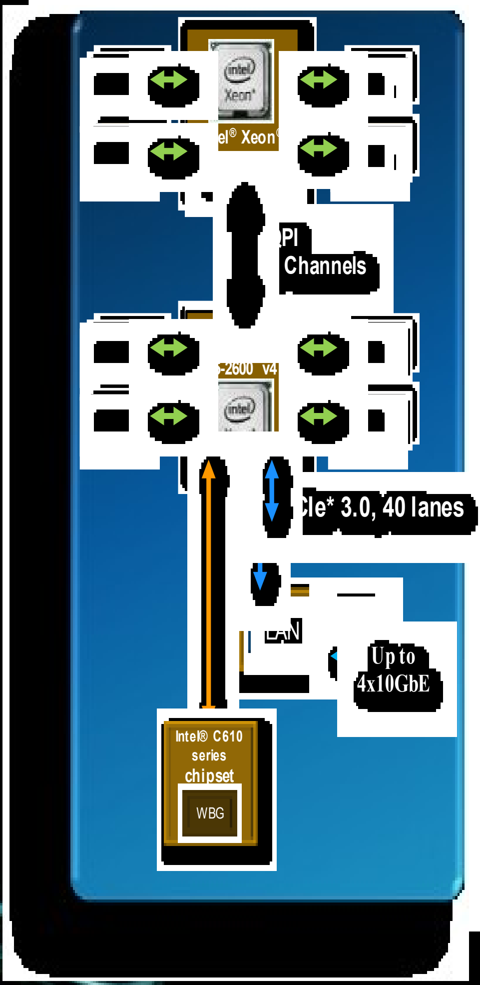

PowerEdge R740XD Server

2 X Intel® Xeon® Gold 6136 3.0G,12C/24T,10.4GT/s 2UPI,24.75M Cache,Turbo,HT (150W)

192G DDR4-2666

PCI-e Configuration Investigation:

2 X Mellanox ConnectX-5 100GE VPI

4 X Kingston/Liqid AIC NVMe PCI-e X8 SSD Drives

Optional SAS/SATA Drives

3G / 3G + Clock Is Critically Important

STARLIGHT℠

Recent DTN WAN Experimental Research

- Motivated By Large Scale Science
- LHC
- LSST
- ESA
- Geophysical Sciences
- Genomics
- BioInformatics
- Precision Medicine
- Etc

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

iCAIR

STARLIGHT℠

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

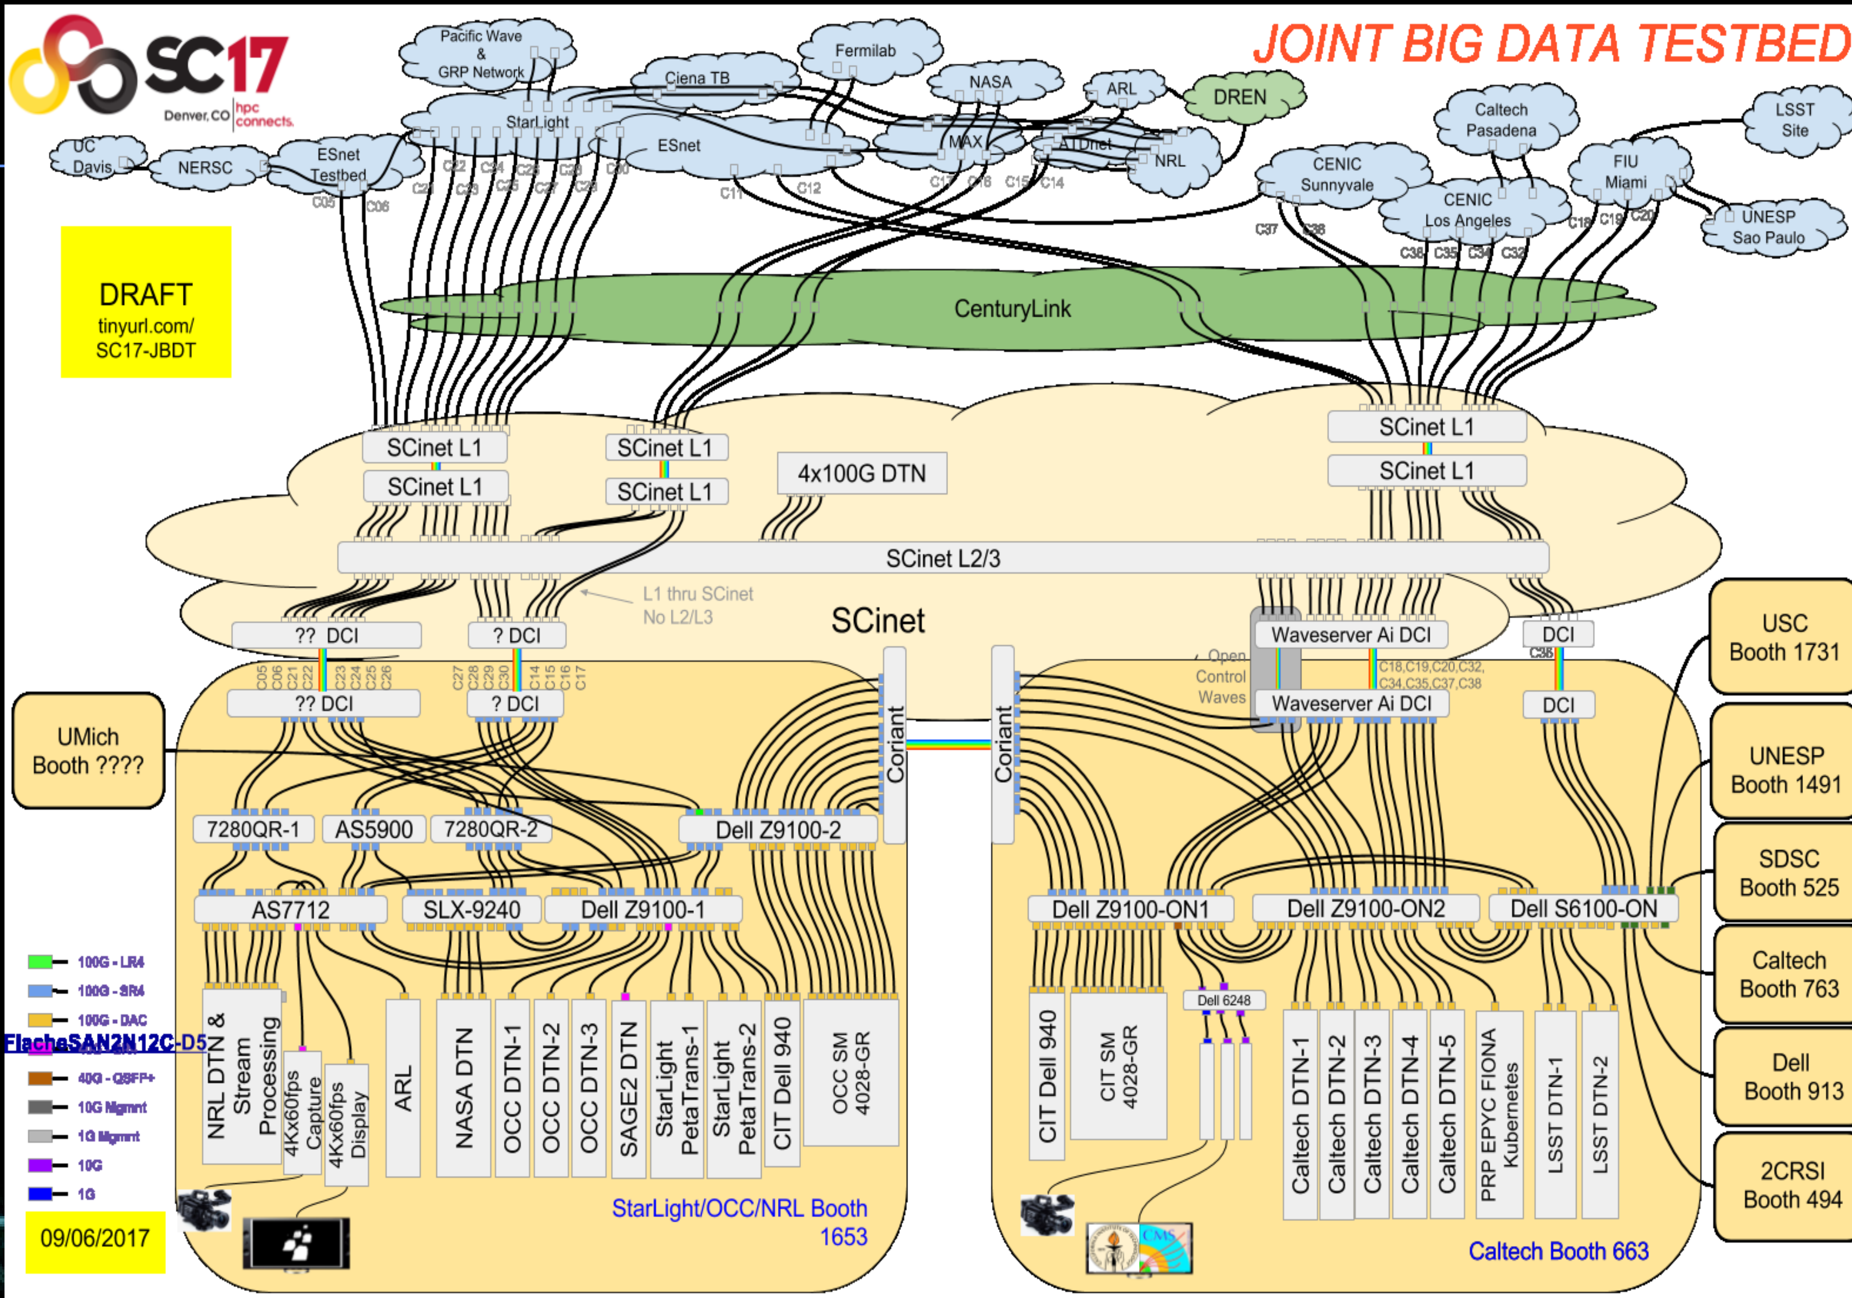

SC17
Denver, CO
hpc connects.
JOINT BIG DATA TESTBED
DRAFT
tinyurl.com/
SC17-JBDT
Pacific Wave & GRP Network
Fermilab
Ciena TB
NASA
ARL
DREN
StarLight
ESnet
MAX
ATDnet
NRL
UC Davis
NERSC
ESnet Testbed
Caltech Pasadena
LSST Site
CENIC Sunnyvale
CENIC Los Angeles
FIU Miami
UNESP Sao Paulo
CenturyLink
SCinet 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
Waveserver Ai DCI
DCI
Open Control Waves
Waveserver Ai DCI
DCI
?? DCI
? DCI
Coriant
Coriant
UMich
Booth ????
USC
Booth 1731
UNESP
Booth 1491
SDSC
Booth 525
Caltech
Booth 763
Dell
Booth 913
2CRSI
Booth 494
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
StarLight/OCC/NRL Booth 1653
Caltech Booth 663
100G - LR4
100G - SR4
100G - DAC
FlachaSAN2N12C-D5
40G - QSFP+
10G Mgmt
1G Mgmt
10G
1G
09/06/2017

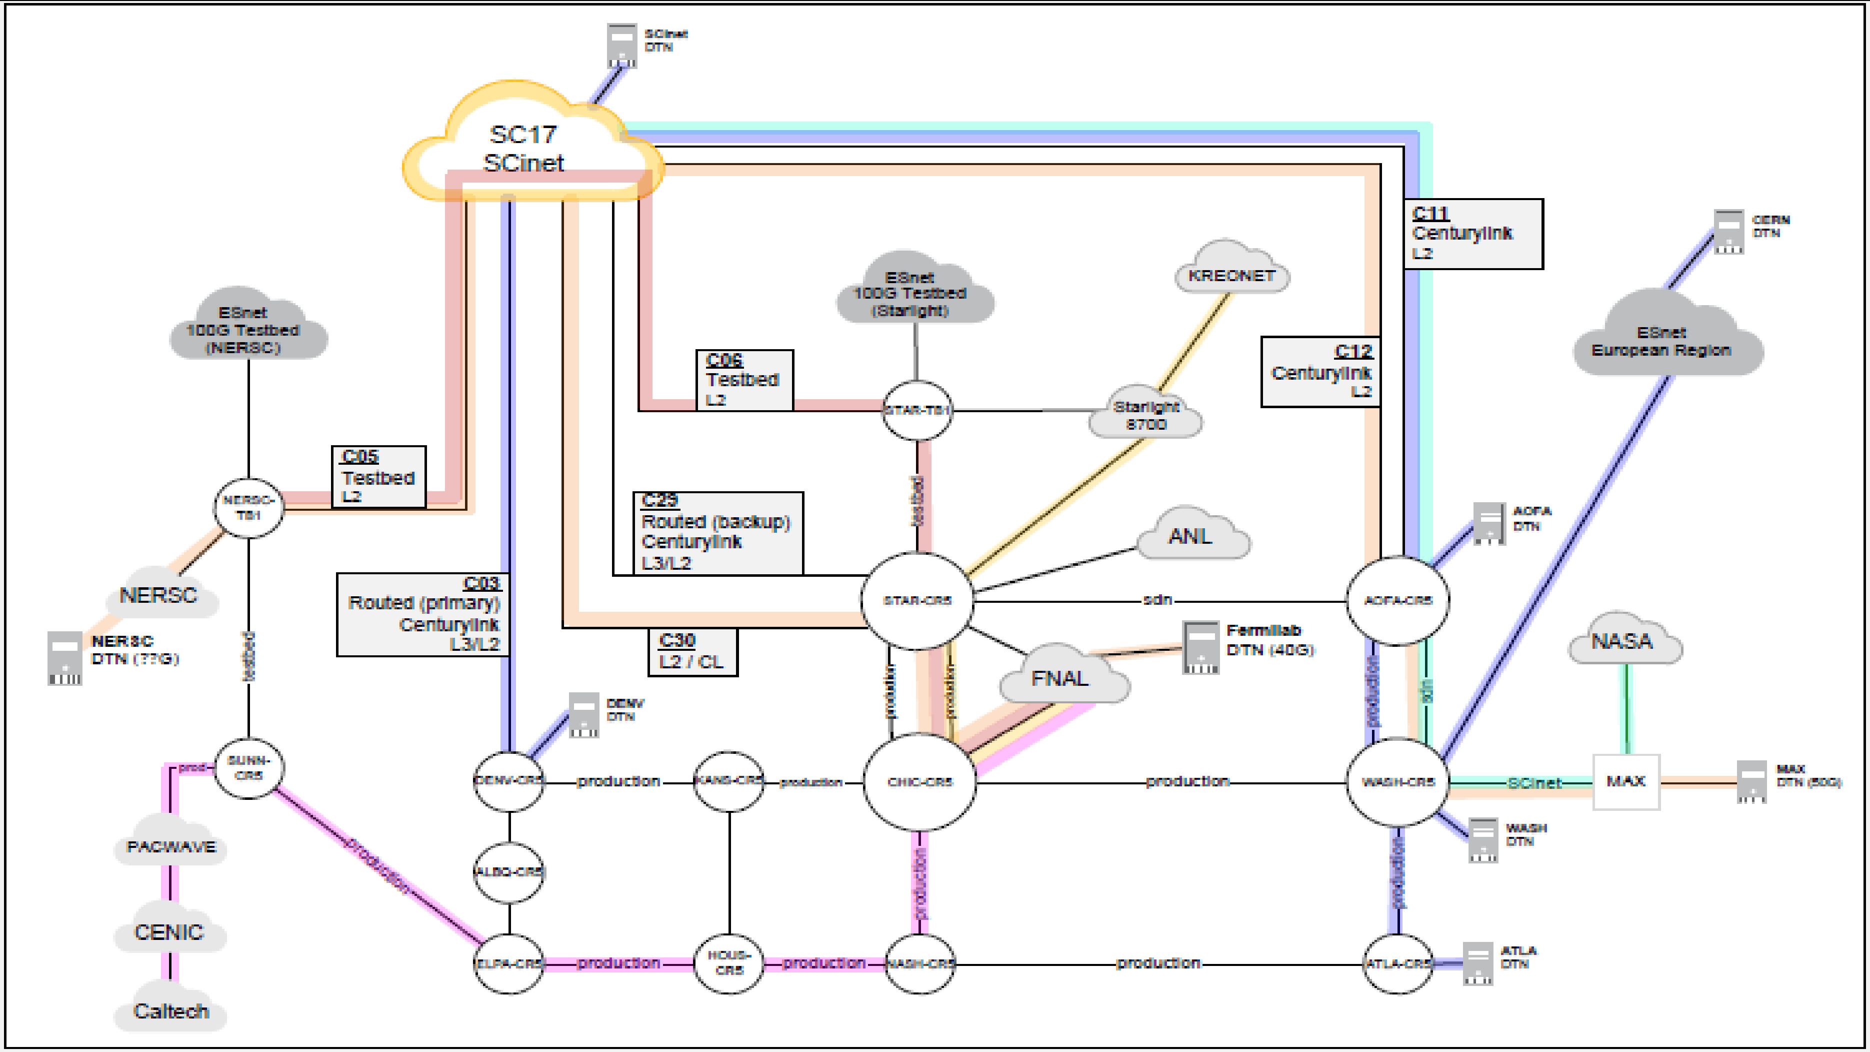

- NASA GODDARD 400G Disk Transfer (100G ESnet path, not active during show)
- SENSE (40G FNAL, 50G MAX)
- FNAL BigData Experess AmoebaNet (2x10G OSCARS)
- BDE FNAL KISTI SDN (2x1G vlans 1662, 1663)
- HEPcloud (10G)(vlan 3602)
- Calibers (1G vlan TBD)

