Global Lambda Exchanges

Dr. Thomas A. DeFanti

Distinguished Professor of Computer Science, University of Illinois at Chicago

Director, Electronic Visualization Laboratory , University of Illinois at Chicago

Principal Investigator, TransLight/StarLight

Research Scientist, California Institute for Telecommunications and Information Technology, University of California, San Diego

Electronic Visualization Laboratory
33 Years of Computer Science and Art

- EVL established in 1973
- Tom DeFanti, Maxine Brown, Dan Sandin, Jason Leigh
- Students in CS, ECE, Art+Design
- >1/3 century of collaboration with artists and scientists to apply new computer science techniques to these disciplines
- Computer Science+ Art→Computer Graphics, Visualization, VR
- Supercomputing+Networking→ Lambda Grids
- Research in:
  - Advanced display systems
  - Visualization and virtual reality
  - Advanced networking
  - Collaboration and human computer interaction
- Funding mainly NSF, ONR, NIH. Also: (NTT), General Motors

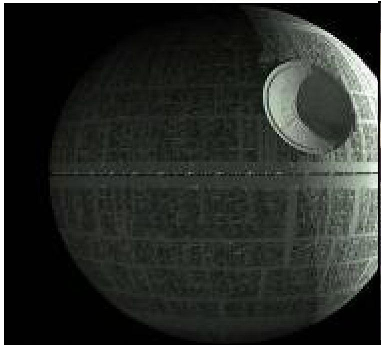
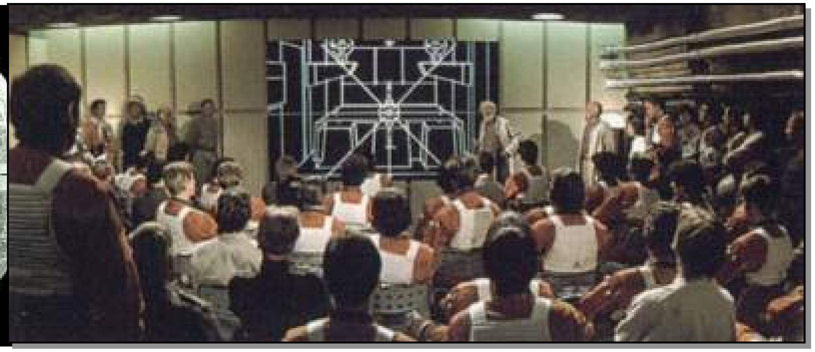
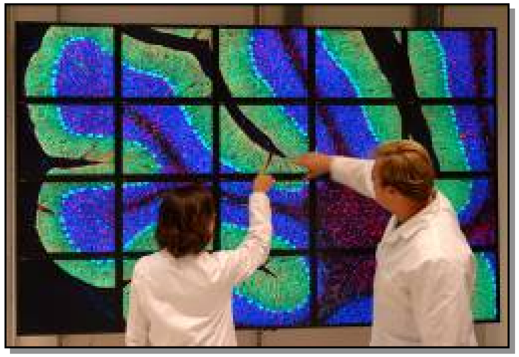


STAR TAP and StarLight

NSF-funded support of STAR TAP (1997-2000) and STAR TAP2/ StarLight (2000-2005), and the High Performance International Internet Services program (Euro-Link, TransPAC, MIRnet and AMPATH).

