

FutureNets: Designing Architecture, Technology, and Facilities for Multi-Service, Deterministic, Differentiated-Attribute Networks

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

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

Northwestern University

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

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

Conducting Research on LEARN Workshop

University of Houston

February 18, 2011



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 Science**
- **Advanced Applications, Middleware, Large-Scale Infrastructure, NG Optical Networks and Testbeds, Public Policy Studies and Forums Related to NG Networks**
- **Three Major Areas of Activity: a) Basic Research b) Design and Implementation of Prototypes c) Operations of Specialized Communication Facilities (e.g., StarLight)**



Advanced Communications Research Topics

- Many Current Topics Could Be Considered “Grand Challenges” In Communications
- Scaling the Internet from A Service For 1-2 Billion Individuals (Current) to 4-6 Billion (Future) and Beyond
- Improving the Current Internet (Creating a “Better Internet,” Removing Limitations, Adding Capabilities, Increasing Security, Reliability, etc.)
- Migrating Services from Layer 3 Only to Multi-Layer Services, Including L2.5, L2, L1, e.g., Lightpaths
- Empowering *Edge* Processes, Applications, and Users
- *Creating a Fundamentally New Architecture That Allows for Accomplishing All of These Goals*



Motivation for New Communications Architecture

- Traditional Networking Architecture and Technology Are Oriented to Supporting A Relatively Few Communications Modalities e.g., Voice, Video, Common Data, for a Very Long Time (Many Years...).
- Traditional Networking Infrastructure Is Too Rigid To Accommodate Changes Quickly
- Traditional Services Are Essentially Based on 19th Century Utility Models of Service and Infrastructure, Which --
 - *Severely Restrict the Inherent Potential of Digital Technology*
 - *Cannot Meet Many Emerging Requirements for 21st Century Services*
- A Fundamentally New Architectural Model is Required
- A New Architecture Replaces The Traditional Network With a New Communication Services Foundation – a Highly Distributed Facility That Can Support Multiple Networks With Different Characteristics Each Supporting Multiple Highly Differentiated Services




Selected Macro Network Research Themes

- **Transition From Legacy Networks To Networks That Take Full Advantage of IT Architecture and Technology**
- **Extremely Large Capacity Networks (Multi-Tbps Streams)**
- **High Degrees of Communication Services Customization**
- **Highly Programmable Networks**
- **Network Facilities As Enabling Platforms for Any Type of Services**
- **Extremely High Levels of Network Virtualization**
- **Highly Distributed Processes**
- **Open Communication Exchanges**
- **Service Exchanges at All Layers (vs Only Traditional L3)**
- **Innovations Based On New Extended Environments**




Paradigm Shift – Ubiquitous Services Based on Large Scale Distributed Facility vs Isolated Services Based on Separate Component Resources

**Traditional Provider Services:
Invisible, Static Resources,
Centralized Management,
Highly Layered**



**Limited Services, Functionality,
Flexibility, Expandability**

**Distributed Programmable Resources,
Dynamic Services,
Visible & Accessible Resources,
Integrated As Required, Non-Layered**

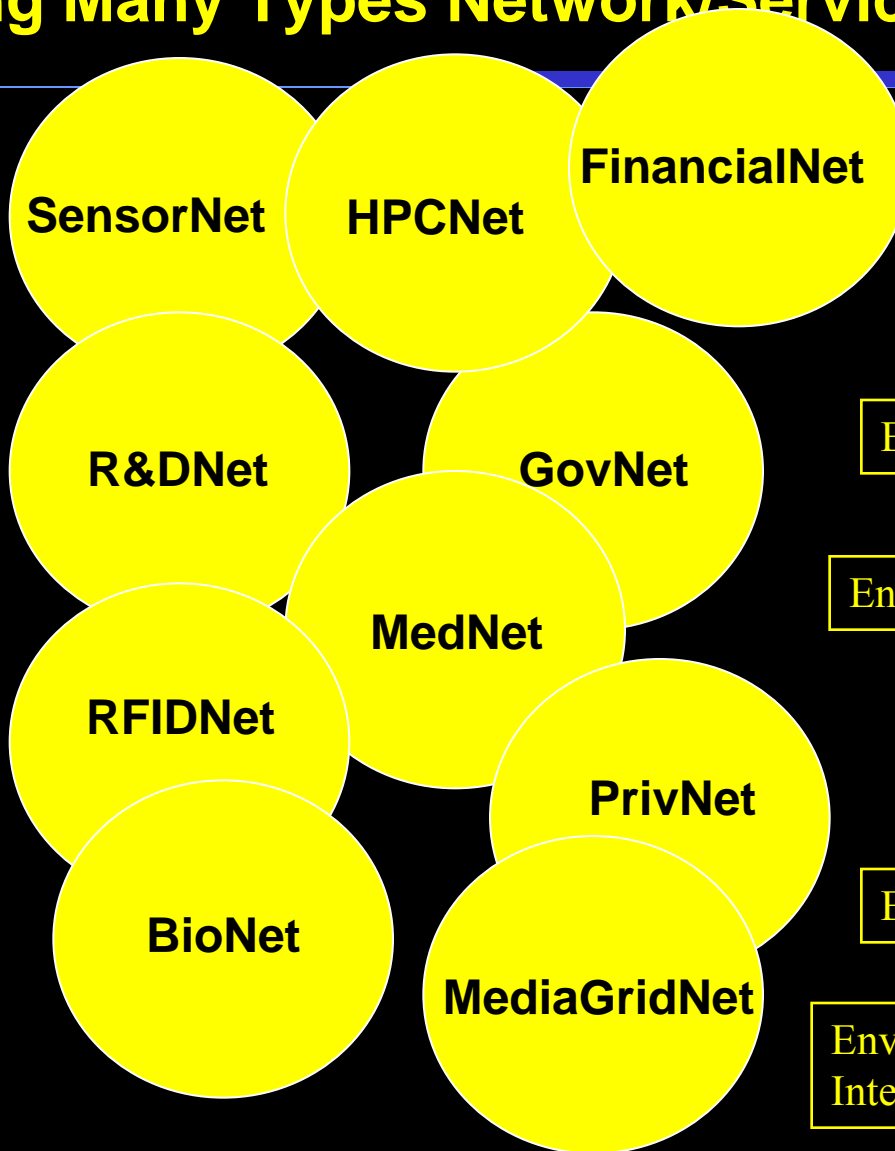


**Unlimited Services, Functionality,
Flexibility, Expandability**

Releasing the Fully Potential of Digital Technologies

STARLIGHTSM

A Next Generation Architecture: *Distributed Facility* Enabling Many Types Network/Services



Environment: VO

Environment: Real Org1

Environment: Intelligent
Power Grid Control

Environment: RFIDNet

Environment: Bio Org

Environment:
Large Scale System Control

Environment: Global App

Environment: Financial Org

Environment: Sensors

Environment: Real Org

Environment: Real Org2

Environment: Gov Agency

Environment:
Control Plane

Environment: Lab

Environment:
International Gaming Fabric



Creating Data-Intensive Science & Engineering e-Science Community Resources




GriPhyN



Data Intensive Science

LHC




ATLAS



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

NEES

ALMA



NEON

National Ecological Observatory Network

Sloan Digital Sky Survey



LIGO



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)

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

SLD 3 TB /year
(1998)

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

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

L3 5 TB /year (1993)

E791 50 TB /year
(1991)




EMC 400GB /year
(1981)



Source: Fermi Lab

RLIGHTSM



Large Hadron Collider at CERN



Fermi National Accelerator Laboratory






**35 Mile LAN PHY
to StarLight**



STARLIGHTSM

Magnetic Fusion Energy

New Sources
Of Power



Source: DOE


Source: DOE

STARLIGHTSM


ITER (Formally- International Thermonuclear Experimental Reactor)

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

ITER, currently under construction in the South of France, aims to demonstrate that fusion is an energy source of the future.



Spallation Neutron Source (SNS) at ORNL



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



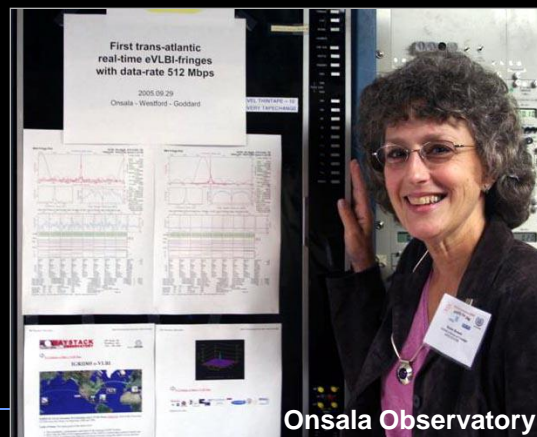
Source: DOE

Real-Time Global e-Very Long Baseline Interferometry

DRAGON (Dynamic Resource Allocation via GMPLS Optical Networks)



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



Onsala Observatory

- 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

eVLBI JIVE-Arecibo Project




Square Kilometer Array



Today: Prototypes – Future: Australia or South Africa

STARLIGHTSM

Square Kilometer Array



Increasing Accuracy in Hurricane Forecasts Ensemble Runs With Increased Resolution

5.75 Day Forecast of Hurricane Isidore

NASA fvGCM — 0.625x0.5 degrees

NASA fvGCM — 0.36x0.25 degrees

Source:
Smarr

Precip [inches/hr] : Sea Level Pressure [mb] : 1000 MB Winds [m/s]

Precip [inches/hr] : Sea Level Pressure [mb] : 1000 MB Winds [m/s]

2002 SEP 26 18:00Z

2002 SEP 26 18:00Z

Operational Forecast

Higher Resolution Research Forecast

Resolution of National Weather Service

NASA Goddard Using Ames Altix

InterCenter
Networking
is
Bottleneck


Intense
Rain-
Bands

Resolved
Eye Wall

4x
Resolution
Improvement

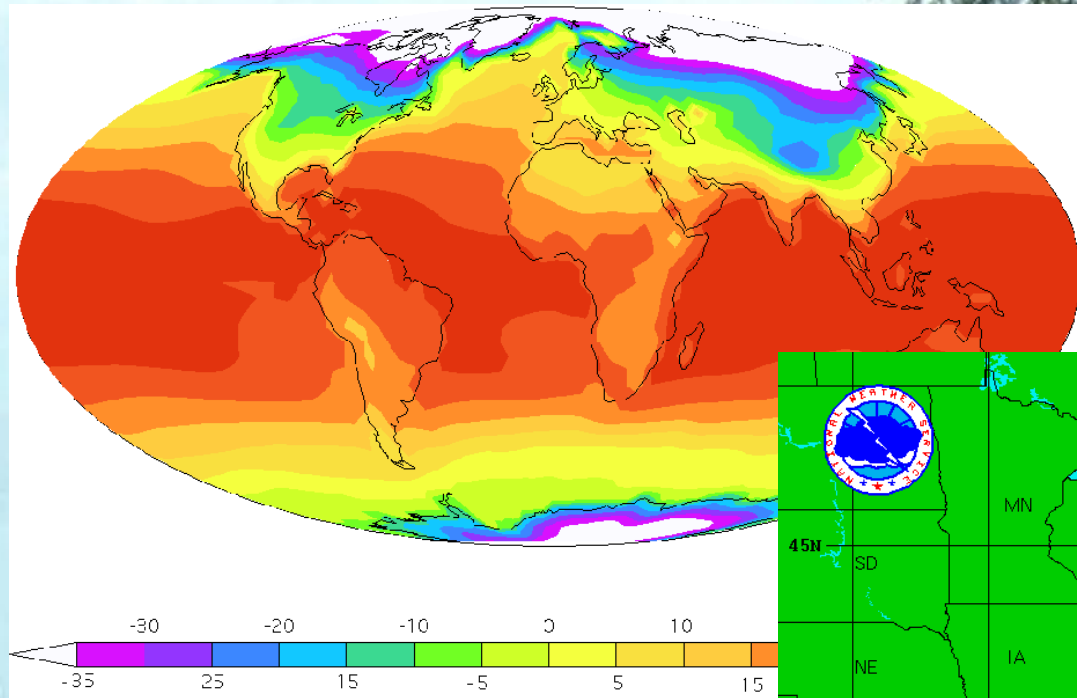
2.5 5.0 7.5 10. 12. 15. 18. 20. 22. 25.

2.5 5.0 7.5 10. 12. 15. 18. 20. 22. 25.

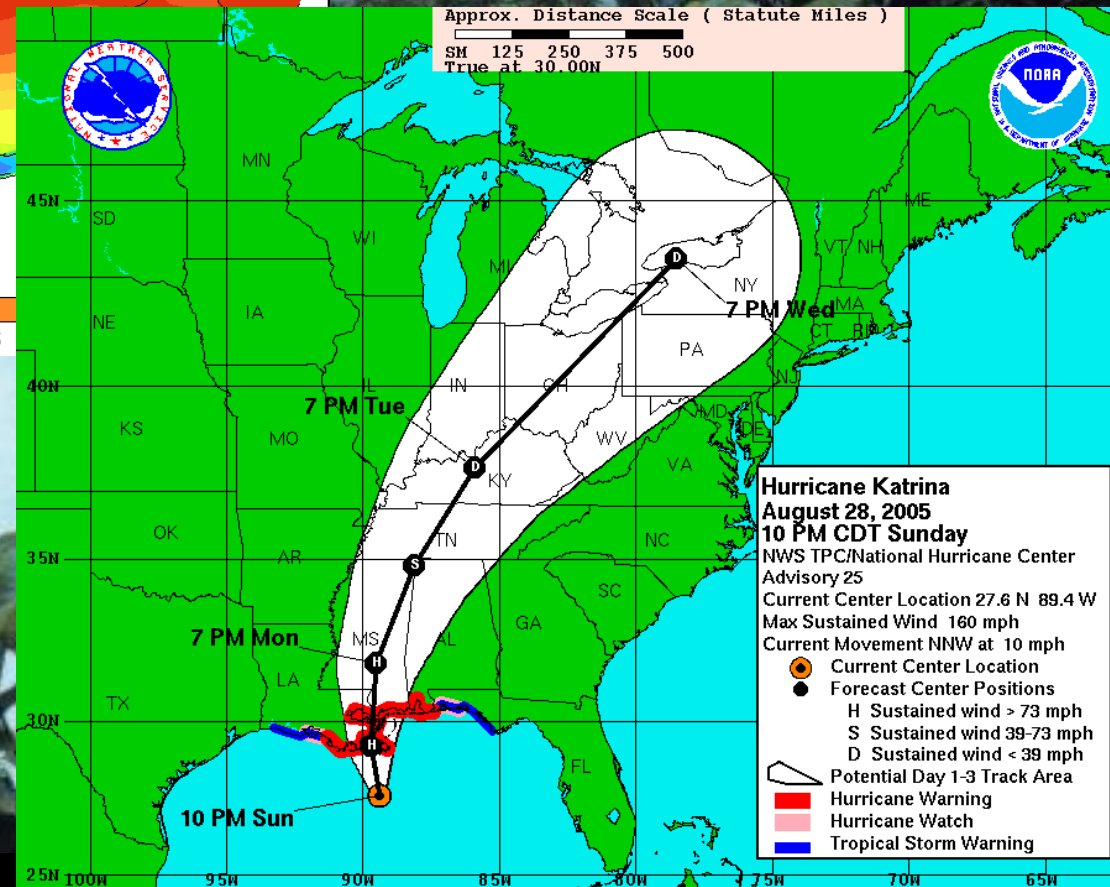


LIGHTSM

Climate Modeling




Source: DOE






USGS Images 10,000 Times More Data than Landsat7

Landsat7 Imagery
100 Foot Resolution
Draped on elevation data



New USGS Aerial Imagery
At 6-inch Resolution

Today's Aerial Imaging is >500,000 Times More Detailed than Landsat7



The Crisis Response Room of the Future



SHD Streaming TV -- Immersive Virtual Reality -- 100 Megapixel Displays

Source: EVL, UIC

STARLIGHTSM

Brain Imaging Collaboration -- UCSD & Osaka Univ. Using Real-Time Instrument Steering and HDTV

Most Powerful Electron
Microscope in the World -
- Osaka, Japan



HDTV



NCMIR
NATIONAL CENTER for
MICROSCOPY and
IMAGING RESEARCH
at San Diego, an NIH supported resource center




OSAKA UNIVERSITY



KDDI Labs USA


Southern California OptIPuter



Source: Mark Ellisman, UCSD

STARLIGHTSM

OptIPuter JuxtaView Software for Viewing High Resolution BioImages on Tiled Displays

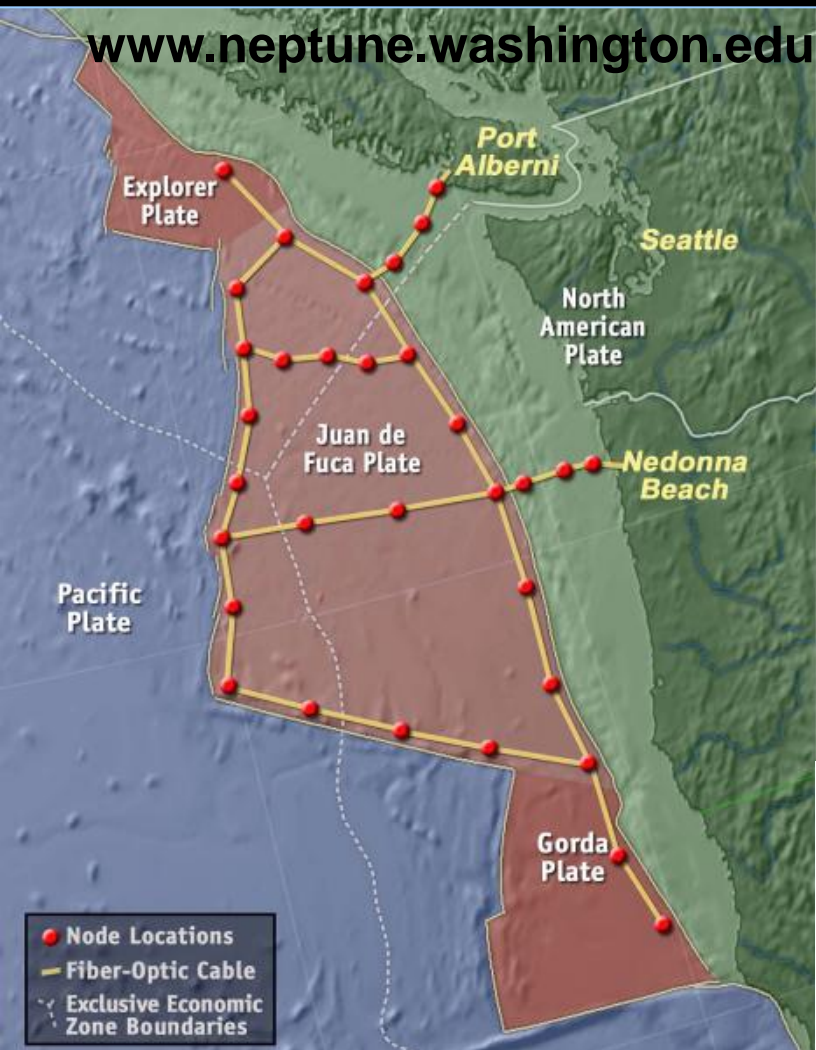


 30 Million Pixel Display
NCMIR Lab UCSD

Source: David Lee, Jason  R L I G H TSM
Leigh, EVL, UIC

New OptIPuter Driver: Gigabit Fibers on the Ocean Floor

A Working Prototype Cyberinfrastructure for NSF's ORION




- NSF ITR with Principal Investigators
 - John Orcutt & Larry Smarr - UCSD
 - John Delaney & Ed Lazowska –UW
 - Mark Abbott – OSU
- Collaborators at:
 - MBARI, WHOI, NCSA, UIC, CalPoly, UVic, CANARIE, Microsoft,



LOOKING (Laboratory for the Ocean Observatory Knowledge
Integration Grid) –
Adding Web Services to LambdaGrids

STARLIGHTSM

MARS New Gen Cable Observatory Testbed - Capturing Real-Time Basic Environmental Data



MARS Installation
Oct 2005 - Jan 2006



**Tele-Operated
Crawlers**



Source:
Jim
Bellingham
, MBARI

CAMERA: Community Cyberinfrastructure for Advanced Marine Microbial Ecology Research and Analysis

National LambdaRail
Direct Connect
Computation and Storage Complex

Funded by: Gordon and Betty Moore Foundation



Joint Partnership of:

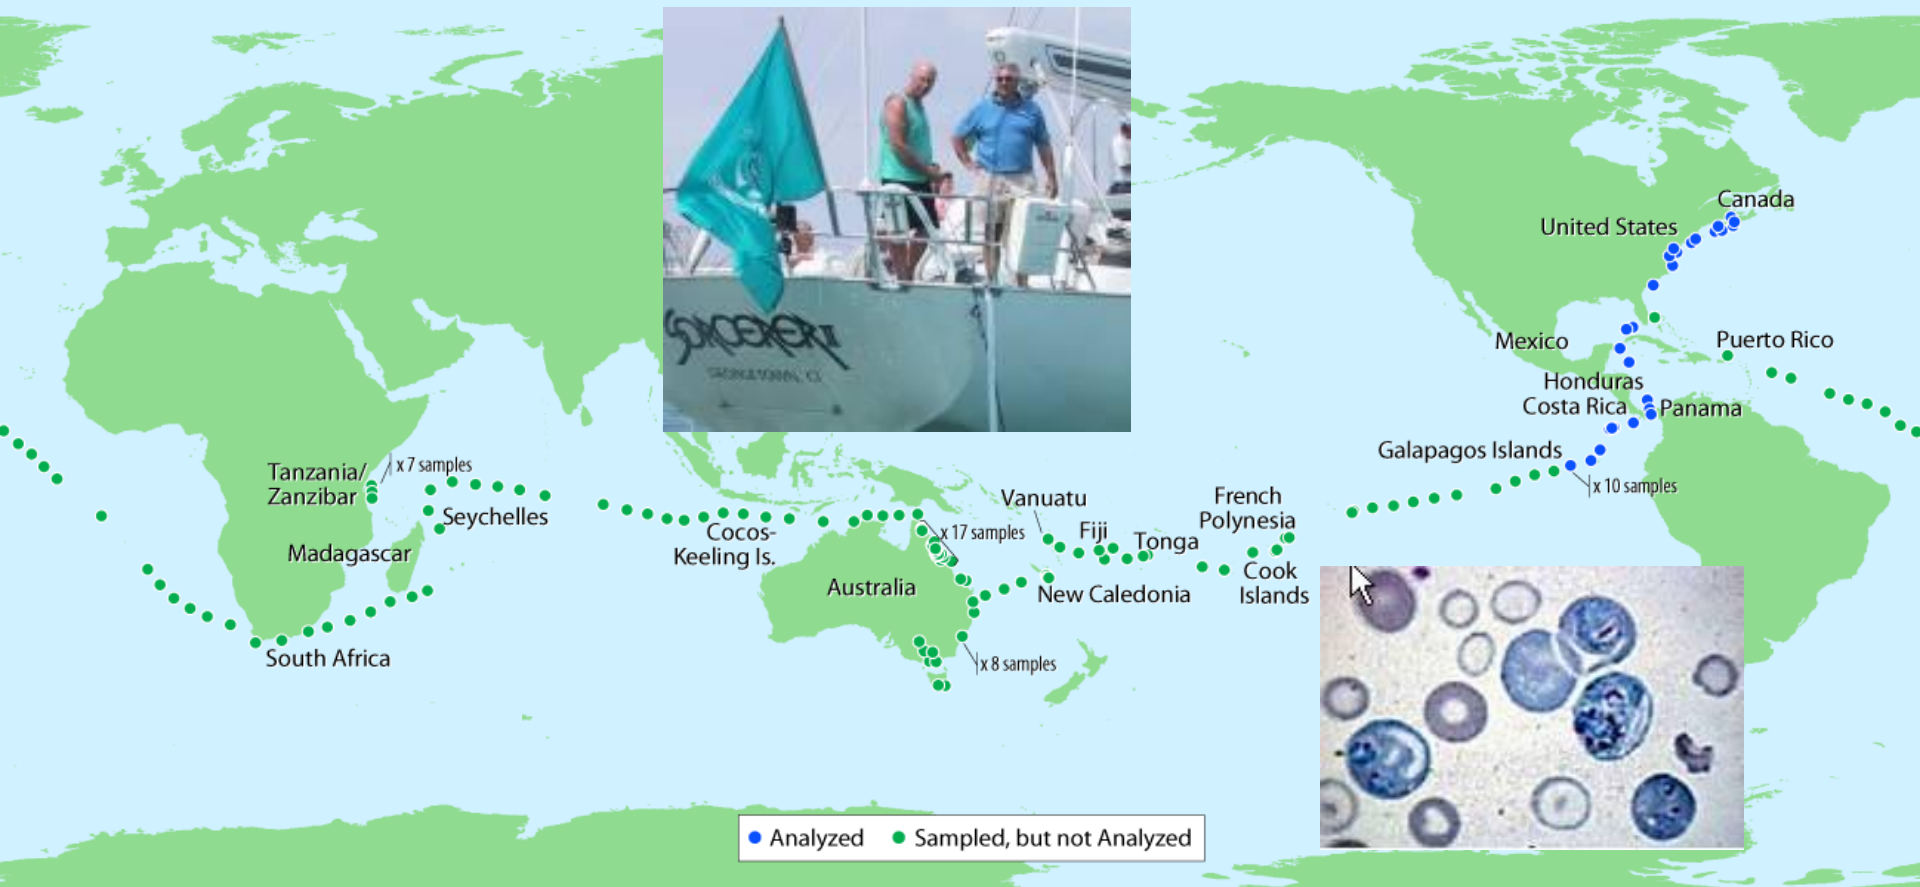


PI Larry Smarr

STARLIGHTSM

Marine Genome Sequencing Project

Measuring the Genetic Diversity of Ocean Microbes



**CAMERA will include
All Sorcerer II Metagenomic Data**


Source: Larry Smarr, C. Venter



J. Craig Venter
INSTITUTE



STARLIGHTSM

Craig Venter Announces Creation of the First Synthetic Life Form



Creation of a Bacterial Cell Controlled by a Chemically Synthesized Genome
Daniel G. Gibson,¹ John I. Glass,¹ Carole Lartigue,¹ Vladimir N. Noskov,¹ Ray-Yuan Chuang,¹ Mikkel A. Algire,¹ Gwynedd A. Benders,² Michael G. Montague,¹ Li Ma,¹ Monzia M. Moodie,¹ Chuck Merryman,¹ Sanjay Vashee,¹ Radha Krishnakumar,¹ Nacyra Assad-Garcia,¹ Cynthia Andrews-Pfannkoch,¹ Evgeniya A. Denisova,¹ Lei Young,¹ Zhi-Qing Qi,¹ Thomas H. Segall-Shapiro,¹ Christopher H. Calvey,¹ Prashanth P. Parmar,¹ Clyde A. Hutchison, III,² Hamilton O. Smith,² J. Craig Venter^{1,2,*}
Published in Science, May 20, 2010
Science DOI: 10.1126/science.1190719

OptlPortal Scalable Display Systems (Source: Smarr, DeFanti OptlPuter)



Advanced Visualization Enabled By Optical Lightpaths



The 200-million-pixel HIPerWall at Calit2 on the UC Irvine campus is now part of the OptiPortal Collaboratory.



Source: UCSD

STARLIGHTSM




Earth and Planetary Sciences are an OptIPuter Large Data Object Visualization Driver


EVL Varrier Autostereo 3D Image



SIO 18-Mpixel IBM OptIPuter Viz Cluster



SIO NIVE 3-Mpixel Panorama



Source: Smarr, Defanti: OptIPuter

LIGHTSM

3D Modeling and Simulation


Distributed Simulation Analysis

- Sandia National Laboratories, USA
- High Performance Computing Center Stuttgart (HLRS), Germany

Thanks to the Computer Services for Academic Research Centre (Manchester, UK), the Centre of Virtual Environments at University of Salford (UK), Tsukuba Advanced Computing Center (Japan) and Pittsburgh Supercomputing Center (USA) for additional supercomputing resources.

This application emphasizes distributed parallel supercomputing and a collaborative virtual-reality computation steering environment applied to Grand Challenge problems.


Source: Tom Defanti Dan Sandia, EVL






CORNEA

- Next step up in res/power: KAUST-Saudi Arabia
 - 24 Sony 4K projectors 100 Million Pixels/eye
 - 240,000 lumens!
 - Mechdyne/Iowa State
- Design based on original 1991 EVL CAVE





Source: Tom DeFanti, UCSD, KAUST

STARLIGHTSM



KAUST CORNEA




STARLIGHTSM

Digital Media (iGrid 2000, Yokohama Japan USA, Canada, Japan, Singapore, Netherlands, Sweden, CERN, Spain, Mexico, Korea)

GiDVN: Global Internet Digital Video Network

- Digital Video Working Group, Coordinating Committee for International Research Networks
- CERN, Switzerland
- APAN, Japan; KDD, Japan
- APAN-KR, Korea; Seoul National University, Korea
- SURFnet, The Netherlands
- DFSCA-UNAM, Mexico
- SingAREN, Singapore
- Universitat Politecnica de Catalunya, Spain
- Royal Institute of Technology, Sweden
- Int'l Center for Advanced Internet Research (iCAIR), Northwestern, USA



GiDVN projects have enhanced media capabilities for the next-generation Internet, enabling new applications to interoperate throughout the world.

www.icaair.org/inet2000

STARLIGHTSM

4K Media

4K Digital Media Ultra High Definition Digital Communications

Digital communications using SHD transmits extra-high-quality, digital, full-color, full motion images.