ESnet

SC17 demos – ESnet

Zach Harlan, ESnet 9/26/2017

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

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

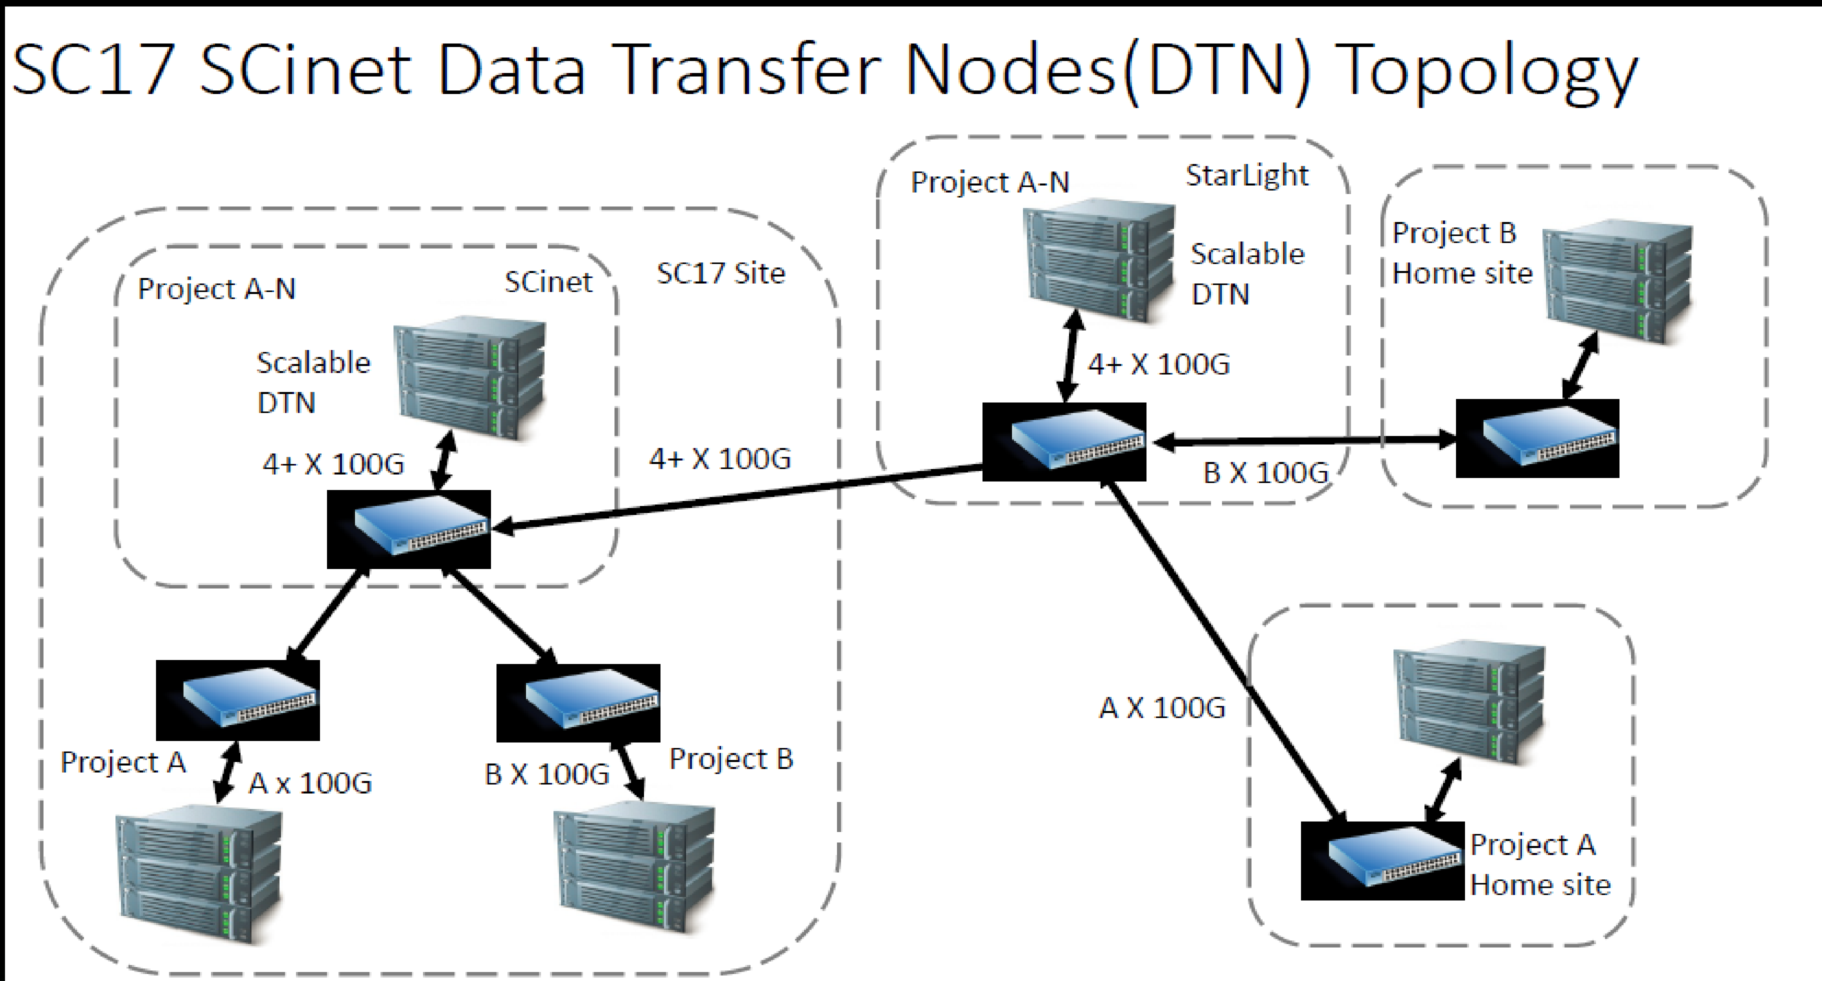

Source; Jim Chen, iCAIR

PetaTrans: Petascale Sciences Data Transfer

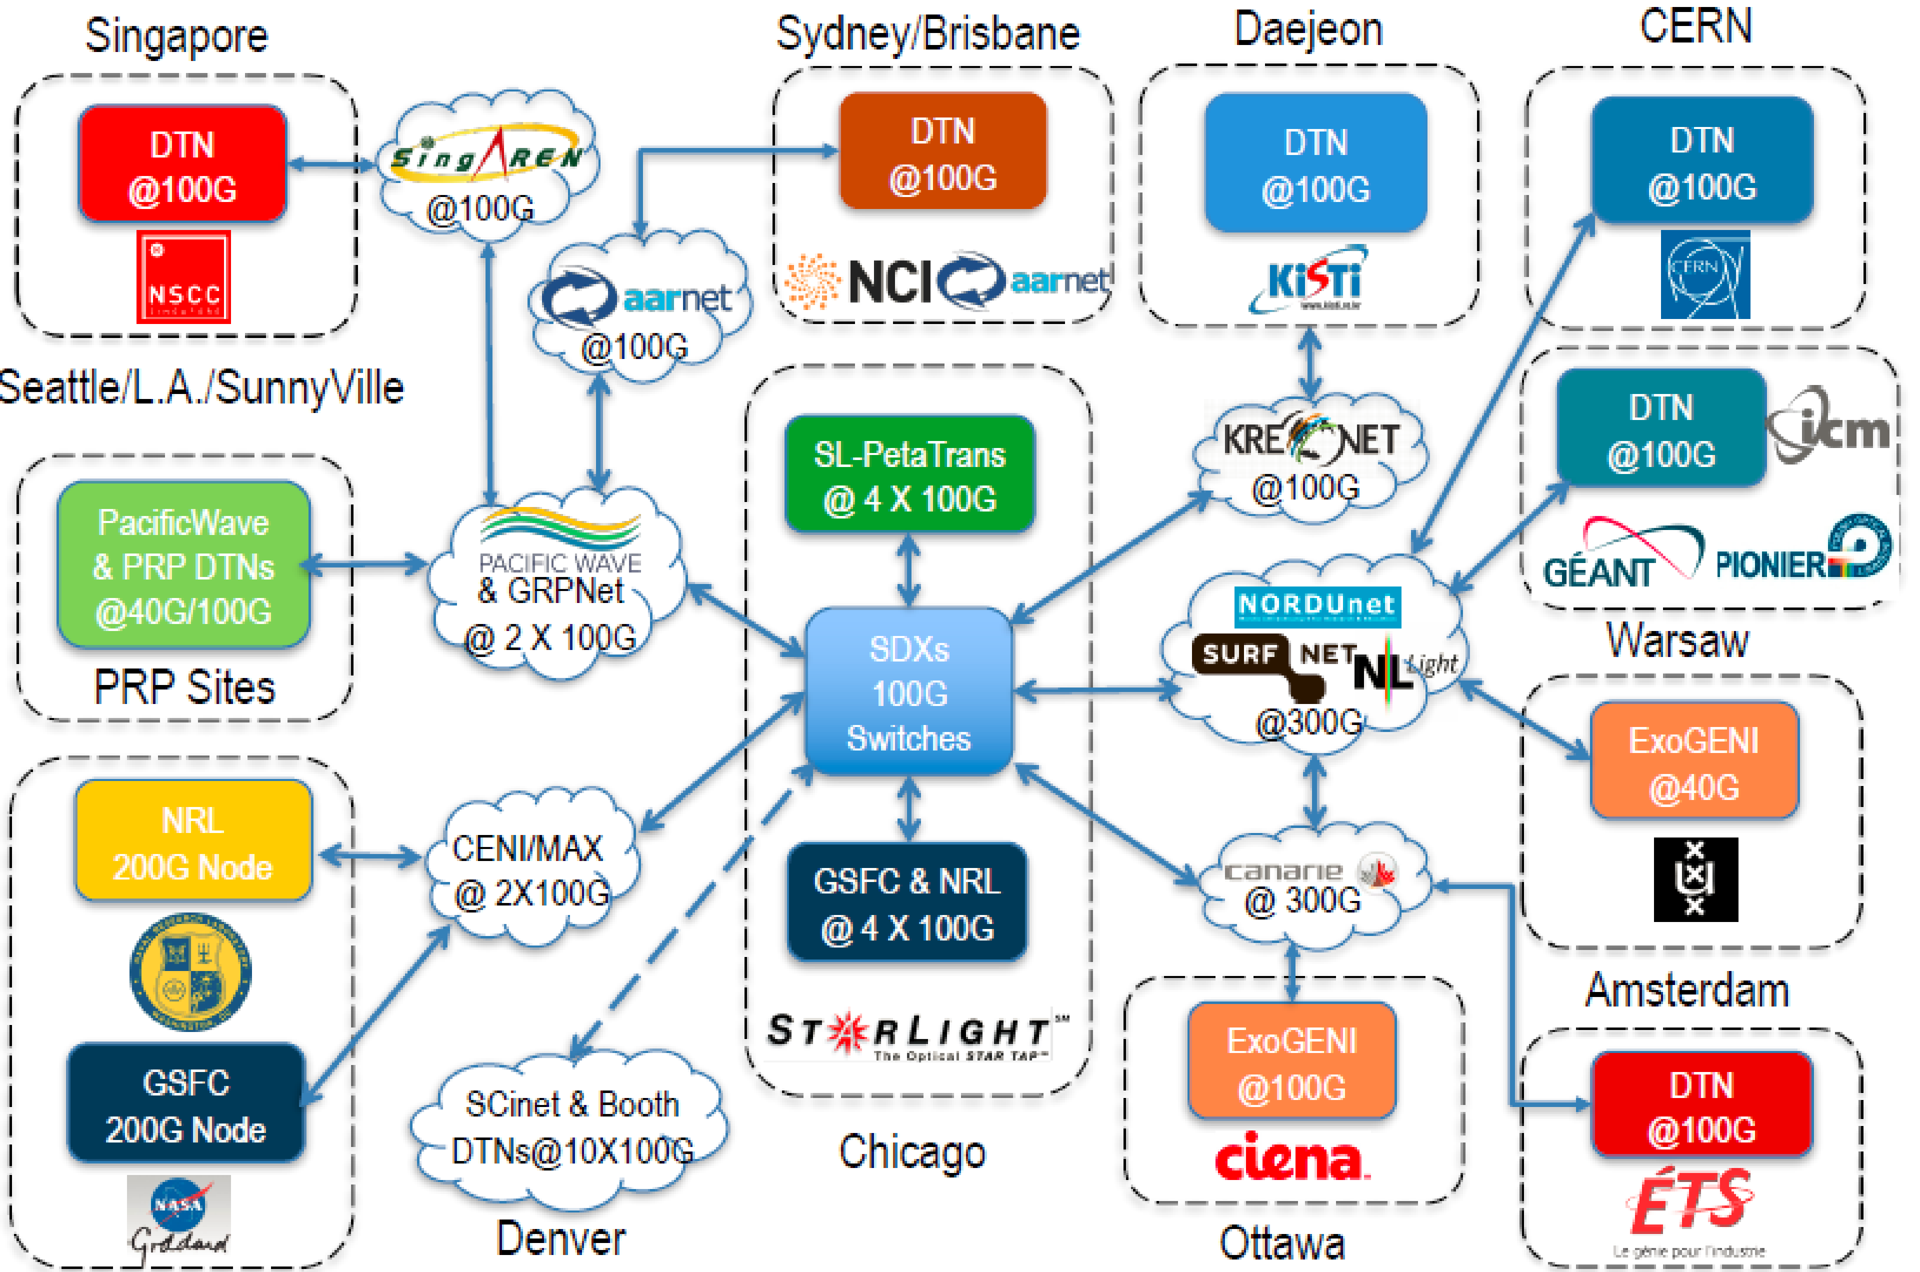

Persistence 100G DTN Services Beyond SC17