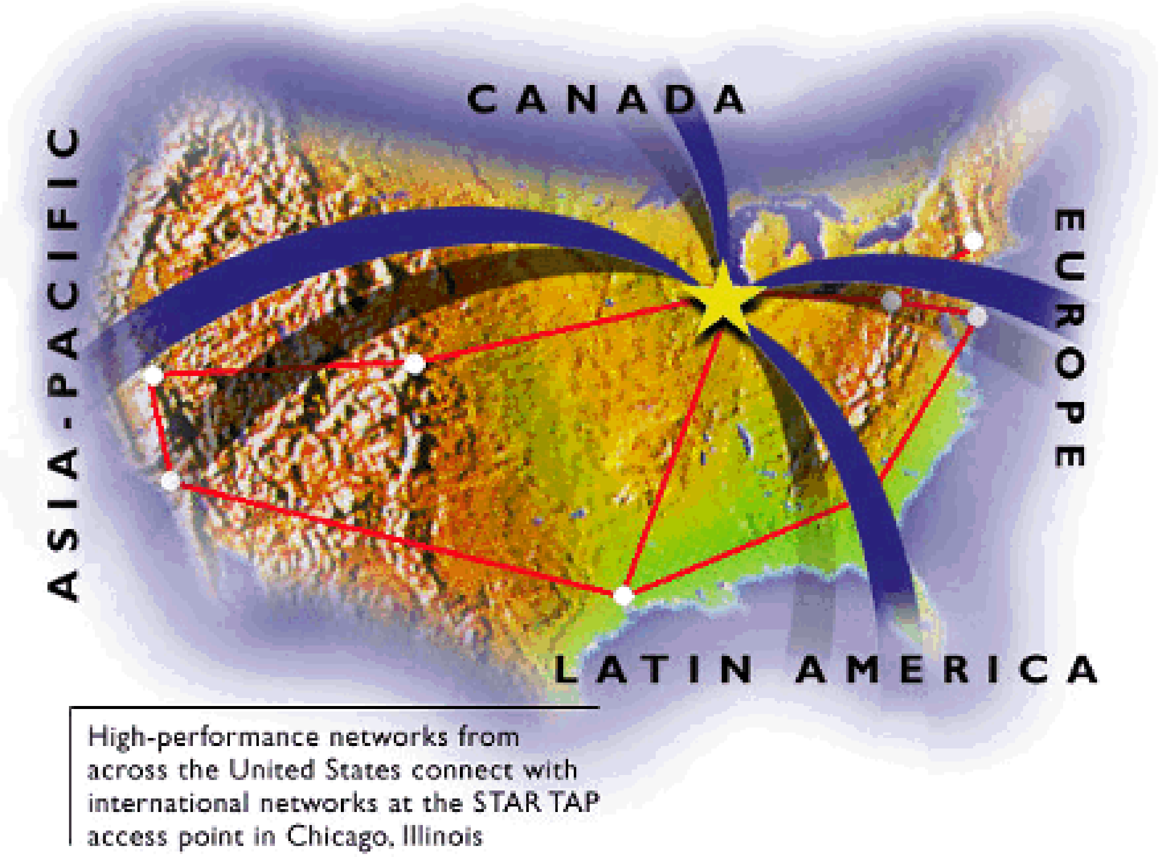

StarLight: A 1 Gigabit and 10 Gigabit Exchange

StarLight hosts optical switching, electronic switching and electronic routing for United States national and international Research and Education networks

StarLight opened in 2001

Abbott Hall, Northwestern University's Chicago downtown campus

TRANSLIGHT℠ glif it² iGrid 2005 evl

iGrid 1998 at SC'98
November 7-13, 1998, Orlando, Florida, USA

- **10 countries:** Australia, Canada, CERN, Germany, Japan, Netherlands, Russia, Singapore, Taiwan, USA
- **22 demonstrations** featured technical innovations and application advancements requiring high-speed networks, with emphasis on remote instrumentation control, tele-immersion, real-time client server systems, multimedia, tele-teaching, digital video, distributed computing, and high-throughput, high-priority data transfers. See: www.startap.net/igrid98