4k pixels horizontal, 2k vertical

4 * HDTV – 24 * DVD

4K Video is approximately 4X standard HD

HD = 720x1280 or 1080x1920 pixels

4K = 3840x2160 pixels



www.onlab.ntt.co.jp/en/mn/shd

STSMRLIGHT

Apps

Clusters

Dynamically
Allocated
Lightpaths

Switch Fabrics

Physical
Monitoring

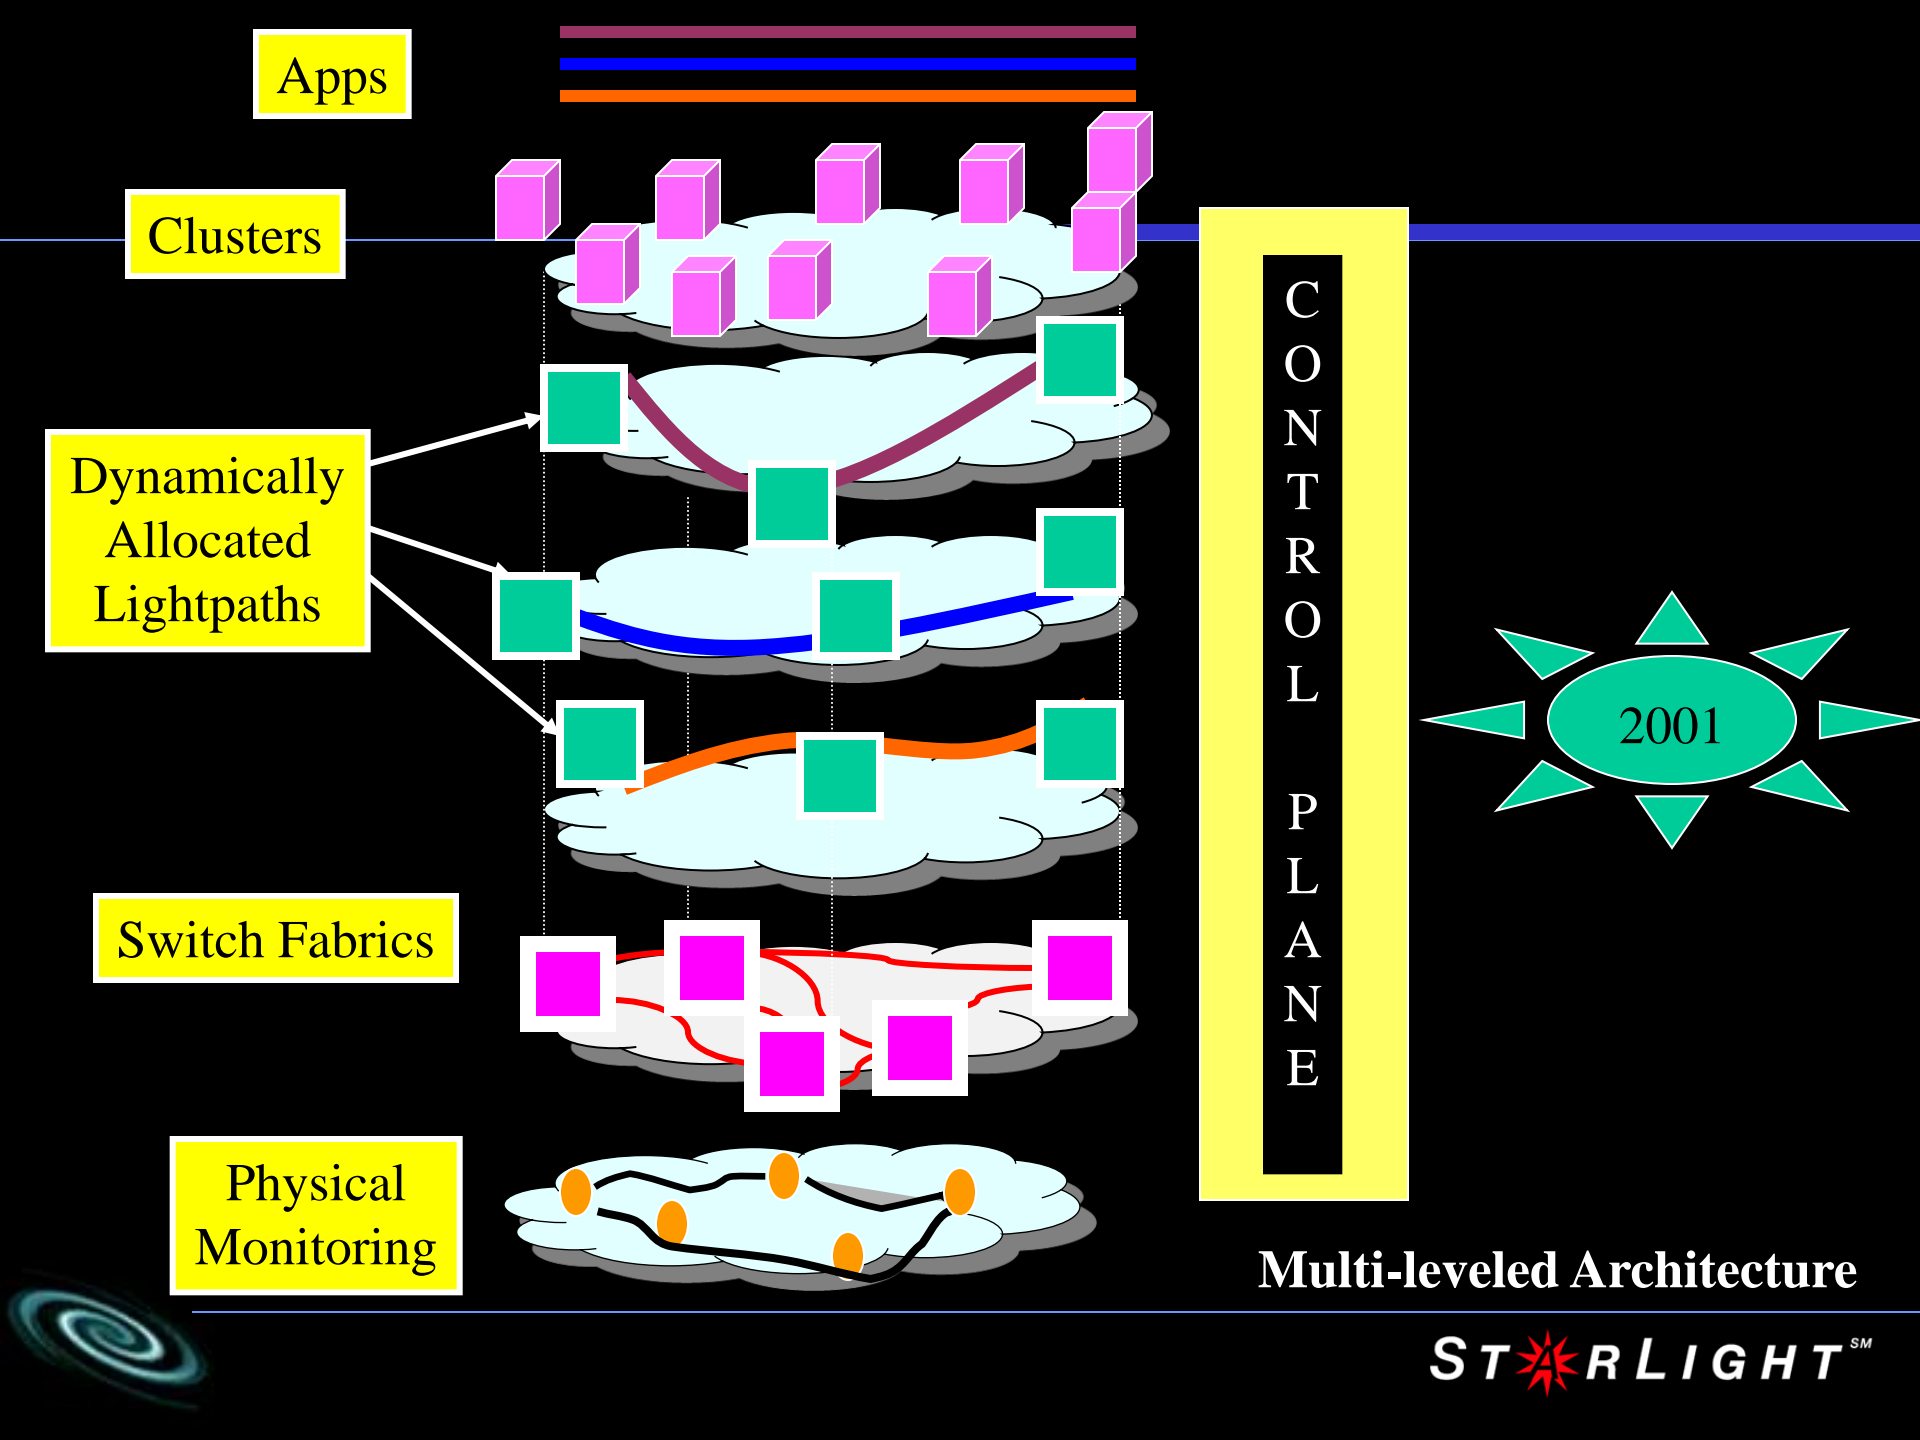
C
O
N
T
R
O
L

P
L
A
N
E

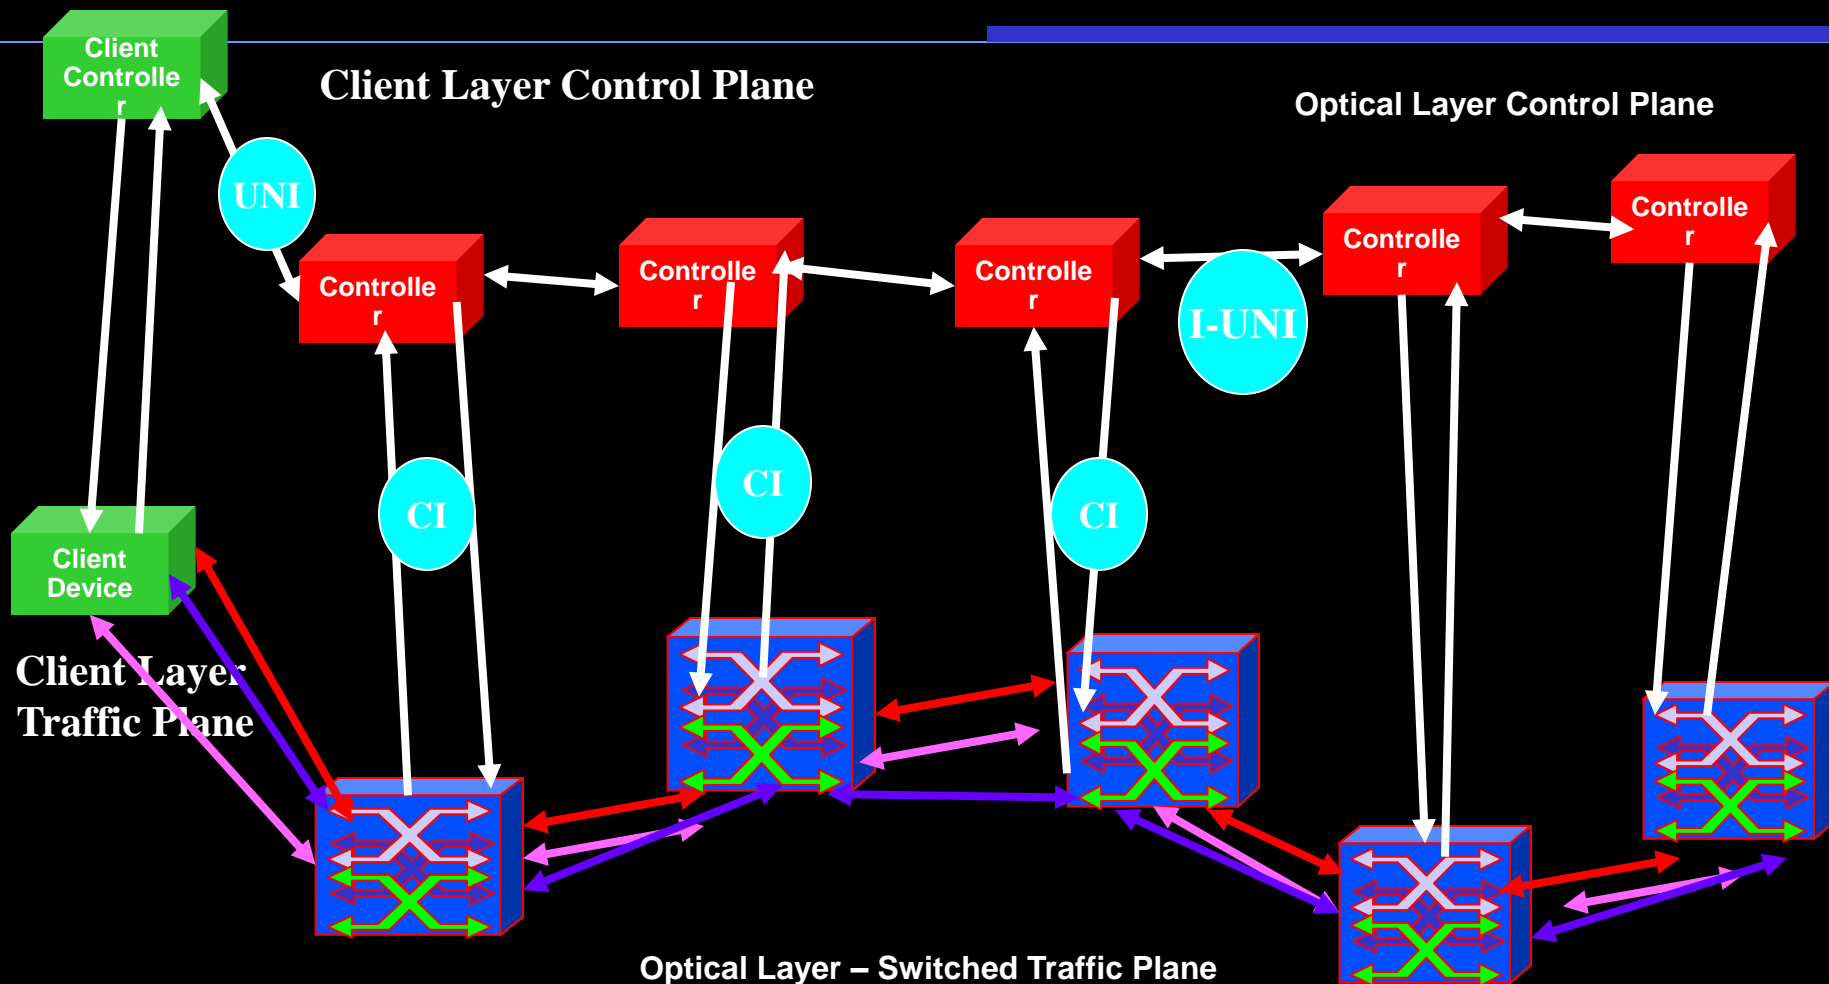
2001

Multi-level Architecture

STARLIGHTSM

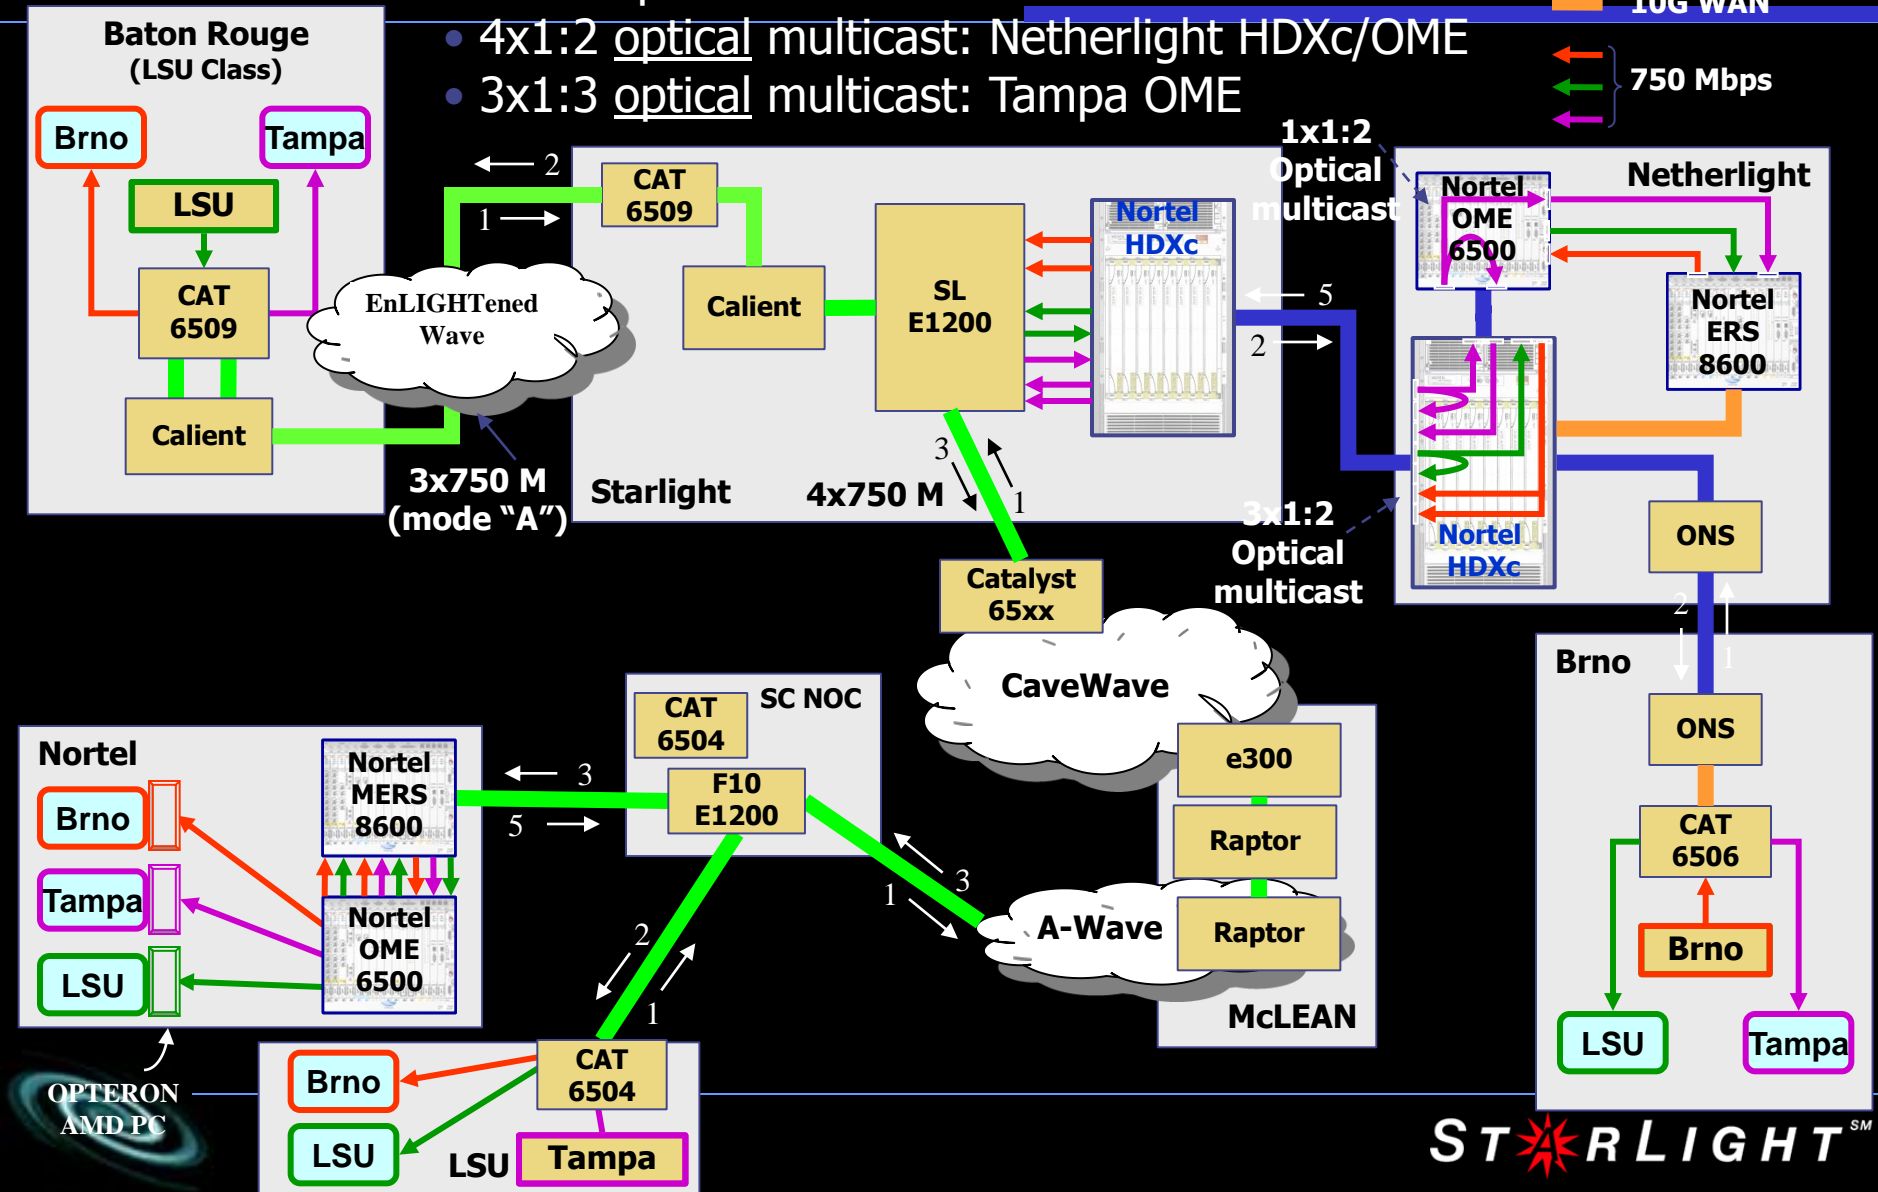



Optical Layer Control Plane



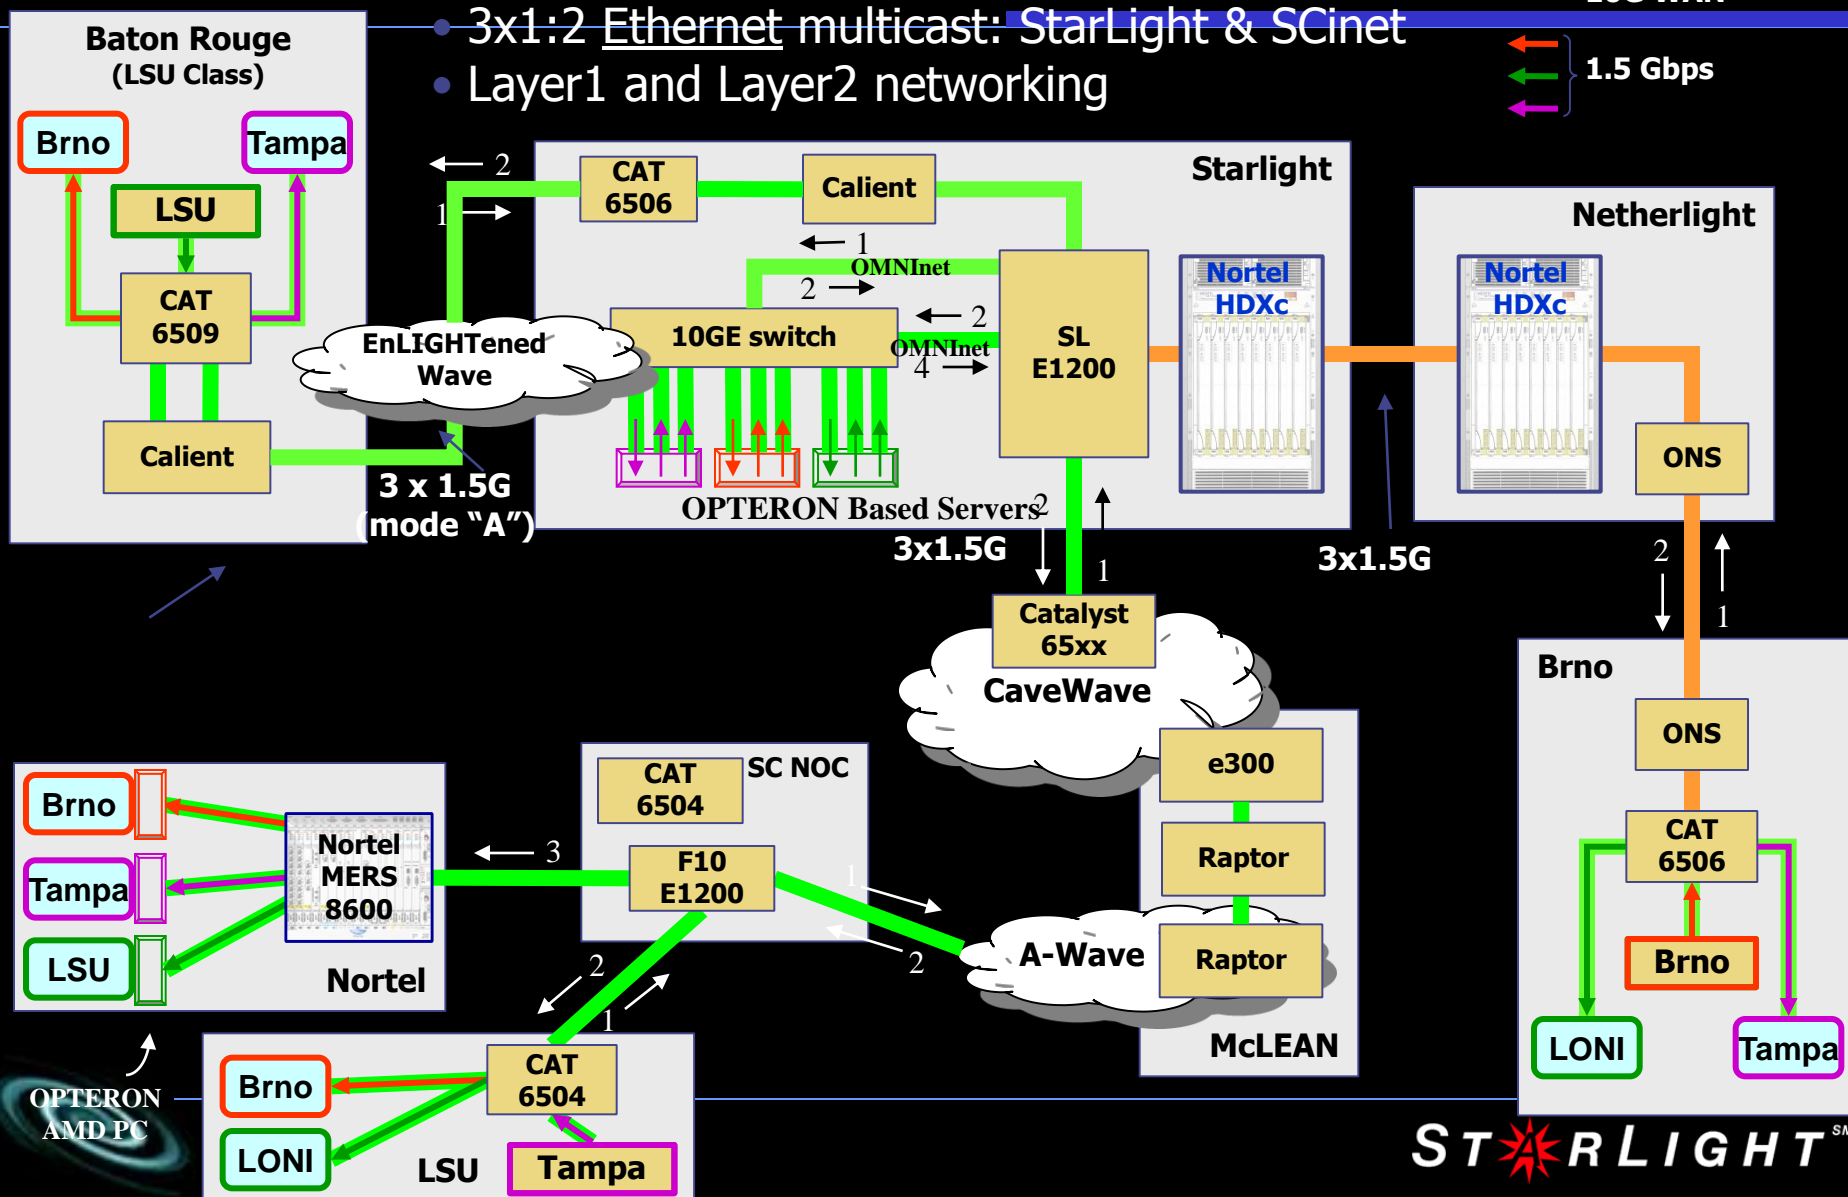

Topology for L1 Dynamic Optical Multicast SC06 Demo


- 750 Mbps Hi-Def video
- 4x1:2 optical multicast: Netherlight HDXc/OME
- 3x1:3 optical multicast: Tampa OME



Topology for L2 Multicast SC06 Demo (L2 Stream Duplication)

- 1.5 Gbps Hi-Def video
- 3x1:2 Ethernet multicast: StarLight & SCinet
- Layer1 and Layer2 networking





Introduction to High Performance Computing

Dr. Thomas Sterling
January 16, 2007



Week One: Class One
Introduction



**The First Public Demonstration
Of HPDMnet GDOM
Was Staged
At
GLIF 2007 in Prague,
Czech Republic
Sept 17-18, 2007**

High Performance Digital Media GLIF

- A Consortium of Research Centers From Around the World Has Formed a Cooperative Partnership To Explore the Key Issues Related to the Challenges and Opportunities Related to Using Lightpaths for High Performance Digital Media (HPDM)
- At the Annual Global LambdaGrid Workshop in Prague, Demonstrations Have Been Designed to Show the Current Project Status
- Multiple Sites Require High Performance/High Volume/High Definition Digital Media Streaming Simultaneously Among All Locations (Multi-Point to Multi-Point)



Music Grid Canada - Prague

GrabFileEditCaptureWindowHelp

Current Status

Reservations

/O=MCNC/OU=GCNS/OU=mcnc.org

/C=JP/O=AIST GTRC/CN=Hidemoto

/O=MCNC/OU=GCNS/OU=mcnc.org

/O=MCNC/OU=GCNS/OU=mcnc.org

/O=MCNC/OU=GCNS/OU=mcnc.org

/O=MCNC/OU=GCNS/OU=mcnc.org

/O=MCNC/OU=GCNS/OU=mcnc.org

/O=Louisiana State University/OU=

/O=Louisiana State University/OU=

/O=Louisiana State University/OU=

/O=Louisiana State University/OU=

/O=Louisiana State University/OU=

/O=Louisiana State University/OU=

/O=Louisiana State University/OU=

/O=Louisiana State University/OU=

/C=JP/O=AIST GTRC/CN=Atsuko Ta

glif02.cesnet.cz:1 (glambda)

1,10,33: icmp_seq=3654 ttl=62 time=277 ms

1,10,33: icmp_seq=3655 ttl=62 time=277 ms

1,10,33: icmp_seq=3656 ttl=62 time=277 ms

1,10,33: icmp_seq=3657 ttl=62 time=277 ms

1,10,33: icmp_seq=3658 ttl=62 time=277 ms

1,10,33: icmp_seq=3659 ttl=62 time=277 ms

1,10,33: icmp_seq=3660 ttl=62 time=277 ms

1,10,33: icmp_seq=3661 ttl=62 time=277 ms

1,10,33: icmp_seq=3662 ttl=62 time=277 ms

1,10,33: icmp_seq=3663 ttl=62 time=277 ms

1,10,33: icmp_seq=3664 ttl=62 time=277 ms

pověda

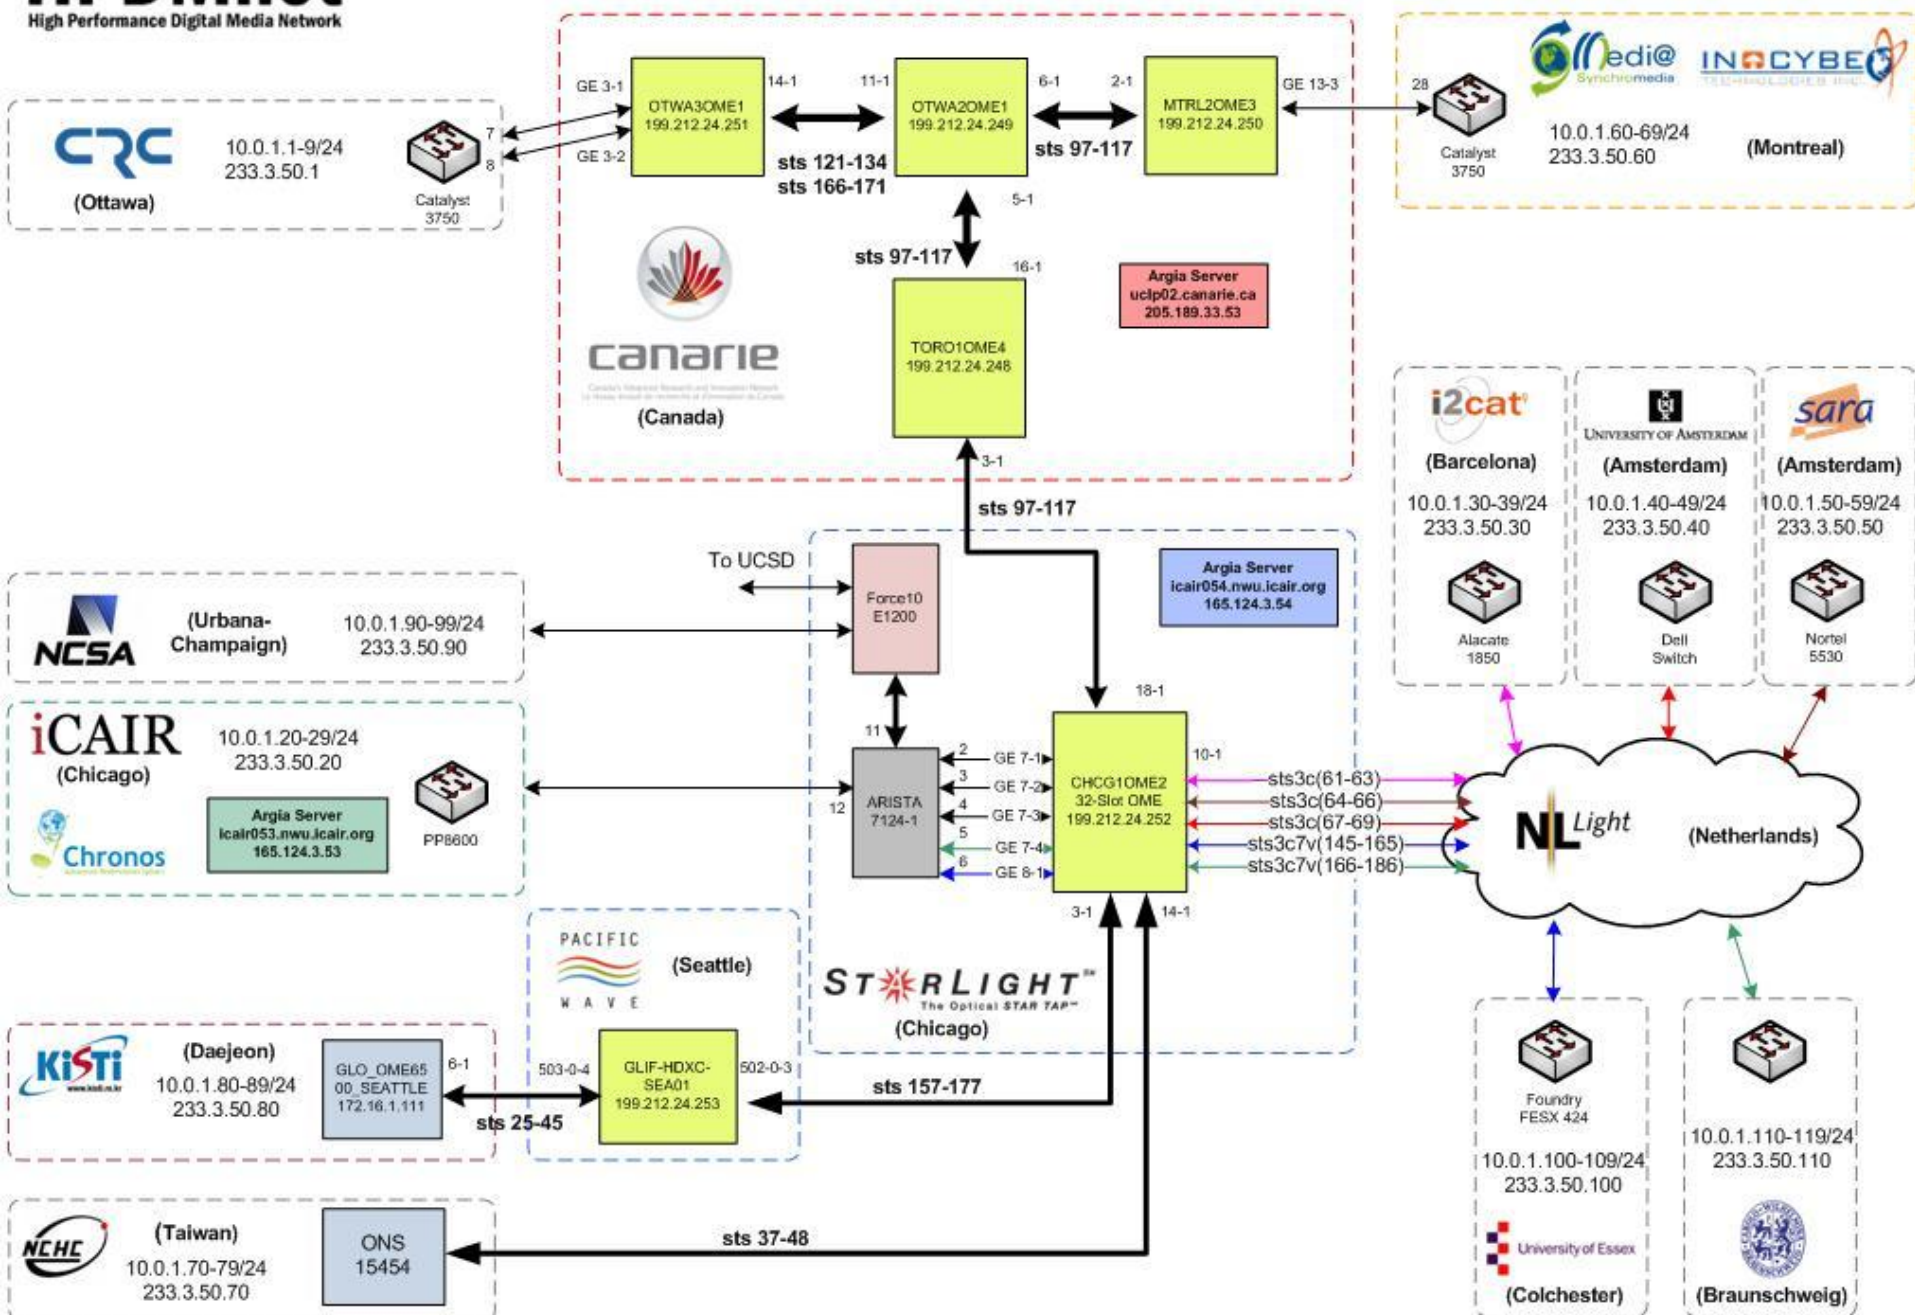
Time Table (Networks)