A Disaggregated SCinet Optical Layer

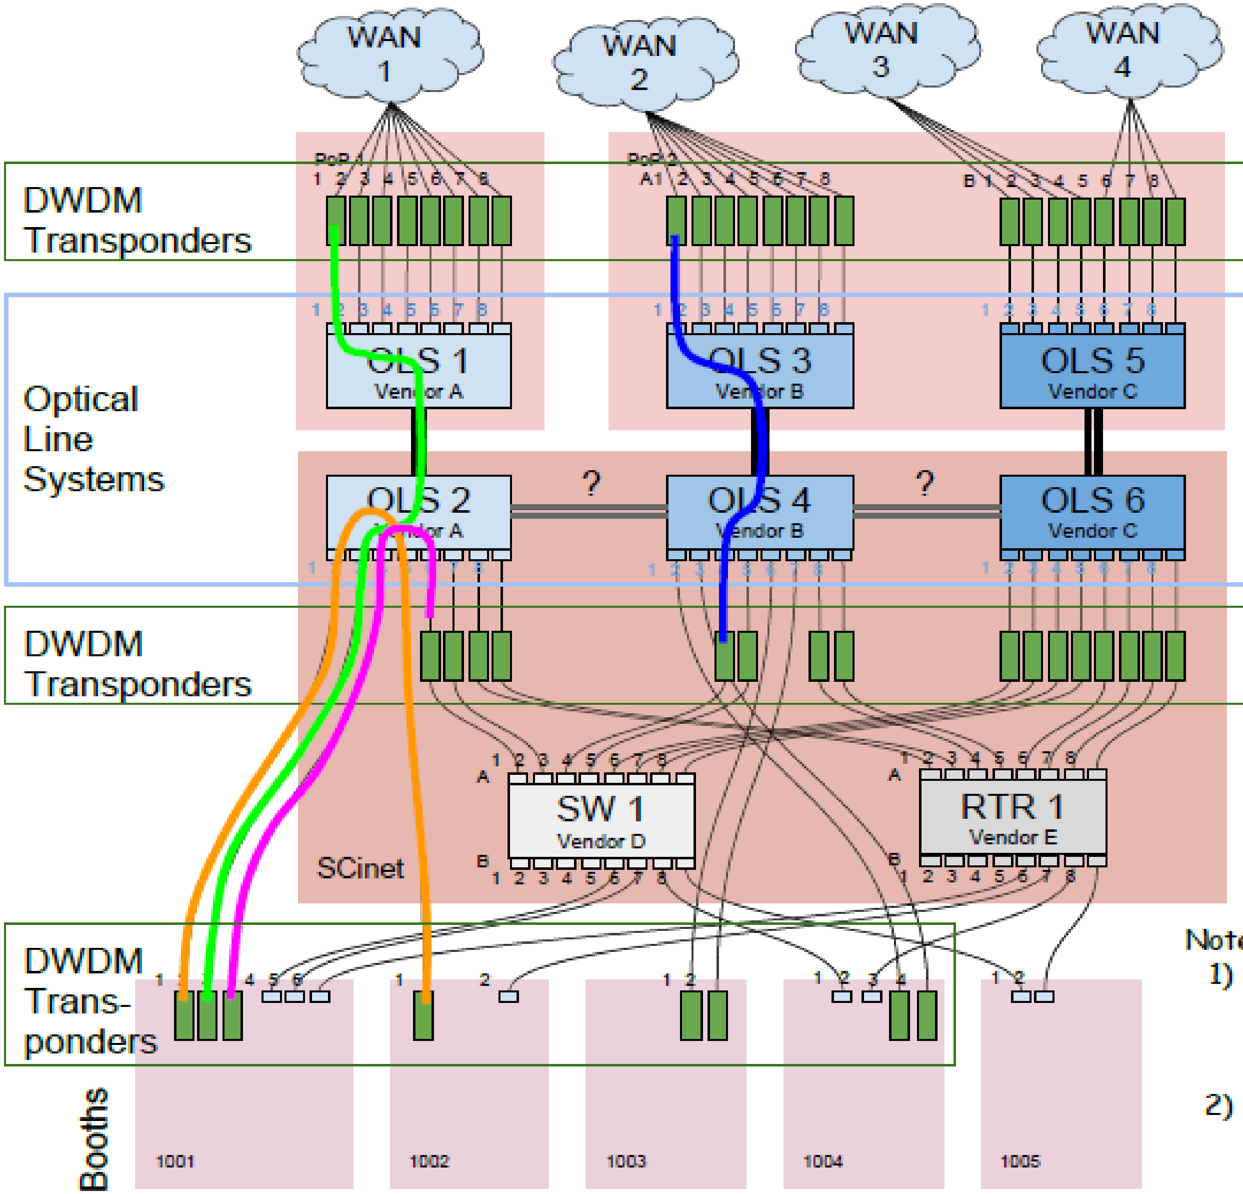

Reconfiguration options

- 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

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

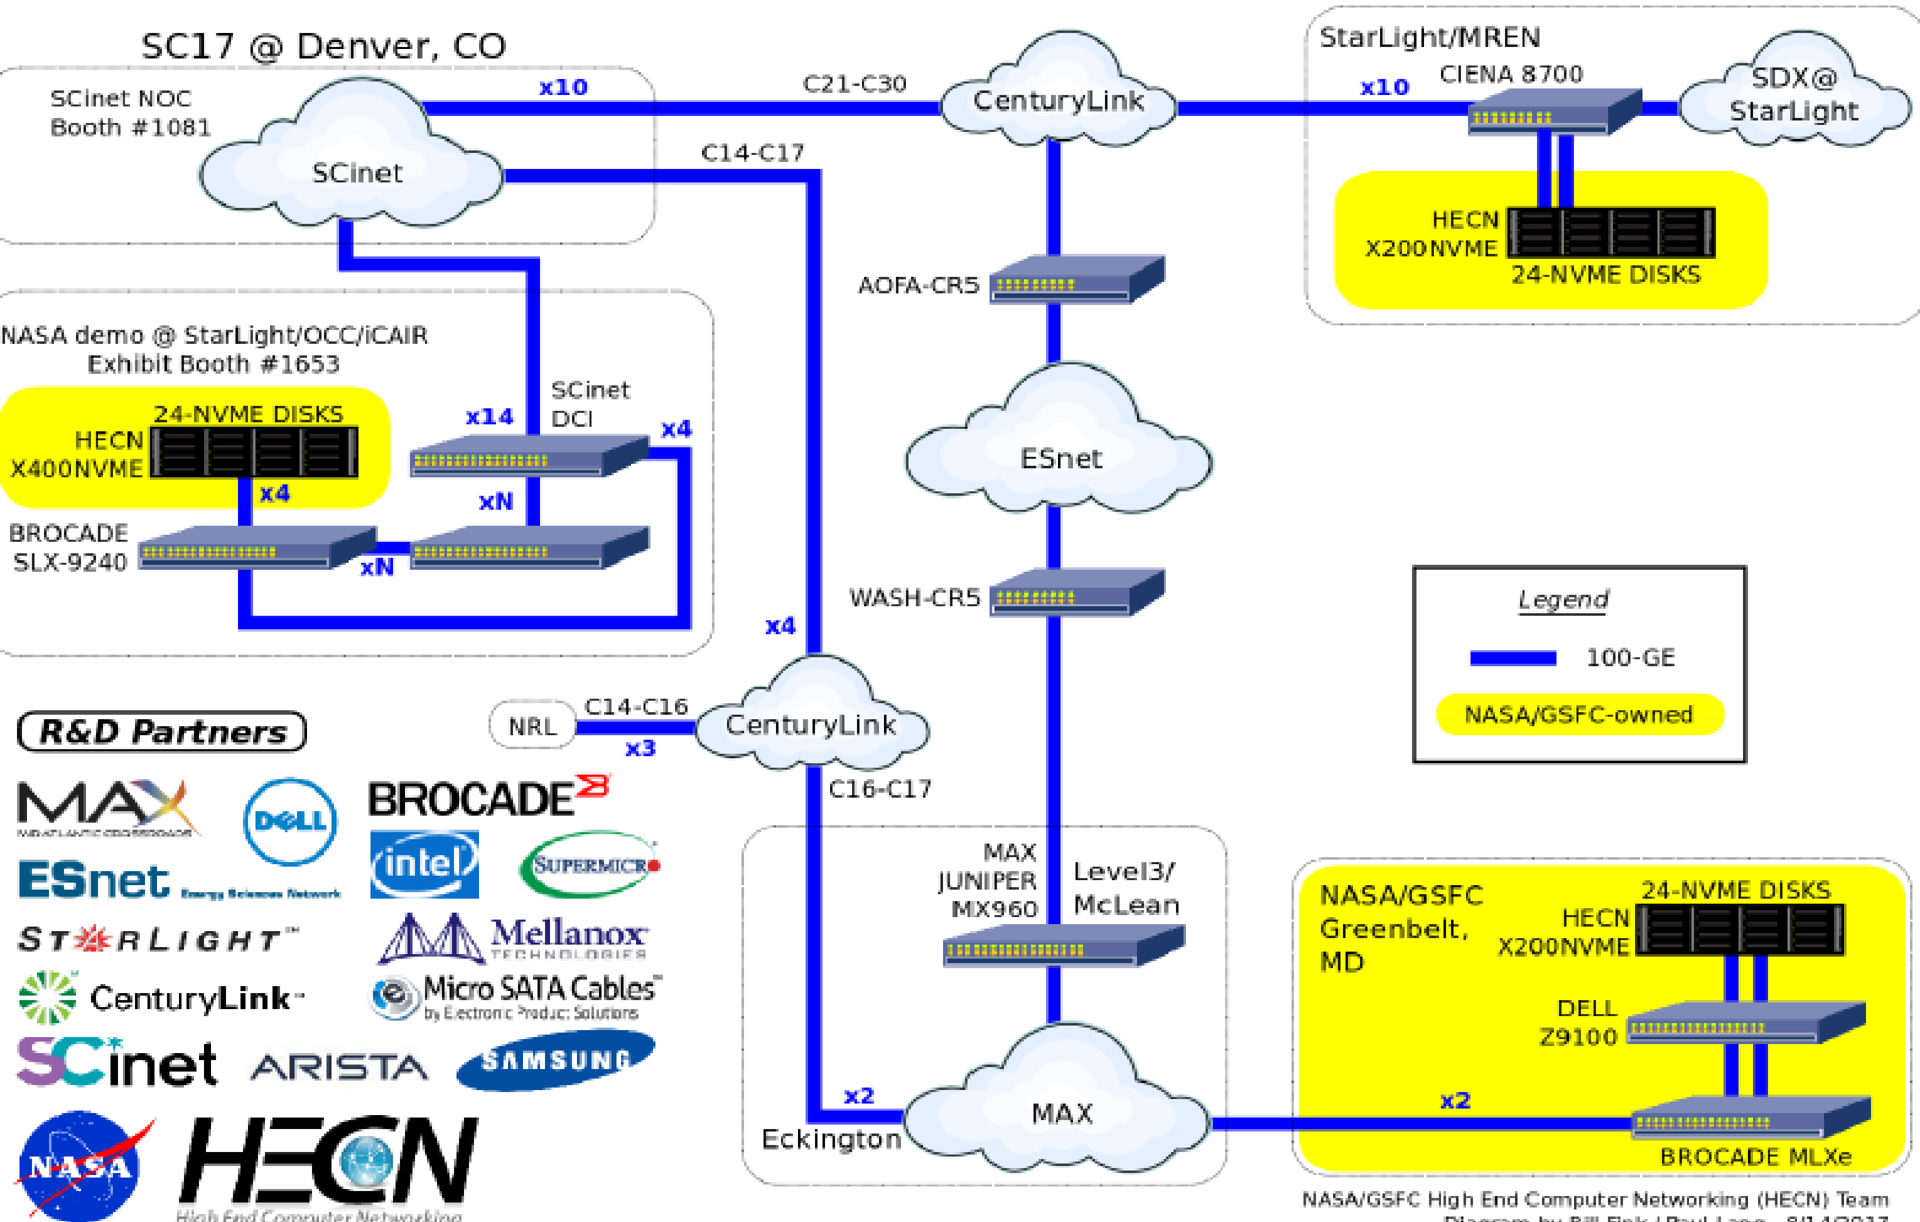

Dynamic Distributed Data Processing

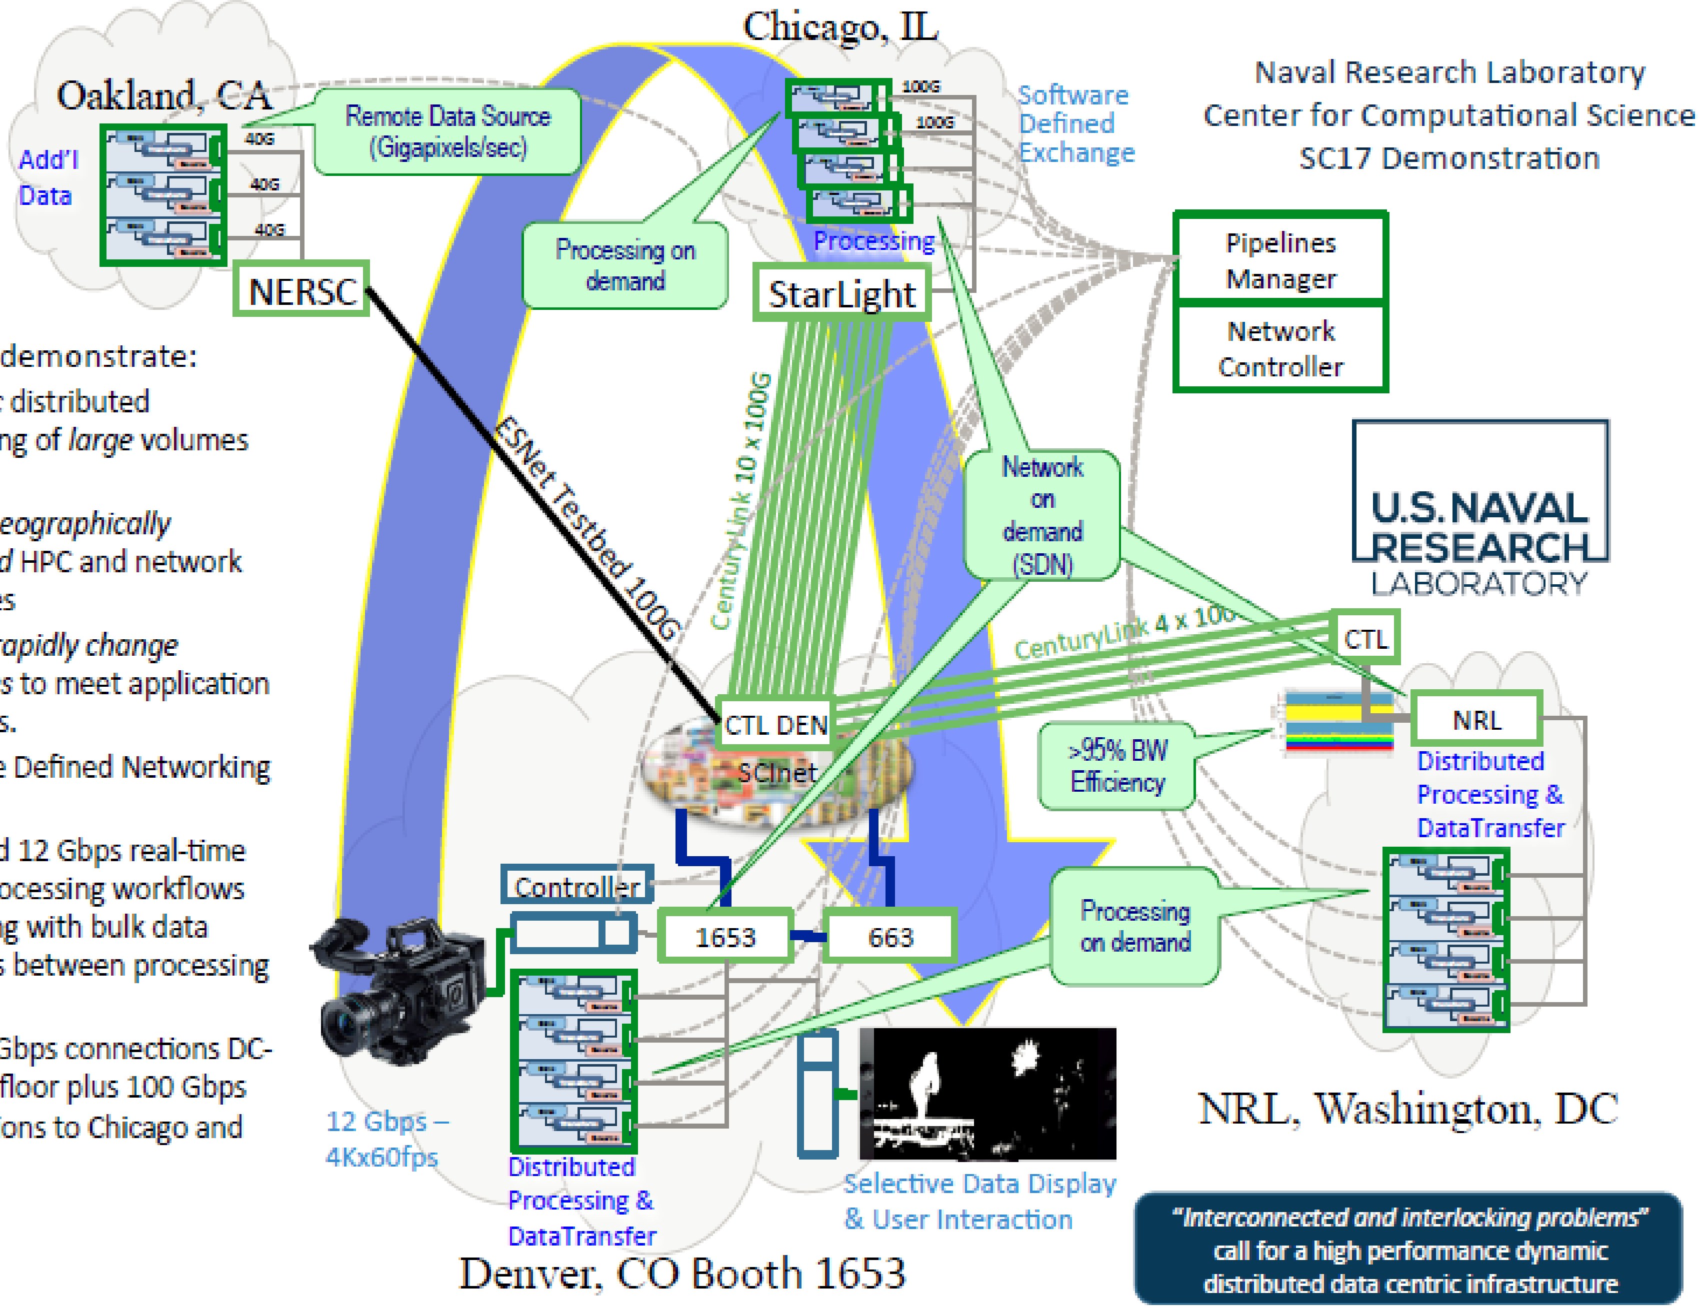

NRL will demonstrate:

- *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

High Speed Research Network

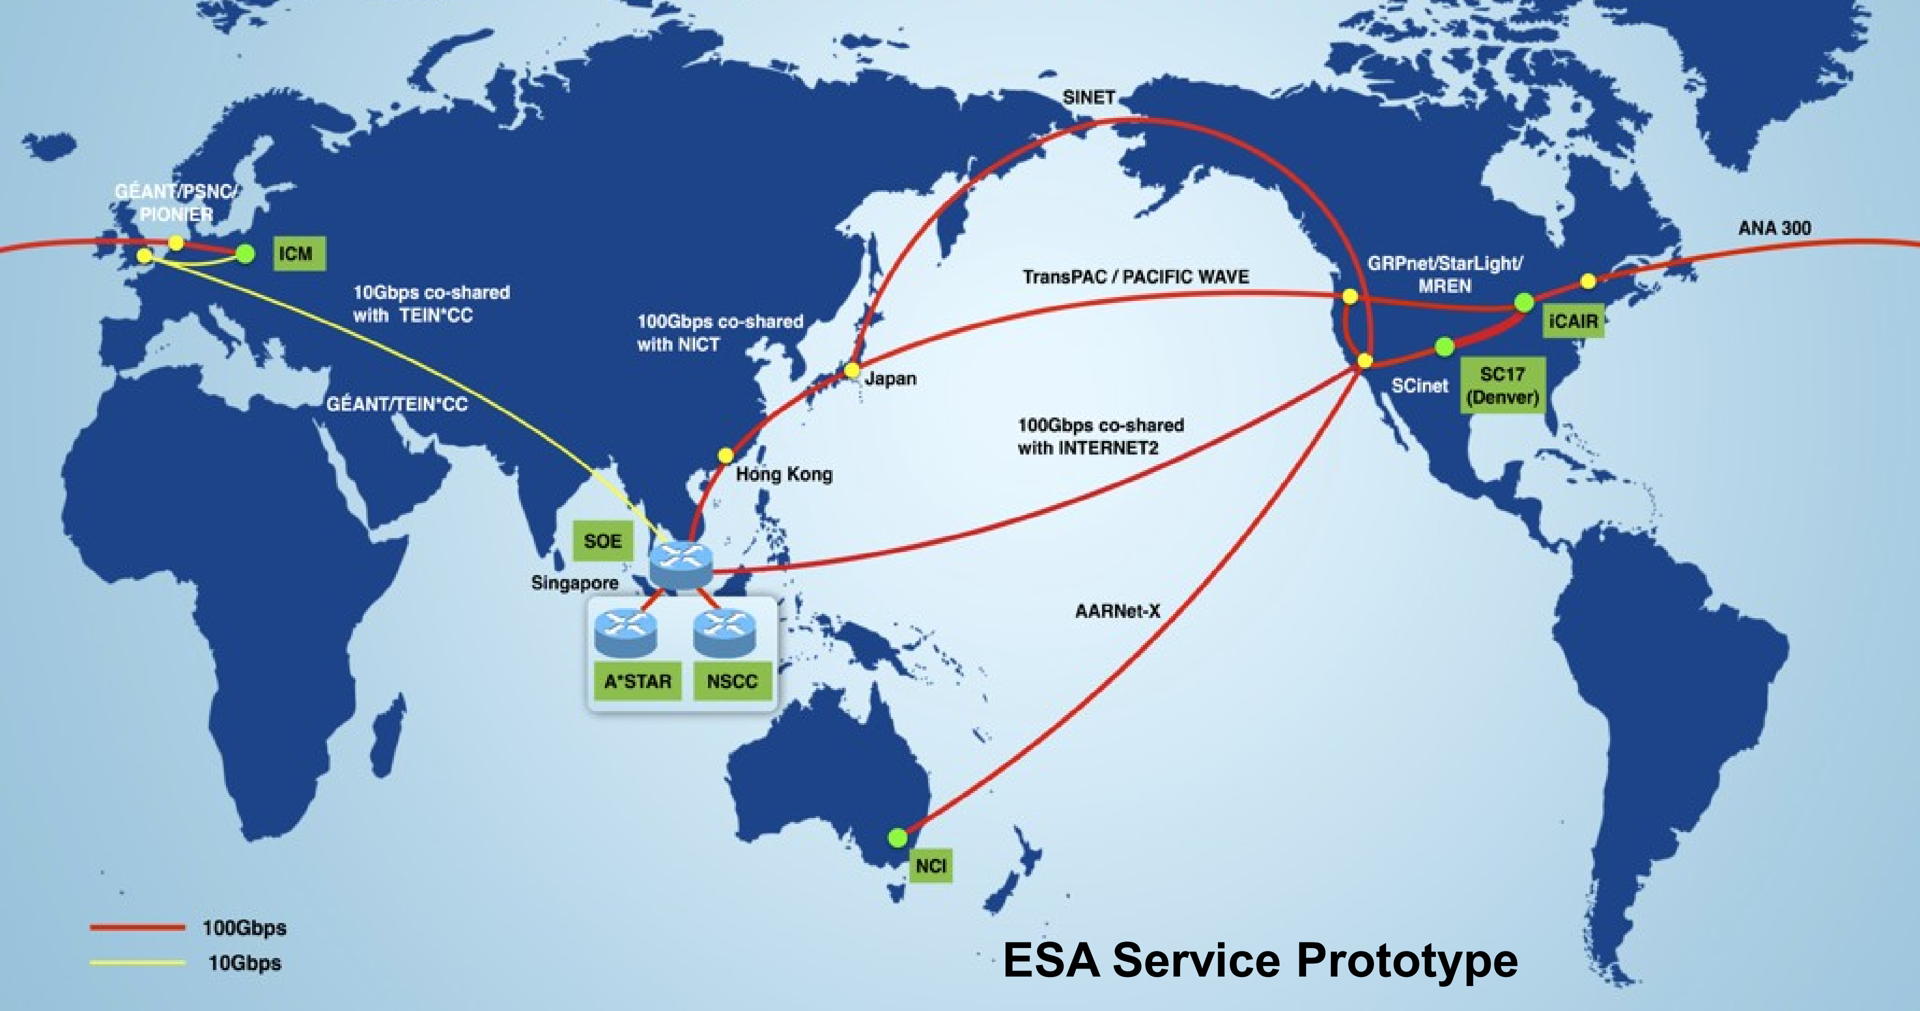

ESA Service Prototype

StarLight SDX Geoscience Research Workflow

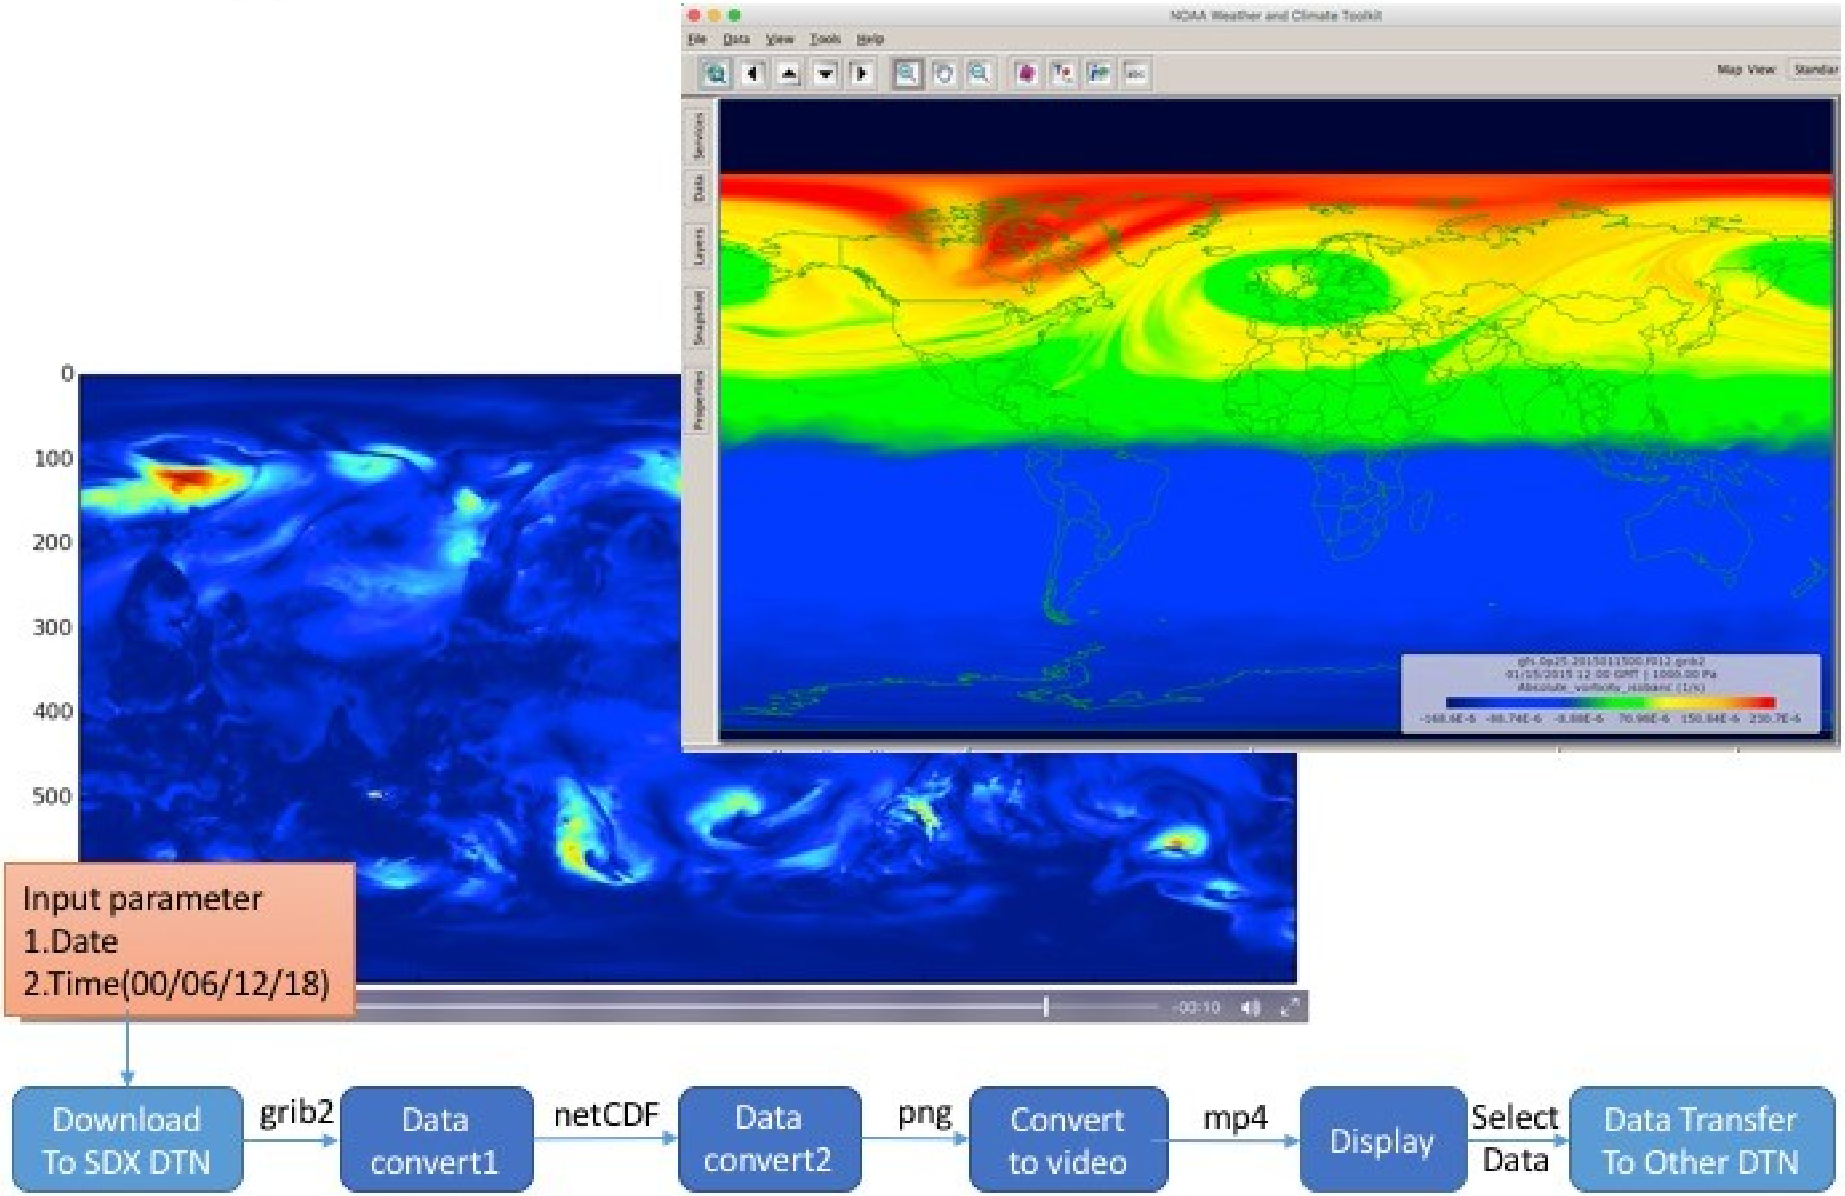

100 Gbps DTN Optical Testbed

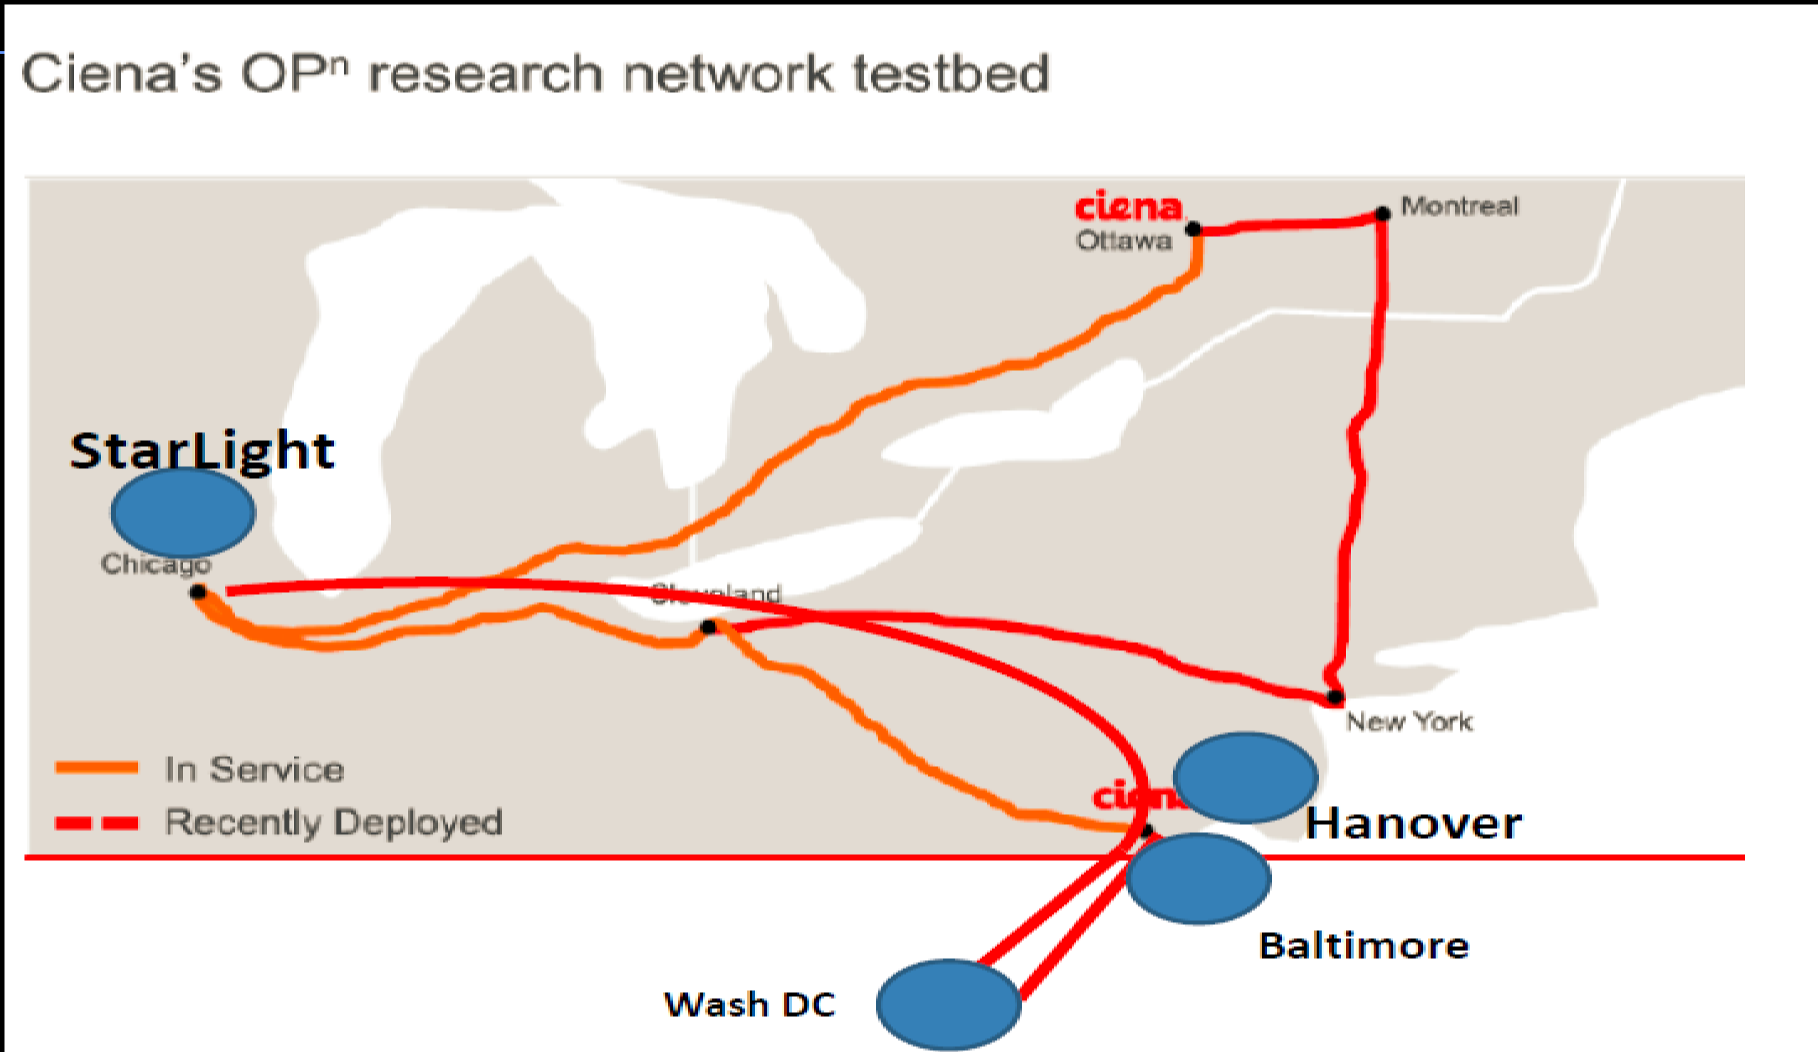

Additional WAN DTN Testing

- Preparation – Tests Conducted By Se-Young Yu (iCAIR)
- DTNs:
- @iCAIR : Intel(R) Xeon(R) Gold 6136 CPU @ 3.00GHz, Mellanox ConnectX-5 100G NIC
- @PACWAVE - LA : Intel(R) Xeon(R) CPU E5-2667 v4 @ 3.20GHz, Mellanox ConnectX-5 100G NIC