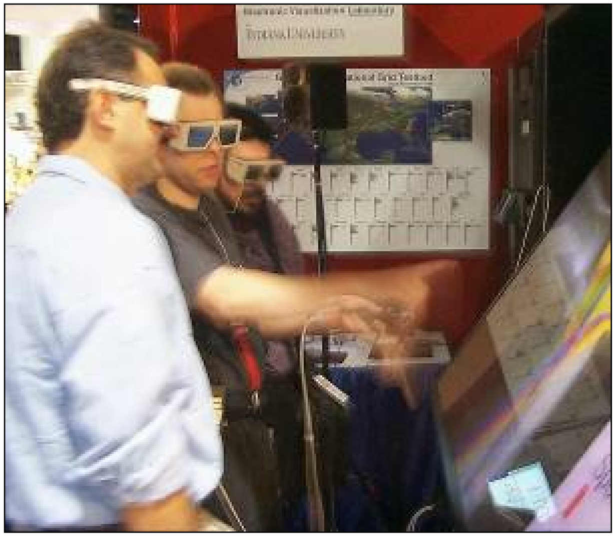

TRANSLIGHT SM glif it2 iGrid 2005 evl

iGrid 2000 at INET 2000
July 18-21, 2000, Yokohama, Japan

- **14 countries:** Canada, CERN, Germany, Greece, Japan, Korea, Mexico, Netherlands, Singapore, Spain, Sweden, Taiwan, United Kingdom, USA
- **24 demonstrations** featuring technical innovations in tele-immersion, large datasets, distributed computing, remote instrumentation, collaboration, streaming media, human/computer interfaces, digital video and high-definition television, and grid architecture development, and application advancements in science, engineering, cultural heritage, distance education, media communications, and art and architecture. See: www.startap.net/igrid2000
- **100Mb transpacific bandwidth carefully managed**

TRANSLIGHT℠ glif it² iGrid 2005 evl

iGrid 2002
September 24-26, 2002, Amsterdam, The Netherlands

- **28 demonstrations from 16 countries:** Australia, Canada, CERN/Switzerland, France, Finland, Germany, Greece, Italy, Japan, Netherlands, Singapore, Spain, Sweden, Taiwan, the United Kingdom and the USA.
- Applications demonstrated: art, bioinformatics, chemistry, cosmology, cultural heritage, education, high-definition media streaming, manufacturing, medicine, neuroscience, physics. See: www.startap.net/igrid2002



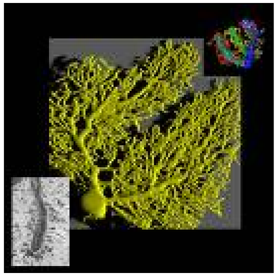


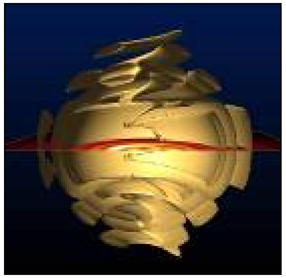

- **Grid technologies** demonstrated: Major emphasis on grid middleware, data management grids, data replication grids, visualization grids, data/visualization grids, computational grids, access grids, grid portals
- **25Gb transatlantic bandwidth (100Mb/attendee, 250x iGrid2000!)**

TRANSLIGHT℠ glif it² iGrid 2005 evl

iGrid 2005
September 26-30, 2005, San Diego, California

- **Networking enabled by the Global Lambda Integrated Facility (GLIF) – the international virtual organization creating a global LambdaGrid laboratory**
- **More than 150Gb GLIF transoceanic bandwidth alone; 100Gb of bandwidth into the Calit2 building on the UCSD campus!**
- **49 demonstrations showcasing global experiments in e-Science and next-generation shared open-source LambdaGrid services**
- **20 countries: Australia, Brazil, Canada, CERN, China, Czech Republic, Germany, Hungary, Italy, Japan, Korea, Mexico, Netherlands, Poland, Russia, Spain, Sweden, Taiwan, UK, USA. See: www.startap.net/igrid2005**




TRANSLIGHT℠ glif it² iGrid 2005 evl

iGrid 2005:
Demonstrating Emerging LambdaGrid Services

- Data Transport
- High-Definition Video & Digital Cinema Streaming
- Distributed High-Performance Computing
- Lambda Control
- Lambda Security
- Scientific Instruments
- Visualization and Virtual Reality
- e- Science

Source: Maxine Brown, EVL UIC

iGrid2005 Data Flows Multiplied Normal Flows by Five Fold!