	2:30	2:32	2:34	2:36	2:38	2:40	2:42
TKB-KAN (2.0)							
TKB-FUK (2.0)							
TKB-X1N (1.0)							
AKB-OSA (1.0)							
AKB-KHN (1.0)							
AKB-X1S (3.0)							
RA1-X1U (1.0)							
RAH-X1U (1.0)							
BTH-X1U (1.0)							
AKB-AKH (1.0)							
X1N-X1S (5.0)							
X1N-X1U (4.0)							
X1S-X1U (5.0)							
KAN-FUK (2.0)							
KAN-X1N (2.0)							
OSA-KHN (1.0)							
OSA-X1S (1.0)							
KHN-X1S (1.0)							
RA1-RAH (1.0)							
RA1-BTH (1.0)							
RAH-BT2 (1.0)							
RAH-BTH (1.0)							
BT2-BTH (1.0)							
TKB-X2N (2.0)							
KAN-X2N (1.0)							
FUK-X1N (2.0)							
FUK-X2N (2.0)							
AKB-X2S (1.0)							
OSA-X2S (2.0)							

00:11:31

HPDMnet

High Performance Digital Media Network



RSV

Server: RSV2_CENIC, Service Name: Bassett-Hand

Get Channels Start RSV

주소(D) http://198.164.40.210:8080/savoir/ 이동

HSVO

SAVOIR - HSVO Project
Web Services Development Team
Release Web Page


For the latest demo
[Click Here](#)

Not Knowing How To Use The
[Download User Guide](#)

[SAVOIR GALLERY](#)

iCAIR 3D HD HPDMnet Demonstration At GLIF Deajong, South Korea Oct 2009






Testbed Demonstrations With National Science Foundation at the Annual Conference of The American Association for the Advancement of Science February 2009

Using An Optical Fiber Extension from StarLight/GLIF



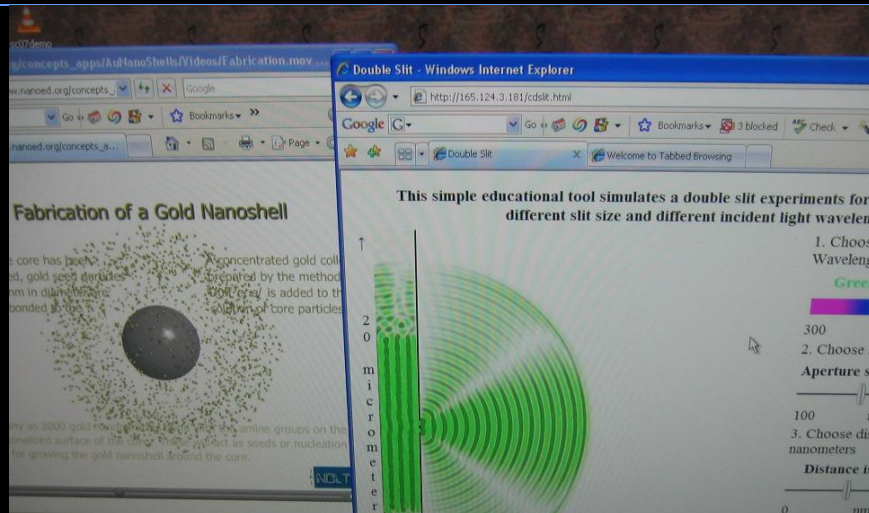
TransLight / StarLight, University of Illinois at Chicago,
Northwestern University



Virtual Instruments for Science

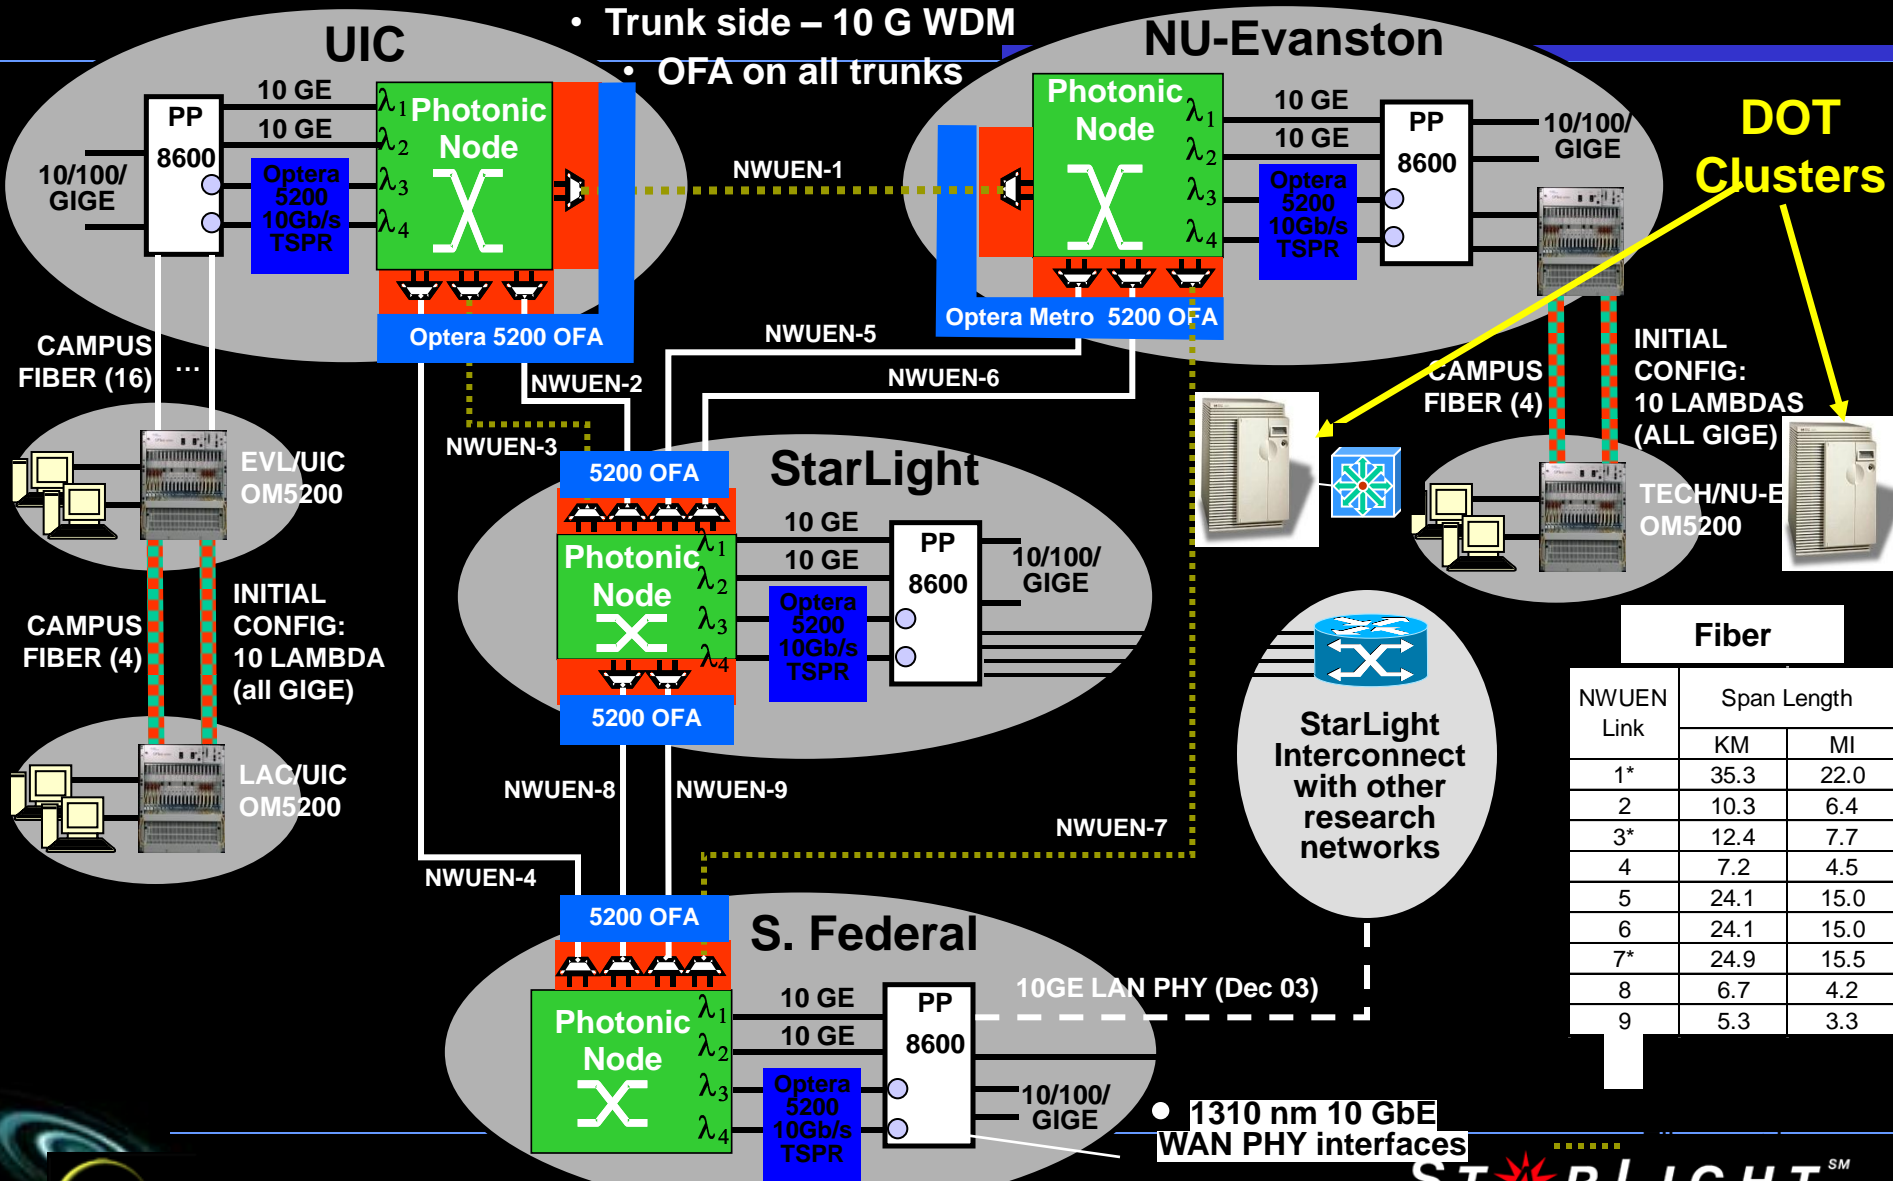


Cooperative Project iCAIR, NCLT



OMNinet Network Configuration Phase 2 (Extended Via Demonstrations Nationally and Internationally)

- 8x8x8λ Scalable photonic switch
- Trunk side – 10 G WDM
- OFA on all trunks






Default configuration:
Tribes can be moved as needed
Could have 2 facing L2 SW

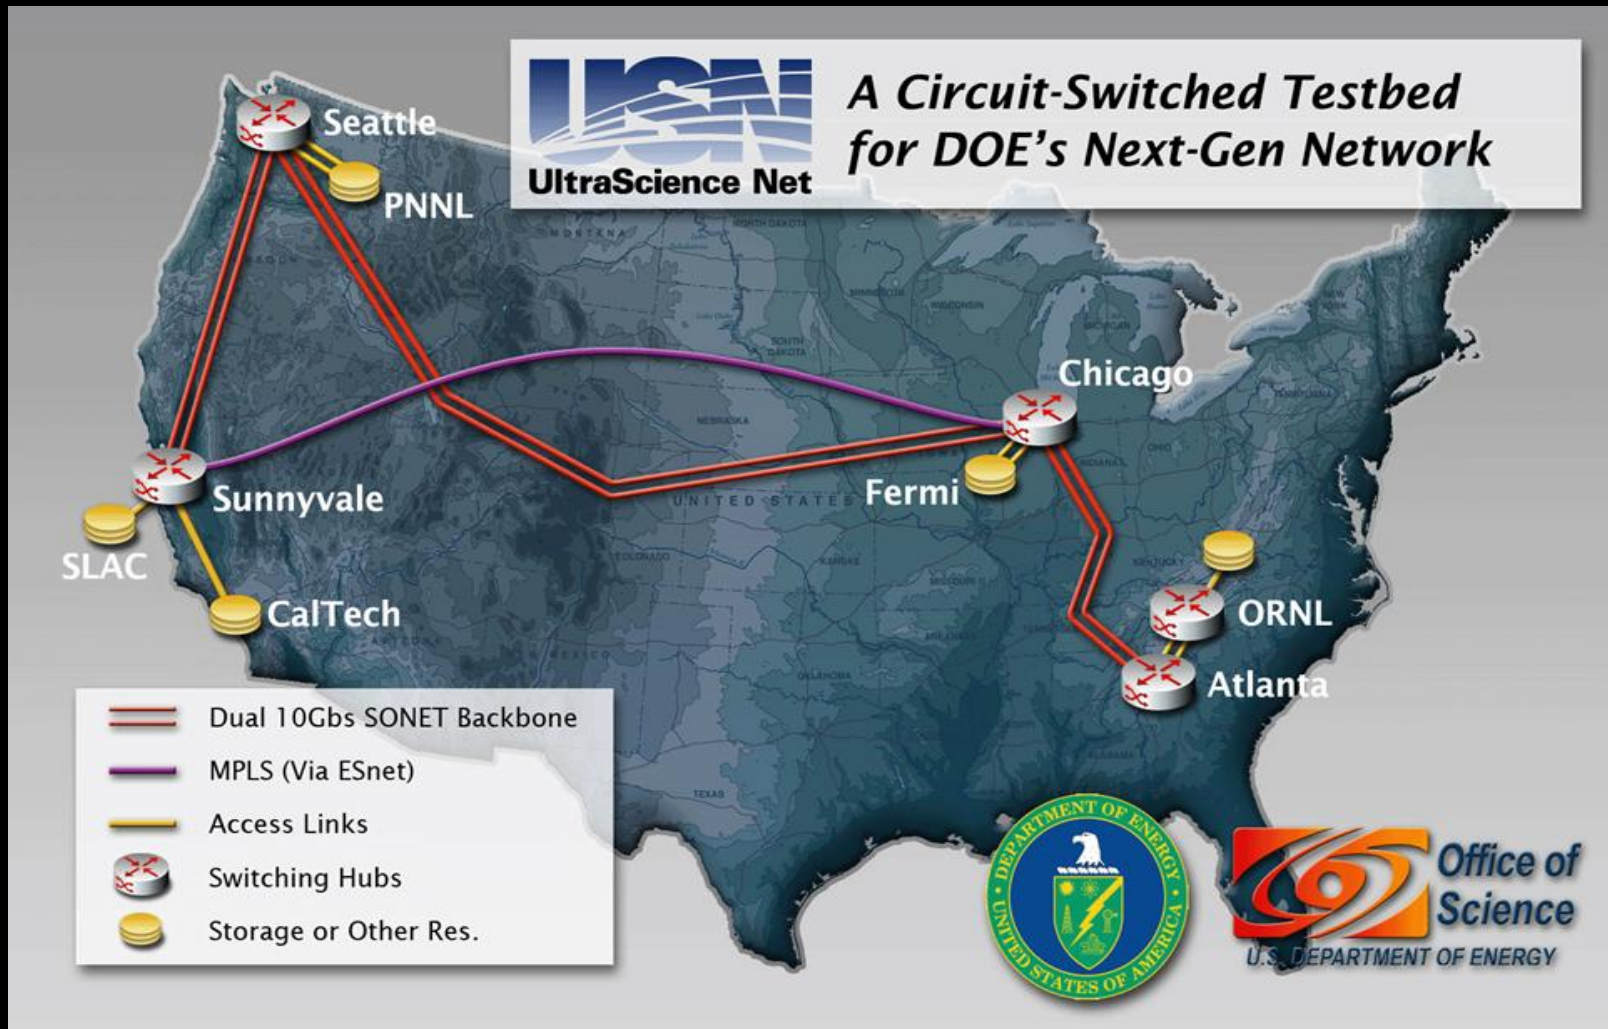
Only TFEC link can support OC-192c (10G Wan) operation

Non -TFEC link used to transport Ge traffic



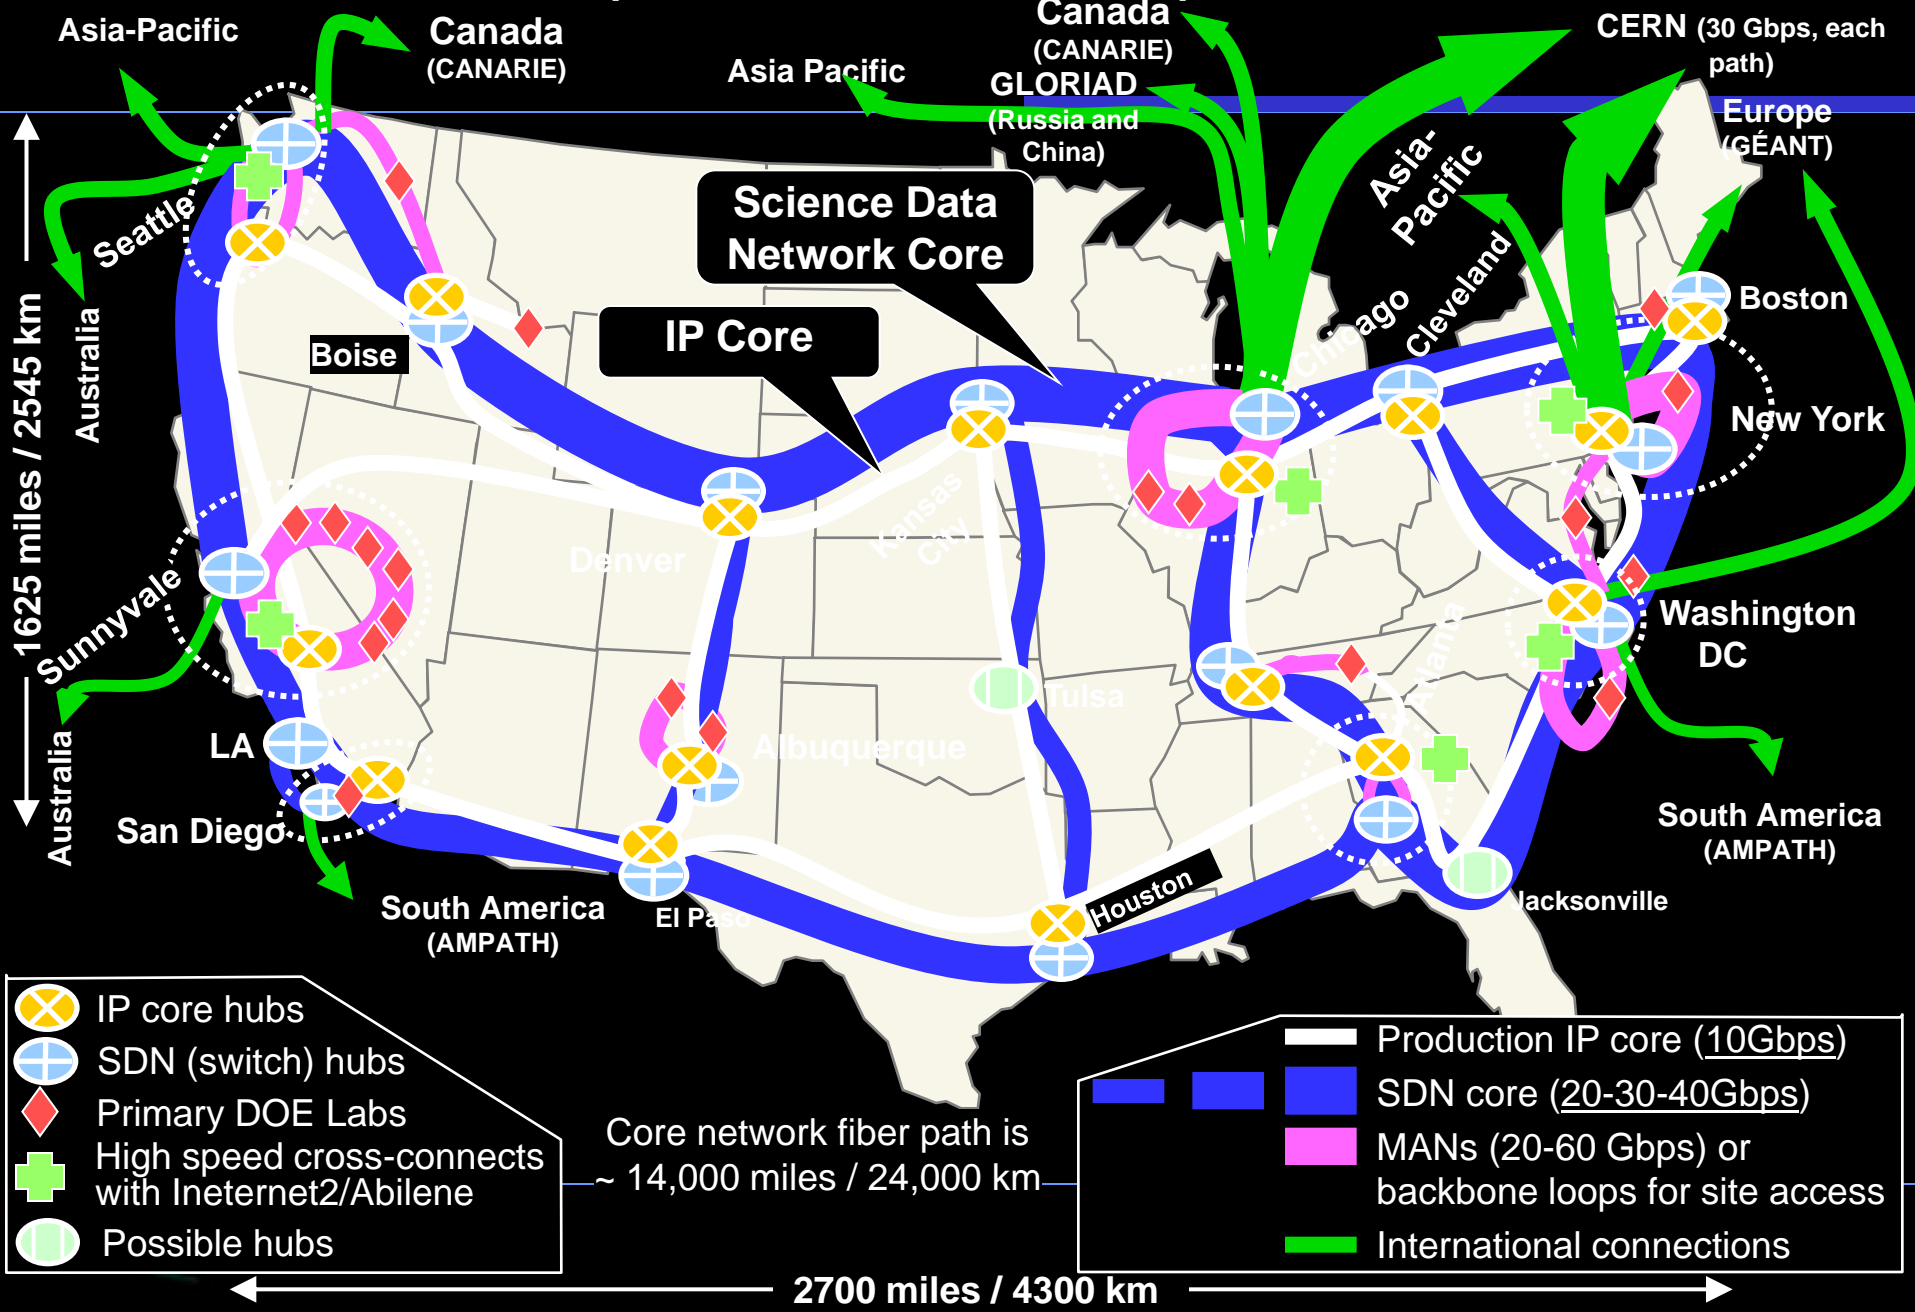
Trib Content	
OC-192 – with TFEC	16
OC-192 – without TFEC	12
Ge	8
OC-48	0

DOE's UltraScience Net at StarLight



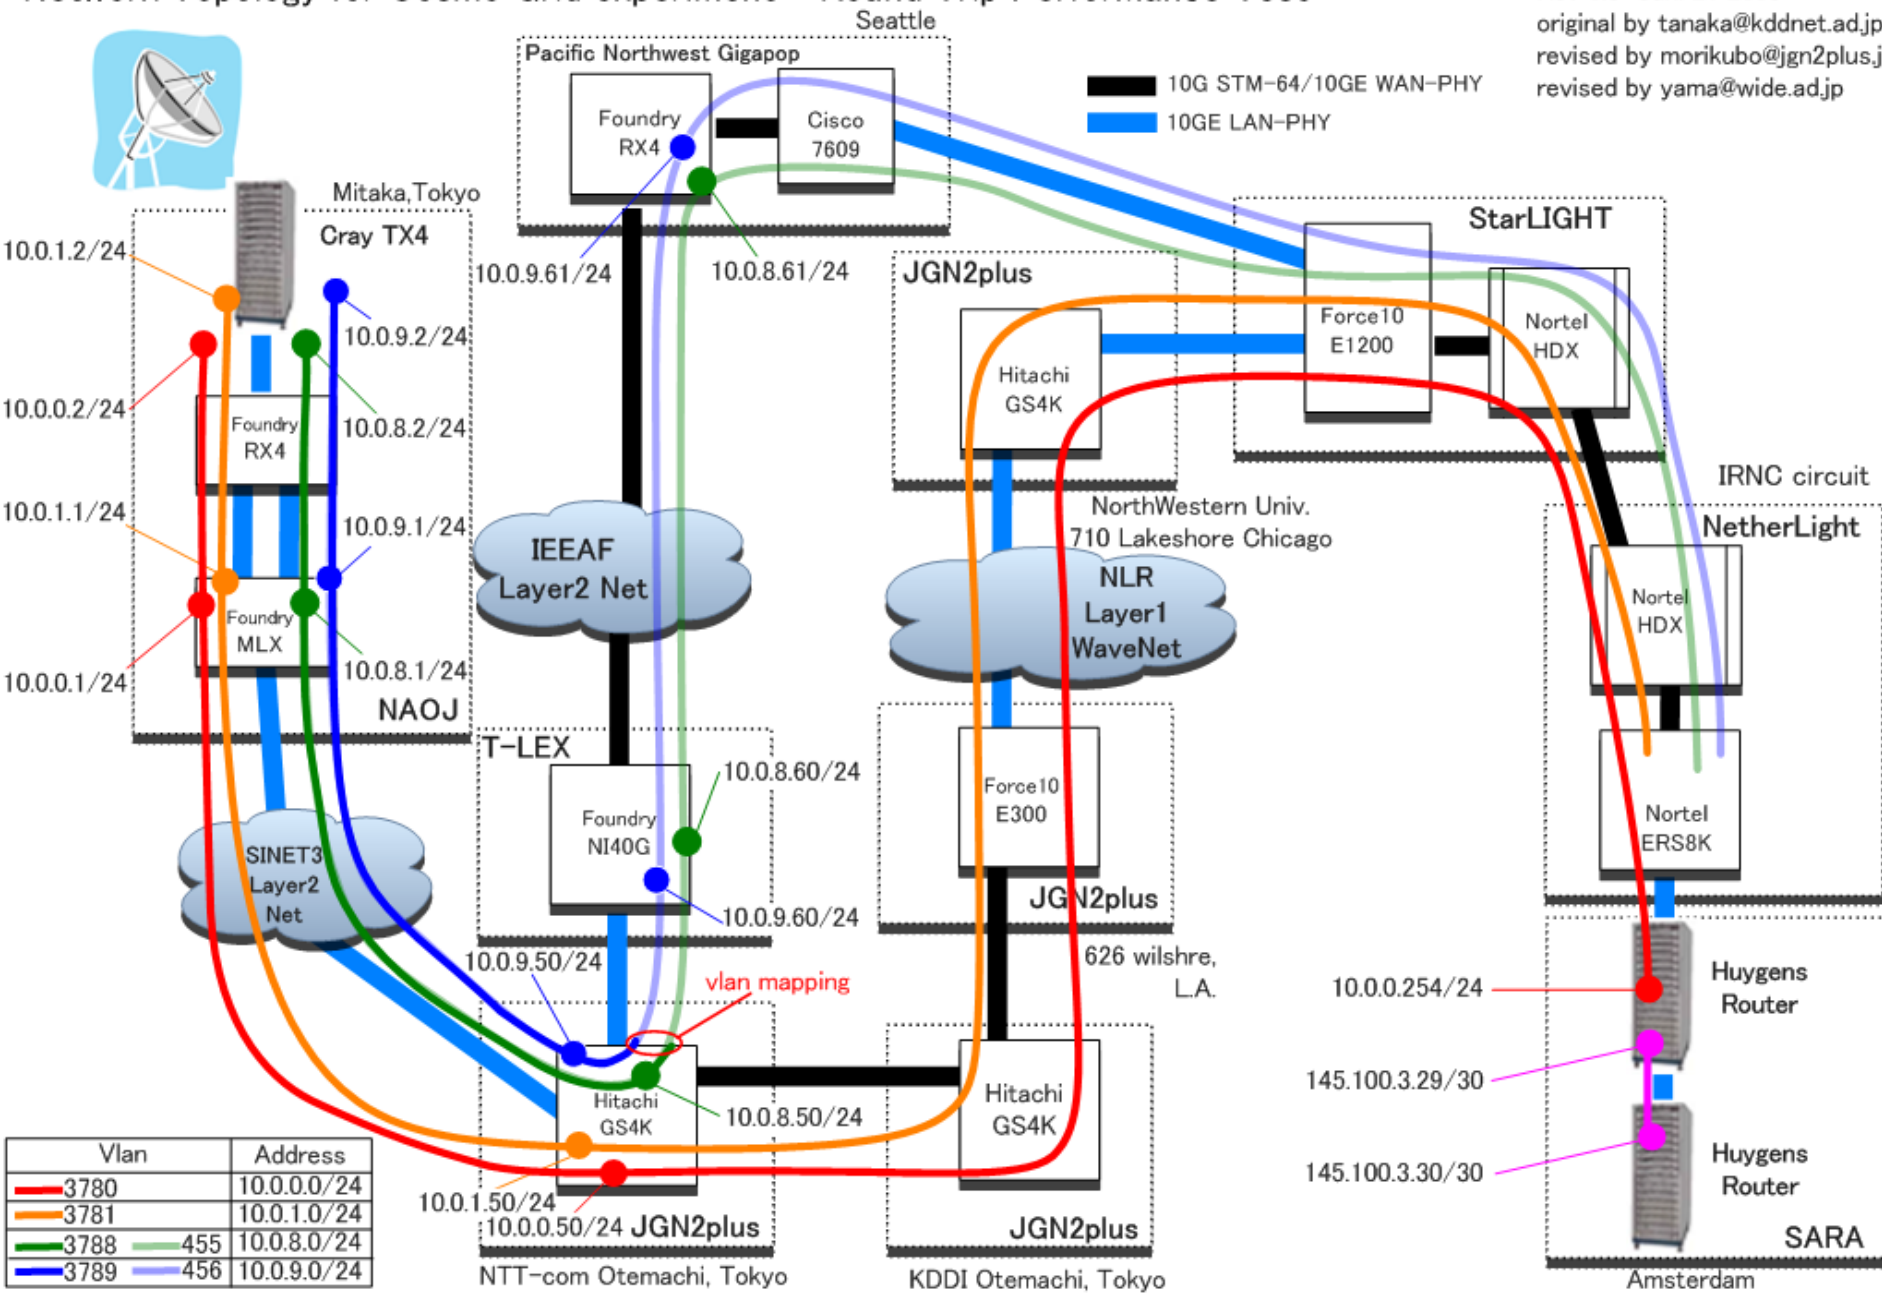
ESnet4 Planed Configuration

40-50 Gbps in 2009-2010. 160-400 Gbps in 2011-2012



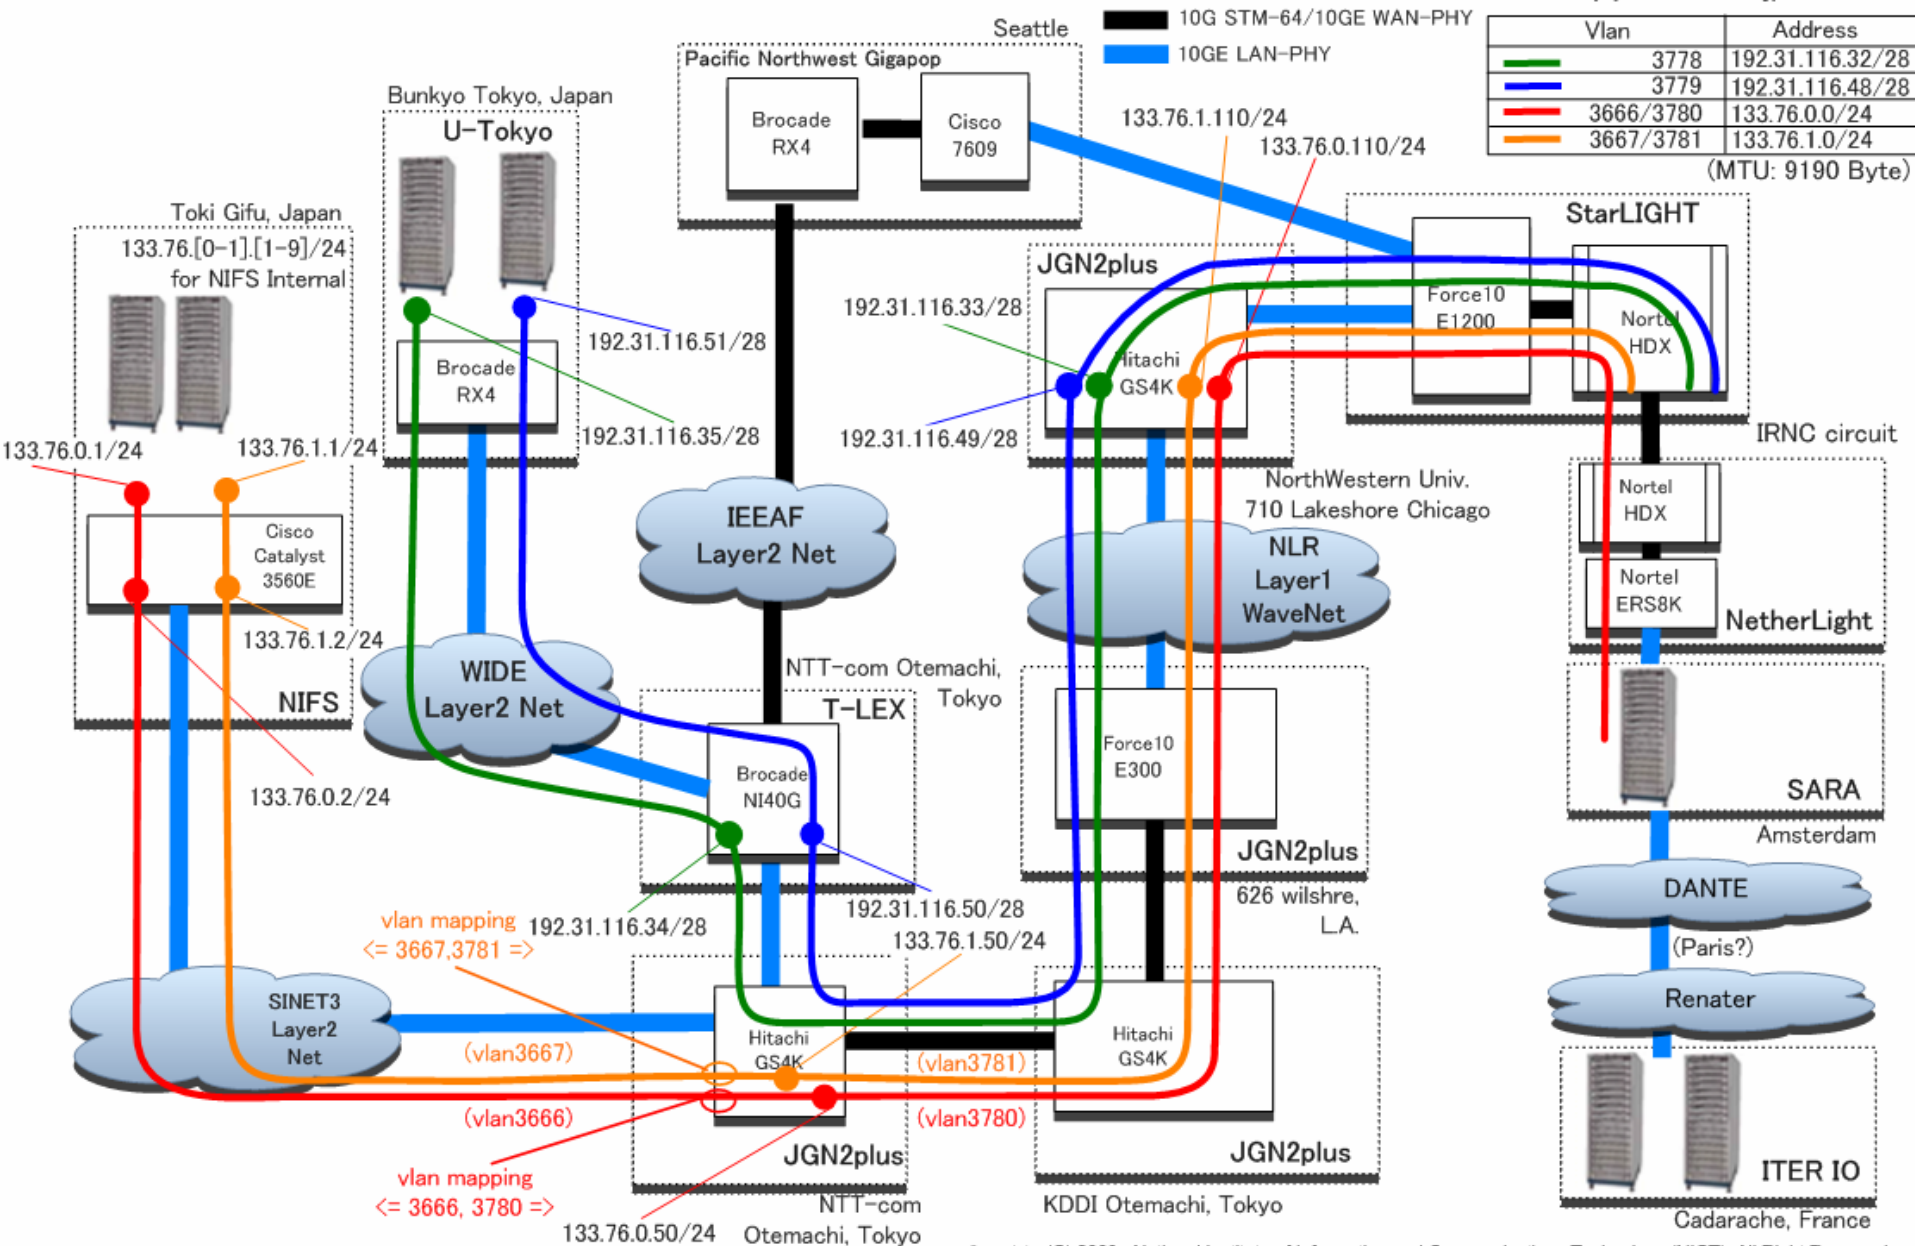
Network Topology for Cosmo Grid experiment – Round Trip Performance Test –