- @UvA : Intel(R) Xeon(R) CPU E5-2630 0 @ 2.30GHz, Mellanox ConnectX-3 40G NIC
- @CERN : Intel(R) Xeon(R) CPU E31220 @ 3.10GHz, Intel 82599ES 10G NIC
- Tuning Parameters :
- BIOS, CPU, NIC, TCP Stack, O/S and MTU Tuning Applied

LHCOPN map

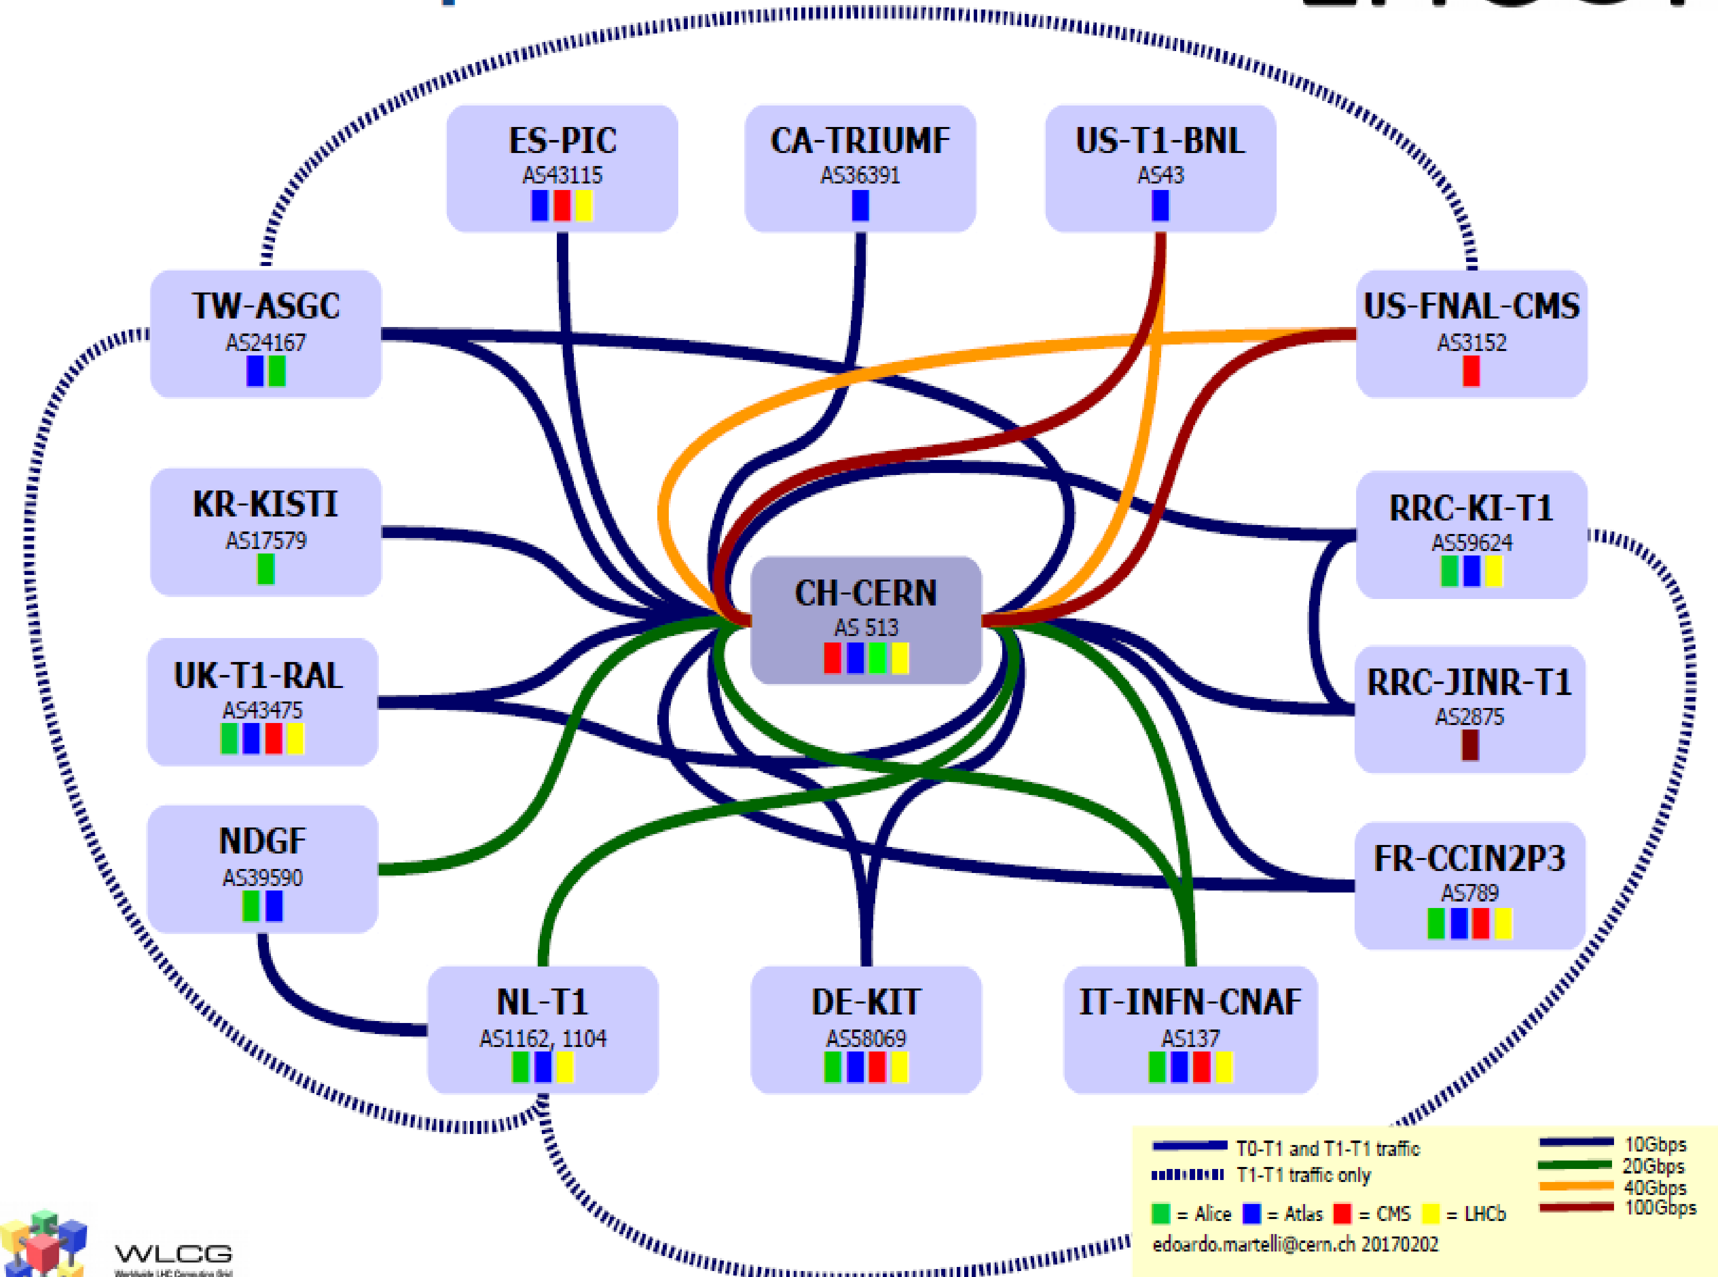

ES-PIC
AS43115
CA-TRIUMF
AS36391
US-T1-BNL
AS43
TW-ASGC
AS24167
US-FNAL-CMS
AS3152
KR-KISTI
AS17579
RRC-KI-T1
AS59624
CH-CERN
AS 513
UK-T1-RAL
AS43475
RRC-JINR-T1
AS2875
NDGF
AS39590
FR-CCIN2P3
AS789
NL-T1
AS1162, 1104
DE-KIT
AS58069
IT-INFN-CNAF
AS137
T0-T1 and T1-T1 traffic
T1-T1 traffic only
= Alice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