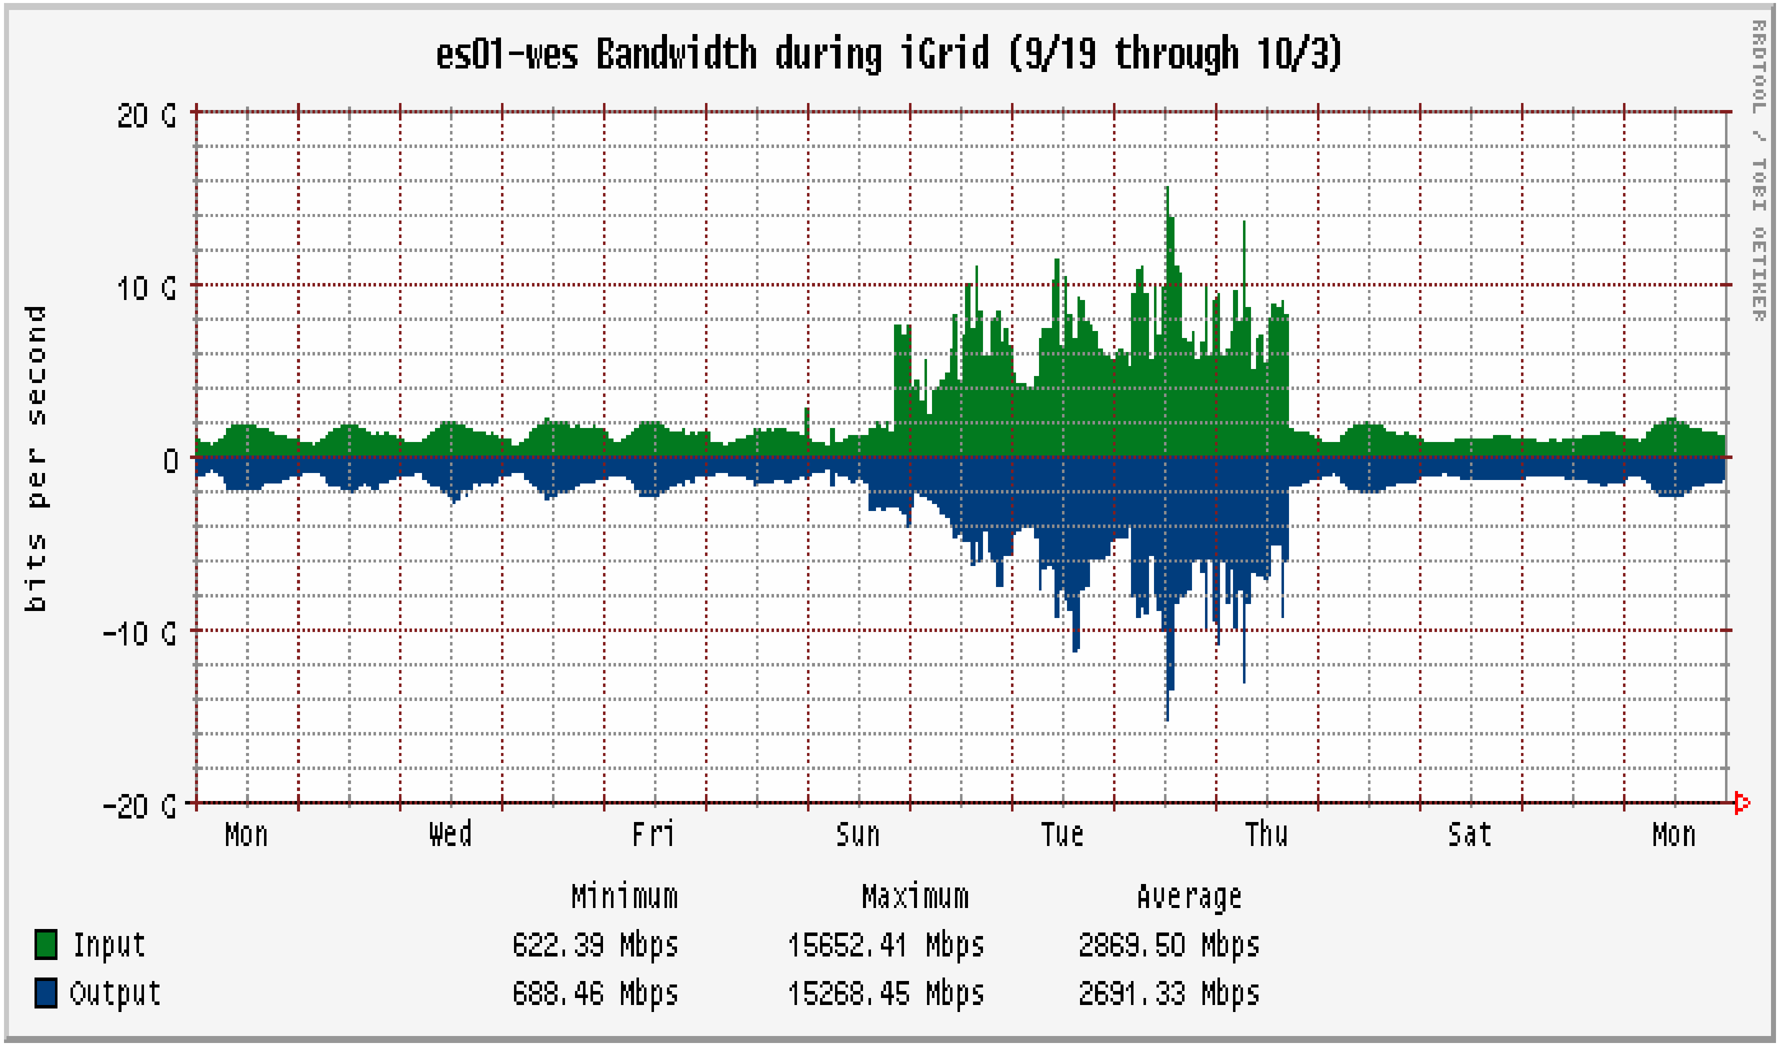

Data Flows Through the Seattle PacificWave International Switch

CENIC 2006 "Innovations in Networking" Award for iGrid 2005

www.igrid2005.org
www.cenic.org

CENIC is the Corporation for Education Network Initiatives in California

TRANSLIGHT[SM]