Rev 0.8 Jan. 21 2008
original by tanaka@kddnet.ad.jp
revised by morikubo@jgn2plus.jp
revised by yama@wide.ad.jp

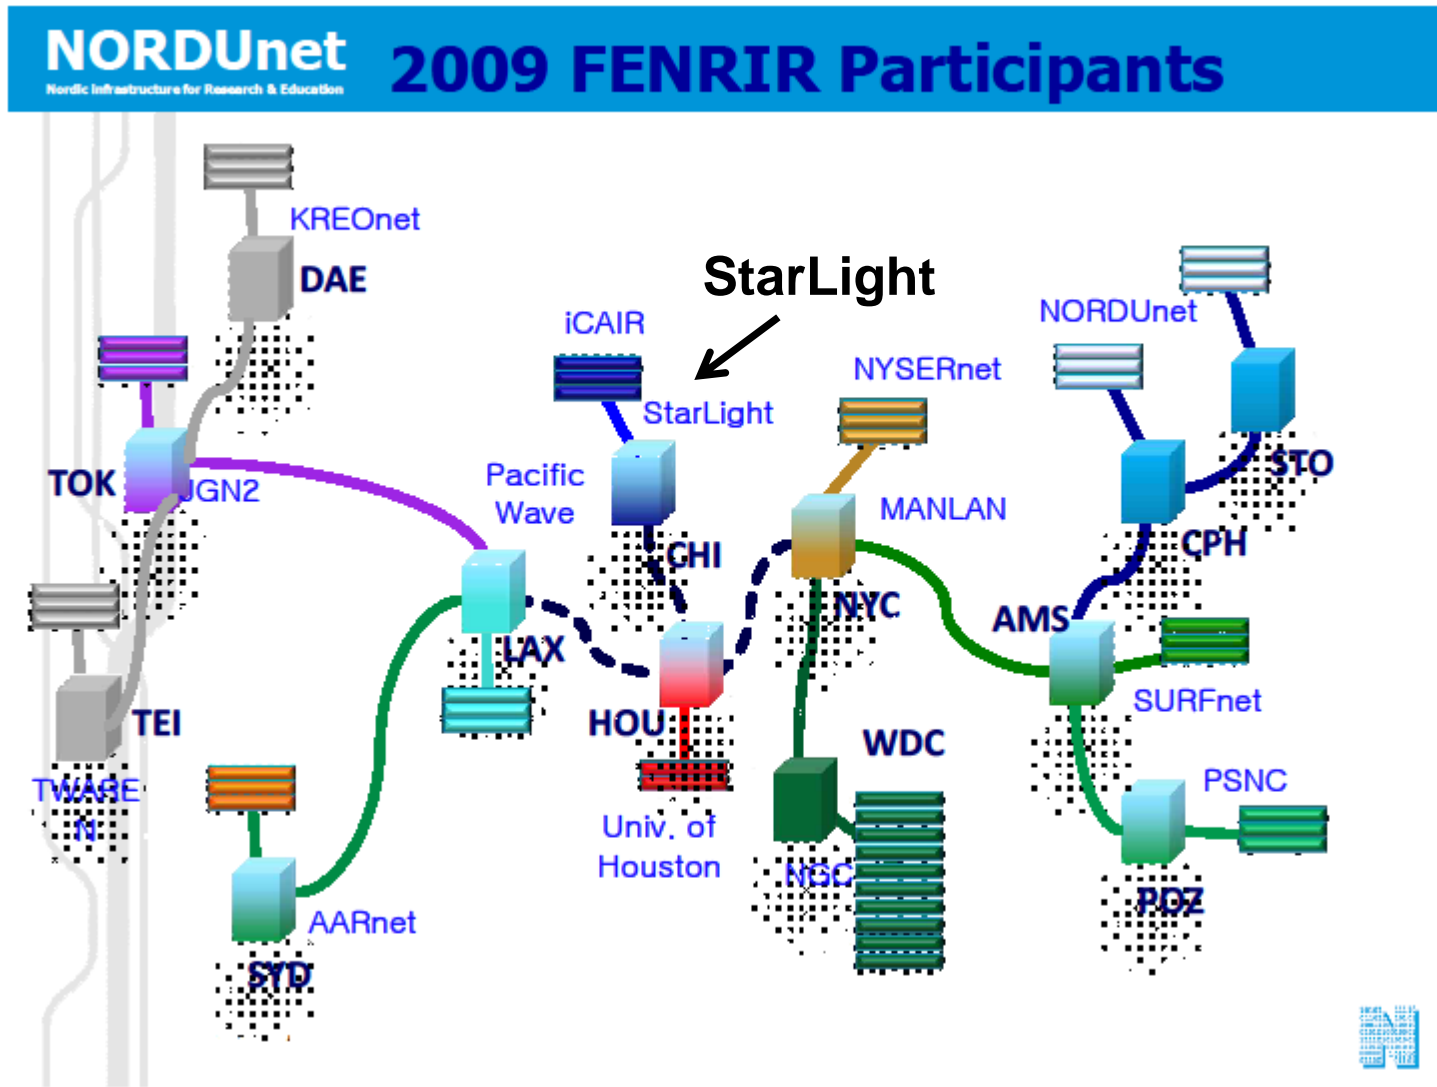


Network Topology for Japan-France High-Speed Data Transfer Experiment

Rev 0.4 Sep. 21 2009
original format by tanaka@kddnet.ad.jp
revised by yama@wide.ad.jp



FENRIR - Federated Experimental Network Resources for International Research




VINI Testbed: Virtual Network Infrastructure





High Performance Flow Switching Network (HPFSnet) Demonstration – ANAGRAN enabled, using the National TeraFlow Network and the National Lambda Rail (NLR)

International Center for
Advanced Internet Research



Chicago


OC192/STM64
10GE LAN
10 G WAN
750-~980 Mbps

STARLIGHTSM

OMNInet

NLR NATIONAL LAMBDA RAIL light the future

National Center for Data Mining



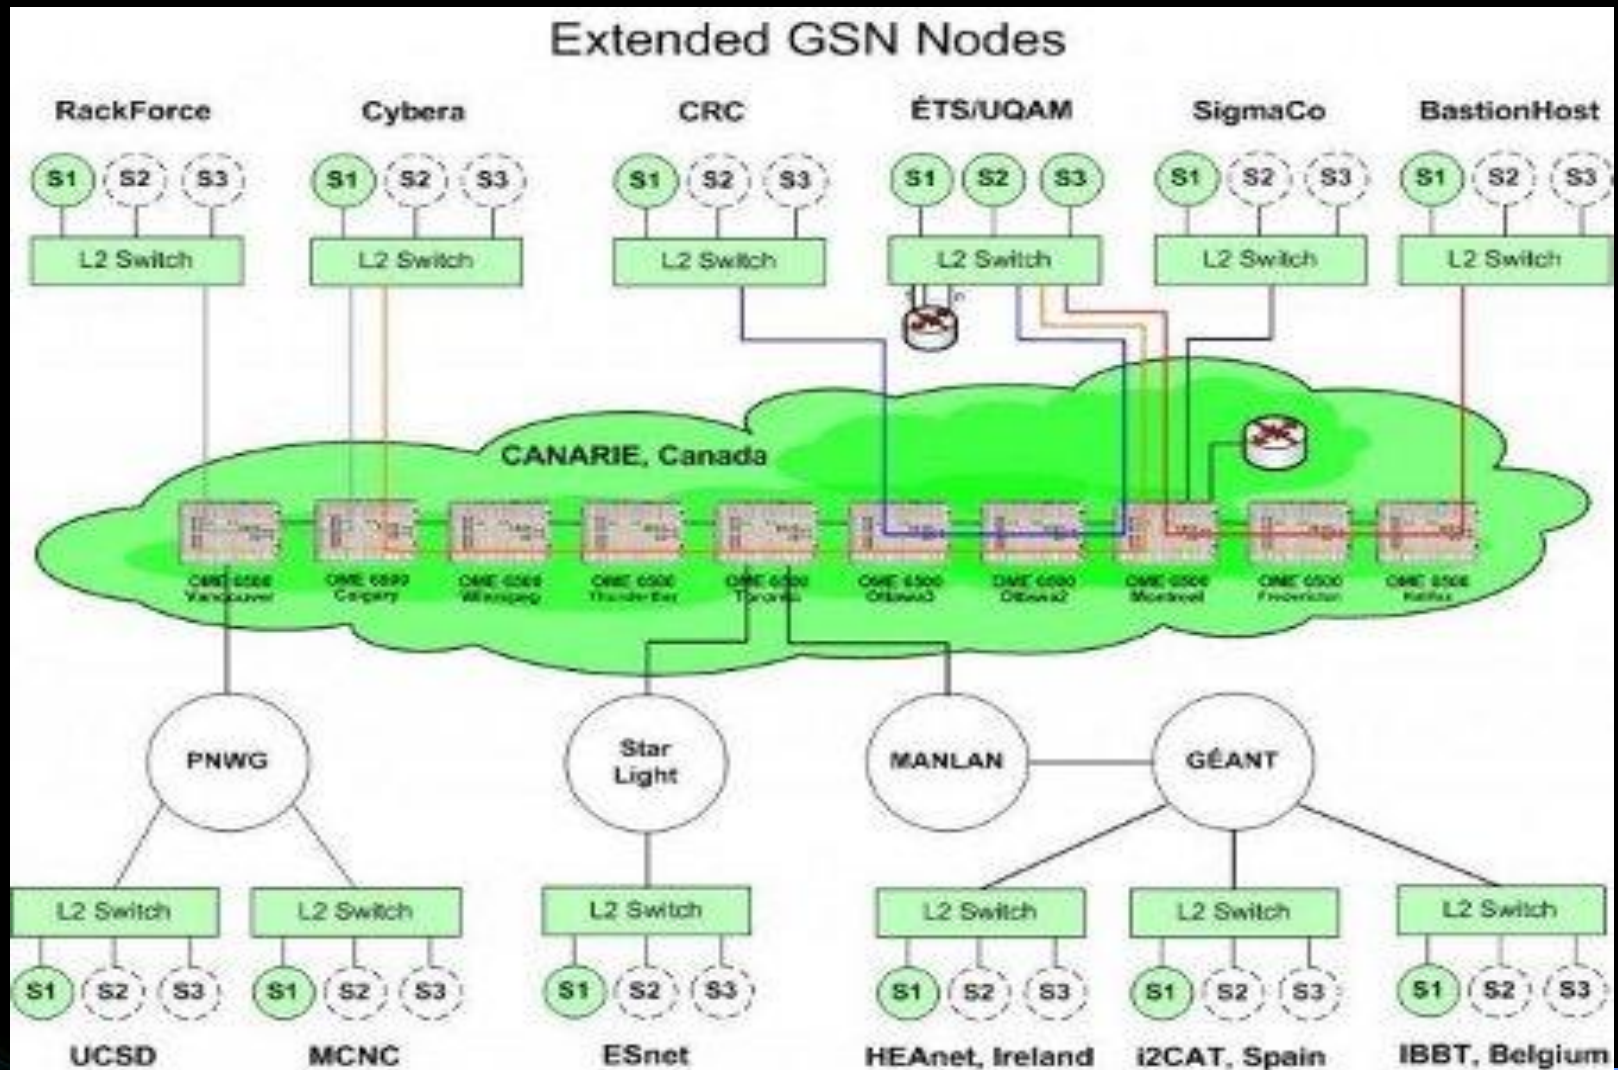
Reno

Showcasing high performance services and capabilities made possible on national networks using sophisticated flow control engineering techniques enabled by Anagran.


- Providing fine-grained control over all network traffic end-to-end
- Allowing targeted delivery of quality of service to support all types of applications, including digital media
- Demonstration Partners: Anagran, International Center for Advanced Internet Research, National Center for Data Mining, National LambdaRail, StarLight International Communications Exchange

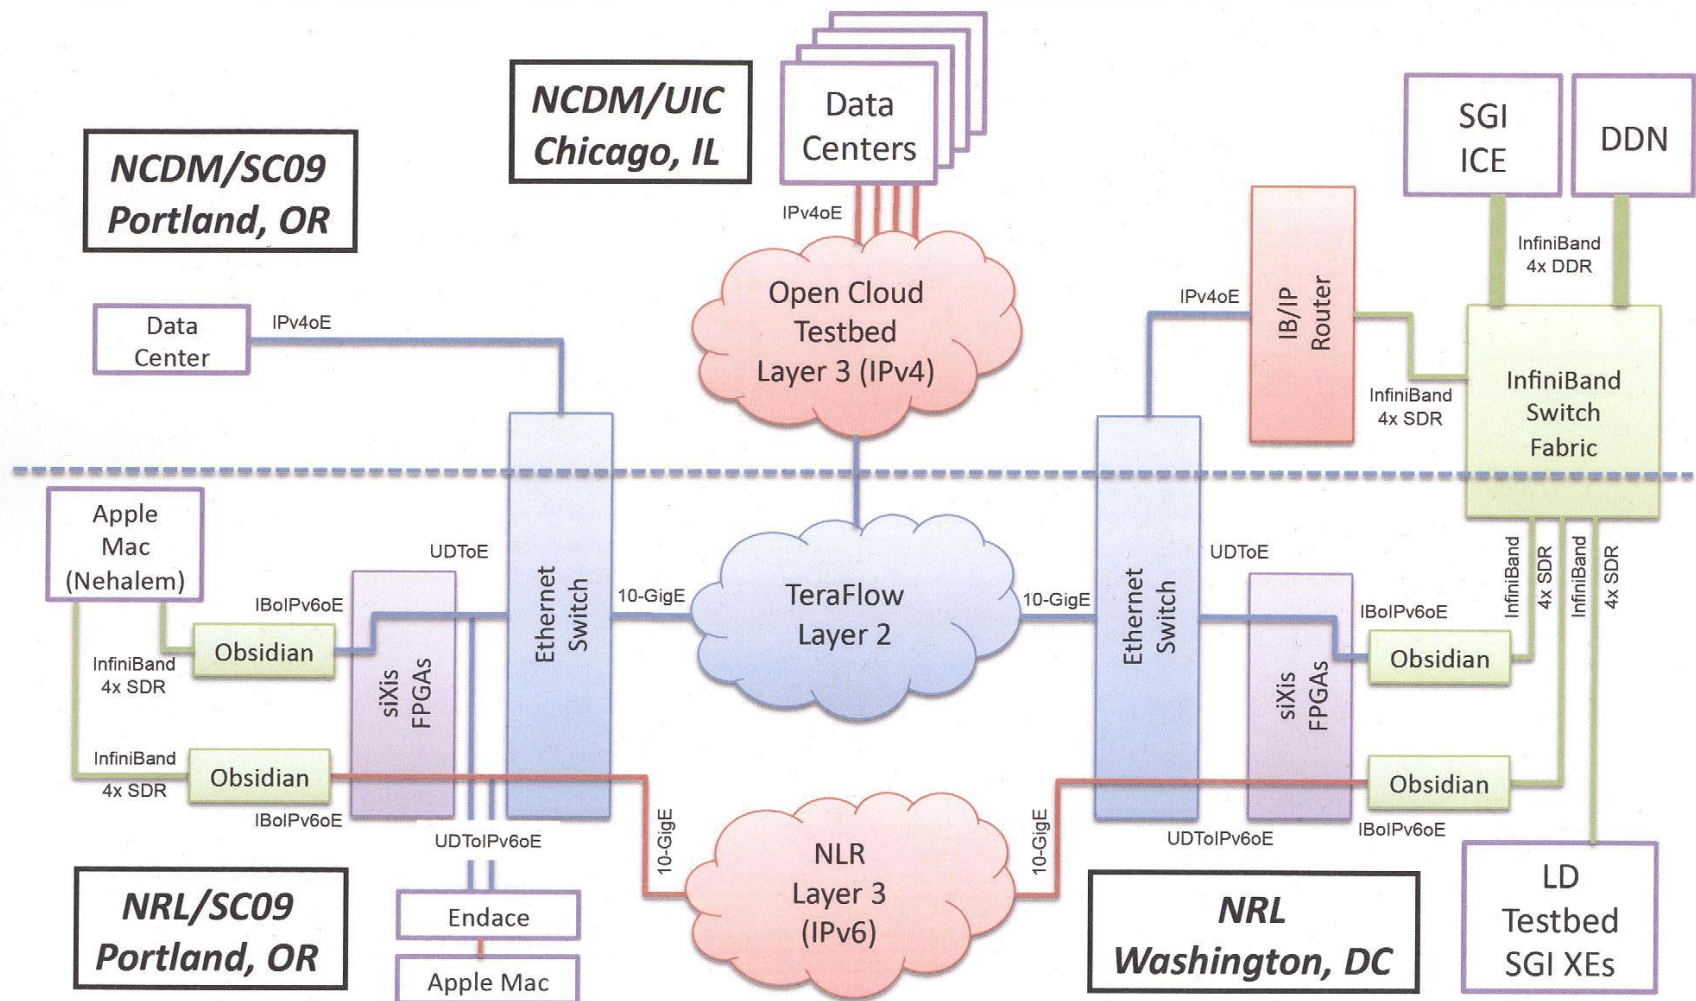
STARLIGHTSM

Green Star Network



iGENI GIST-BEN-KREONET Testbed

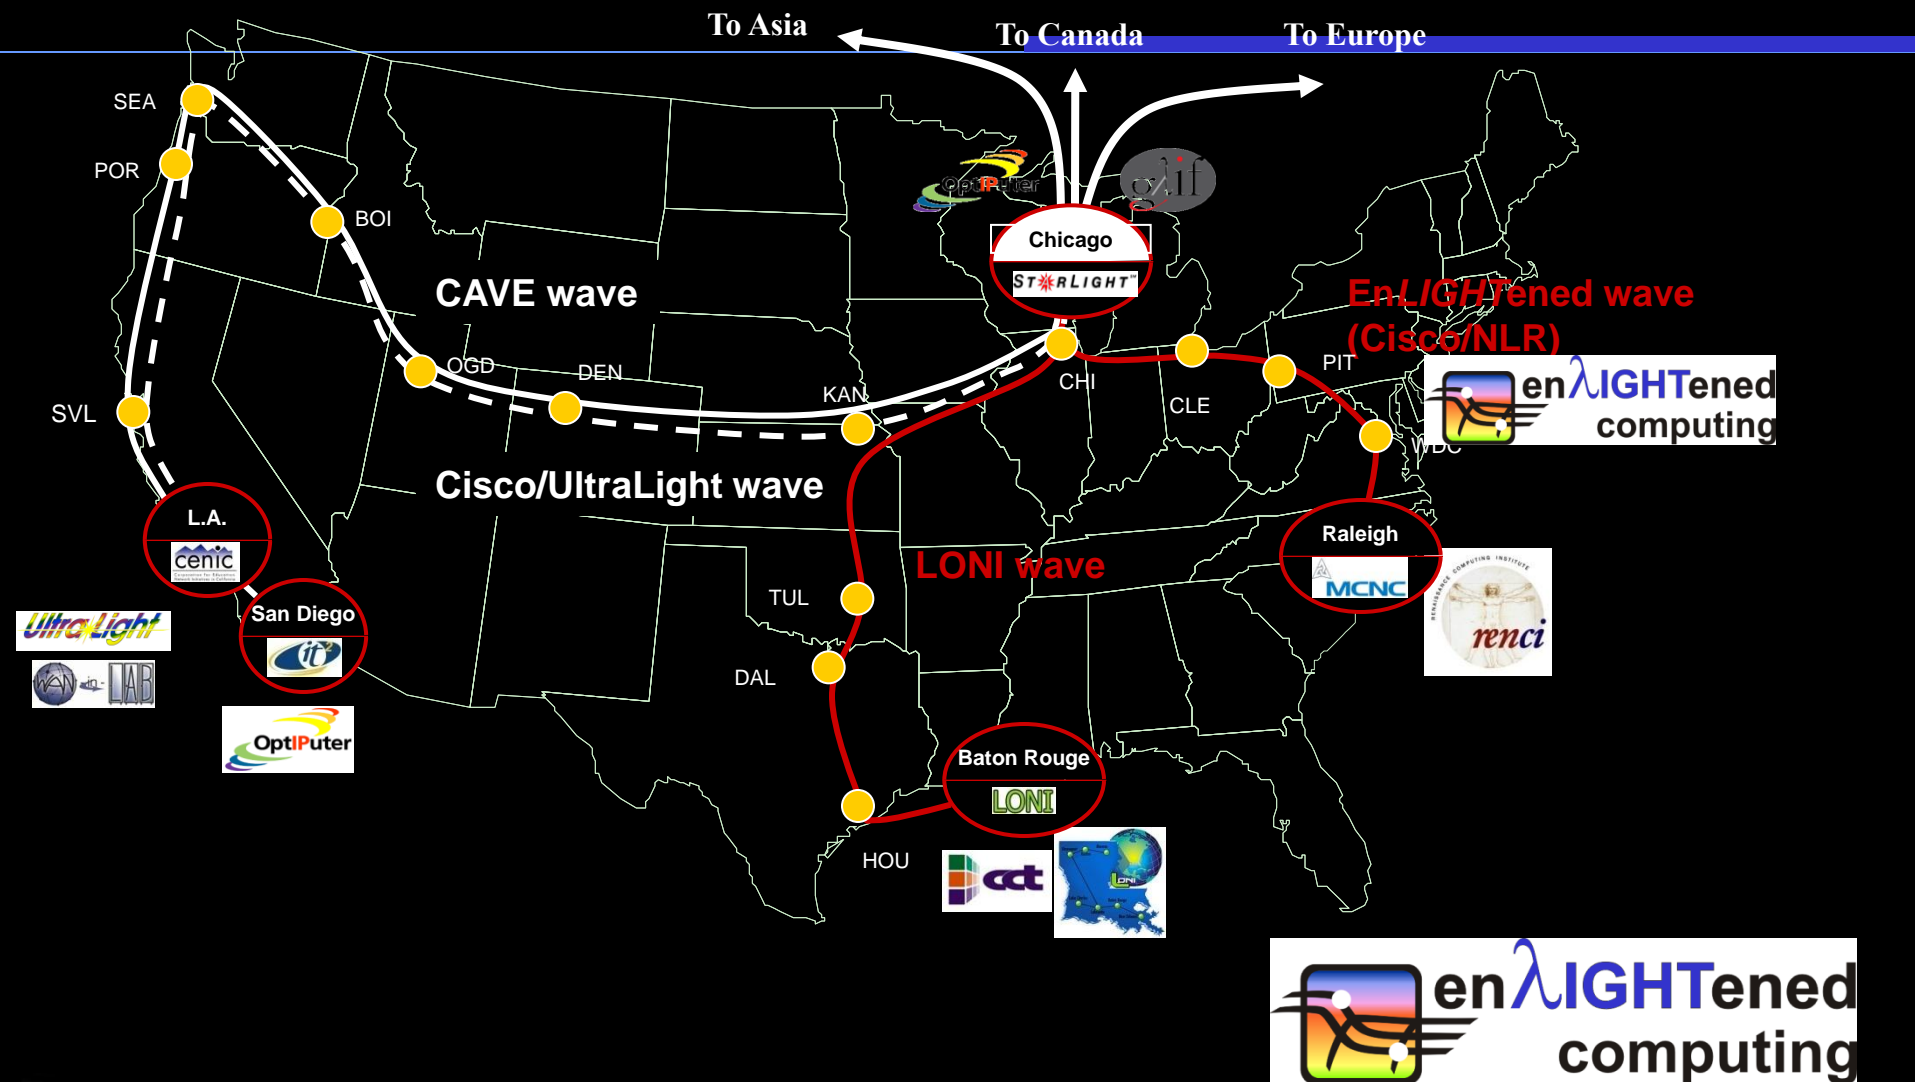




15 November 2009

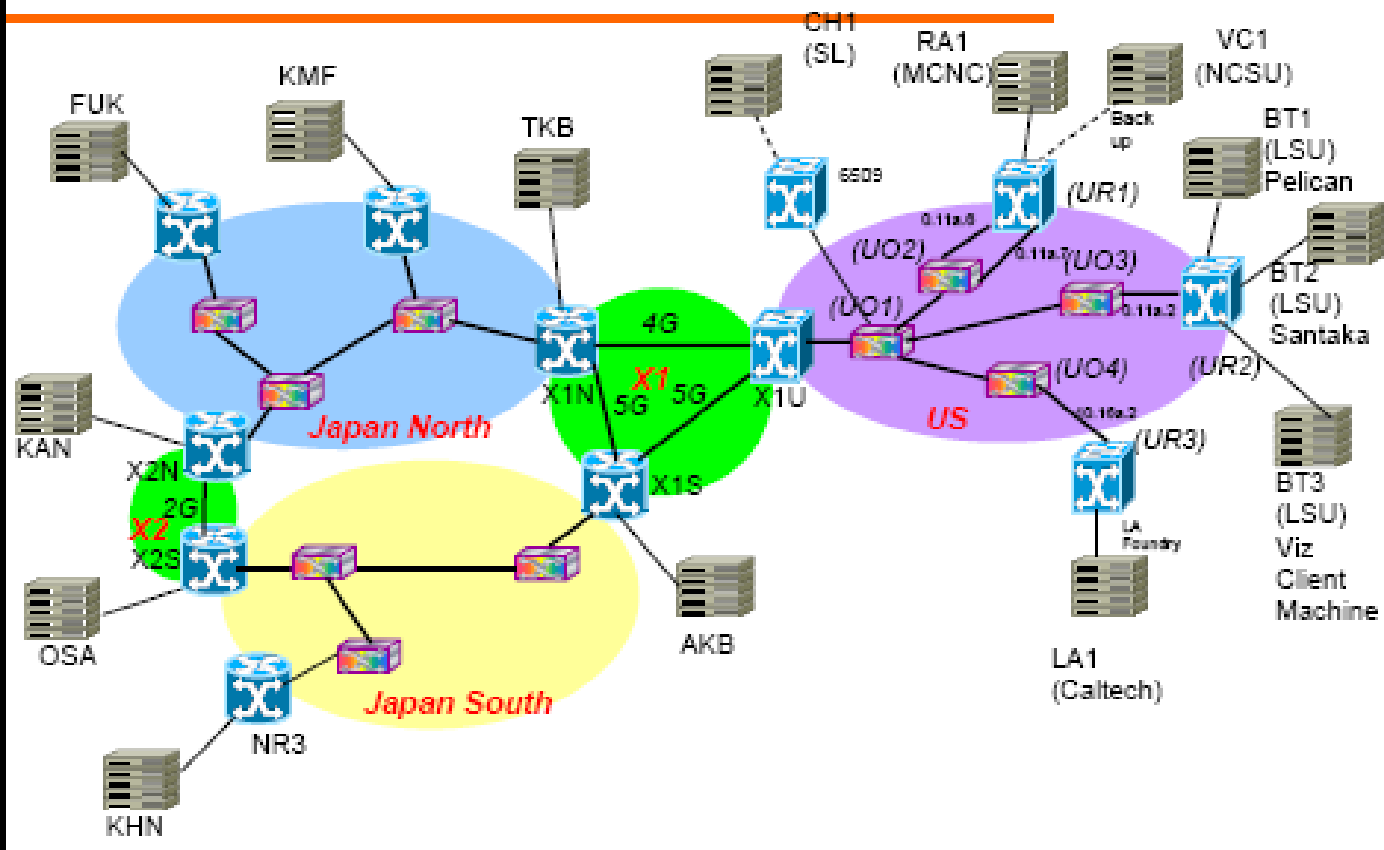


Early Testbed Diagram



EnLIGHTened – G-Lambda Demo at GLIF Conference Tokyo 2006

Resource map of the demo



The OptIPuter LambdaGrid

The diagram illustrates the network topology of The OptIPuter LambdaGrid. It shows connections between various research institutions and network providers. The legend indicates two types of connections: 1 GE Lambda (blue line) and 10 GE Lambda (purple line).