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

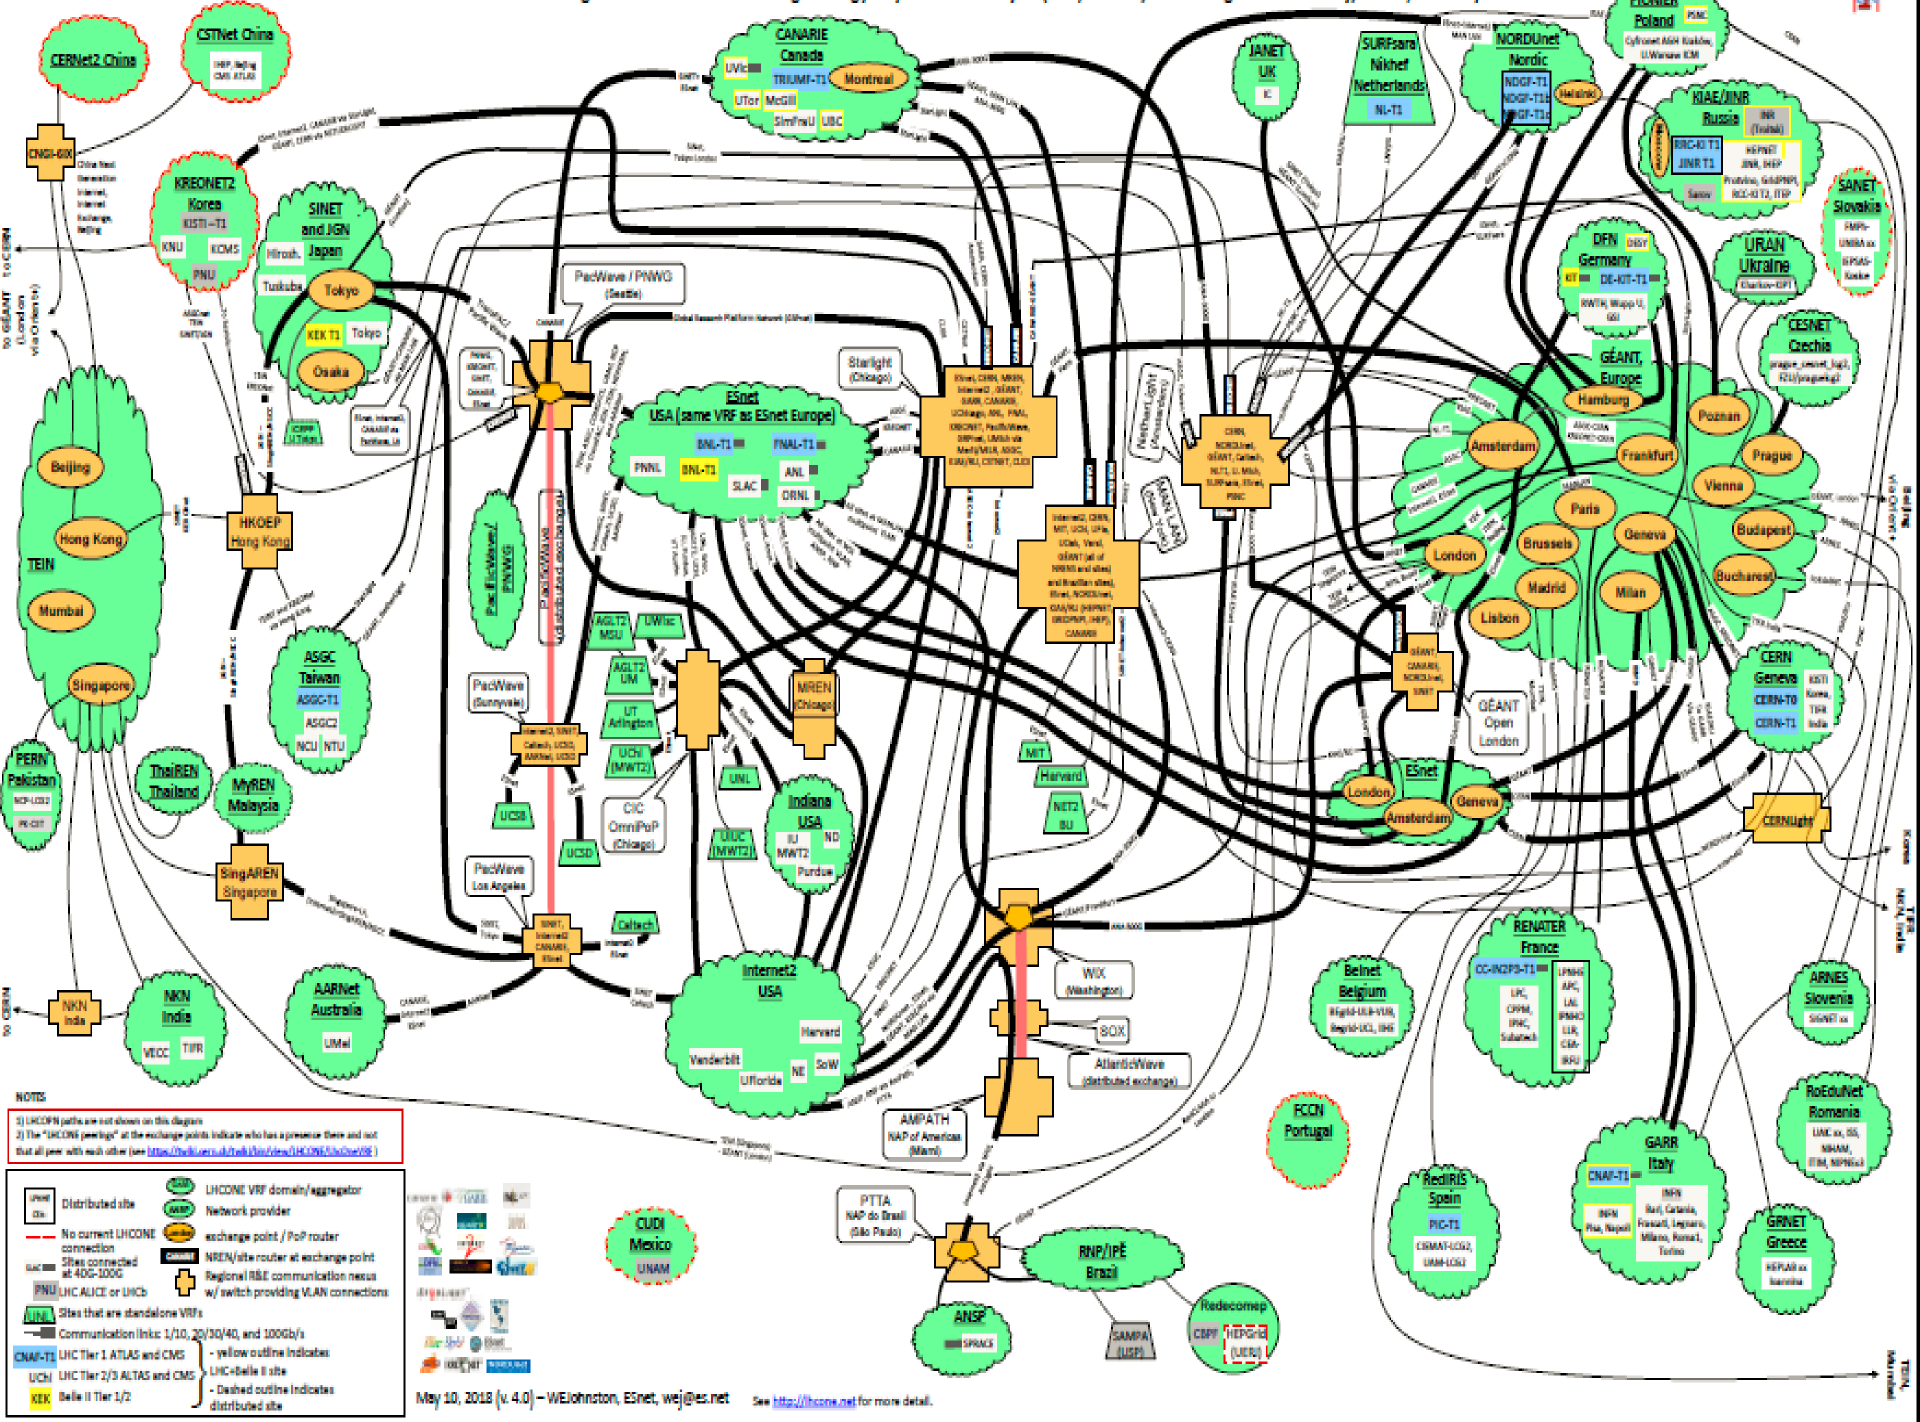

Memory-to-Memory Test Results

Server (X 4 instances) : **iperf3 -s ipaddr -p port**
Client (X 4 instances) : **iperf3 -c ipaddr -p port -Z -fq-rate=*B* -t60s**
Where ***B*** is (link-speed/4).

	PACWAVE - LA (100G)	**UvA (40G)**	**CERN (10G)**
Sending	89.8% (89.8 Gbps)	89.8% (35.92 Gbps)	95.8% (9.58 Gbps)
Receiving	87.9% (87.9 Gbps)	92.5% (37 Gbps)	96.4% (9.64 Gbps)

Disk To Memory Test Results

Server (X 8 instances) : **nuttcp -S -1 -sdz -P port < file**
Client (X 4 instances) : **nuttcp -r -sdz -P port -i1 ipaddr > file**

Starlight DTN	**PACWAVE - LA (100G)**	**UvA (40G)**	**CERN (10G)**
Reading from disk	91.68% (91.68 Gbps)	73.06% (29.22 Gbps)	74.77% (7.47 Gbps)
Writing to disk	53.67% (53.67 Gbps)	99.50%(39.80 Gbps)	77.9% (7.79 Gbps)

Fermilab U.S. DEPARTMENT OF ENERGY | Office of Science

The BigData Express Project (BDE)

BDE Research Team

November 2017

Wenji Wu, Phi Demar et al

STARLIGHT℠

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

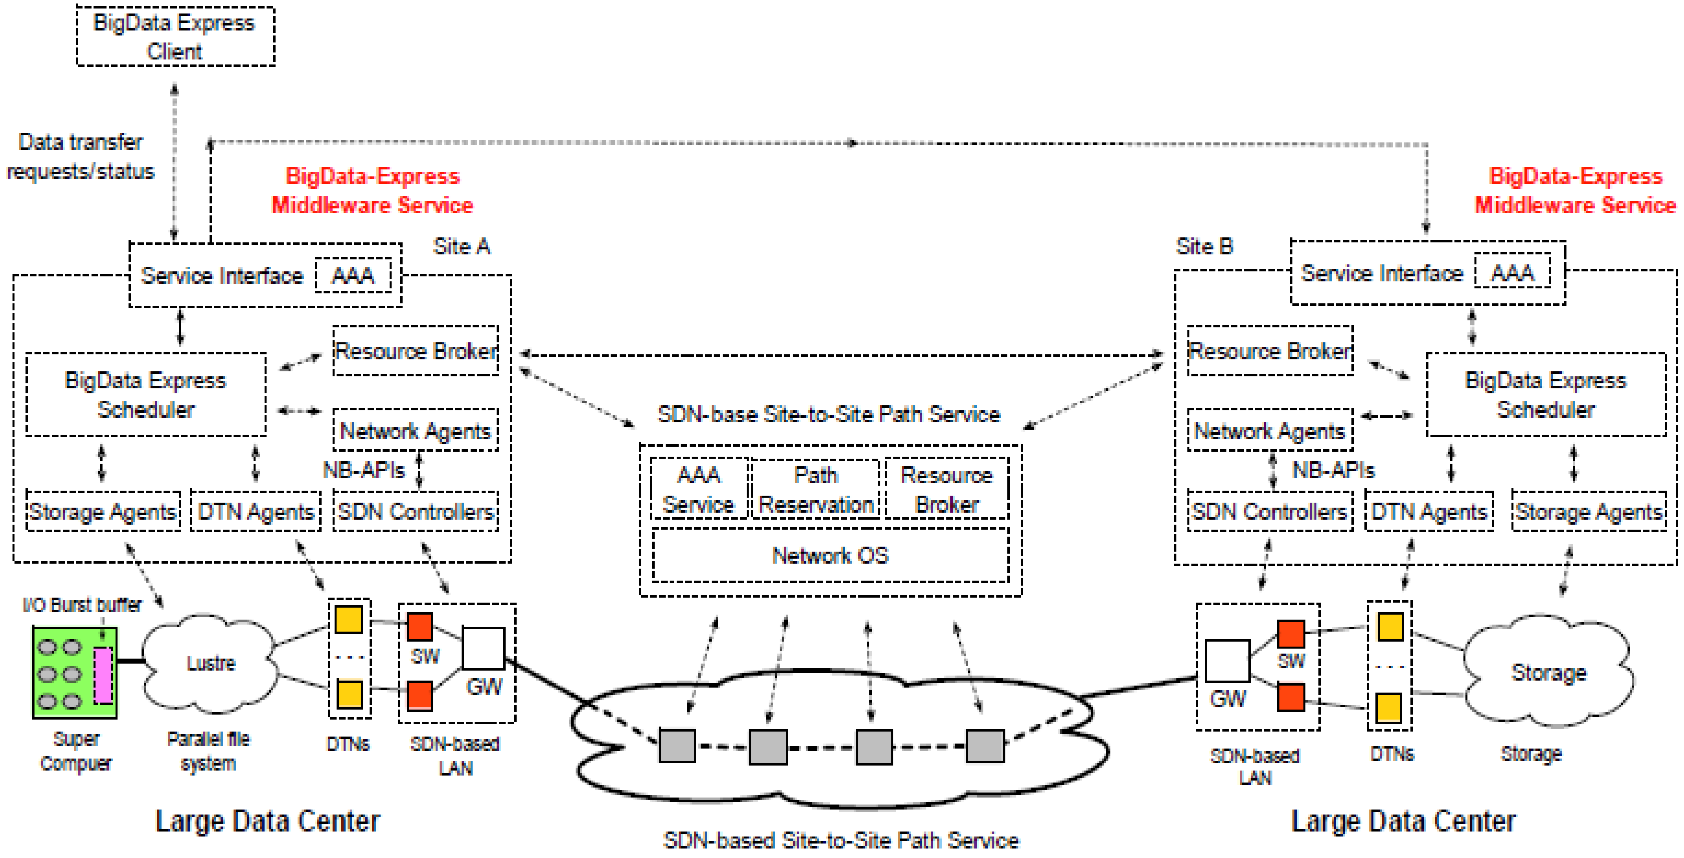

A Cross-Pacific SDN Testbed

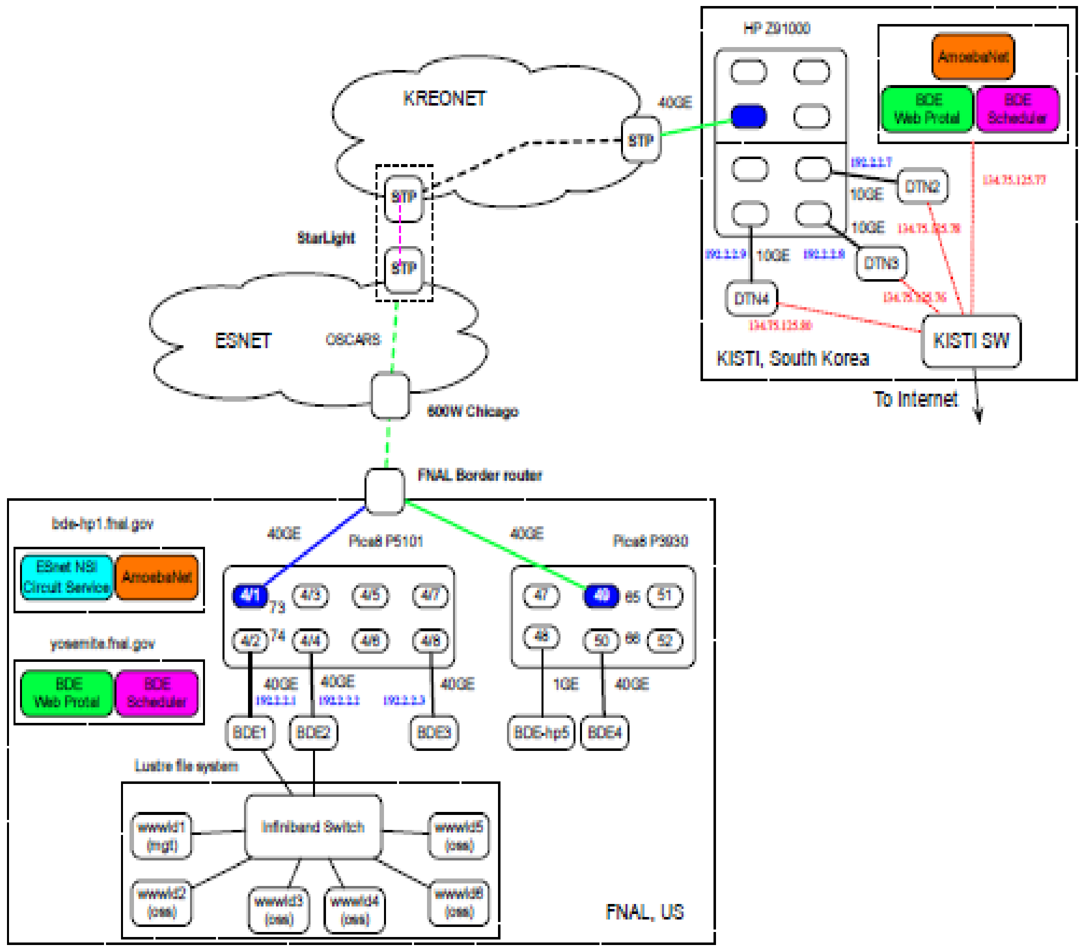

BigData Express SC'17 DEMO

- 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

A DOE/SC/ASCR-sponsored research project

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

Test methodology

- 3rd party data transfer
  - Disk-to-Disk
- mdtmFTP vs. GridFTP
- Test scenarios
  - Run1 – 1 data transfer job, transferring a 300GB file
  - Run2 – 2 parallel data transfer job, with each transferring a 300GB file
  - Run4 – 4 parallel data transfer job, with each transferring a 300GB file
  - Run8 – 8 parallel data transfer job, with each transferring a 300GB file
- Metric
  - Aggregate throughput