StarLight and TransLight Partners 2006

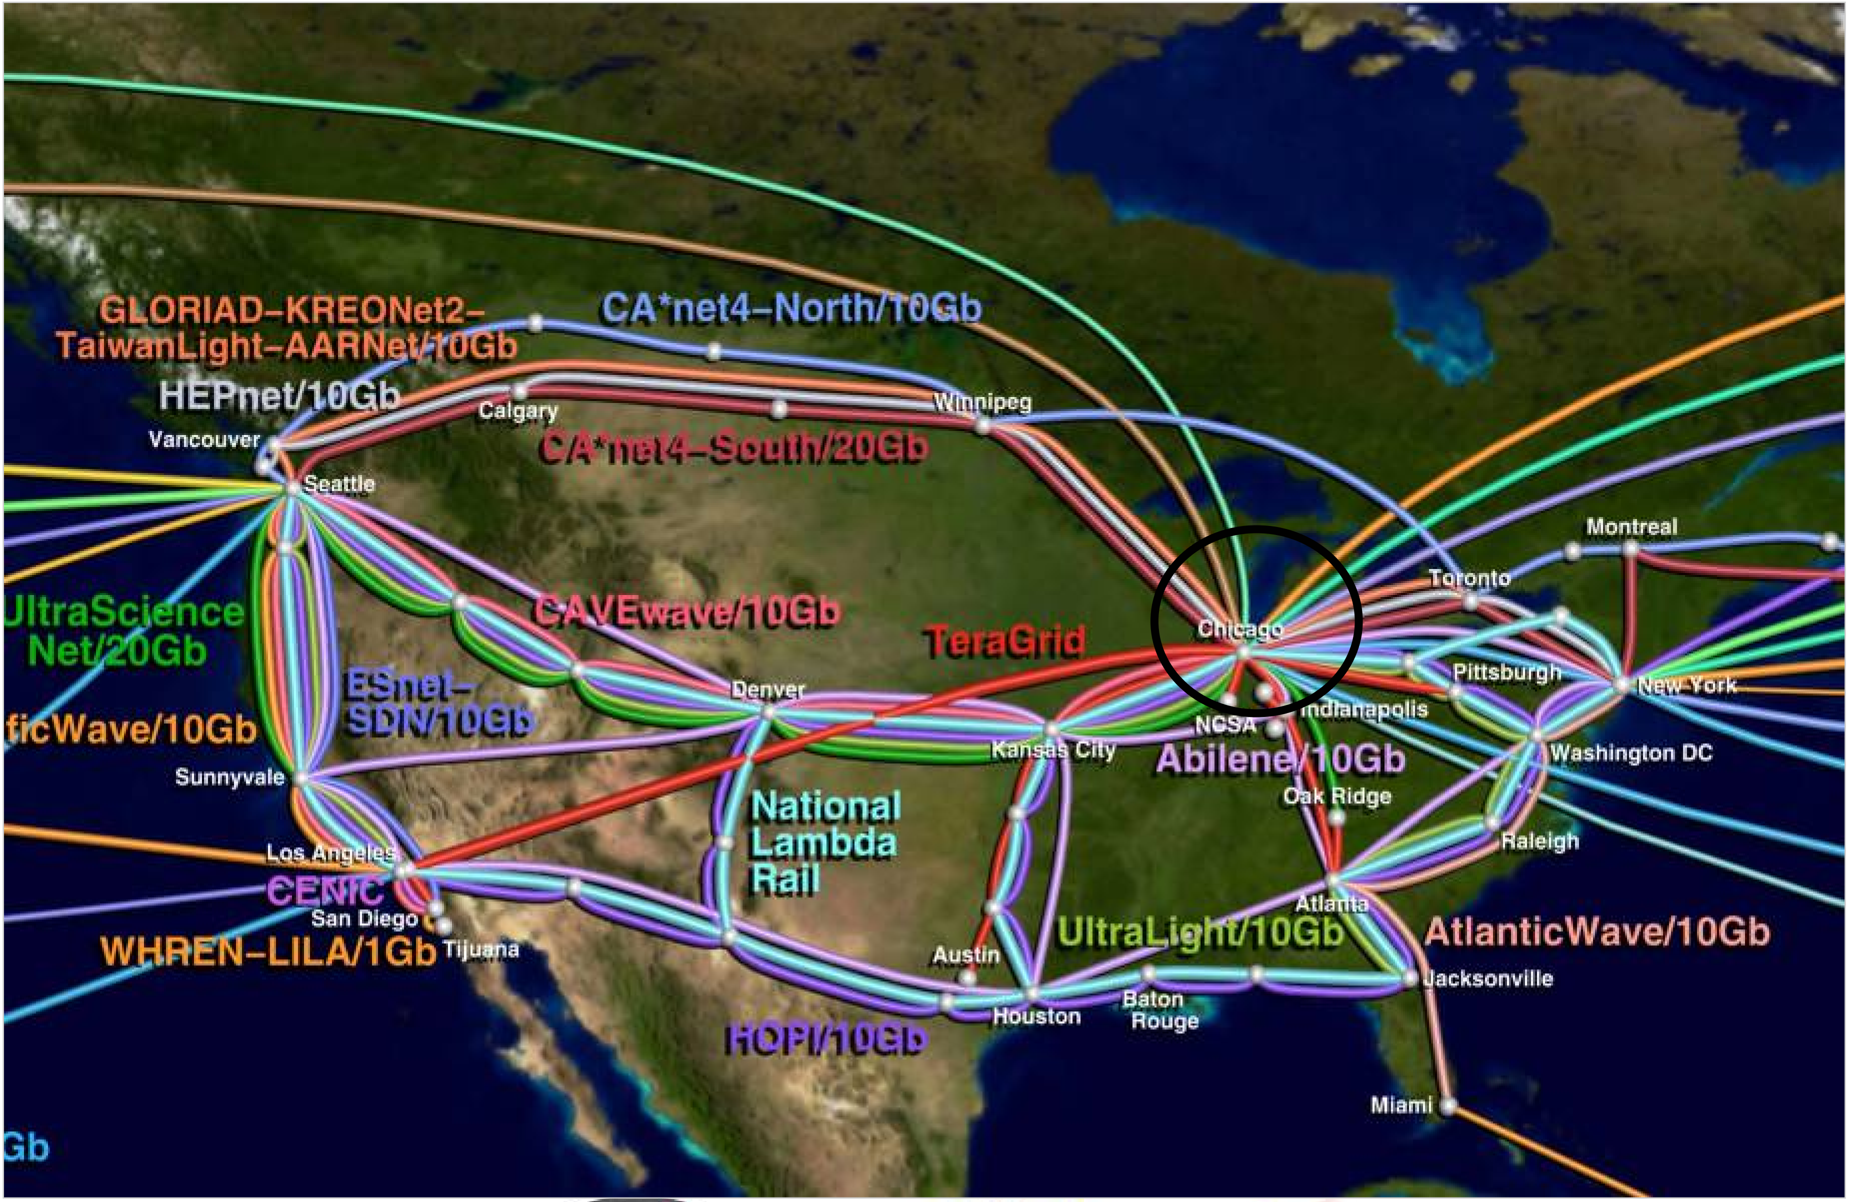

TRANSLIGHT℠ glif it² iGrid 2005 evl

Why Photonics?