Institutions and Network Providers:

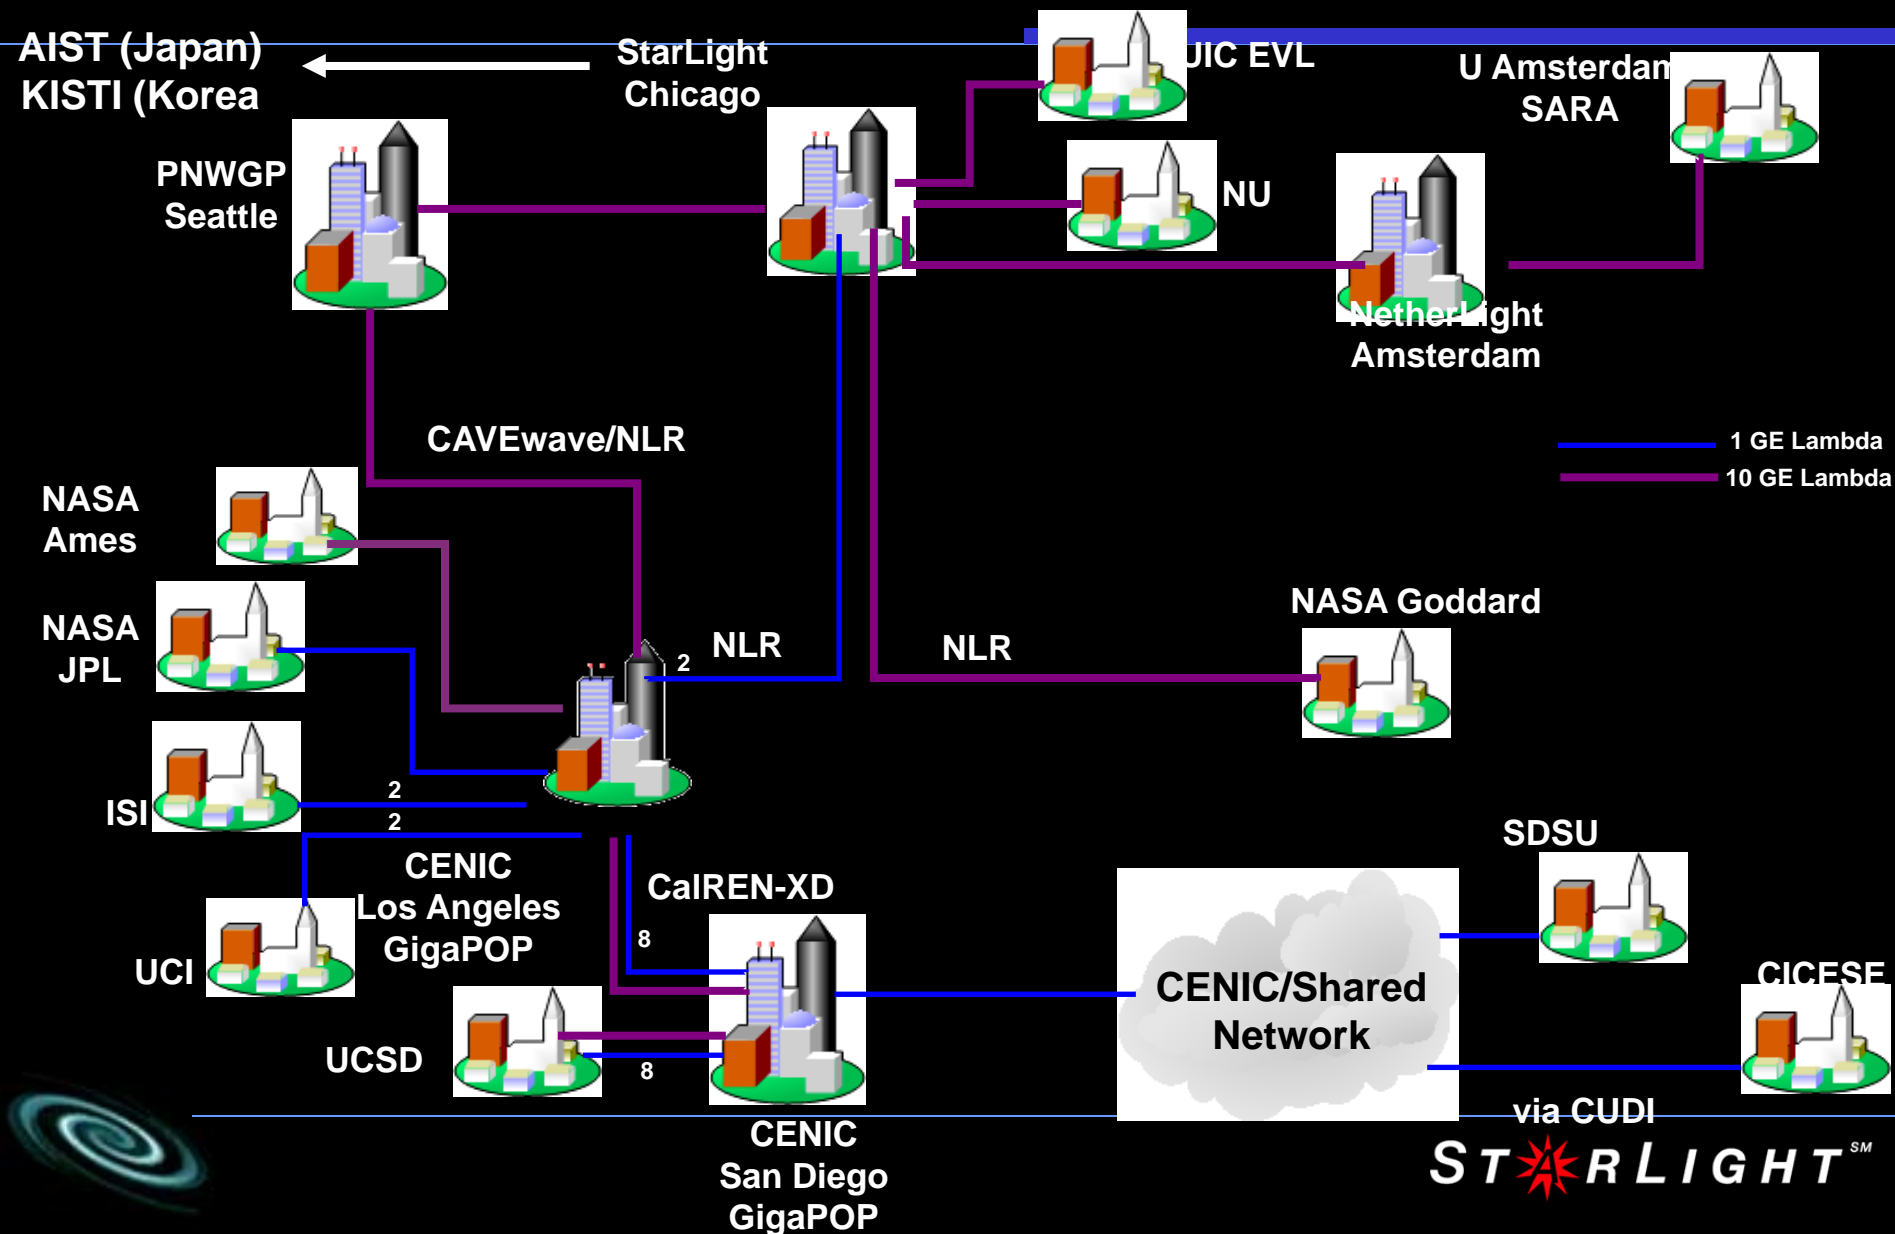
- AIST (Japan)
- KISTI (Korea)
- StarLight Chicago
- JIC EVL
- U Amsterdam SARA
- PNWGP Seattle
- NU
- NetherLight Amsterdam
- CAVEwave/NLR
- NASA Ames
- NASA JPL
- ISI
- CENIC Los Angeles GigaPOP
- UCI
- UCSD
- CaREN-XD
- CENIC San Diego GigaPOP
- CENIC/Shared Network
- NASA Goddard
- SDSU
- GIGASE

Legend:


- Blue line: 1 GE Lambda
- Purple line: 10 GE Lambda

Connections:

- StarLight Chicago connects to AIST (Japan), KISTI (Korea), JIC EVL, U Amsterdam SARA, PNWGP Seattle, NU, NetherLight Amsterdam, CAVewave/NLR, NASA Ames, NASA JPL, ISI, CENIC Los Angeles GigaPOP, UCI, UCSD, CaREN-XD, CENIC San Diego GigaPOP, and CENIC/Shared Network.
- PNWGP Seattle connects to StarLight Chicago and CAVewave/NLR.
- CAVewave/NLR connects to StarLight Chicago, NASA Ames, NASA JPL, ISI, CENIC Los Angeles GigaPOP, and UCSD.
- NASA Ames connects to CAVewave/NLR and NASA JPL.
- NASA JPL connects to CAVewave/NLR and ISI.
- ISI connects to CAVewave/NLR and CENIC Los Angeles GigaPOP.
- CENIC Los Angeles GigaPOP connects to ISI, UCSD, and CaREN-XD.
- UCI connects to CENIC Los Angeles GigaPOP.
- UCSD connects to CENIC Los Angeles GigaPOP and CaREN-XD.
- CaREN-XD connects to CENIC Los Angeles GigaPOP and CENIC San Diego GigaPOP.
- CENIC San Diego GigaPOP connects to CaREN-XD and CENIC/Shared Network.
- CENIC/Shared Network connects to SDSU, GIGASE, and via CUDI to STARLIGHT.
- NASA Goddard connects to StarLight Chicago.
- SDSU connects to CENIC/Shared Network.
- GIGASE connects to CENIC/Shared Network.



FILE



10GE CAVEwave on the National LambdaRail



Global Environment for Network Innovations (GENI)

- **GENI**

- **Supports At-Scale Experimentation on Shared, Heterogeneous, Highly Instrumented Infrastructure**
- **Enables Deep Programmability Throughout the Network,**
- **Promotes innovations in Network Science, Security, Technologies, Services and Applications**
- **Provides Collaborative and Exploratory Environments for Academia, Industry and the Public to Catalyze Groundbreaking Discoveries and Innovation.**



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

- **GENI Is Funded By The National Science Foundation's Directorate for Computer and Information Science and Engineering (CISE)**
- **GENI Is a Virtual Laboratory For Exploring Future Internets At Scale.**
- **GENI Is Similar To Instruments Used By Other Science Disciplines, e.g., Astronomers – Telescopes, HEP - Synchrotrons**
- **GENI Creates Major Opportunities To Understand, Innovate and Transform Global Networks and Their Interactions with Society.**
- **GENI Is Dynamic and Adaptive.**
- **GENI Opens Up New Areas of Research at the Frontiers of Network Science and Engineering, and Increases the Opportunity for Significant Socio-Economic Impact.**



iGENI: The International GENI

- **The iGENI Initiative Will Design, Develop, Implement, and Operate a Major New National and International Distributed Infrastructure.**
- **iGENI Will Place the “G” in GENI Making GENI Truly Global.**
- **iGENI Will Be a Unique Distributed Infrastructure Supporting Research and Development for Next-Generation Network Communication Services and Technologies.**
- **This Infrastructure Will Be Integrated With Current and Planned GENI Resources, and Operated for Use by GENI Researchers Conducting Experiments that Involve Multiple Aggregates At Multiple Sites.**
- **iGENI Infrastructure Will Connect Its Resources With Current GENI National Backbone Transport Resources, With Current and Planned GENI Regional Transport Resources, and With International Research Networks and Projects,**



Initial iGENI Consortium

- **Consortium Partners Include Several Major Network Research Organizations:**
 - International Center for Advanced Internet Research (iCAIR) at Northwestern University,
 - Electronic Visualization Laboratory (EVL) at the University of Illinois at Chicago
 - The California Institute for Telecommunications and Information Technology (Calit2) at the University of California, San Diego
 - Cisco Systems, Inc. Research
 - BBN Technologies GENI Program Office (GPO).
 - The StarLight Consortium
 - RENCI and North Carolina University Partners, e.g. Duke, North Carolina State
 - Other Cluster D Participants (Univ of Mass - Amherst, Columbia, Ohio State, Wayne State, Univ of Houston, Rice, Texas A&M, UT Austin, Oklahoma State)
 - iGENI Initiatives Also Extend To Activities In Other GENI Clusters
 - *iGENI Research Initiatives Have Multiple International Partners*

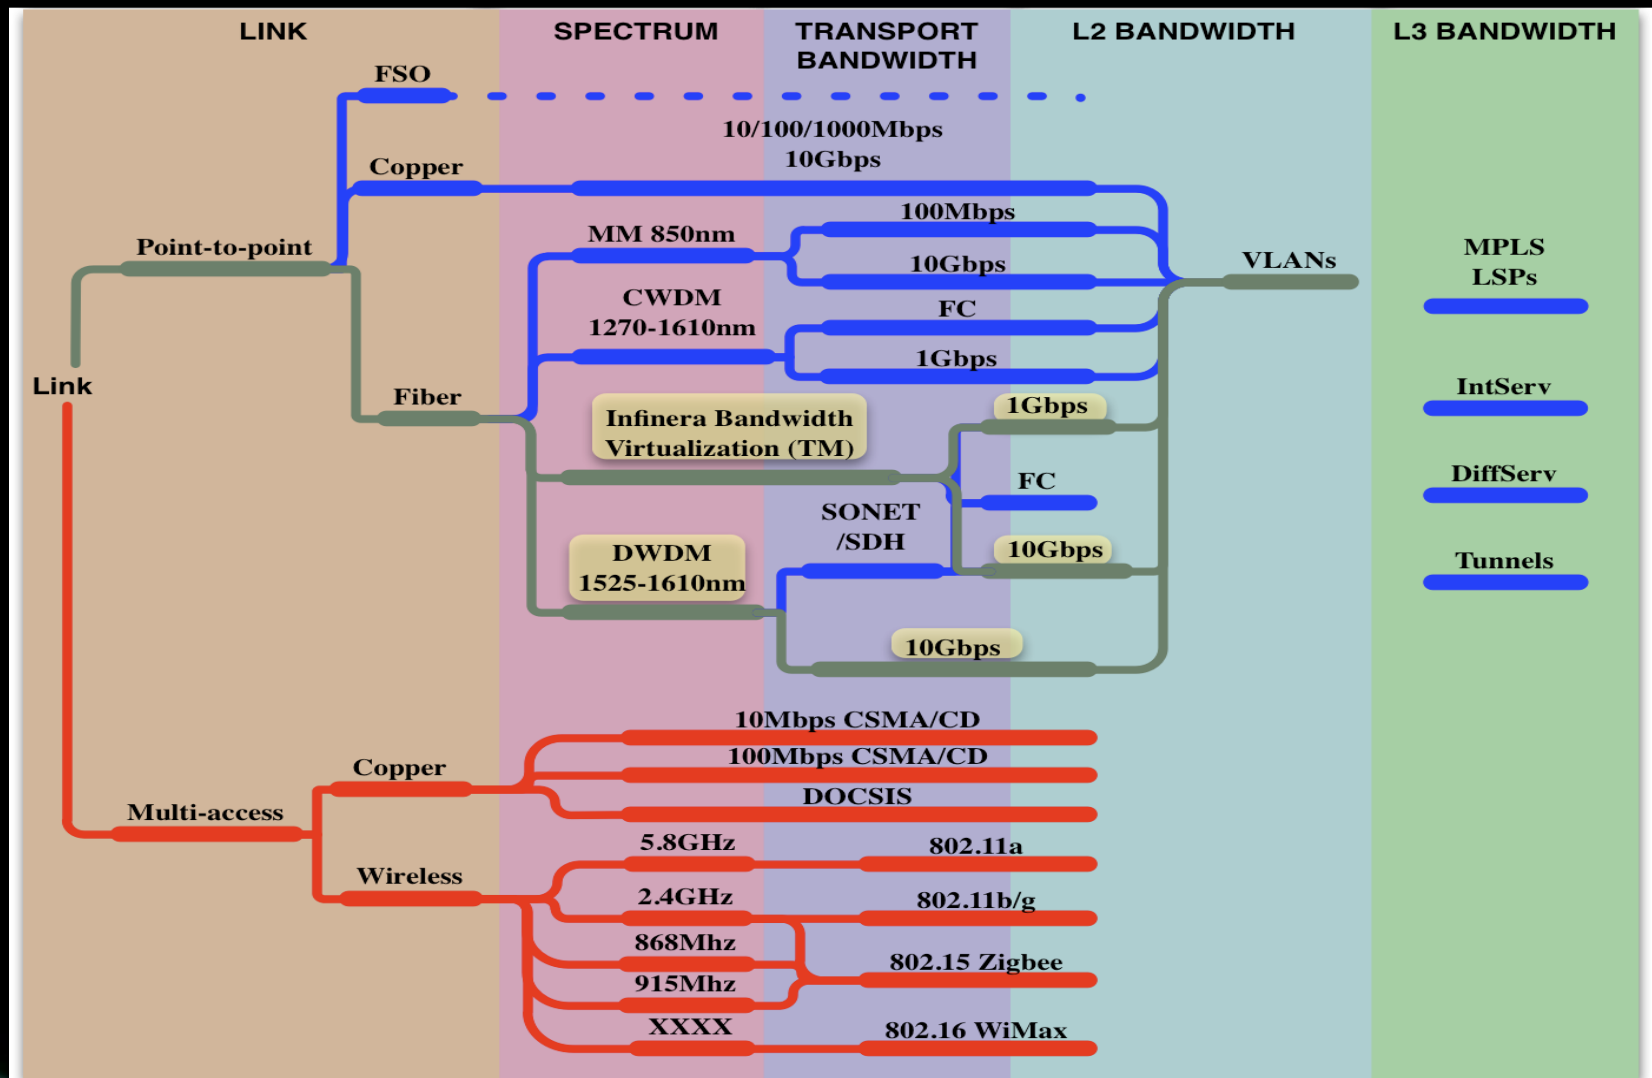


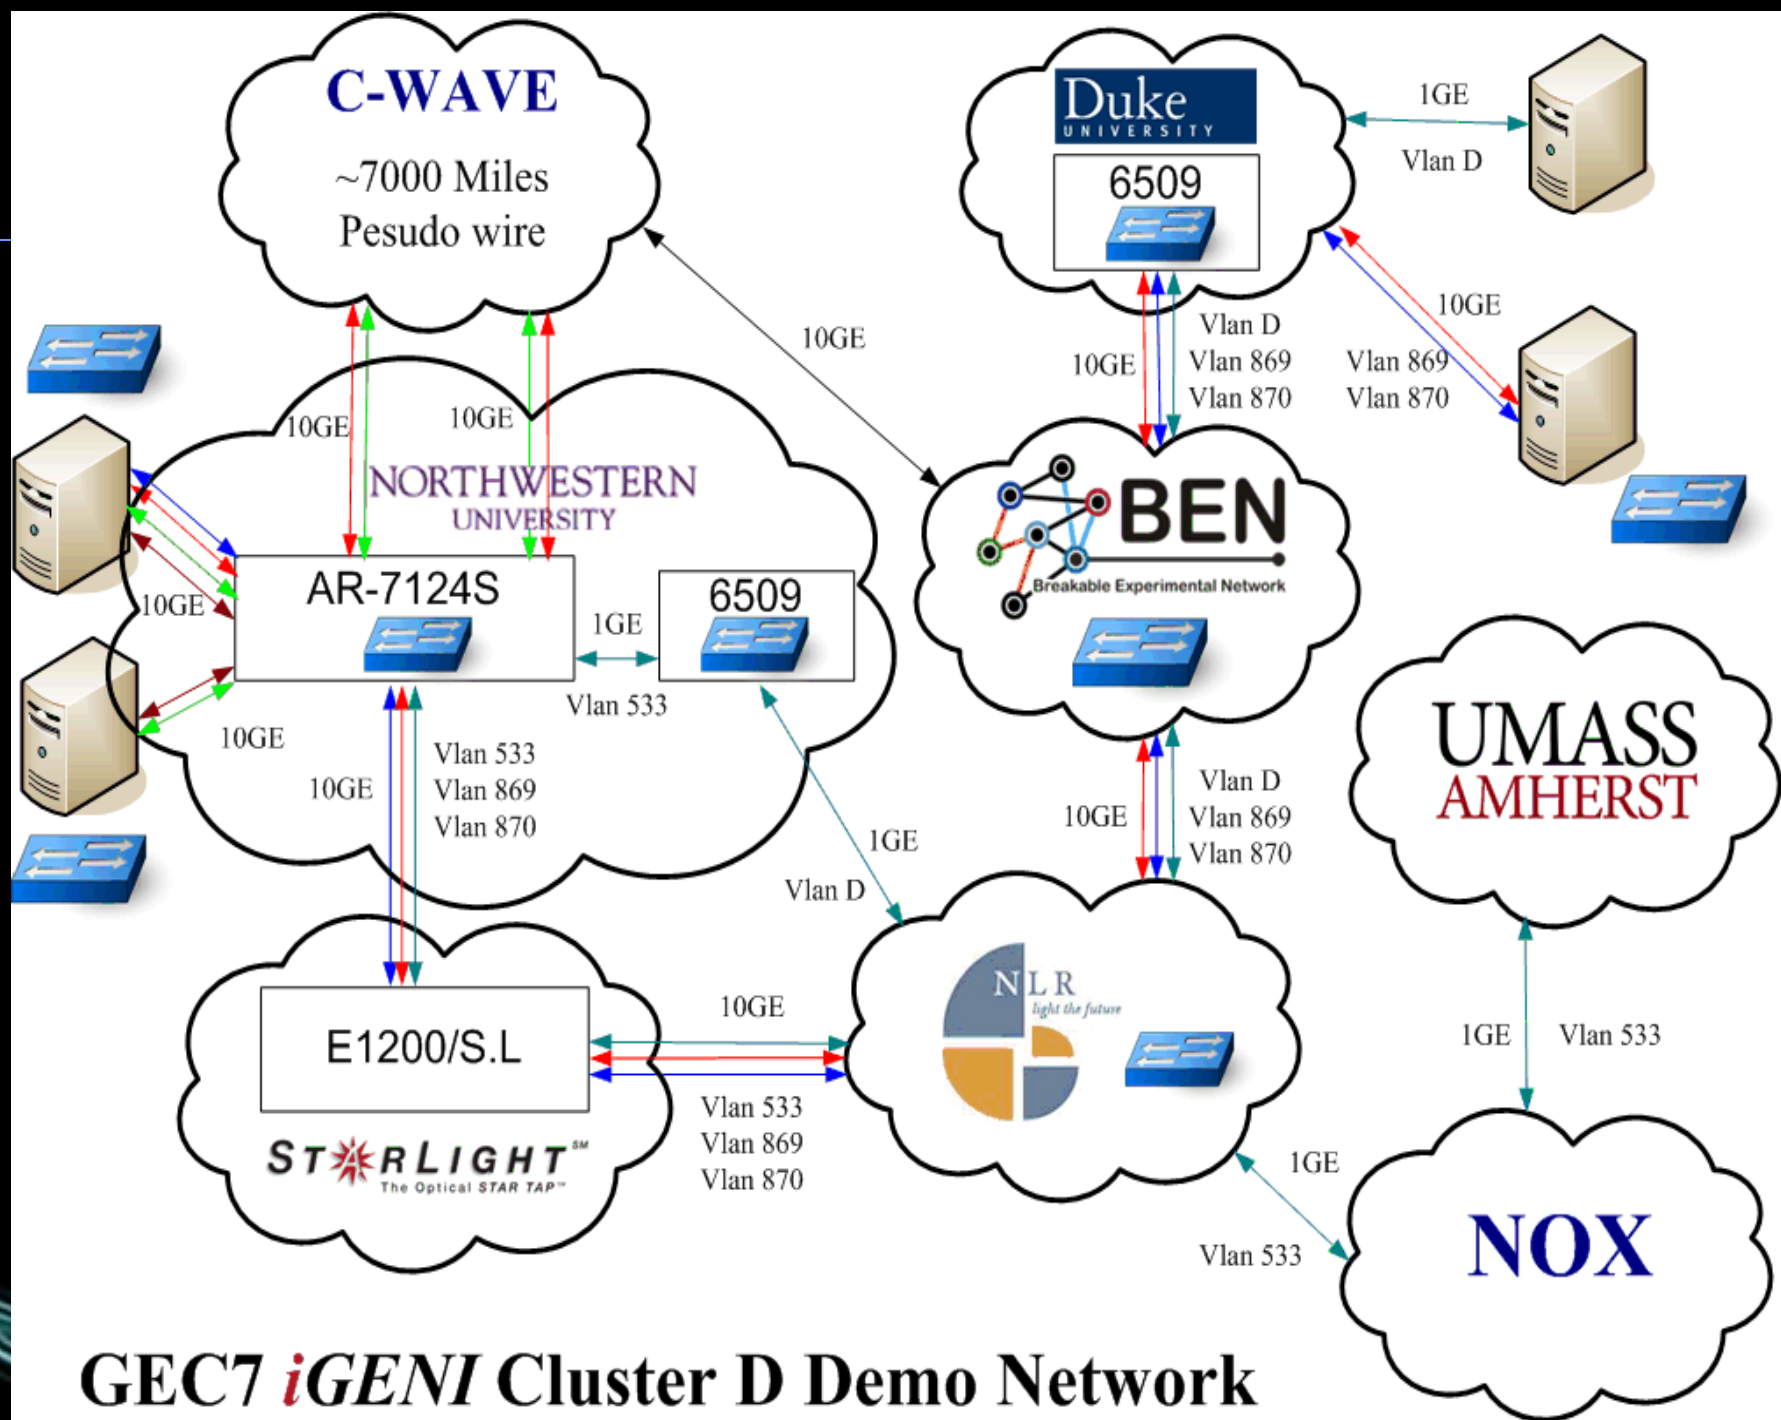
Cluster D Initiatives

- **ORCA/BEN -- Primary Distinction Among Other Control Frames = A Unique Comprehensive Strategy Cross-Layer Provisioning and Experimentation**
- **DOVE -- Diverse Outdoor Mobile Environment**
- **ViSE -- Sensor Virtualization and Slivering in an Outdoor Wide-Area Wireless -- GENI Sensor/Actuator Network Testbed**
- **ERM -- Embedded Real-Time Measurements**
- **KANSAI -- KanseiSensorNet**
- **DICLOUD -- Data Intensive Cloud Control**
- **OKGEMS -- Cyber-Physical System**
- **IMF -- Integrated Measurement Network**
- **LEARN -- Measurement Handler and Network Integration (Programmable Measurements)**
- **BBN ORCA Xen Cluster**
- **iGENI**