Source: Wenji Wu

Server 1 Hardware Configuration

- Two NUMA Nodes
  - 28 cores, Intel(R) Xeon(R) CPU E5-2683 v3 @ 2.00GHz
- 64GB MEM
- One 100GE Mellanox ConnectX-4
- 8 NVME Drives

wenji — root@datanode: /home/mdtmftp_server — ssh -l wenji 165.124.3...

```
/dev/nvme0n1p1 384448652 104925636 259971116  29% /disk0
/dev/nvme1n1p1 384448652 104925636 259971116  29% /disk1
/dev/nvme2n1p1 384448652 104925636 259971116  29% /disk2
/dev/nvme3n1p1 384448652 104925636 259971116  29% /disk3
/dev/nvme4n1p1 384448652 104925636 259971116  29% /disk4
/dev/nvme5n1p1 384448652 104925636 259971116  29% /disk5
/dev/nvme6n1p1 384448652 104925636 259971116  29% /disk6
/dev/nvme7n1p1 384448652 104925636 259971116  29% /disk7
/dev/sdb1      192169564 113235132  69149704  63% /etc/hosts
shm                65536         0     65536   0% /dev/shm
root@datanode:/home/mdtmftp_server#
```

Source: Wenji Wu

Server 2 Hardware Configuration

- Two NUMA Nodes
  - 24 cores, Intel(R) Xeon(R) CPU E5-2687W v4 @ 3.00GHz
- 24GB MEM
- One 100GE Mellanox ConnectX-4
- 8 NVME Drives

```
wenji — root@petatrans3: /home/mdtmftp_server — ssh -l wenji 165.124.33.1...
root@petatrans3: /home/mdtmftp_serve...    root@petatrans3: /home/wenji — ssh -l...
Filesystem      1K-blocks      Used Available Use% Mounted on
none            192169564  70829144 111555692  39% /
tmpfs           132019976         0 132019976   0% /dev
tmpfs           132019976         0 132019976   0% /sys/fs/cgroup
/dev/nvme12n1p1 769018760 104927996 625003824  15% /disk12
/dev/nvme13n1p1 769018760 104928004 625003816  15% /disk13
/dev/nvme14n1p1 769018760 104927980 625003840  15% /disk14
/dev/nvme15n1p1 769018760 104927996 625003824  15% /disk15
/dev/nvme8n1p1  769018760 104927984 625003836  15% /disk8
/dev/nvme9n1p1  769018760 104927992 625003828  15% /disk9
/dev/nvme10n1p1 769018760 104927980 625003840  15% /disk10
/dev/nvme11n1p1 769018760 104927984 625003836  15% /disk11
/dev/sda1       192169564  70829144 111555692  39% /etc/hosts
shm                 65536         0     65536   0% /dev/shm
root@petatrans3:/home/mdtmftp_server#
```

Source: Wenji Wu

Testing environment – 3rd party data transfer

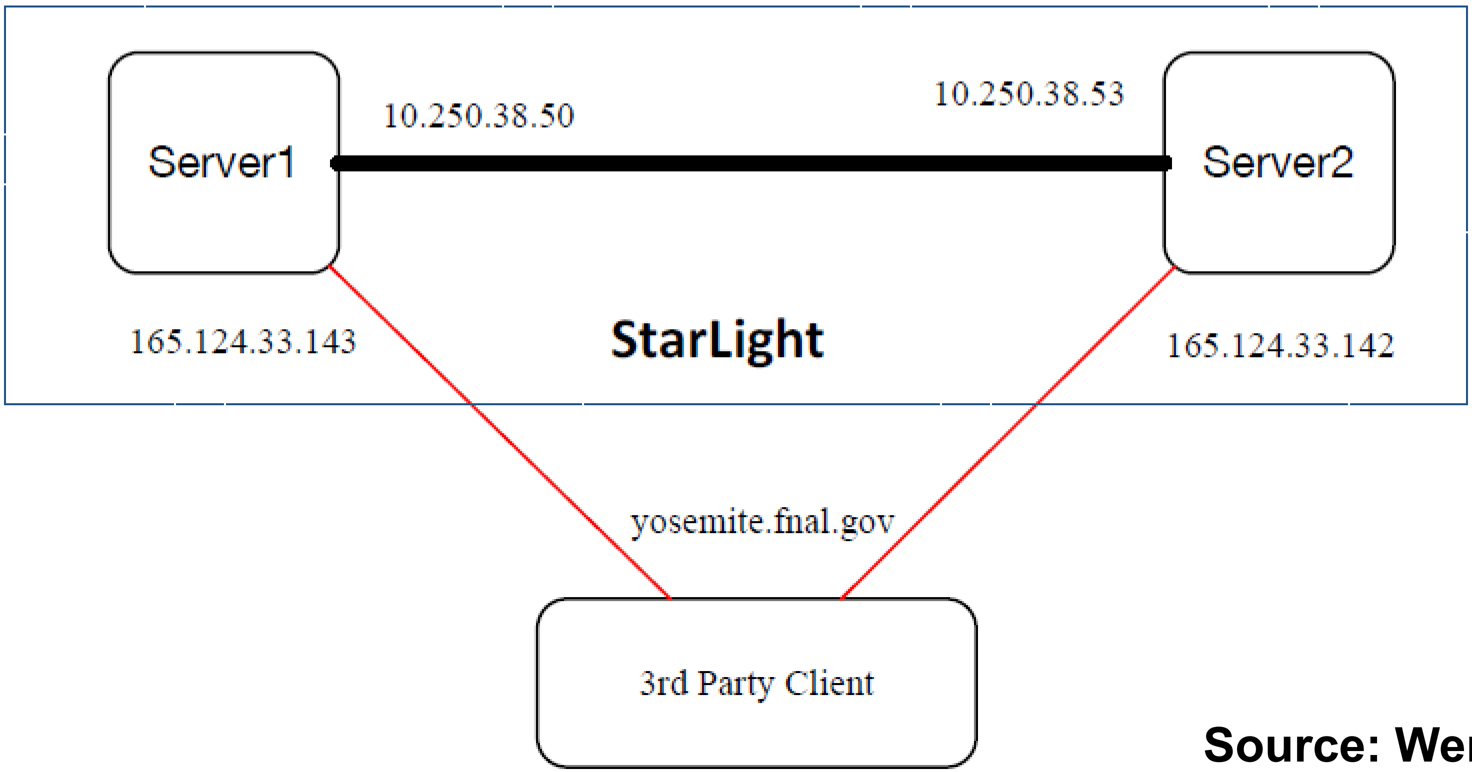

Source: Wenji Wu

Performance – Aggregate throughput

Gb/s	Run1	Run2	Run4	Run8
GridFTP	6.2Gbps	12.24Gbps	20.35Gbps	28.32 Gbps
mdtmFTP	13.27Gbps	23.80Gbps	28.354Gbps	43.94 Gbps

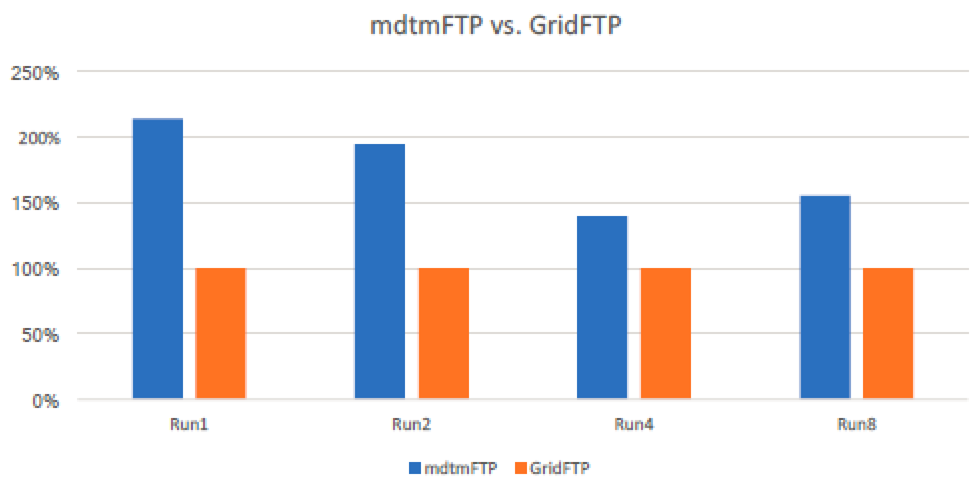

Source: Wenji Wu

Future DTNs

- Enhanced Core Components
- Much Faster Channels
- Optimal Use of Multi Core
- Higher Clock Speeds
- Enhanced NUMA Management
- Enhanced Thread Bindings/Affinity Management
- Enhanced In Network Computing
- Integration Into Network Orchestration
- Enhanced Programmability (e.g., DTN + P4 Switch)
- 400 Gbps/Tbps

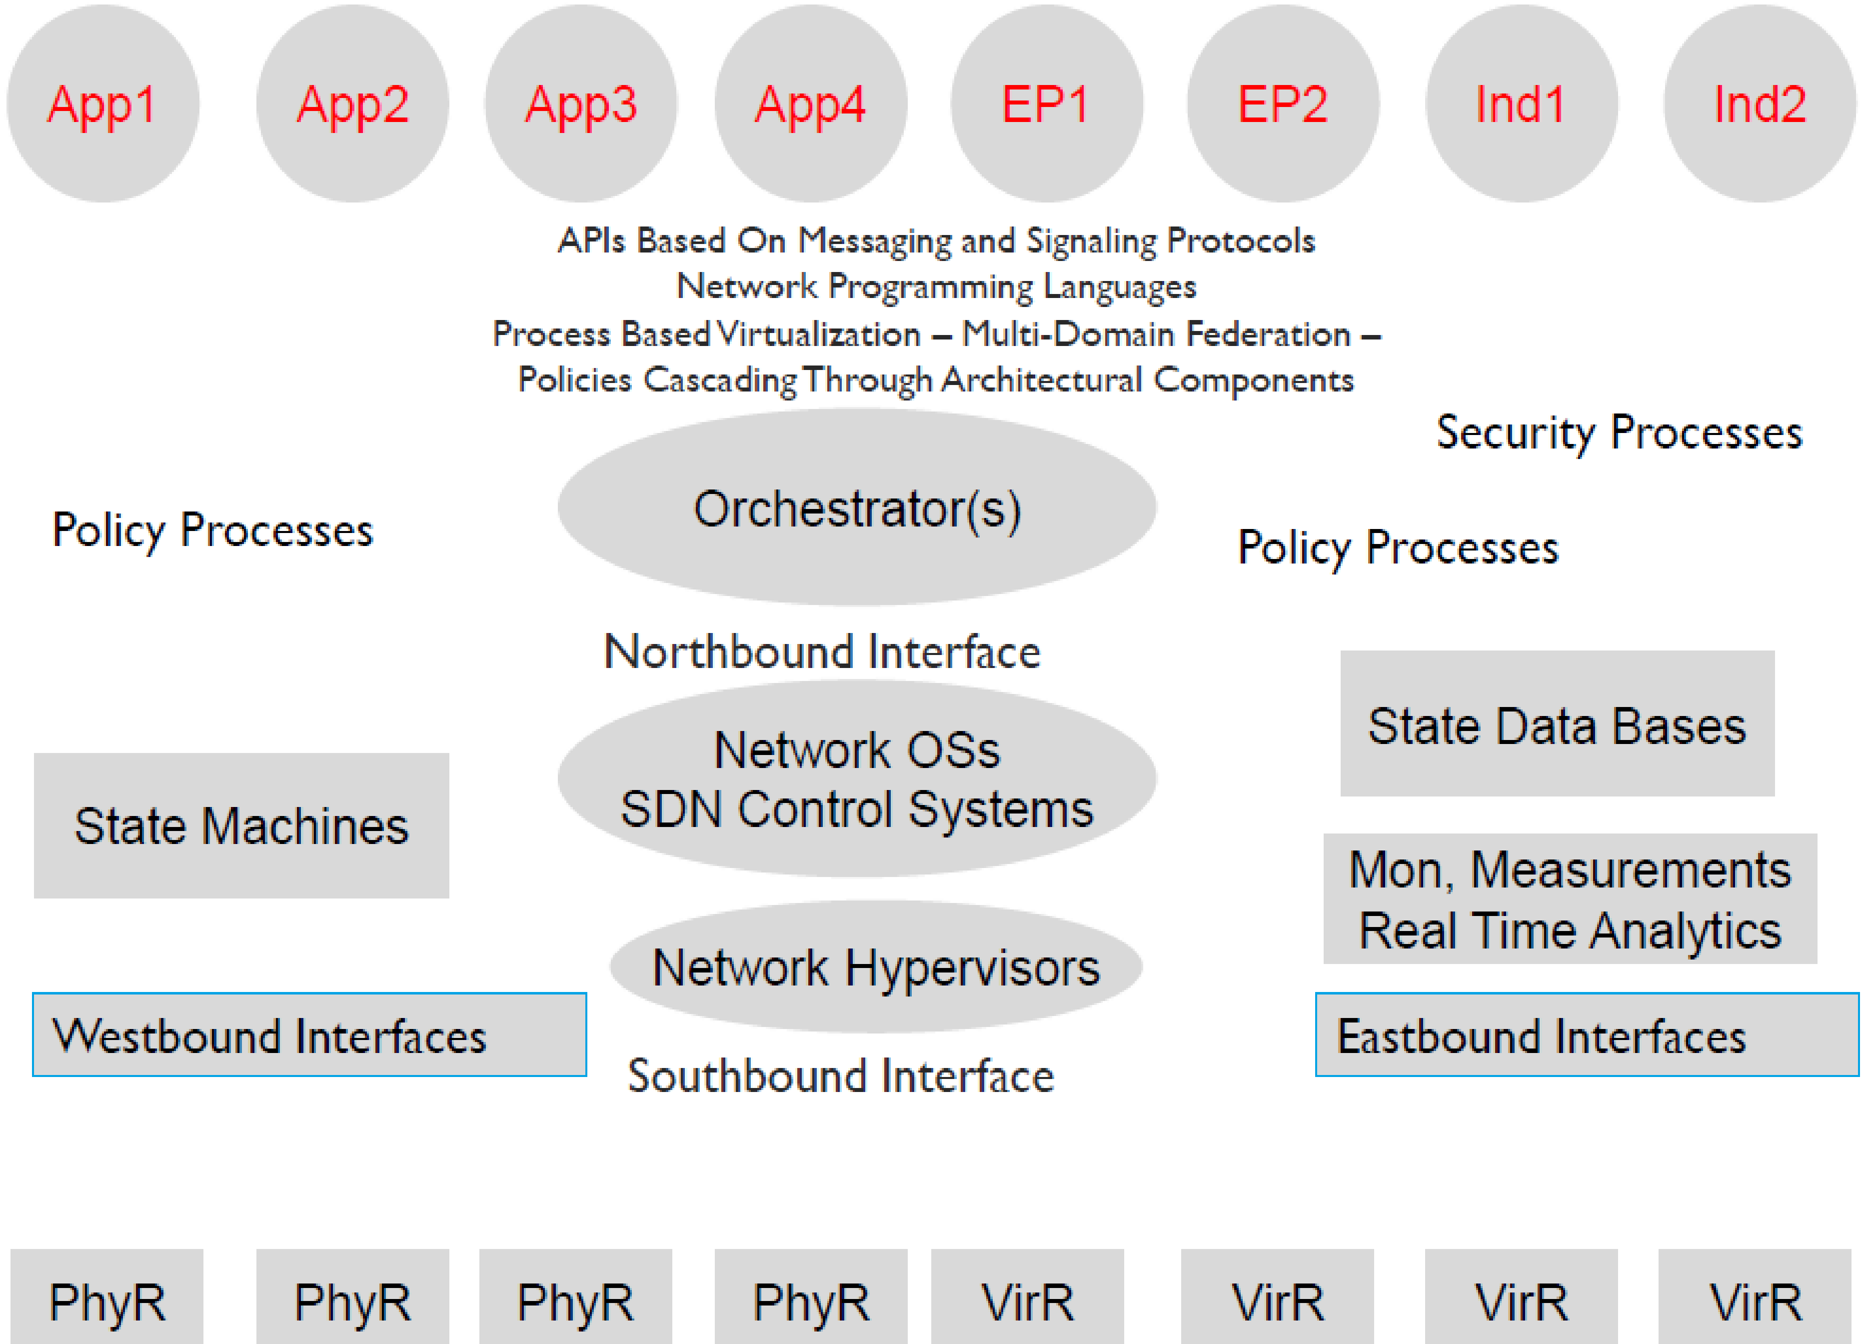
App1
App2
App3
App4
EP1
EP2
Ind1
Ind2
APIs Based On Messaging and Signaling Protocols
Network Programming Languages
Process Based Virtualization – Multi-Domain Federation –
Policies Cascading Through Architectural Components
Security Processes
Policy Processes
Orchestrator(s)
Policy Processes
Northbound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