- **Many of the highest performance e-science applications involve national and international collaboration.**
- **This was the purpose of StarTAP (ATM) and StarLight (GE and 10GE).**
- **The next generation networking infrastructure must interoperate globally!**
- **Colleagues in Japan (such as Aoyama-sensei and Murai-sensei, colleagues at the University of Tokyo, Keio, and NTT Labs) and in America, Canada, Netherlands, Korea, China, UK, Czech Republic and elsewhere, agreed in 2003 to form a loose global initiative to create a global photonic network testbed for the common good.**
- **We call this *GLIF*, the Global Lambda Integrated Facility.**

TRANSLIGHT℠ glif it² iGrid 2005 evl

Some Applications that Need Photonics

- **Interactive collaboration using video (SD, HD, SHD) and/or VR**
  - **Low latency streaming (real-time use)**
  - **High data rates**
  - **Lossy protocols OK**
  - **Multi-channel, multi-cast**
- **Biomedical Imaging**
  - **Very high resolution 2D (tens to hundreds of megapixels)**
  - **Volume visualizations (billions of zones in 3D)**
- **Geoscience Imaging**
  - **Very high resolution 2D (tens to hundreds of megapixels)**
  - **Volume visualizations (billions of zones in 3D)**
- **Digital cinema**
  - **Large data sets**
  - **Security**
- **Metagenomics**
  - **Large computing**
  - **Large data sets**

TRANSLIGHT[SM] glif it² iGrid 2005 evl

High-Resolution Media Users Need Multi-Gb/s Networks

- **e-Science 2D images with hundreds of Mega-pixels**
  - **Microscopes and telescopes**
  - **Remote sensing satellites and aerial photography**
  - **Multi-spectral, not just visible light, so 32 bits/pixel or more**
- **GigaZone 3-D objects with billions of volume elements**
  - **Supercomputer simulations**
  - **Seismic imaging for earthquake R&D and energy exploration**
  - **Medical imaging for diagnosis and R&D**
  - **Zones are often multi-valued (taking dozens of bytes each)**
- **Digital Cinema uses 250Mb/s for theatrical distribution, but up to 14Gb/s for post-production**
- **Interactive analysis and visualization of such data objects is *impossible* today**
- **Focus of the GLIF: deploy new system architectures ASSUMING photonic network availability**

California Institute for Telecommunications and Information Technology (Calit2)

- **New Laboratory Facilities**
  - **Nanotech, BioMEMS, Chips, Radio, Photonics, Grid, Data, Applications**
  - **Virtual Reality, Digital Cinema, HDTV, Synthesis**
- **Over 1000 Researchers in Two Buildings**
  - **Linked via Dedicated Optical Networks**
  - **International Conferences and Testbeds**

Preparing for an World in Which Distance Has Been Eliminated…

The OptIPuter Project

Removing Bandwidth as an Obstacle In Data Intensive Sciences

- An NSF-funded project that focuses on developing technology to enable the real-time collaboration and visualization of very-large time-varying volumetric datasets for the Earth sciences and the biosciences
- OptIPuter is examining a new model of computing whereby ultra-high-speed networks form the backplane of a global computer