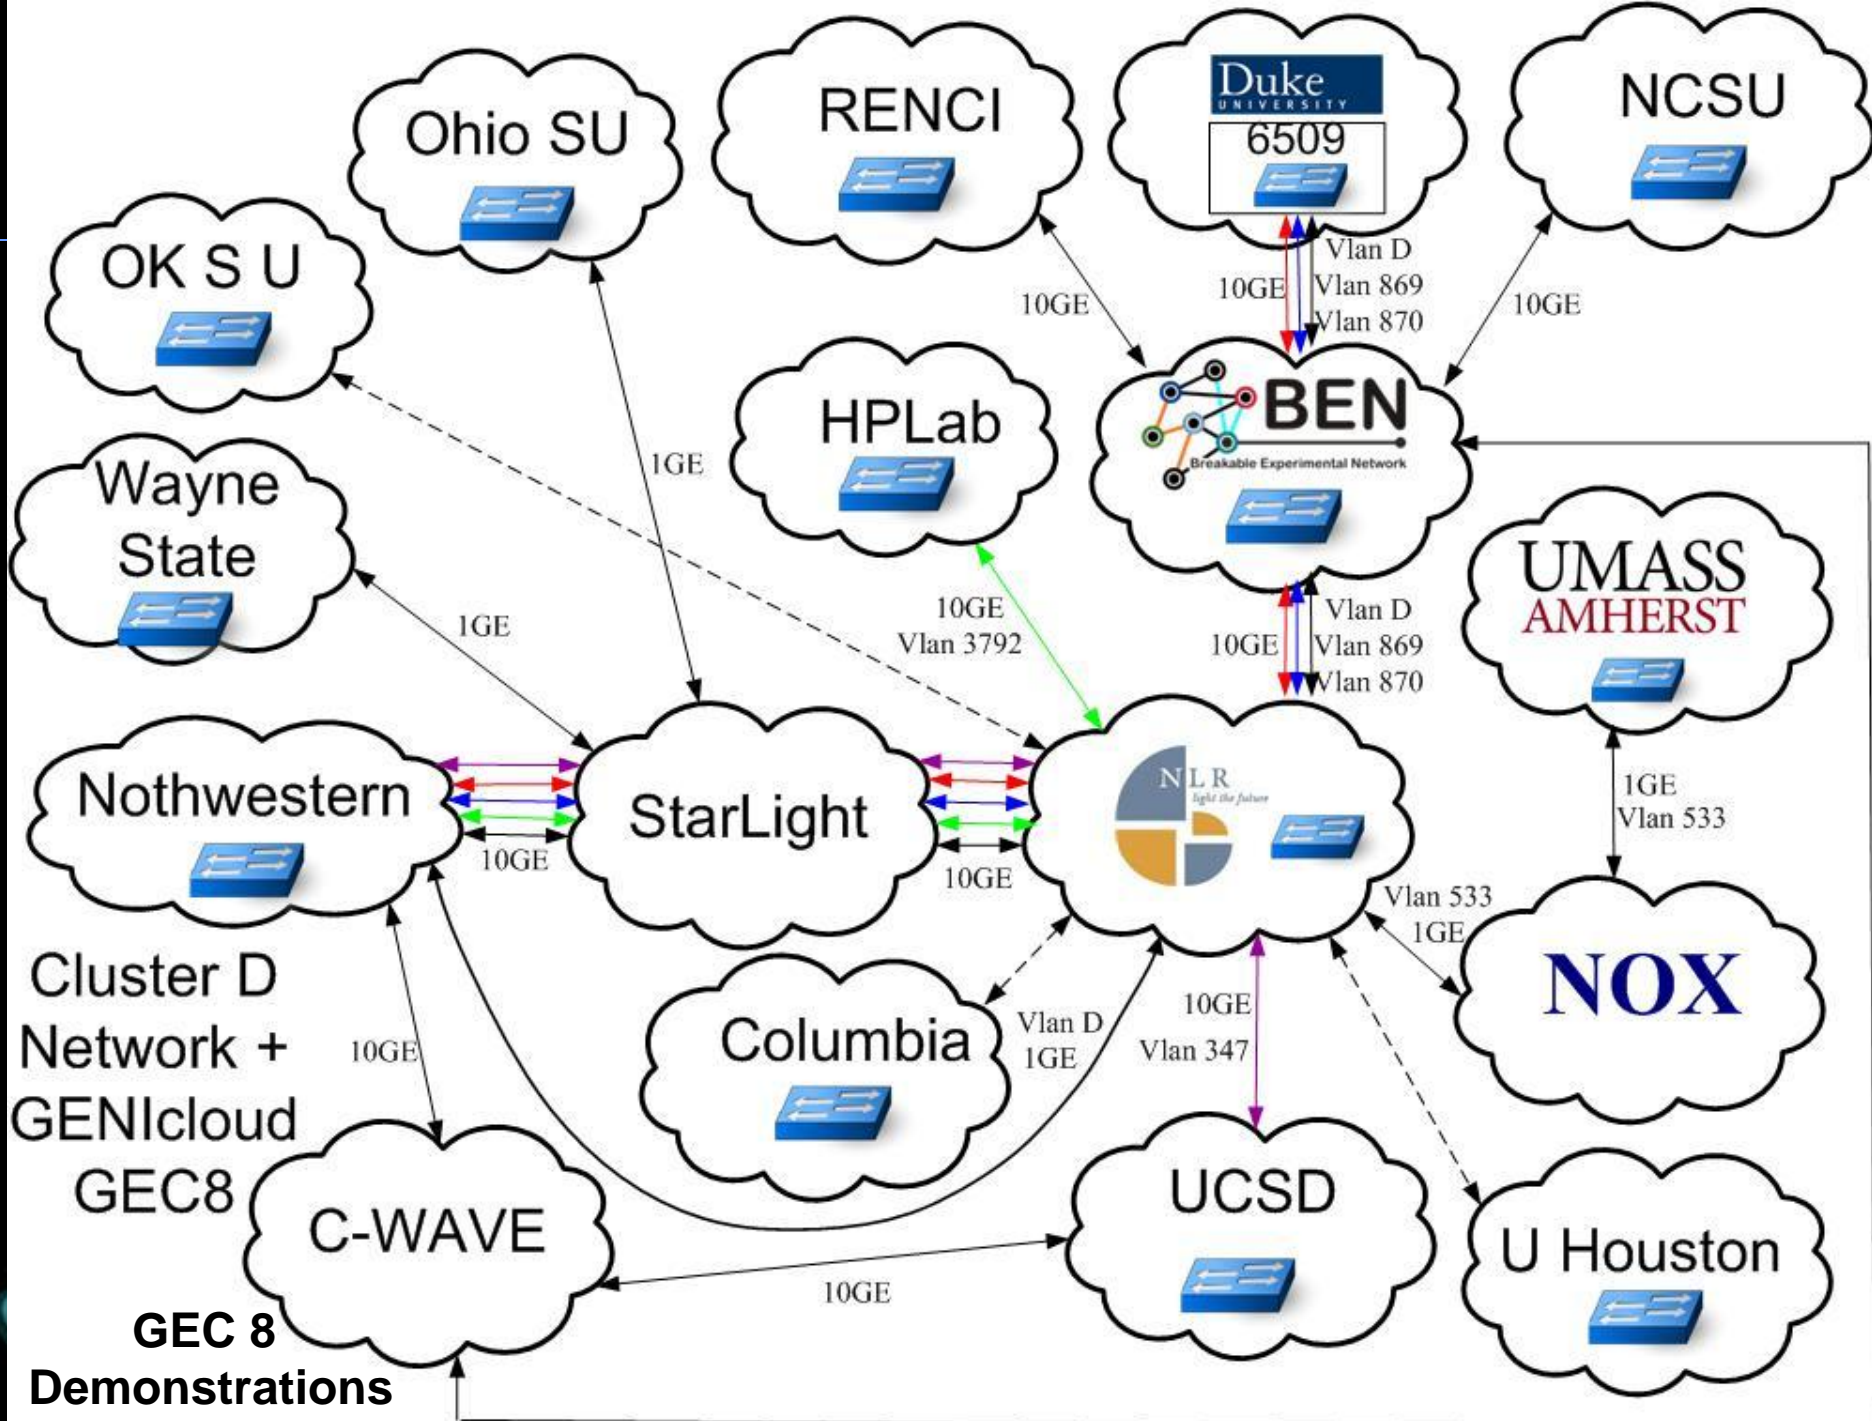


ORCA "Link" Slivering

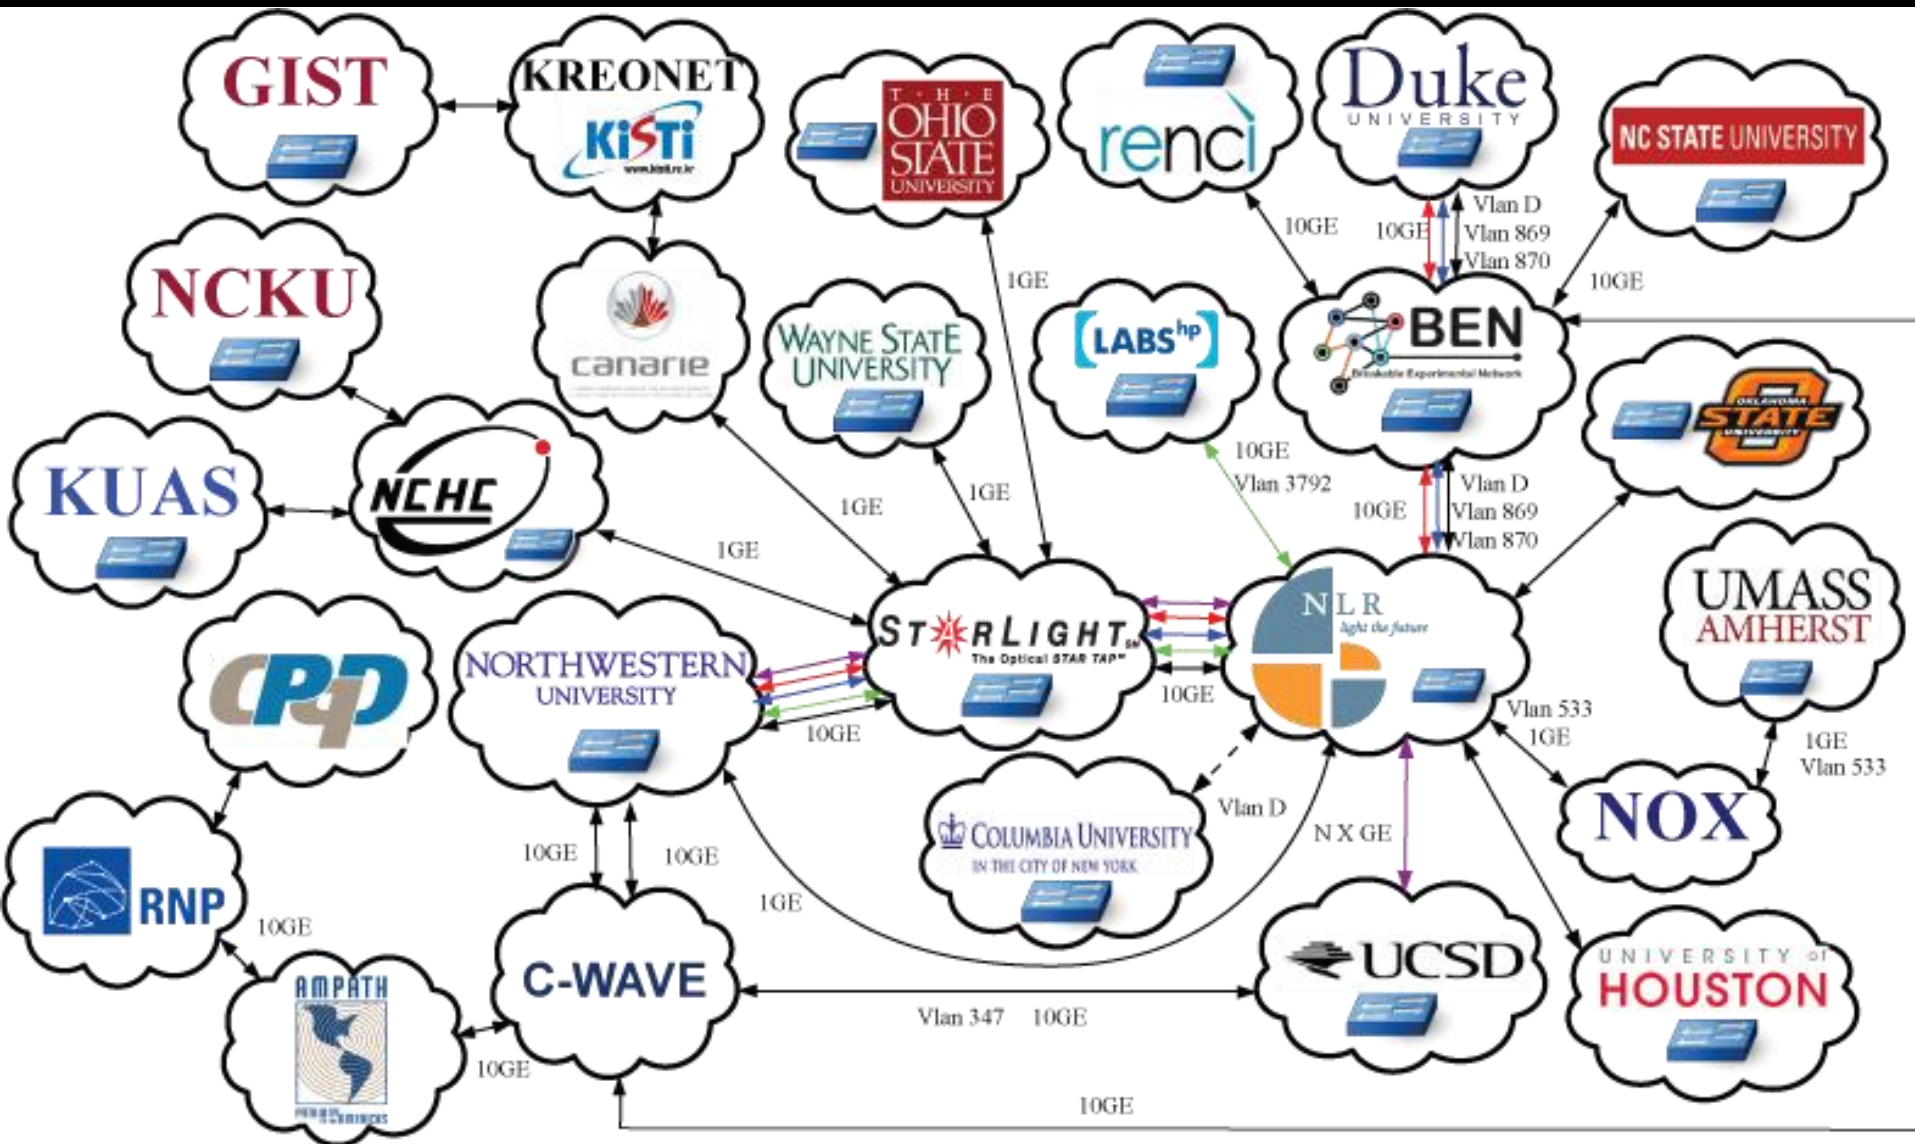




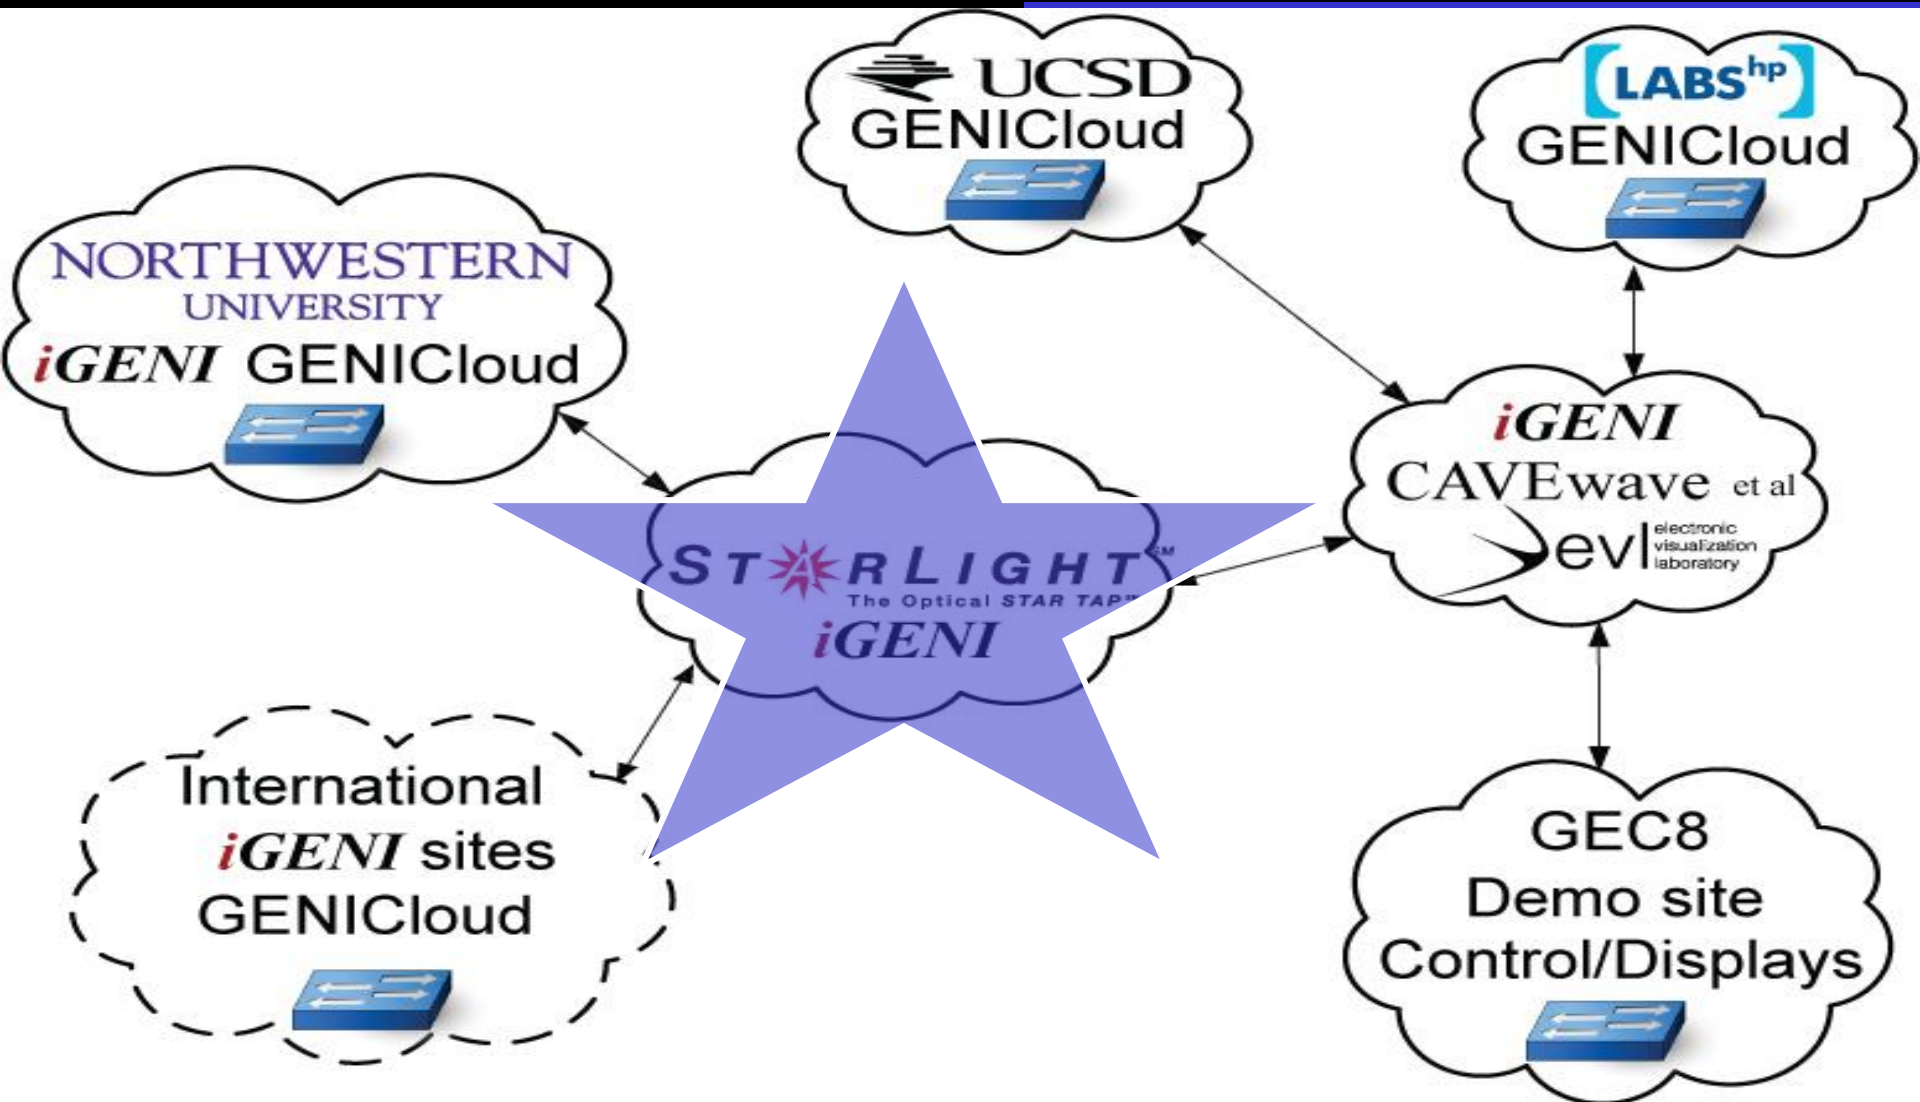
GEC7 iGENI Cluster D Demo Network

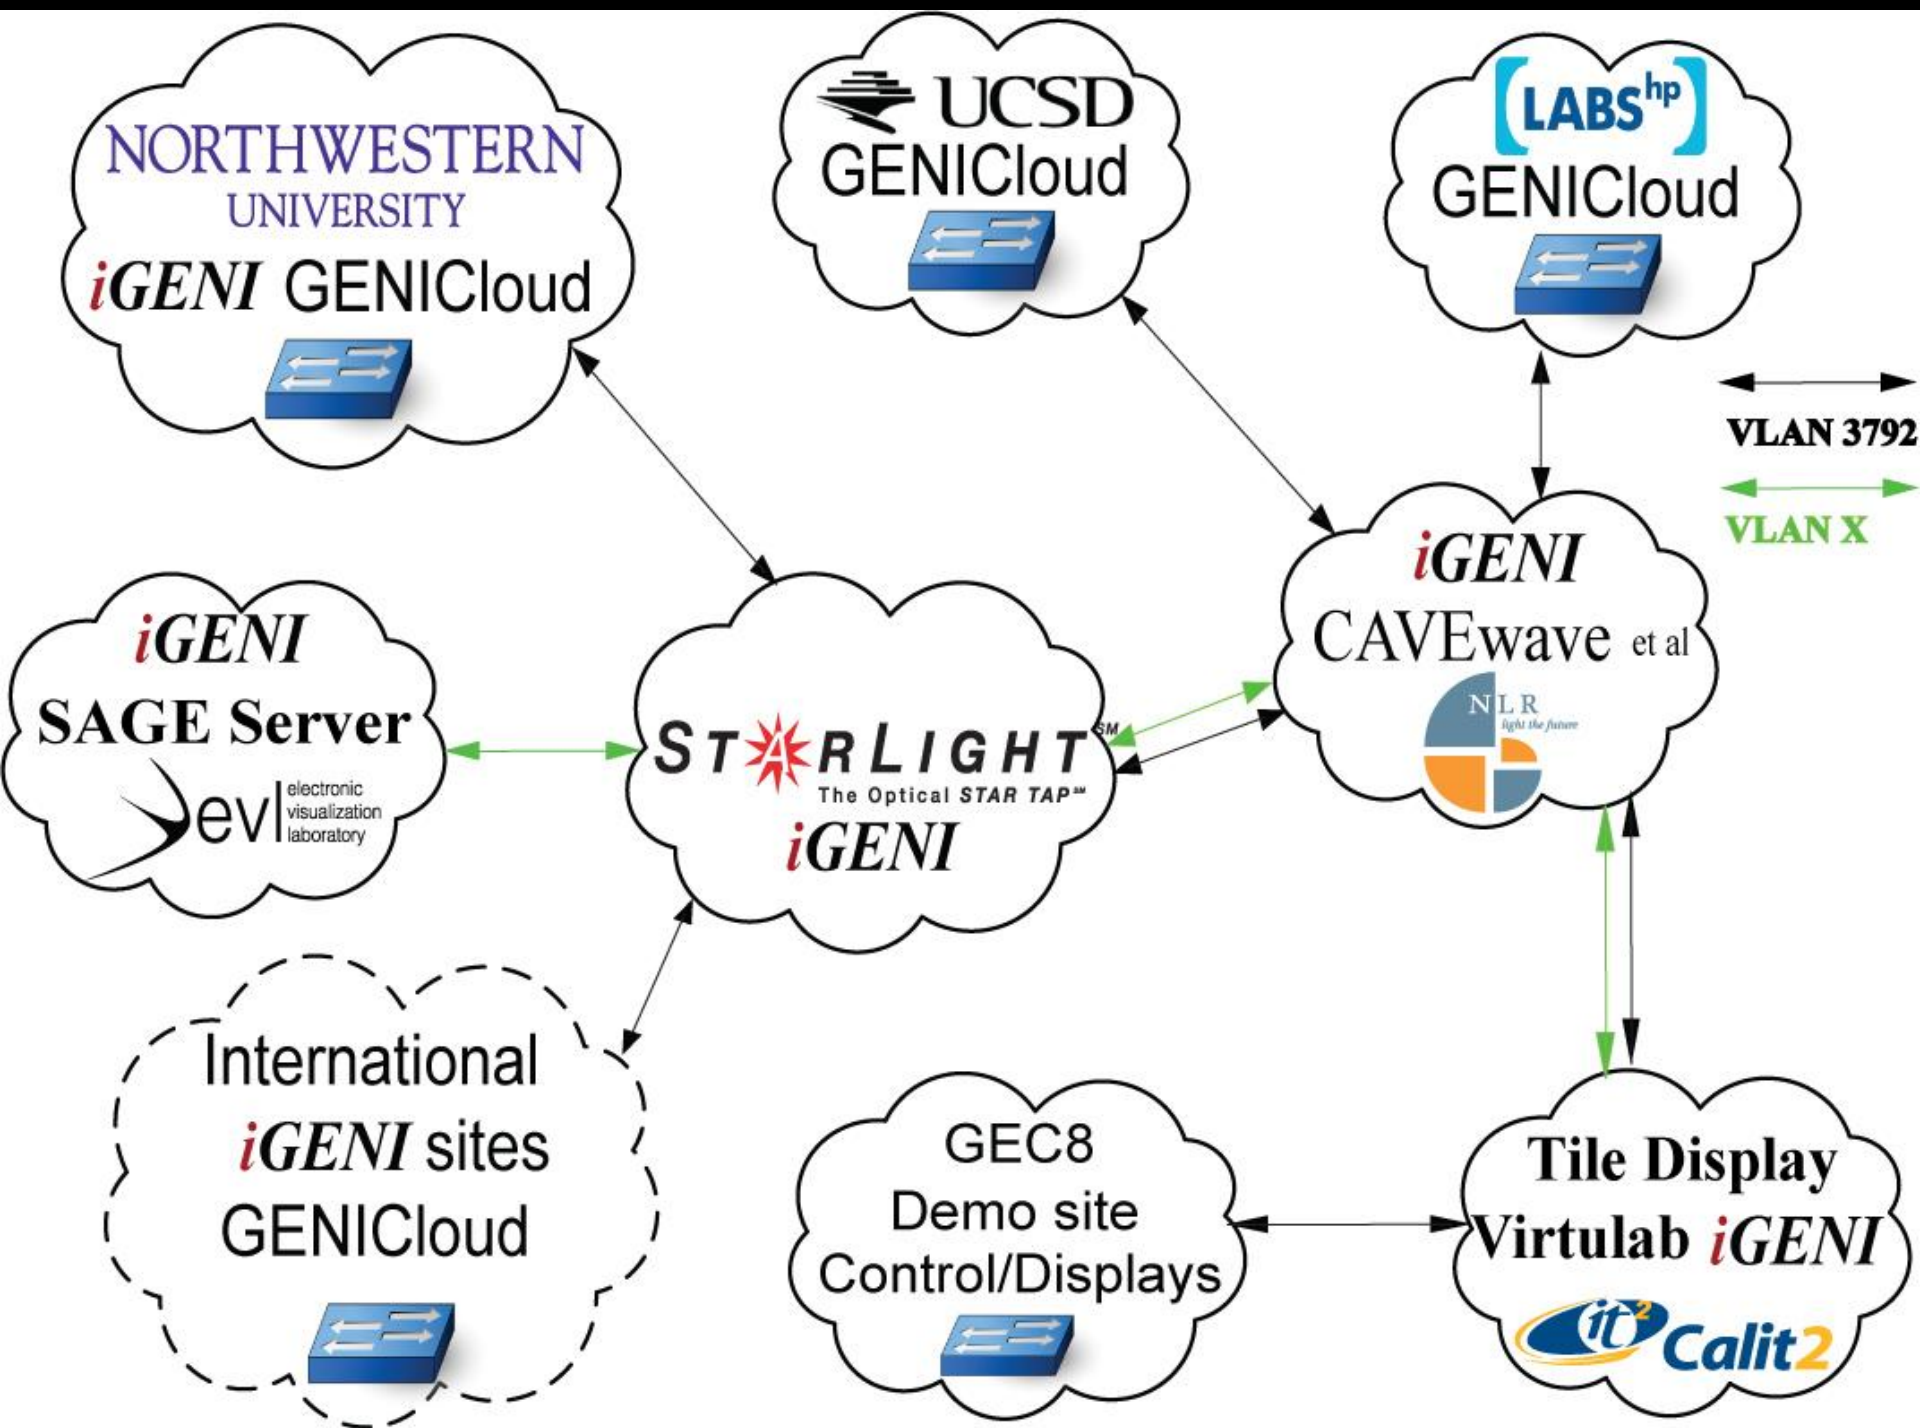


GENI Engineering Conference 9



iGENI Demonstrations at GEC 8







GEC 10 Demonstrations

TransCloud: A Distributed Environment Based On Dynamic Networking

Rick McGeer, HP Labs

Joe Mambretti, Northwestern

Paul Müller, TU Kaiserslautern

Chris Matthews, Chris Pearson, Yvonne Coady, Victoria

Jim Chen, Feh Yeh, Northwestern

Andy Bavier, PlanetWorks

Marco Yuen, Princeton

Jessica Blaine, Alvin Au Young, HP Labs

Alex Snoeren, UC San Diego

March 16, 2010

<http://www.icaair.org>

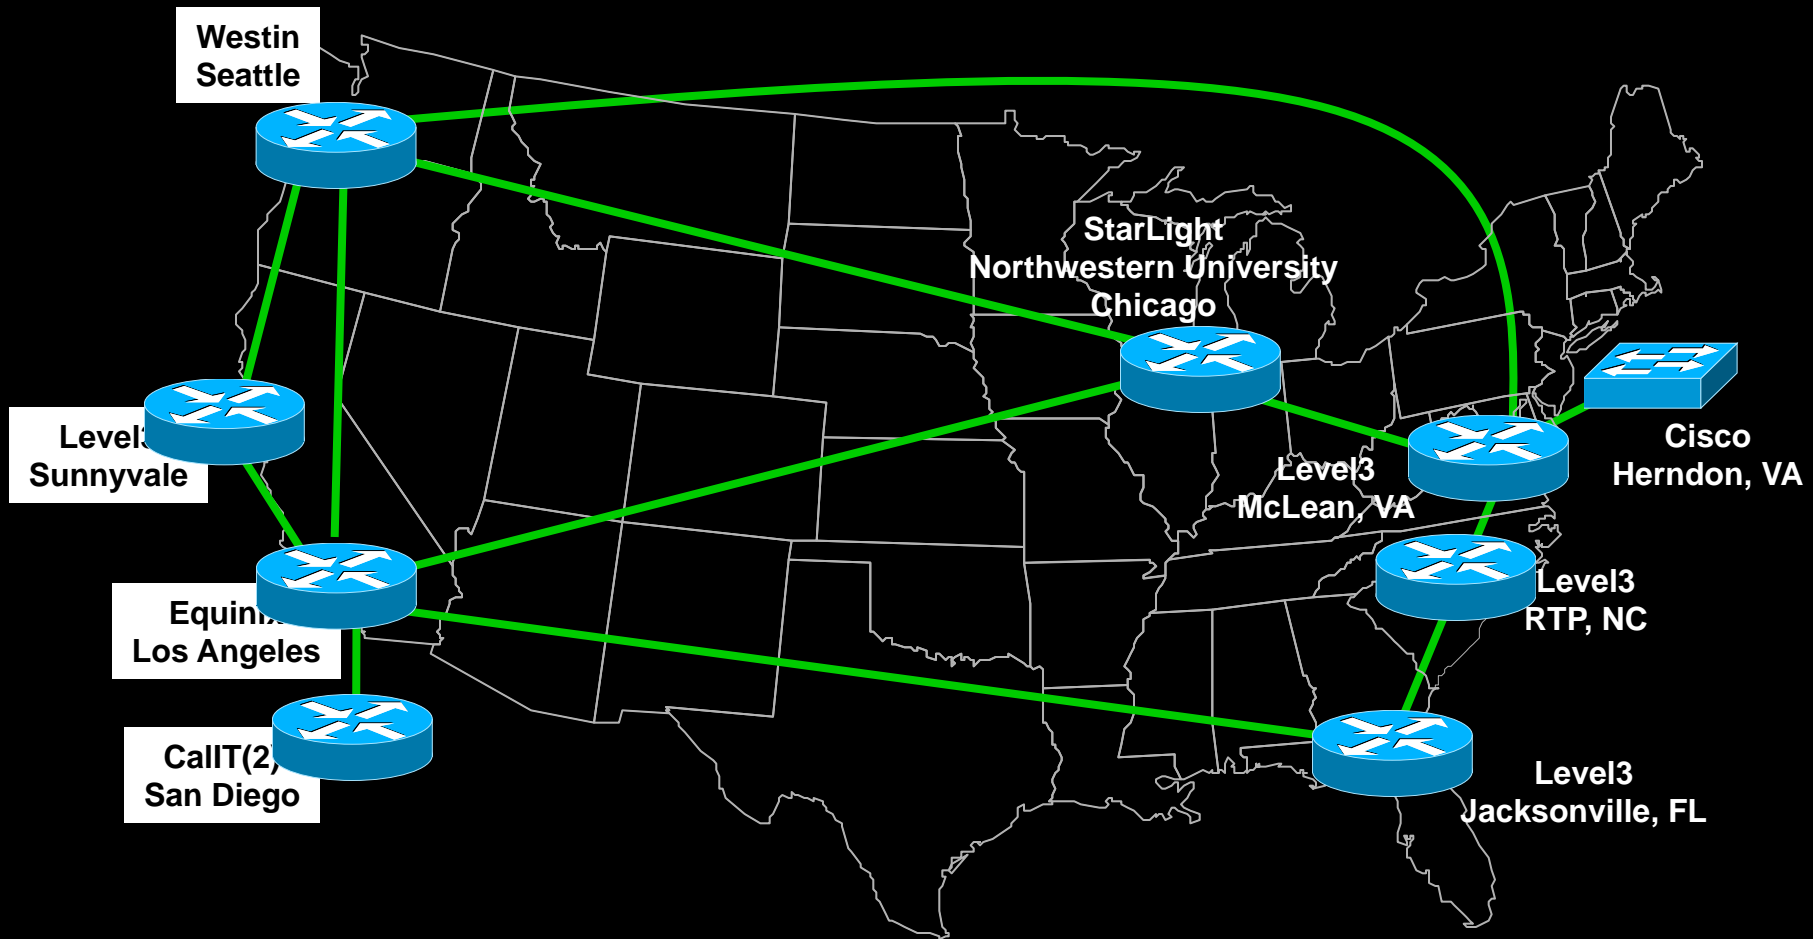
<http://www.geni.net> **STARLIGHTSM**

TransCloud

- **TransCloud = A Cloud Where Services Migrate, Anytime, Anywhere In a World Where Distance Is Eliminated**
 - **Joint Project Between GENICloud, iGENI, et al**
 - **GENICloud Provides Seamless Interoperation of Cloud Resources Across N-Sites, N-Administrative Domains**
 - **iGENI Optimizes Private Networks of Intelligent Devices Capable of Dynamically Provisioned Low-Latency, High-Performance Communications Among Multiple Physically-Distributed Infrastructures and Facilities**



C-Wave Summer 2010



Issues Related To Federated Processes Across Multiple Sites

- Topics Revolve Around Policy Issues Related To Resource Utilization
- Macro APIs
- Fenius Example
- An Experimental Architecture
- Based on a Super Agent That Can Be Used As an API That Can Interact With Individual GLIF GOLE Control Planes
- Edge Processes Can Interact With One API Yet Utilize Multiple Different Control Frameworks



Building On Existing Partnerships, Current & Future International Partners Include Researchers From Many Countries

- **Australia**
- **Brazil**
- **Canada**
- **China**
- **Egypt**
- **Germany**
- **India**
- **Japan**
- **Korea**
- **Taiwan**
- **Spain**
- **Singapore**
- **Netherlands**
- **Spain**
- **New Zealand**
- **Sweden**
- **UK**
- **Et Al**



StarLight – “By Researchers For Researchers”

StarLight is an experimental optical infrastructure and **proving ground for network services** optimized for high-performance applications

GE+2.5+10GE

Exchange

Soon:

Multiple 10GEs

Over Optics –

World’s “Largest”

10GE Exchange

First of a Kind

Enabling Interoperability

At L1, L2, L3



View from StarLight



Abbott Hall, Northwestern University's Chicago downtown campus

StarLight Infrastructure

StarLight is *a large research-friendly co-location facility* with space, power and fiber that is being made available to university and national/international network collaborators as a *point of presence* in Chicago



MREN

Metropolitan Research & Education Network

- An Advanced Network for Advanced Applications
- Designed in 1993; Initial Production in 1994, Managed at L2 & L3
- Created by Consortium of Research Organizations -- over 20
- Partner to STAR TAP/StarLight, I-WIRE, NGI and R&E Net Initiatives, Grid and Globus Initiatives etc.
- Model for Next Generation Internets
- Developed World's First GigaPOP
- Next – the “Optical MREN”
- Soon - Optical ‘TeraPOP’ Services

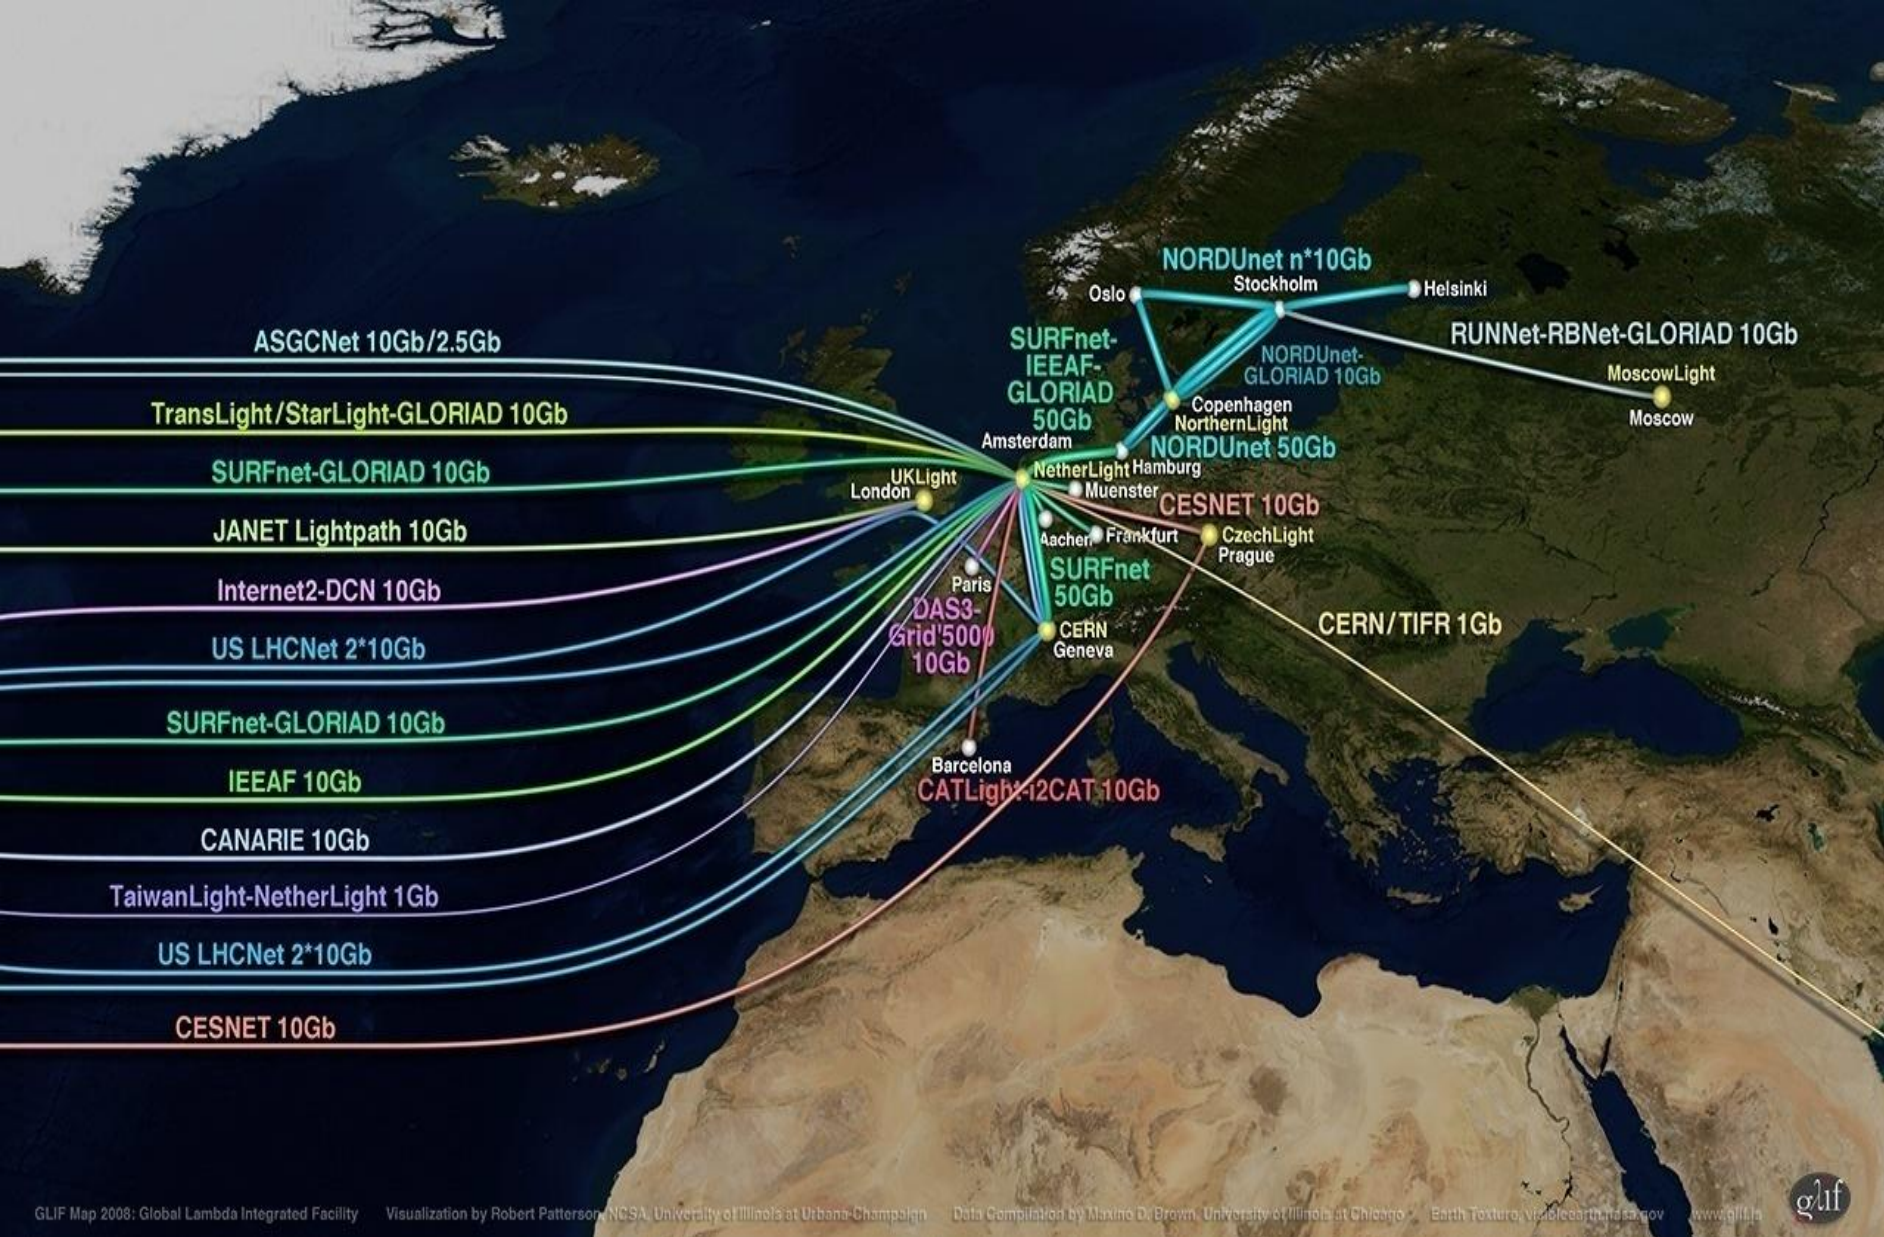
iCAIR: Founding Partner of the Global Lambda Integrated Facility

Available Advanced Network Resources



GLIF is a consortium of institutions, organizations, consortia and country National Research & Education Networks who voluntarily share optical networking resources and expertise to develop the *Global LambdaGrid* for the advancement of scientific collaboration and discovery.