Orchestration And AI & Selected ML Frameworks (Of Many):

- Apache Singa
- Caffe
- H2O
- MLlib (Apache Spark)
- Scikit-Learn (Python)
- Shogun (C++)
- TensorFlow
- Theano (Python)
- Torch (~ Scientific Computing)
- Veles (C++, w/ Some Python)

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

Emerging US SDX Interoperable Fabric

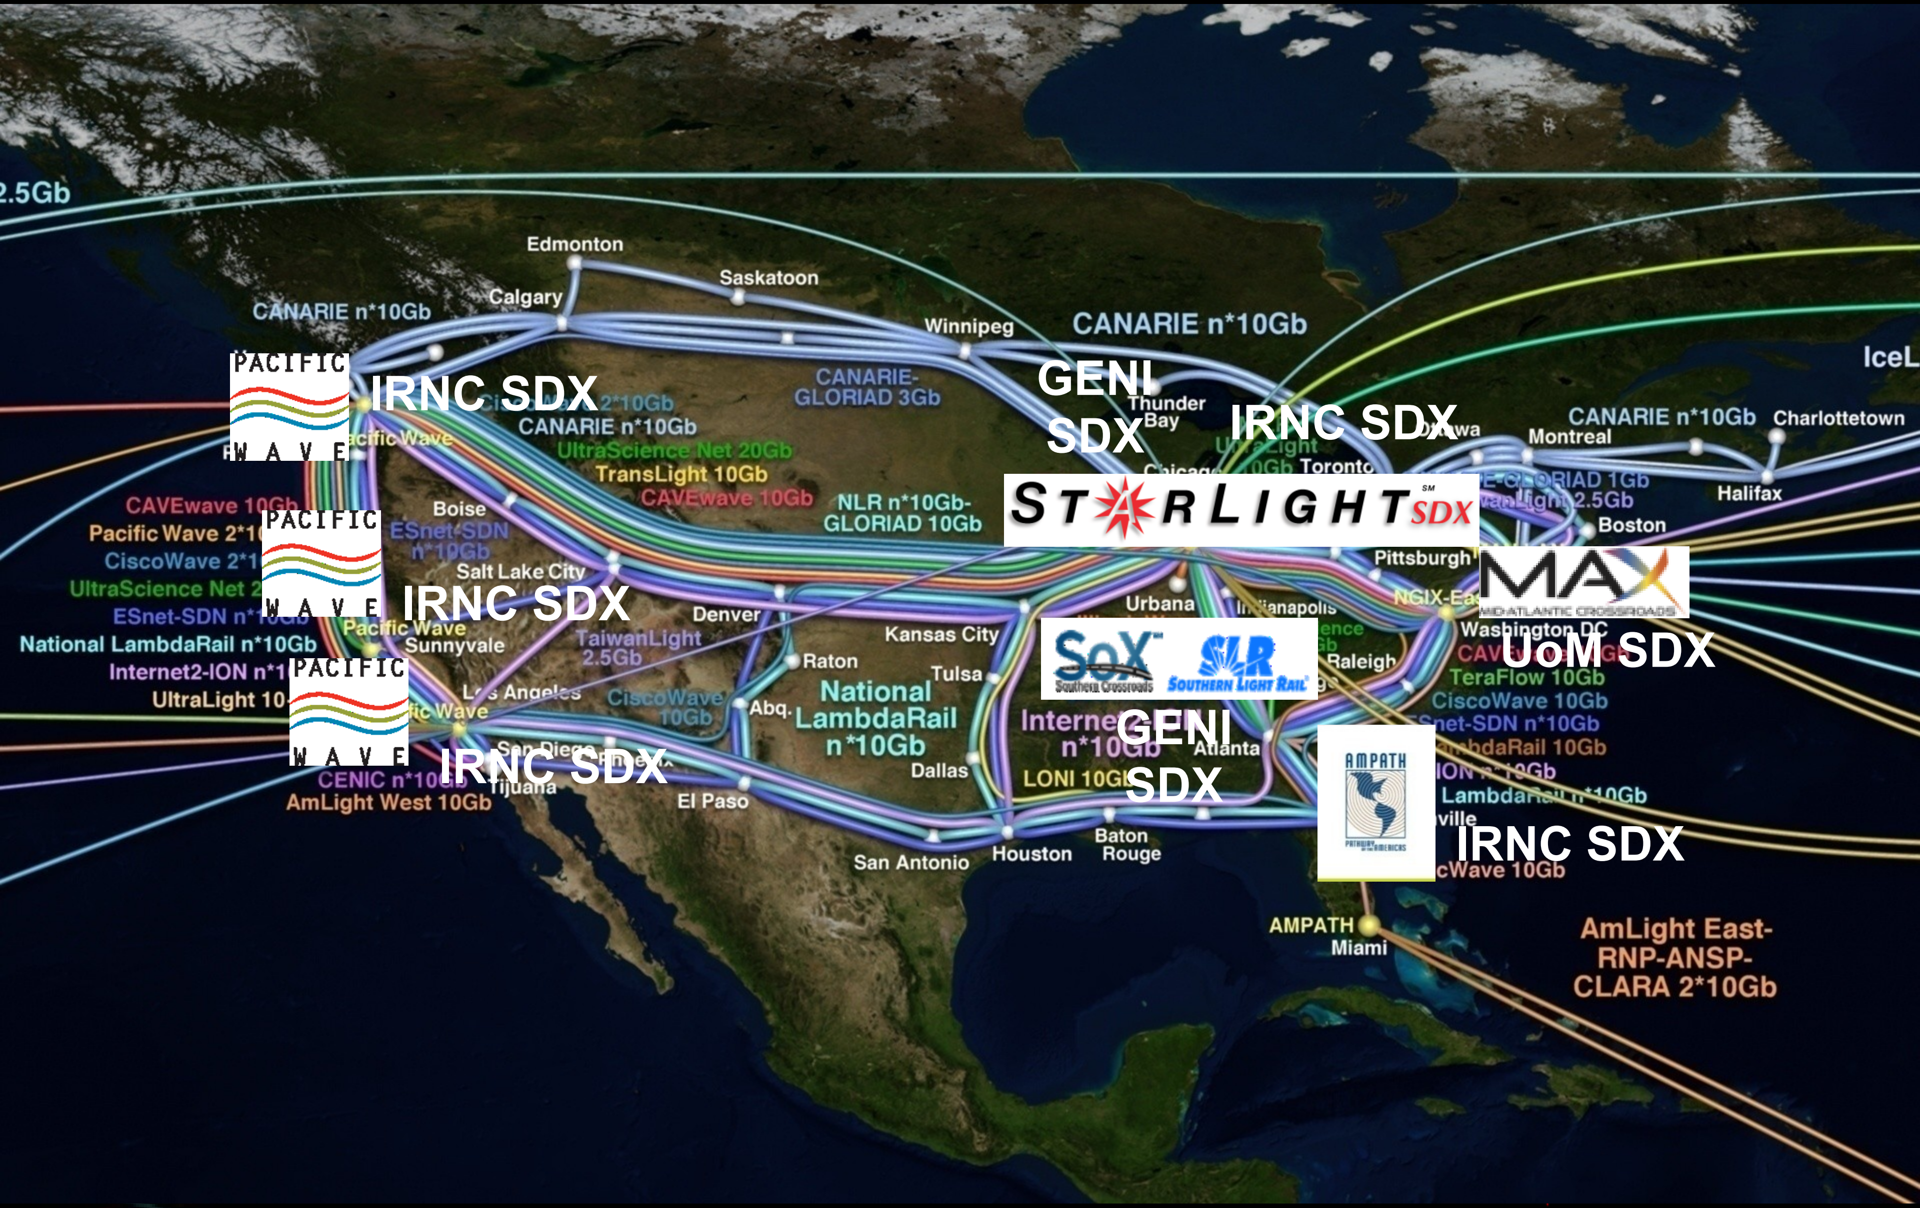

Multi-Institution, Hyper-Converged ScienceDMZ

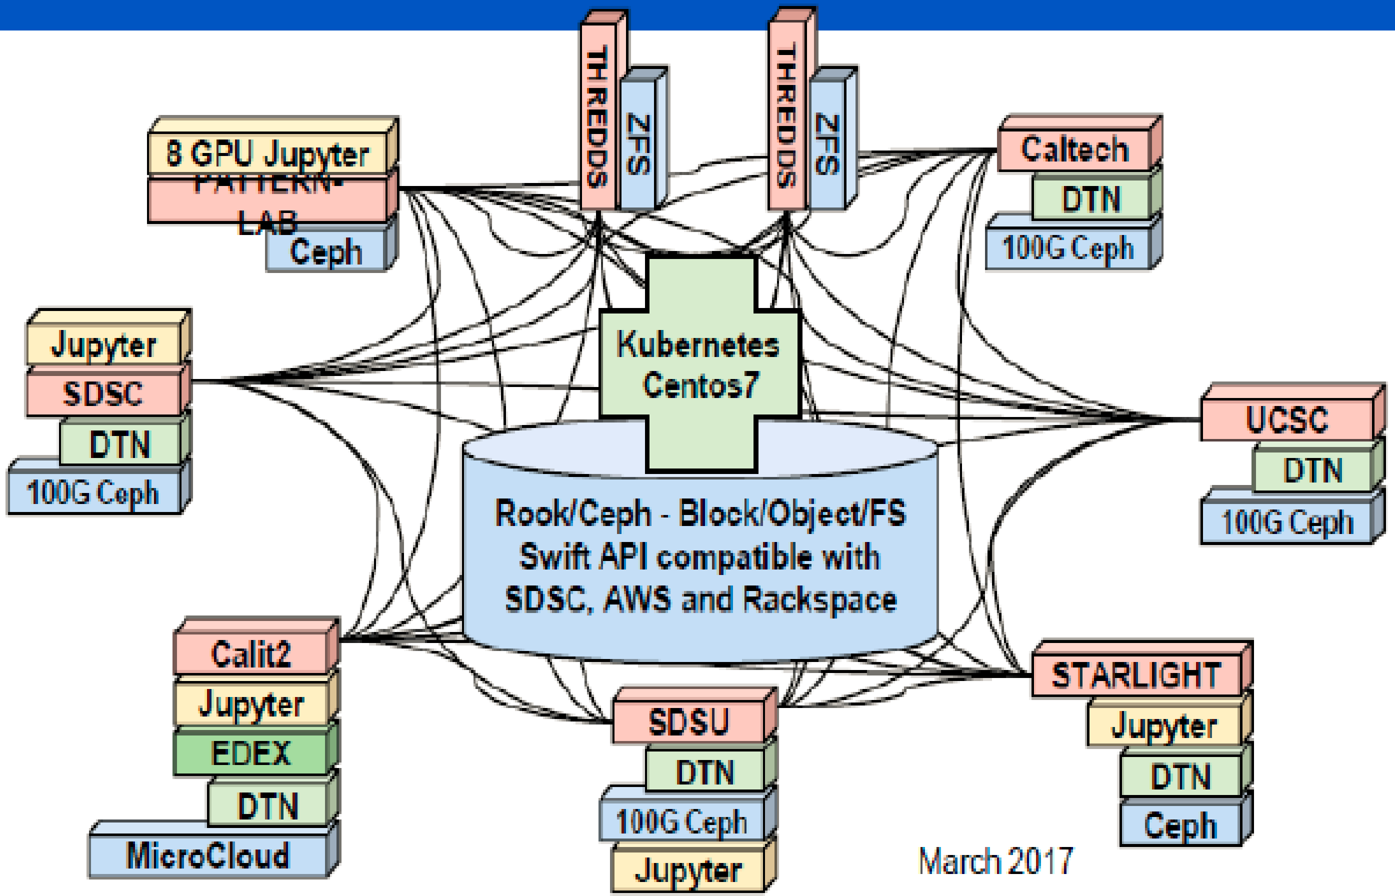

Source; John Graham UCSD

Building the Open Storage Network

Alex Szalay
The Johns Hopkins University

Institute for Data Intensive Engineering and Science
idies

Potential For Close Integration of DTNs & Large Scale Storage

Connections

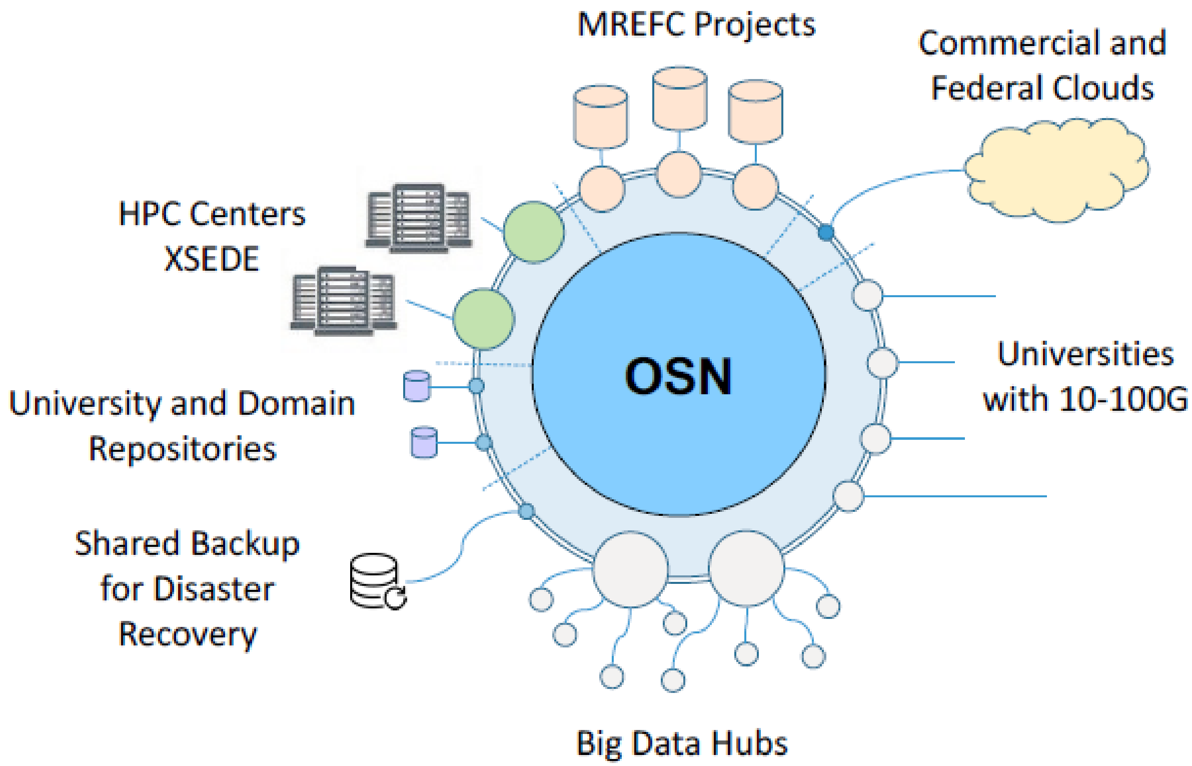

Summary

- Data Intensive Science As Well As Other Services and Applications Can Benefit From DTNs, Which Enable Enhanced Capabilities For High Performance LAN and WAN Data Transport, Including Customized Flow Management
- Key Enabling Capability: Using DTNs Integrated With Specialized WAN Paths, Including L2 Paths To Optimize E2E Data Flows, Including Disk To Disk
- Core Components Can Be Supplemented By Enhancing Software Stacks, e.g., Jupyter, NSI, MEICAN, P4 Programming, BDE, AI/ML/DL, etc
- Today, Many Components Exist To Create An E2E Services For Data Intensive Science
- Major Opportunity=> Creating DaaS Capabilities and Placing Them Into Production, e.g. Using the Global Research Platform (World-Wide Science DMZ)

www.startap.net/starlight

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

iCAIR

STARLIGHT℠