IRNC ProNet: TransLight/StarLight

Announced: July 13, 2010

Tom DeFanti, Maxine Brown, Joe Mambretti, Tajana Rosing

Calit2, University of California, San Diego

Electronic Visualization Lab, University of Illinois at Chicago

**International Center for Advanced Internet Research, Northwestern
University**

**20 years of NSF-Funded High-Performance
International Networking for
Advanced Applications
(1995-2014)**

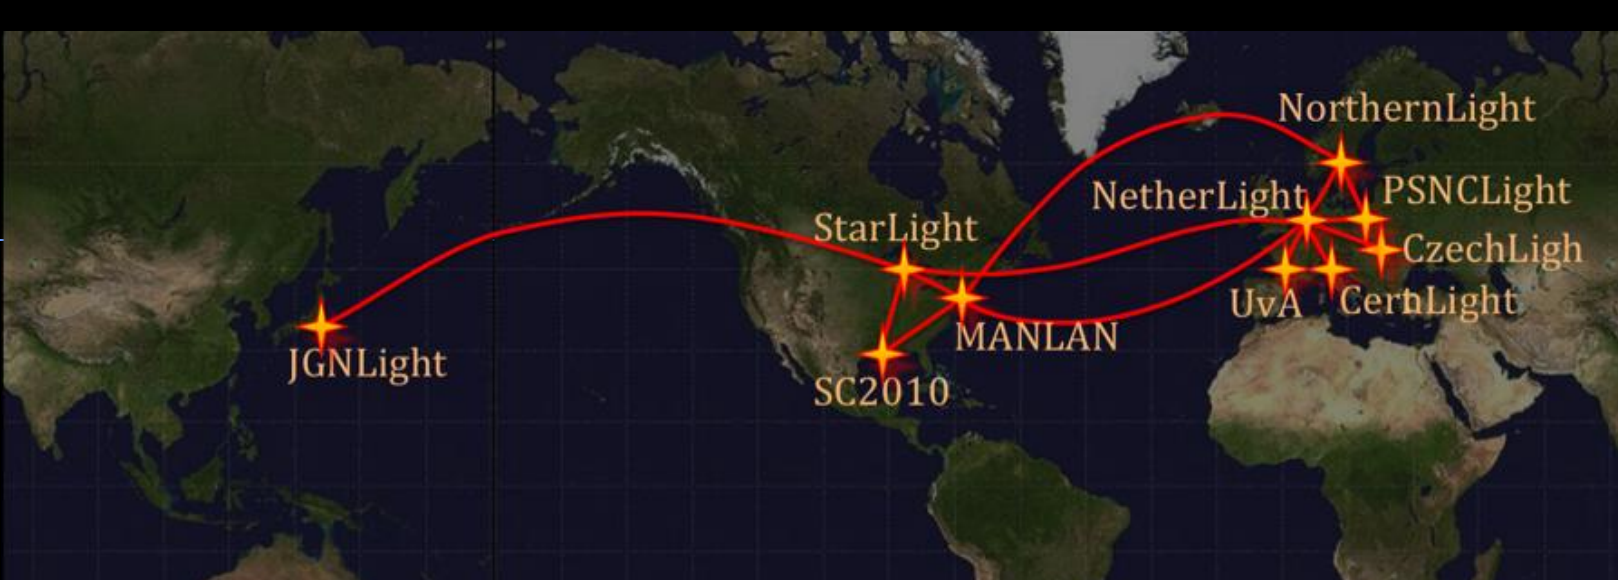


STARLIGHTSM

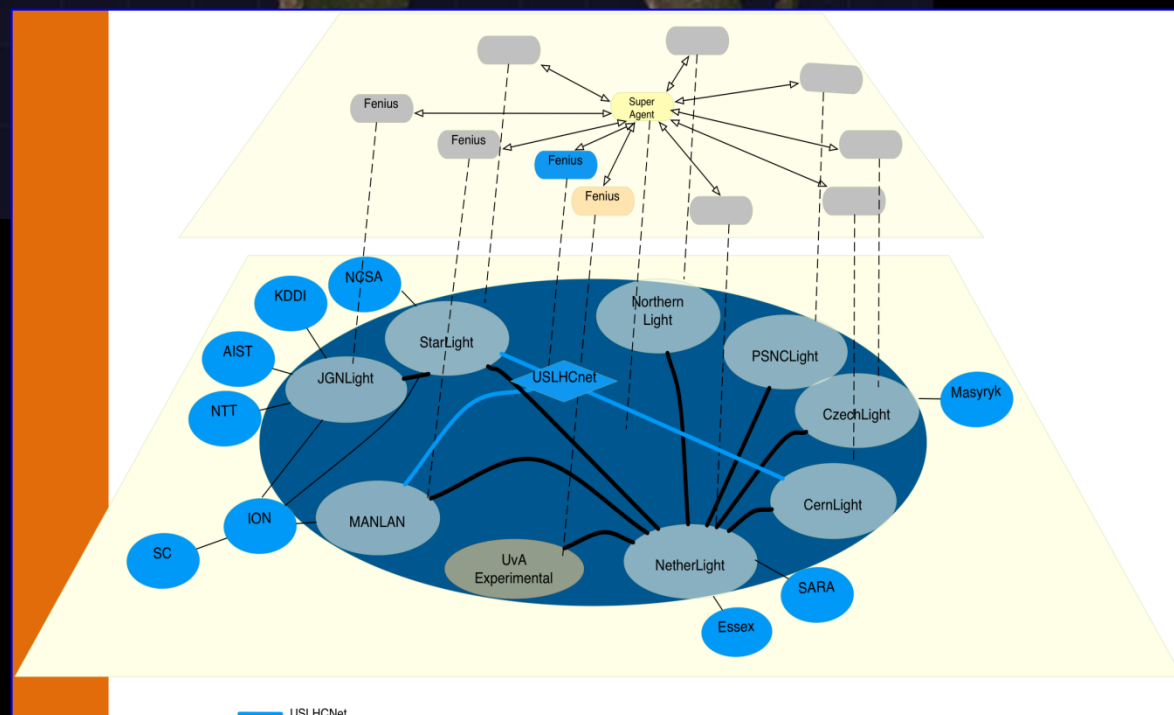
IRNC TL/SL Deliverables

- **Continue Enabling Multi-National Application and Middleware Experiments Through Innovative Services and Technologies On International Networks:**
 - **High-Performance Digital Media Network (HPDMnet)**
 - **iGENI: the GENI-funded international GENI project* ##**
 - **SAGE: connecting people and their data at high-res***
 - **CineGrid: it's all about visual communications**
 - **GreenLight International: less watts/terabyte***
 - **Science Cloud Communication Services Network (SCCSnet)*: the impending disruption**
 - **Build Cooperative National and International Partnerships***
 - **Provide New Services, Including Many with Industrial Partners**
 - **Capitalize On Other Emerging Opportunities***
- ## Now, In Part, A CISE/OCI Partnership!!**

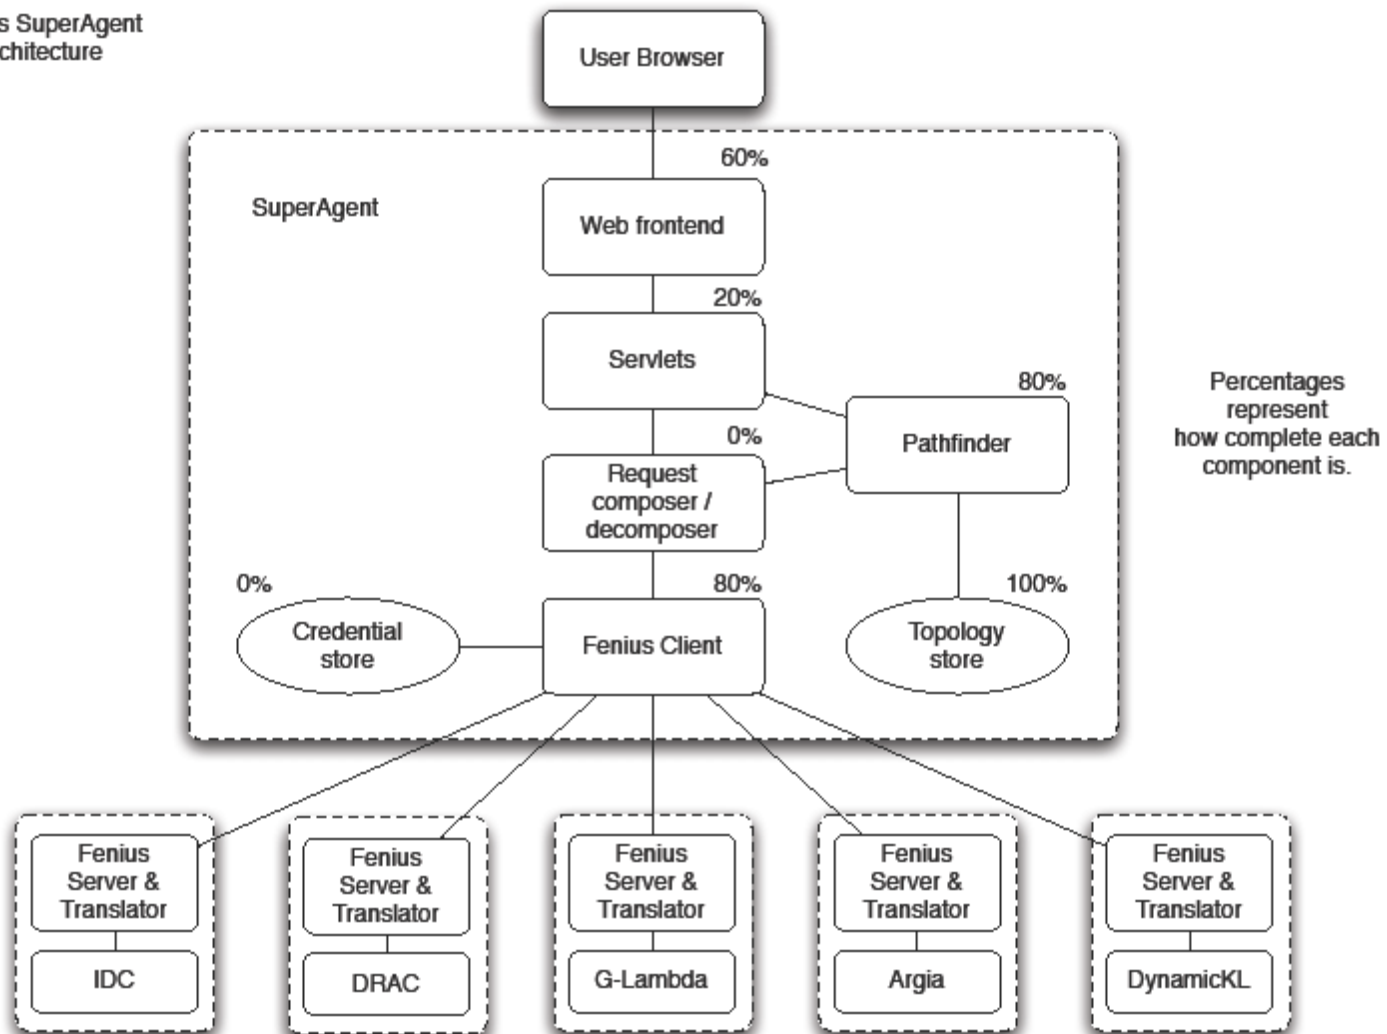
 *Currently also funded by various NSF awards to UCSD/UIC/NUST  R L I G H T SM



Fenius GLIF Demonstrations Global Lambda Grid Workshop SC10

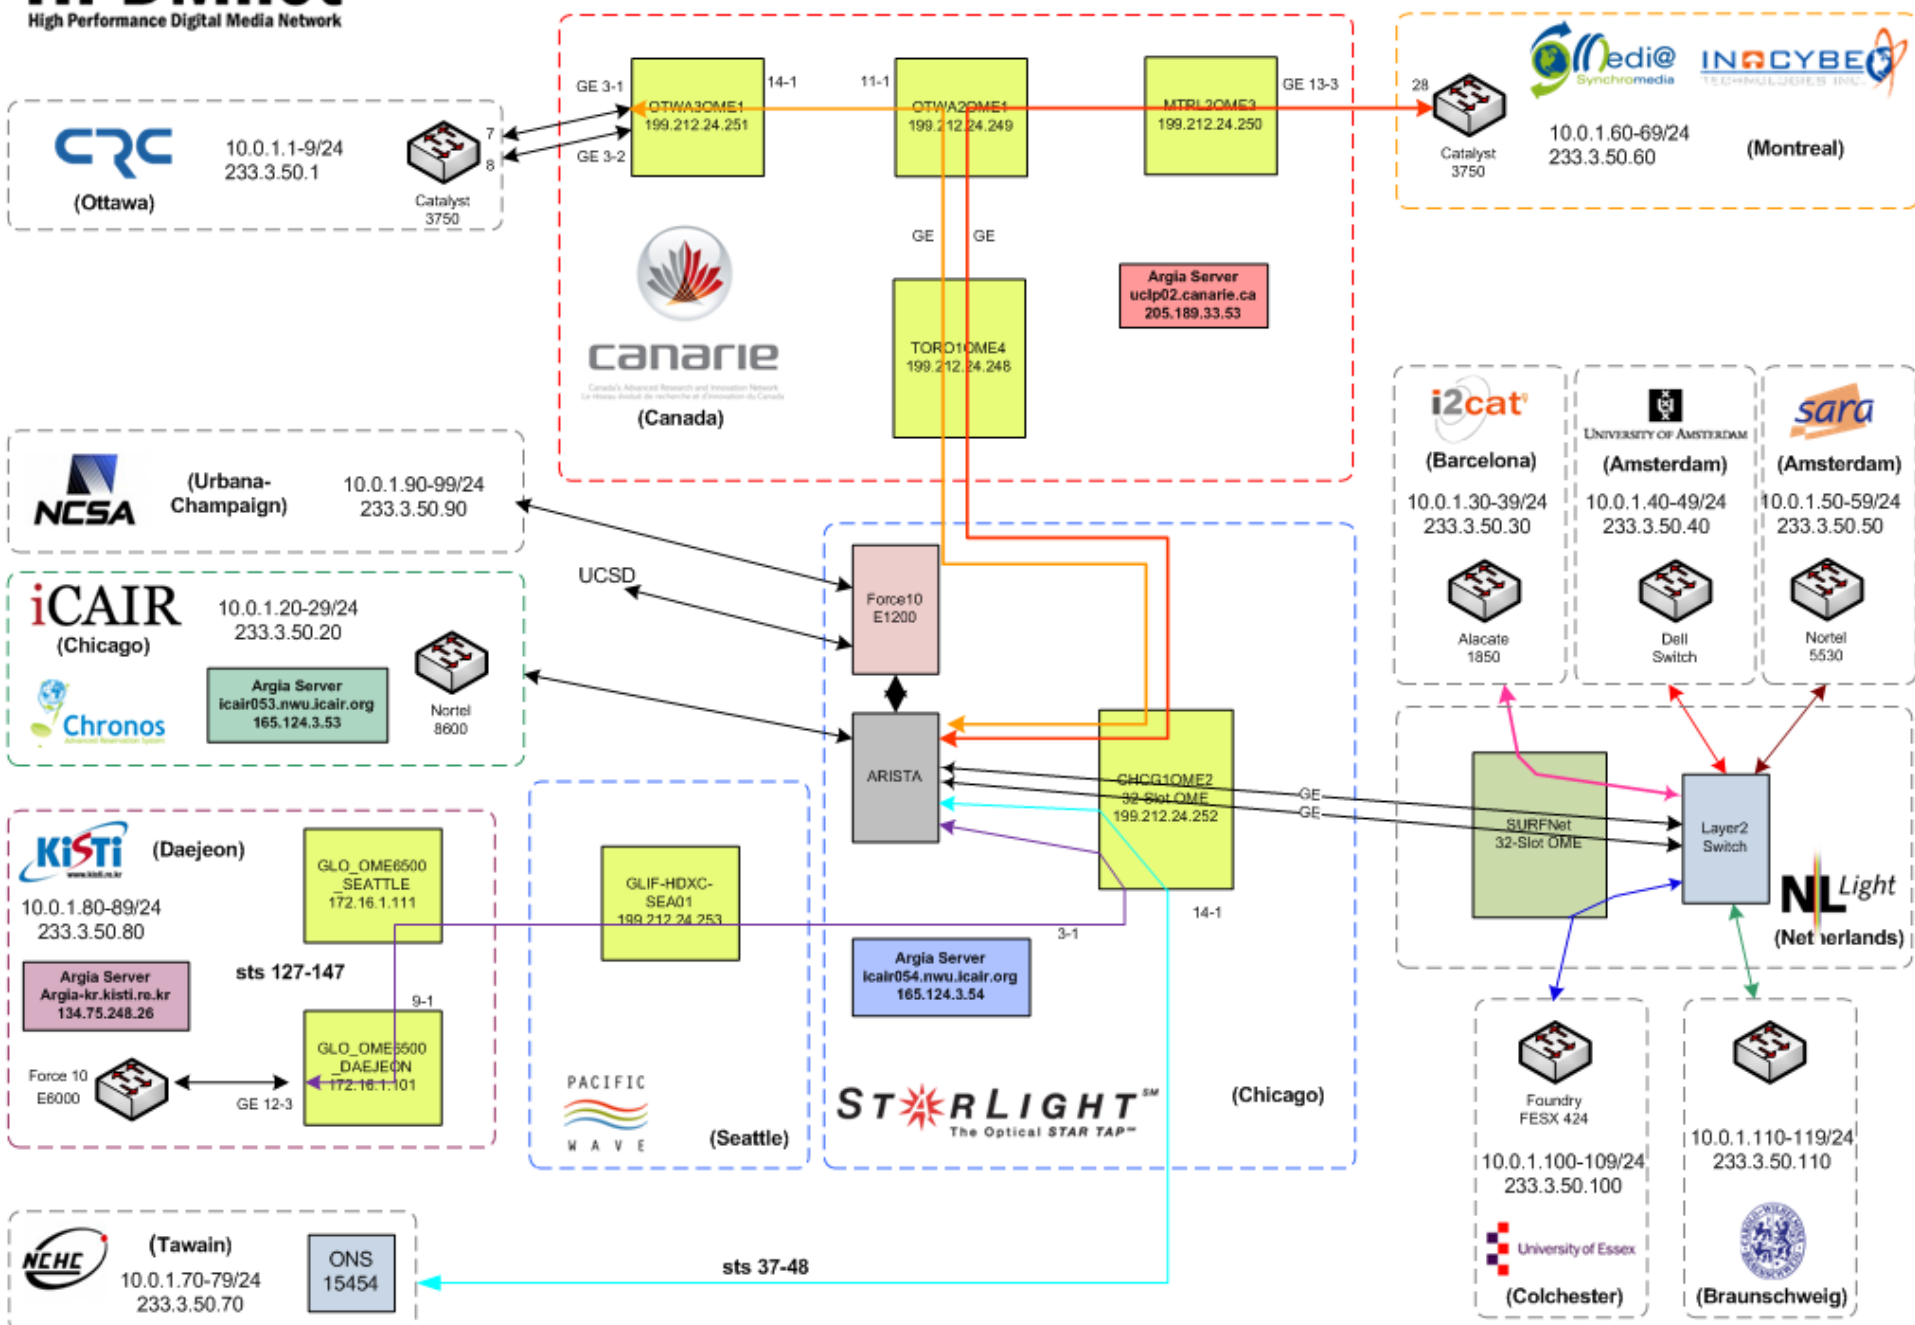


Fenius SuperAgent
architecture



HPDMnet

High Performance Digital Media Network





Projected (minimal) Network Topology 2014



Africa R&E Network Development
Planned by Egyptian GLORIAD
Partners under Chairmanship of
African Ministerial Conference on
Science and Technology (AMCOST)



Legend for Circuits

- GLORIAD Partner contribution
- GLORIAD Partner contribution +
NSF Cost Share under
GLORIAD ProNET award
- ★ Open Exchange Points (GLIF GOEs)

USA-RUSSIA-CHINA-KOREA-NETHERLANDS-CANADA-
DENMARK-FINLAND-ICELAND-NORWAY-SWEDEN

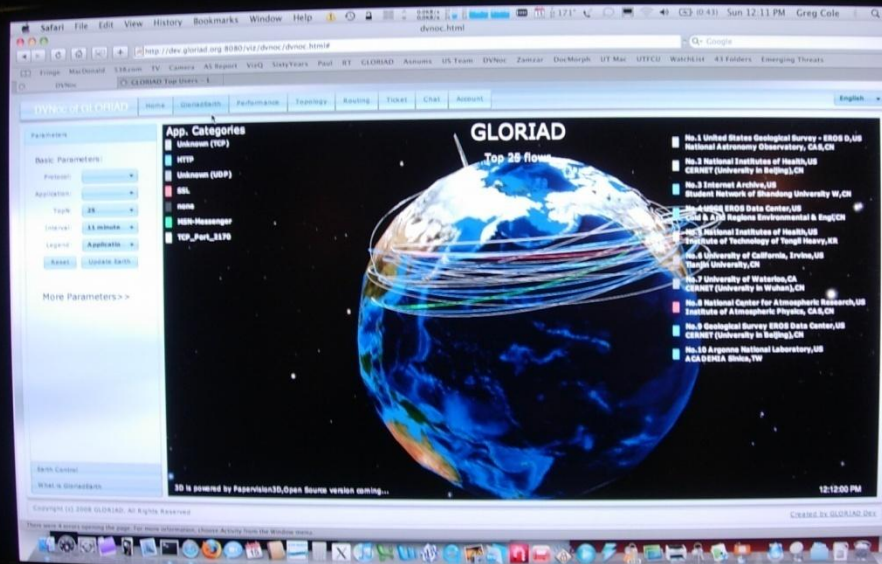
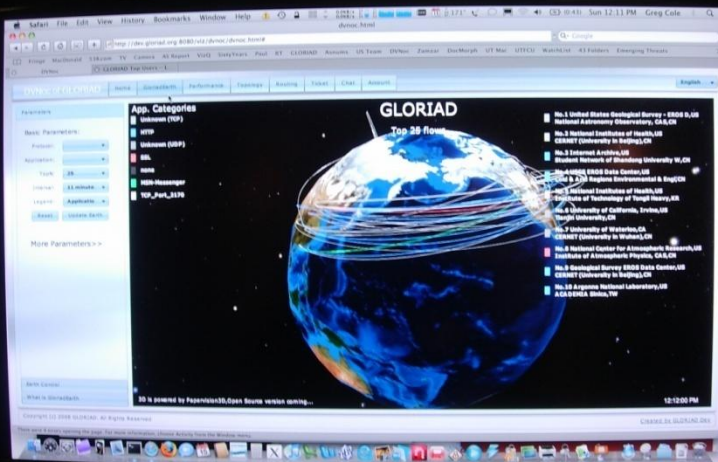
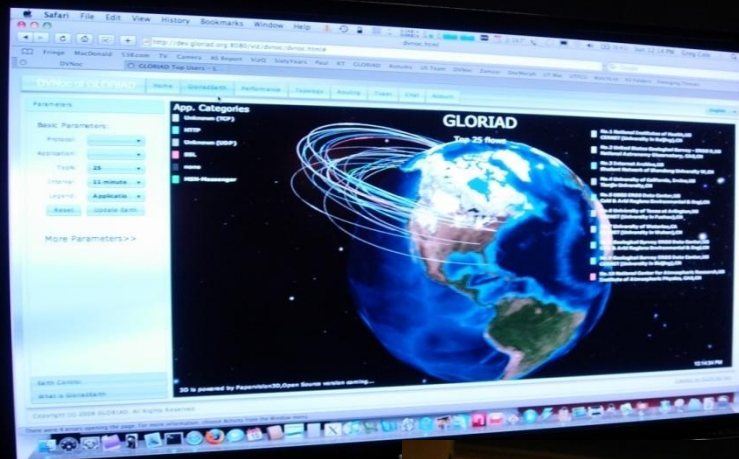
INDIA-EGYPT-SINGAPORE-VIETNAM-GREENLAND

GLORIAD-Taj Expansion



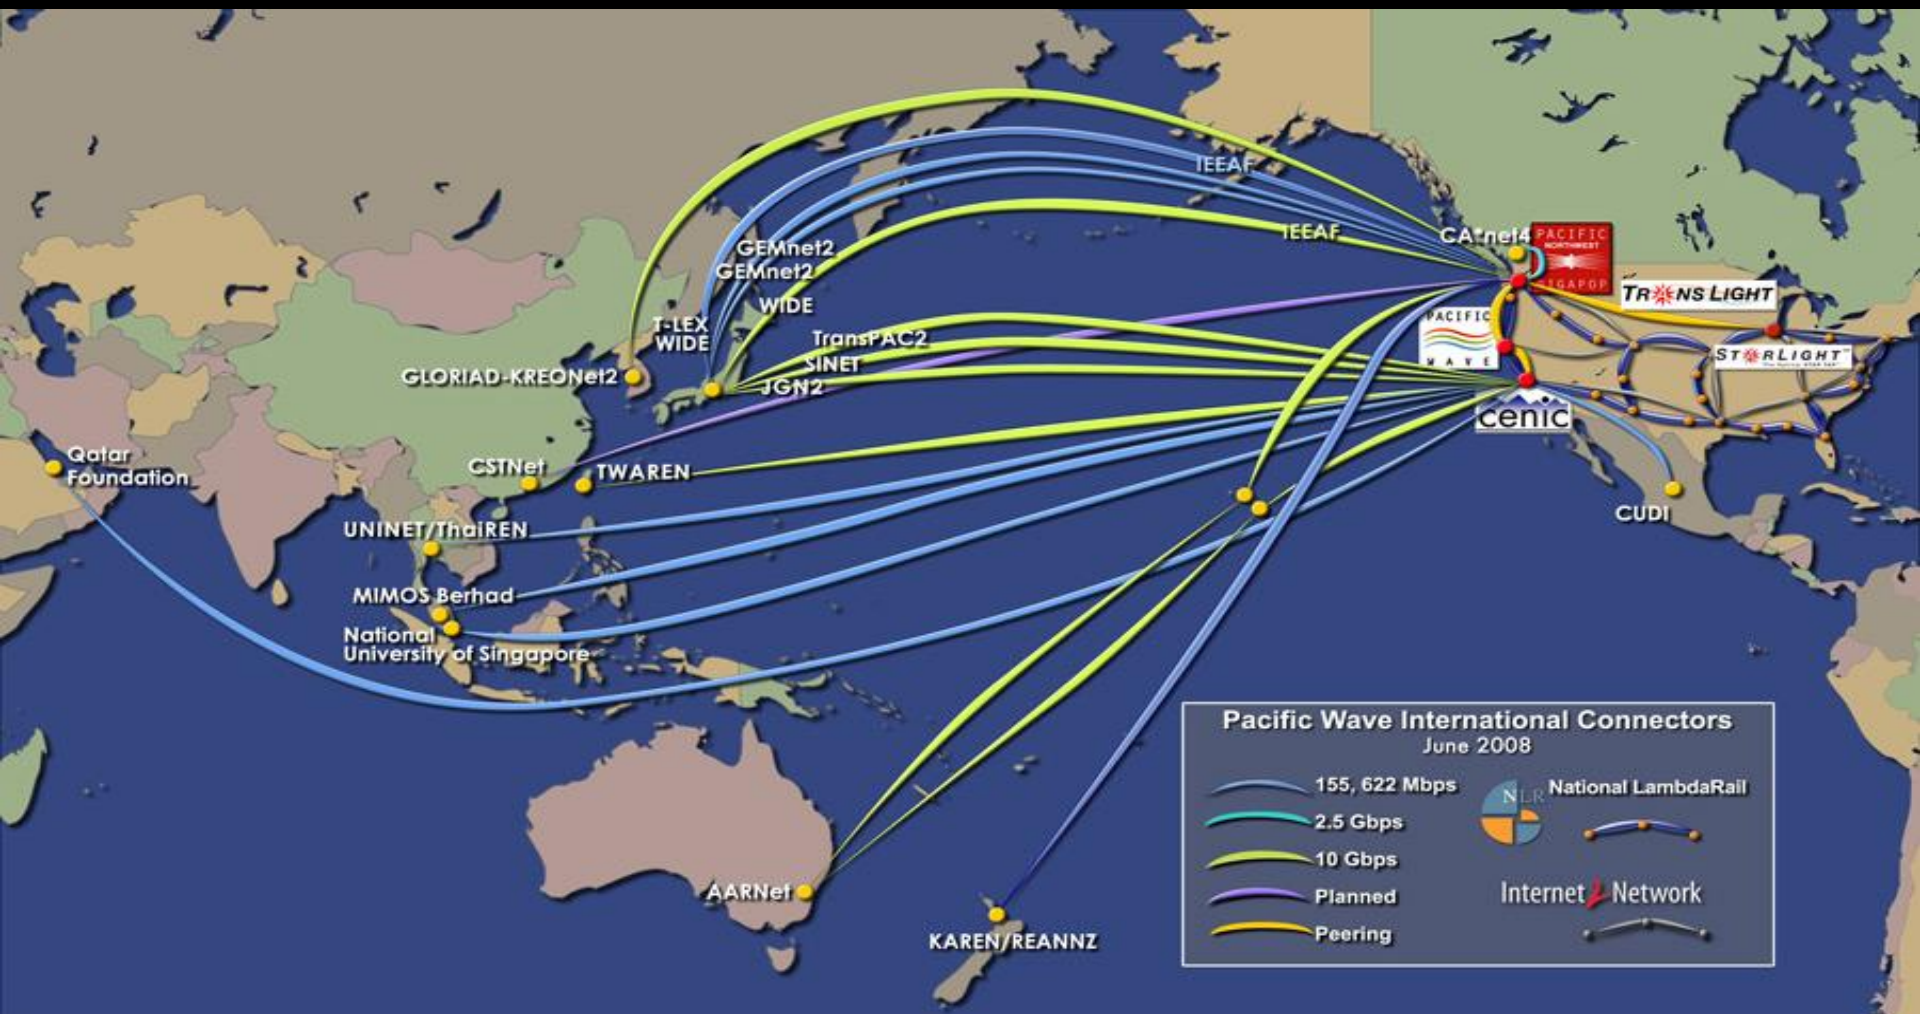
The new Taj expansion
is highlighted in orange
on this map

Global Ring Network for Advanced Applications Development



TransLight/Pacific Wave

10GE Wave Facilitates US West Coast Connectivity



Seattle, Sunnyvale and Los Angeles) to interconnect international and US research and education networks

Testbeds Preparing for Transition To 100 Gbps Services

- **Capacity**
 - Support for Capacity Much Beyond Aggregation of Millions of Small Flows
 - Support for Extremely Large Individual Stream (Including End-To-End)
- **Communications for Data Intensive Petascale Science**
- **Communications for Specialized Distributed Environments**
- **Environments Directly Controlled By Edge Processes (Application Specific Network Services)**
- **Highly Controllable Science Workflows**
- **Science Clouds (vs Consumer and Enterprise)**
- **Many New Applications and Services That Cannot Be Supported Today**



100 Gbps Services: Routing

- 100 Gbps Routing
- Available Today Based on Proprietary Technology
- Optimal Network Designs Place Such Devices At the Network Edge vs Network Core



100 Gbps Services: Client Side (b)

- **100 GigE Physical Layer Standard (PHY) Objectives**
 - **Preserve 802.3 / Ethernet Frame Format Based On 802.3 MAC**
 - **Preserve Min/Max Frame Size of 802.3 Standard**
 - **Provide PHY Specifications For Single-Mode Optical Fiber, Multi-Mode Optical Fiber (MMF), Copper Cables, Backplanes.**
 - **Support Bit Error Ratio (BER) Better Than or Equal to 10^{-12} at the MAC/PLS Service Interface**
 - **Provide Appropriate Support for Optical Transport Network (OTN) Standard**




100 Gbps Services: WAN Side/Line Side (a)

- **100 Gbps Optical Switching**
 - **Standard: ITU G.709 v3 (ODU4 100G)**
 - **ODU4/OTU4 Format -- Designed to Transport 100GbE (OTU4 = the ODU4 With FEC Included)**
 - **Formal Final Approval Took Place In Dec 2009**
 - **Beta Products Available Today**
 - **Ref: Demonstrations at SC10**
 - **1st Commercial Products Available End of Q2- Beginning Q3 2011**

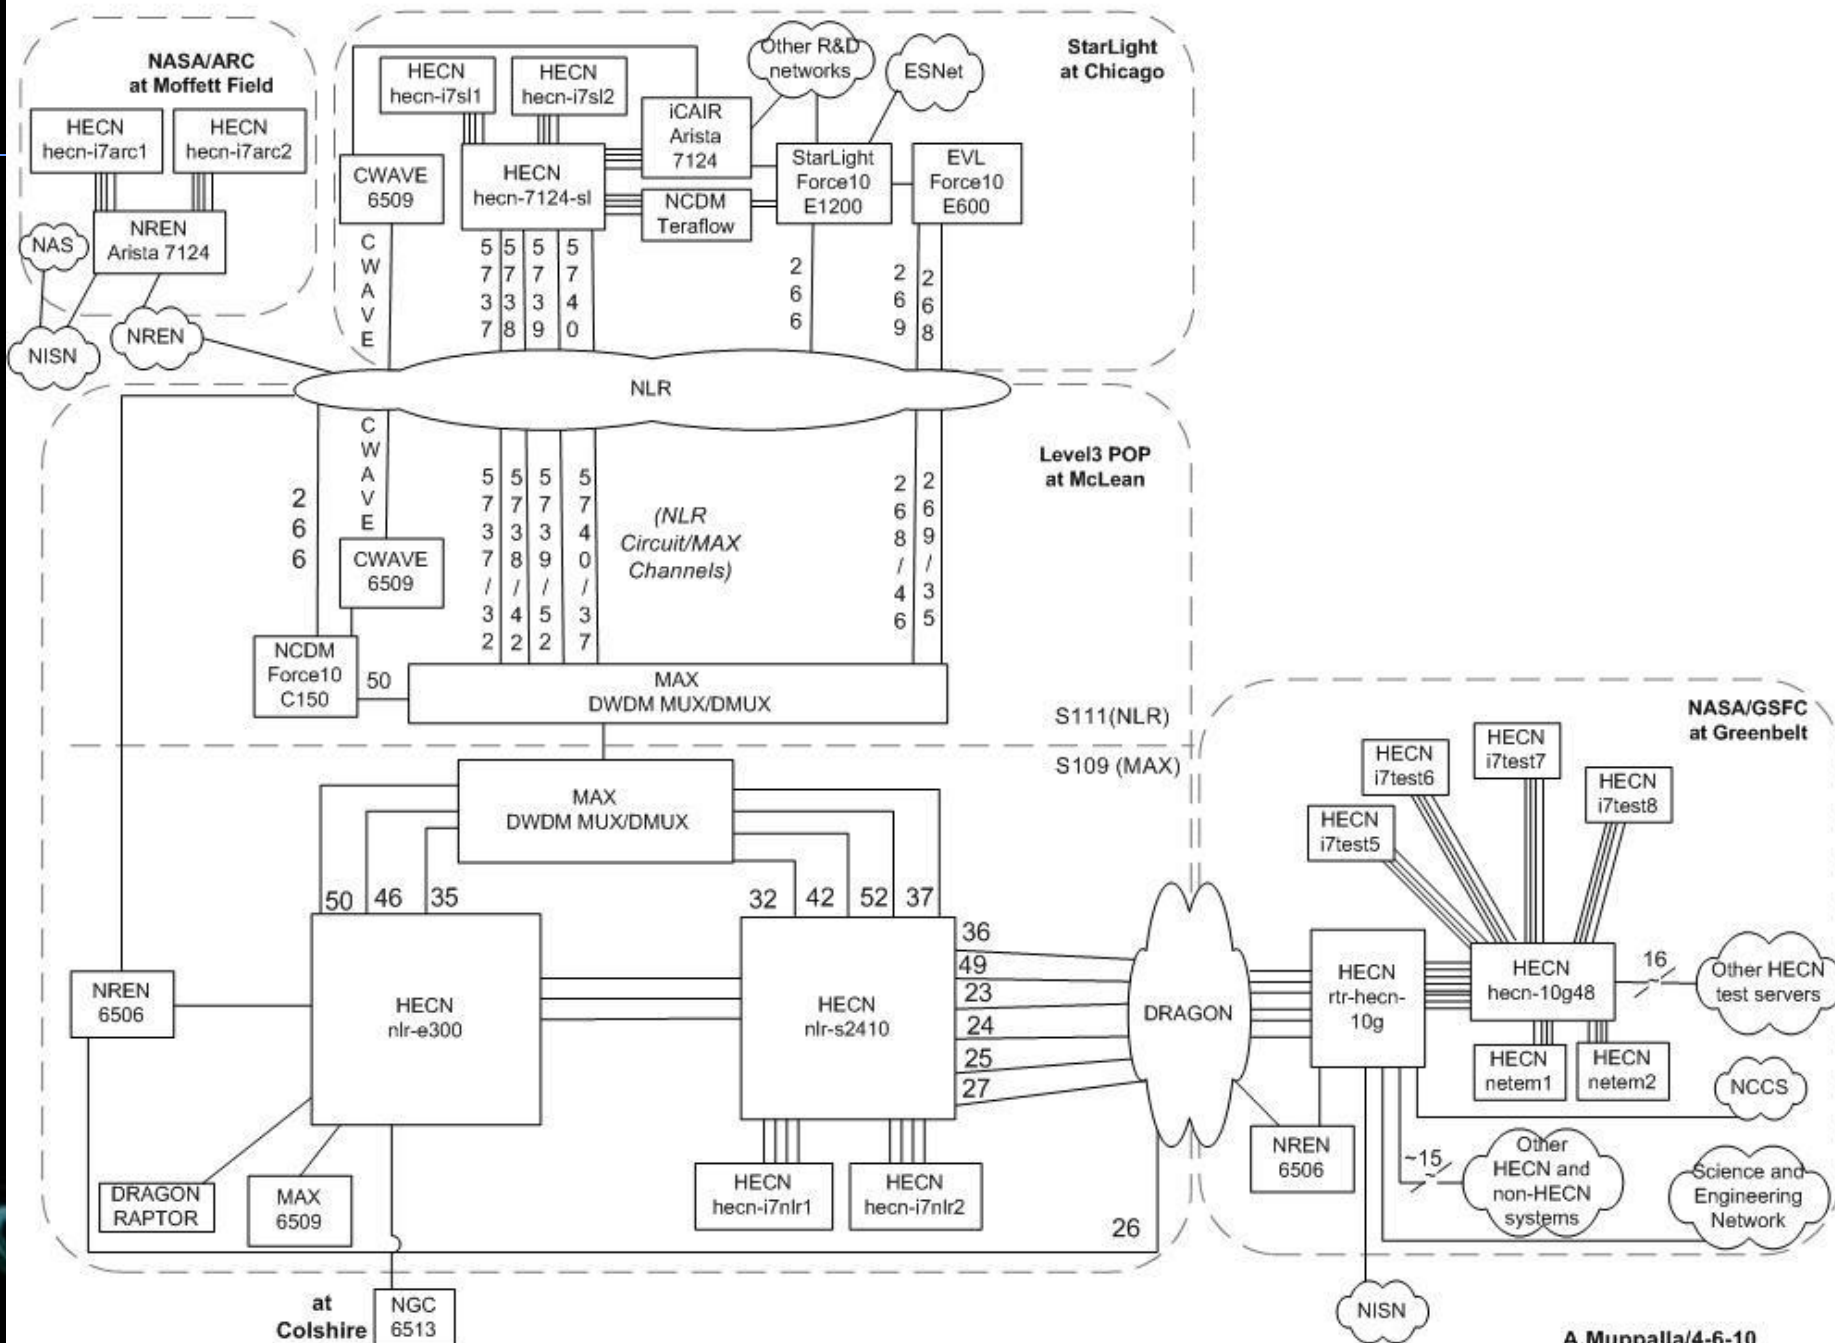


DOE ESnet Advanced Networking Initiative: 100 Gbps

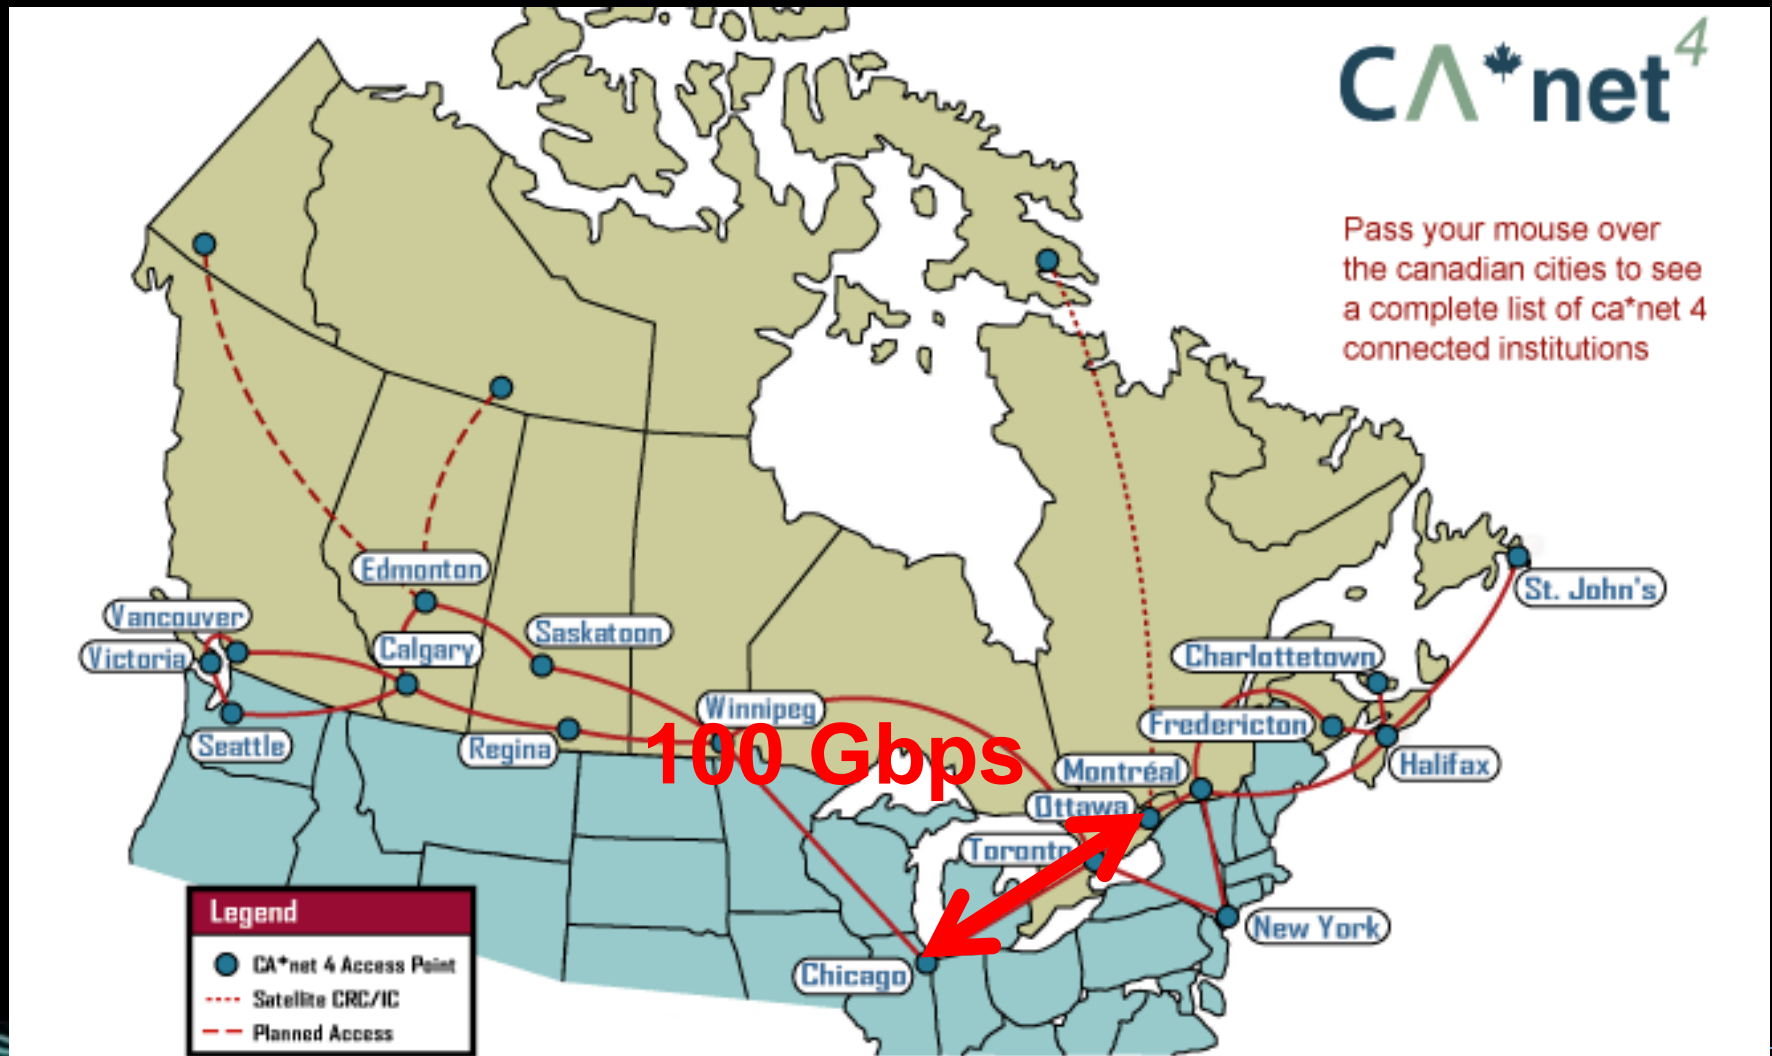


GSFC/High End Computer Network (HECN) and Partners 10GE and 10G Lambda Connections Through McLean

Note: The non-GSFC/HECN systems shown typically have other connections that are not shown in this diagram, as the focus is primarily GSFC/HECN connections




CA*net/Ciena/StarLight/iCAIR 100 Gbps Testbed



Source: CANARIE

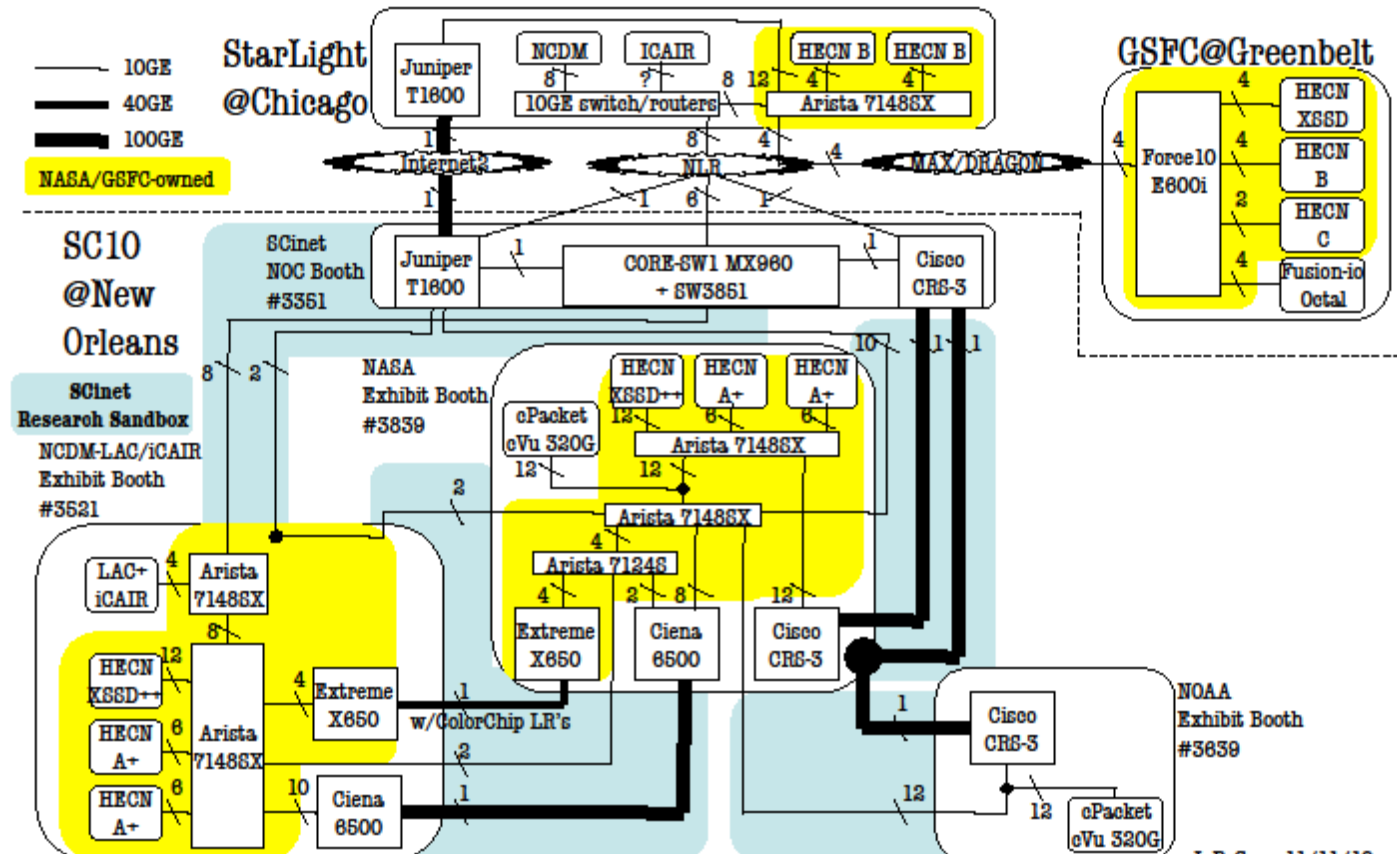
STARLIGHTSM

A Collaborative Initiative Among NASA, NLR, NOAA, Northwestern/iCAIR, SCinet & UIC/LAC



Using 100G Network Technology in Support of Petascale Science

A Collaborative Initiative Among NASA, NLR, NOAA, Northwestern/iCAIR, SCinet & UIC/LAC
Also Using Internet2's Multi-Vendor 100GigE Infrastructure Between StarLight and SC10



11/29/10

J. P. Gary

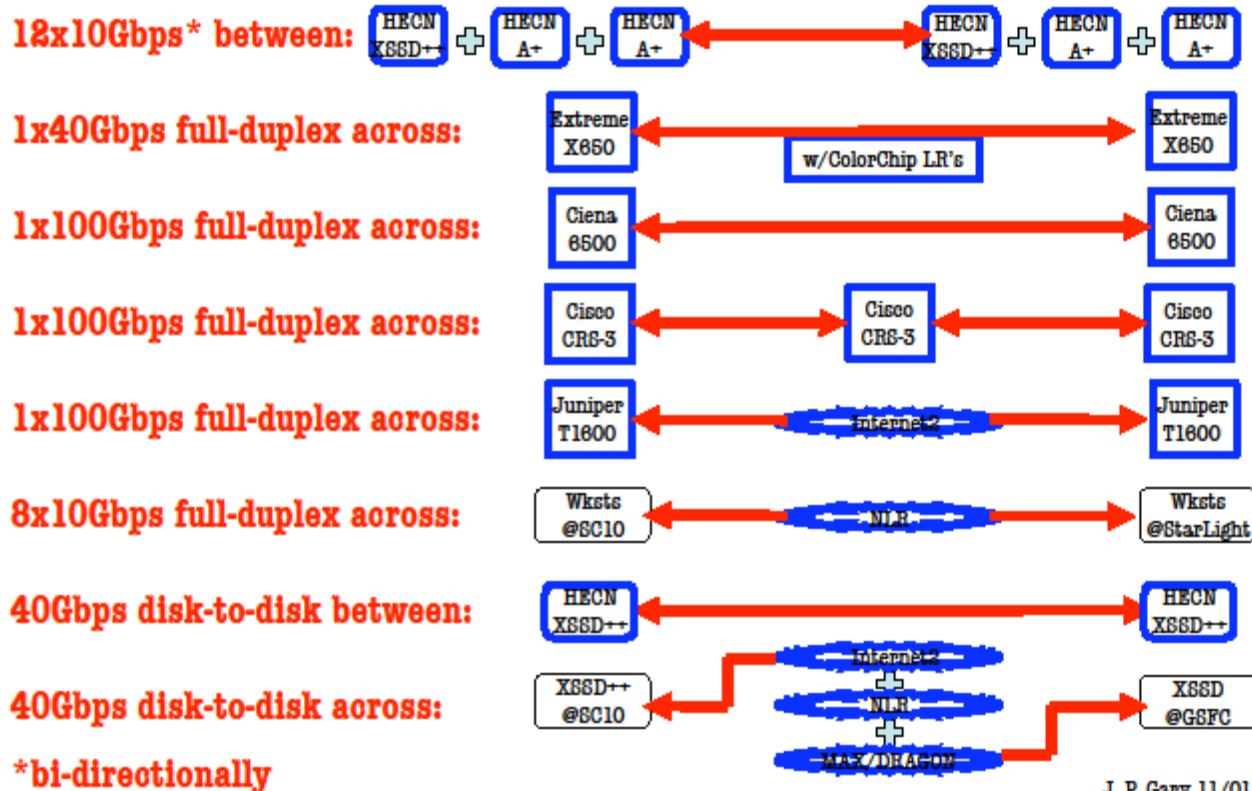
J. P. Gary 11/11/10

7

Using 100G Network Technology in Support of Petascale Science

A Collaborative Initiative Among NASA, NLR, NOAA, Northwestern/iCAIR, SCinet & UIC/LAC
Also Using Internet2's Multi-Vendor 100GigE Infrastructure Between StarLight and SC10

Demo Summary



11/29/10

J. P. Gary

J. P. Gary 11/01/10

6

StarWave Facility With StarLight

- **A National Science Foundation Funded Initiative**
- **Led by iCAIR**
- **Currently In Design Phase**
- **Scheduled for Summer 2011**
- **StarWave Will Provide Multi-100 Gbps Services for Data Intensive Science**
- **StarWave Will Integrated State-of-the-Art Architecture and Technology**
- **More Than Mere “Additional Capacity”**
- **Ref: Subsequent iCAIR Presentation on Multi-100 Gbps Services**



Terabit Networks for Extreme Scale Science

- **Workshop - February 16-17, 2011
Rockville, MD**
- **Sponsored by Advanced Scientific Computing
Research**
- **Department of Energy Office of Science**
- **Requirements for Significantly More Capacity**
- **New Architecture At All Levels**
- **New Technologies**

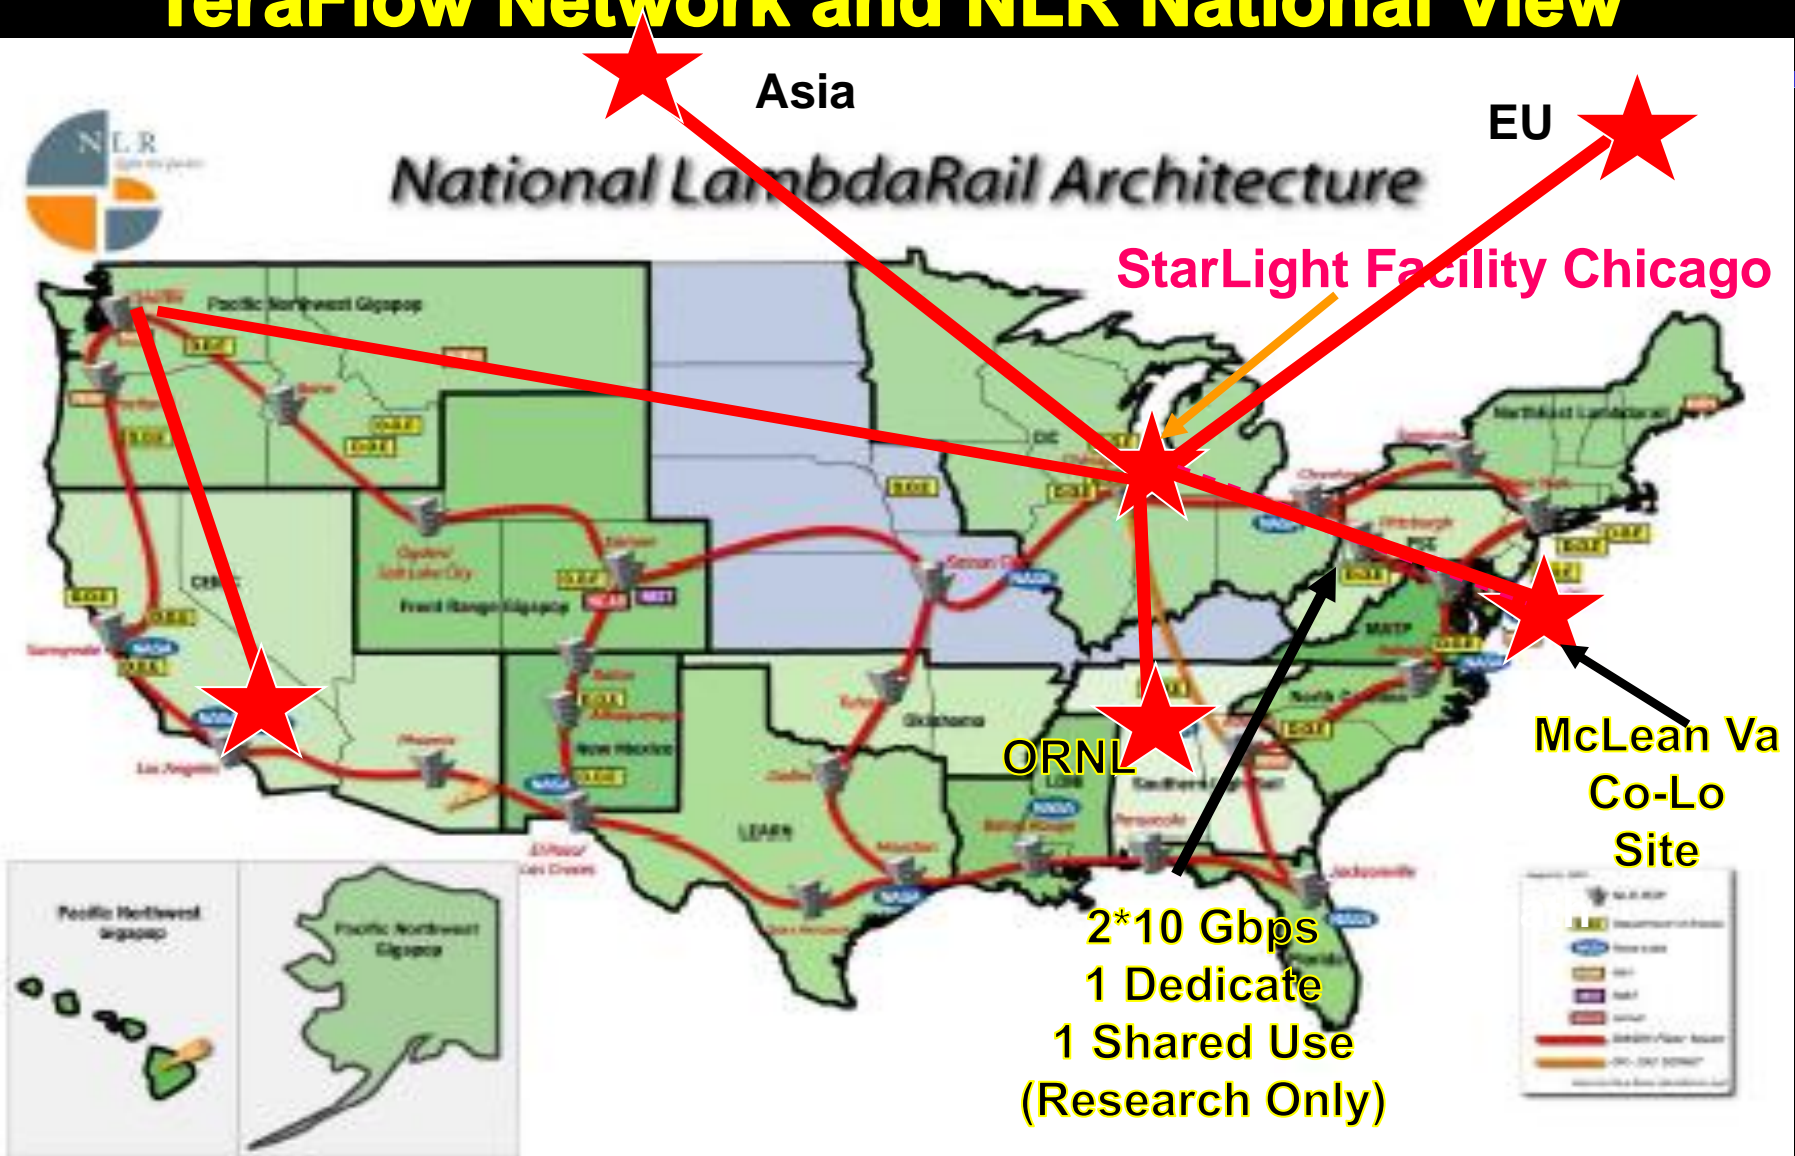




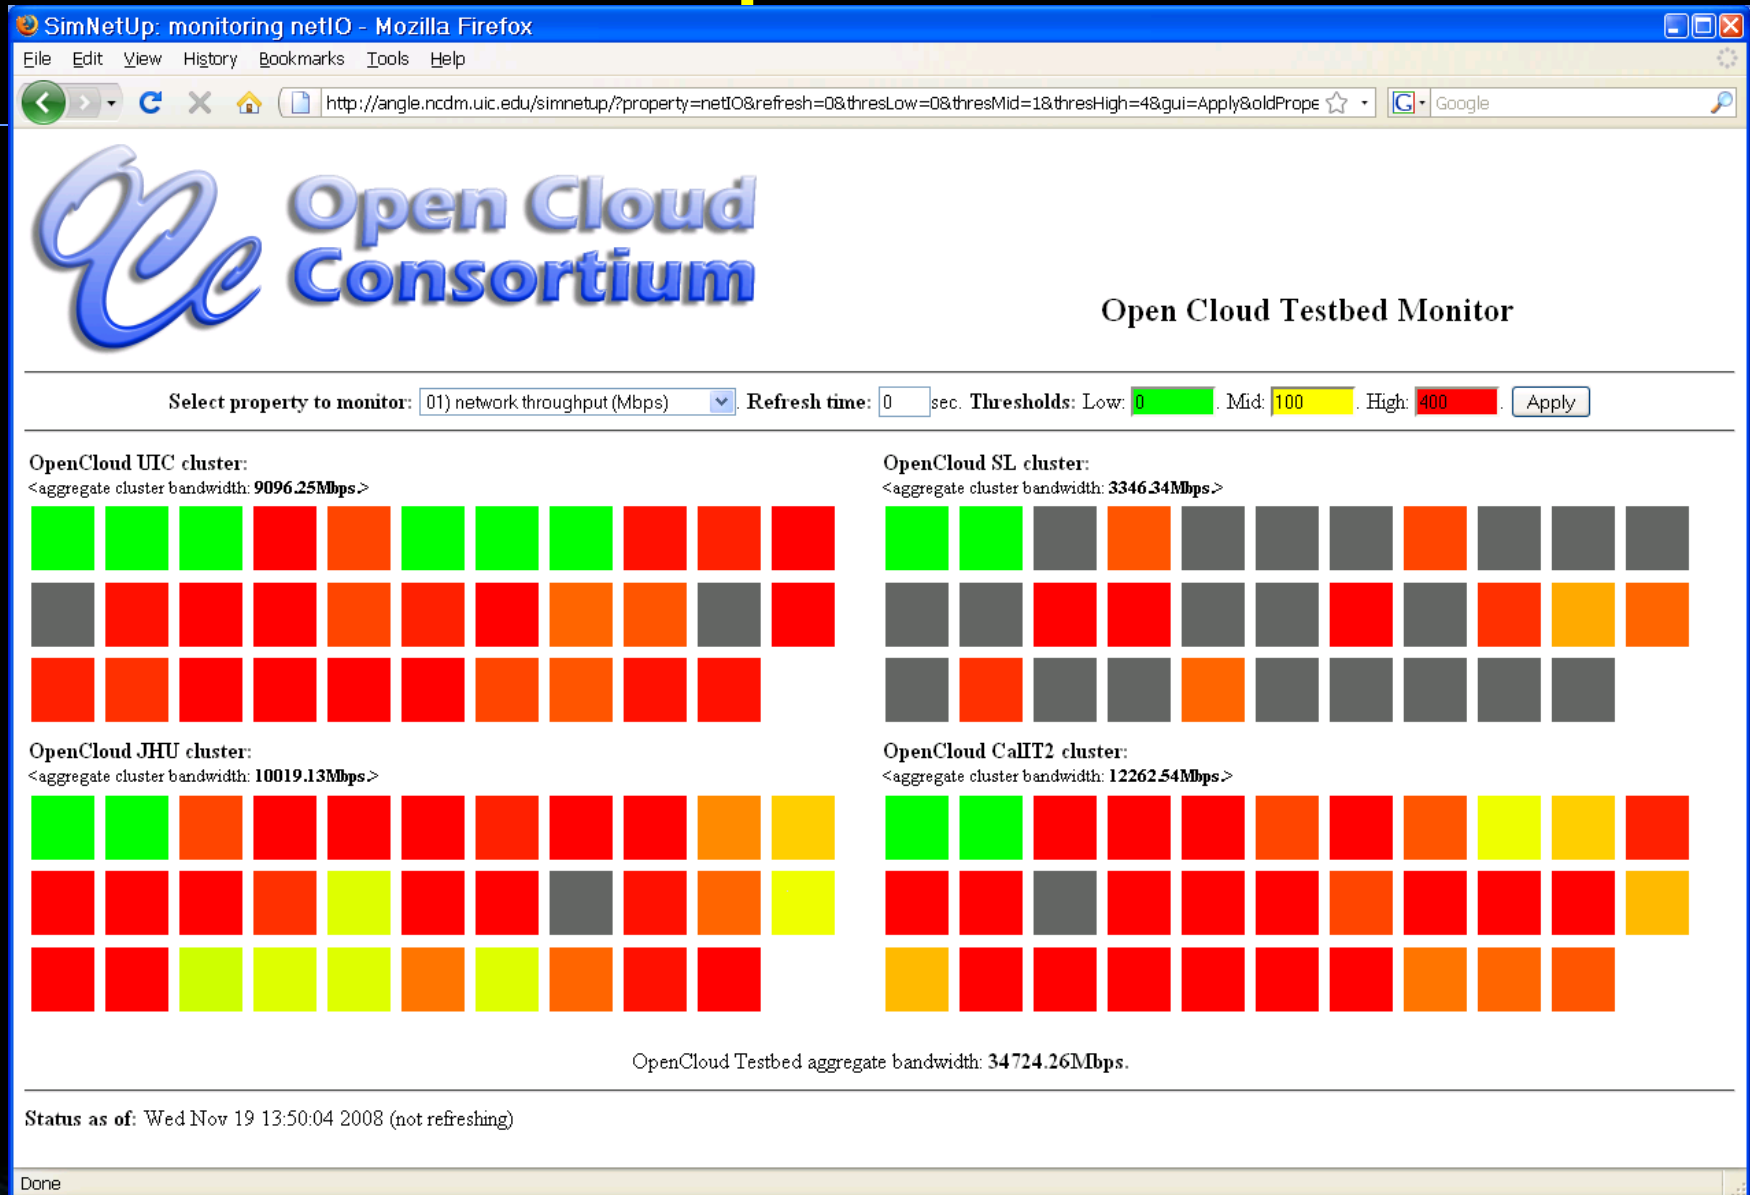
Open Cloud Consortium



TeraFlow Network and NLR National View



Terasort on Open Cloud Testbed



Source: NCDM, UIC

STARLIGHTSM

**IEEE
Communications
March 2006
Special Issue on
“An Optical
Control Plane for
the Grid
Community”
Ref: iCAIR Article
on ODIN – Optical
Dynamic
Intelligent
Networking**




Future Generation Computer Systems 2009

Special Issue Enabling Science With
Optical Communications: The Global
Lambda Integrated Facility

Larry Smarr, Maxine Brown,
and Cees de Laat (guest editors)

Future Generation Computer Systems,
Volume 25, Issue 2, Elsevier, February
2009, pp. 137-141

“OptIPuter: Enabling Advanced
Applications With Novel Optical Control
Planes and Backplanes,” iCAIR



Pascale Vicat-Blanc Primet
Tomohiro Kudoh
Joe Mambretti (Eds.)

2

Networks for Grid Applications

Second International Conference, GridNets 2008
Beijing, China, October 2008
Revised Selected Papers



 Springer

 **R LIGHT**SM

Anastasios Doulamis
Joe Mambretti
Ioannis Tomkos
Theodora Varvarigou (Eds.)




25

Networks for Grid Applications

Third International ICST Conference, GridNets 2009
Athens, Greece, September 2009
Revised Selected Papers



 Springer



www.startap.net/starlight

**Thanks to Deniz Gurkan for this Event
and To Our Many Partner Organizations:
Universities,
National and International Research Labs,
International and National Networking Organizations,
Technology Corporations, Federal Agencies
(Especially the National Science Foundation)
and Other Supporters**



STARLIGHTSM