NIH Biomedical Informatics Research Network

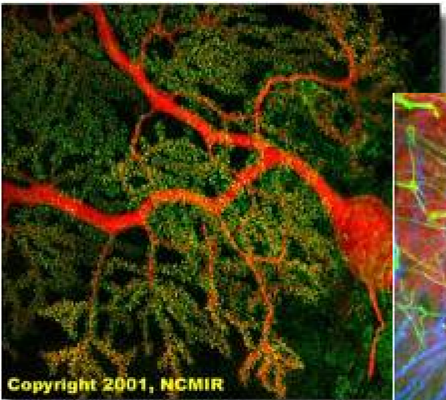

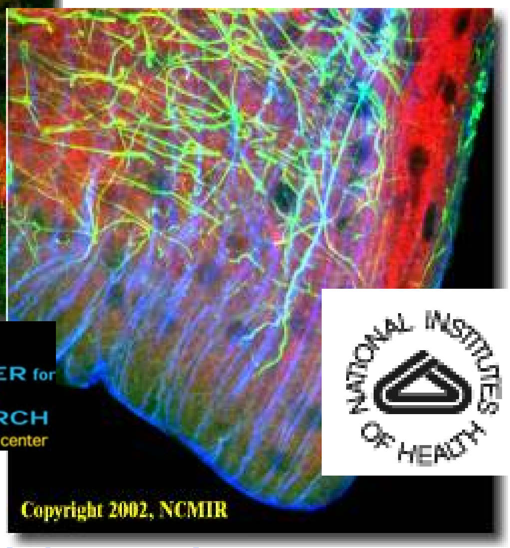

NSF EarthScope and ORION

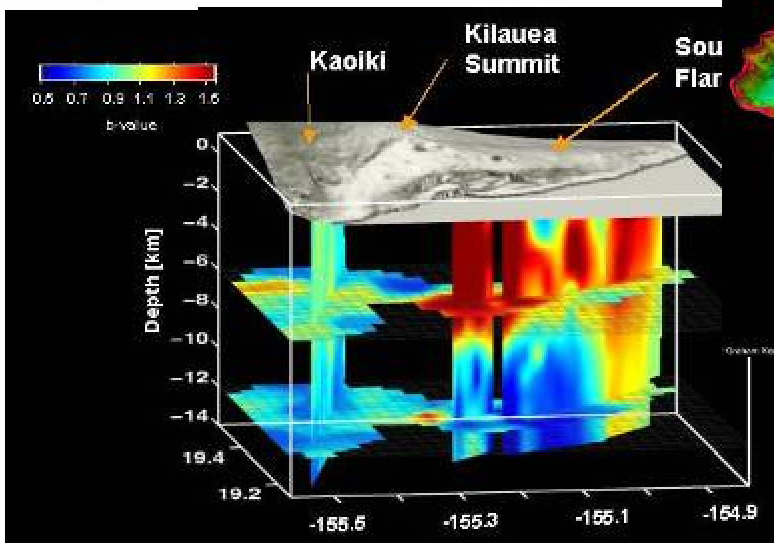

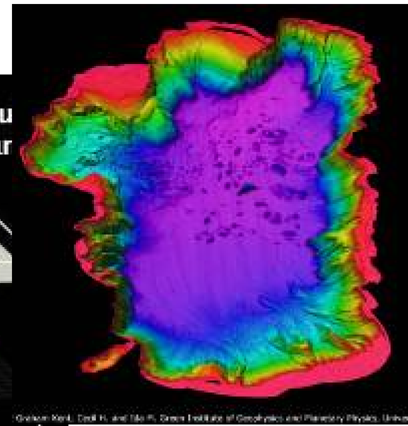

http://ncmir.ucsd.edu/gallery.html

www.optiputer.net

siovizcenter.ucsd.edu/library/gallery/shoot1/index.shtml

TRANSLIGHT℠ glif iGrid 2005 evl

The OptIPuter Tiled Displays and Lambda Grid Enable Persistent Collaboration Spaces

- Hardware installations assembled at each site
- Unified software at each site (Rocks Viz Roll w/ stable integration of SAGE)
- Refined TeraVision for Streaming HDTV (video conferencing and microscope outputs)
- Controls for launching images from application portals

Goal: Use these systems for conducting collaborative experiments

www.optiputer.net

TRANSLIGHT℠ glif it² iGrid 2005 evl

Biomedical Imaging

Source: Steven T. Peltier

JuxtaView showing ~600 megapixel montage dataset from Amsterdam

HDTV stream from a light microscope at NCMIR

HDTV camera feed shows the conference room at NCMIR

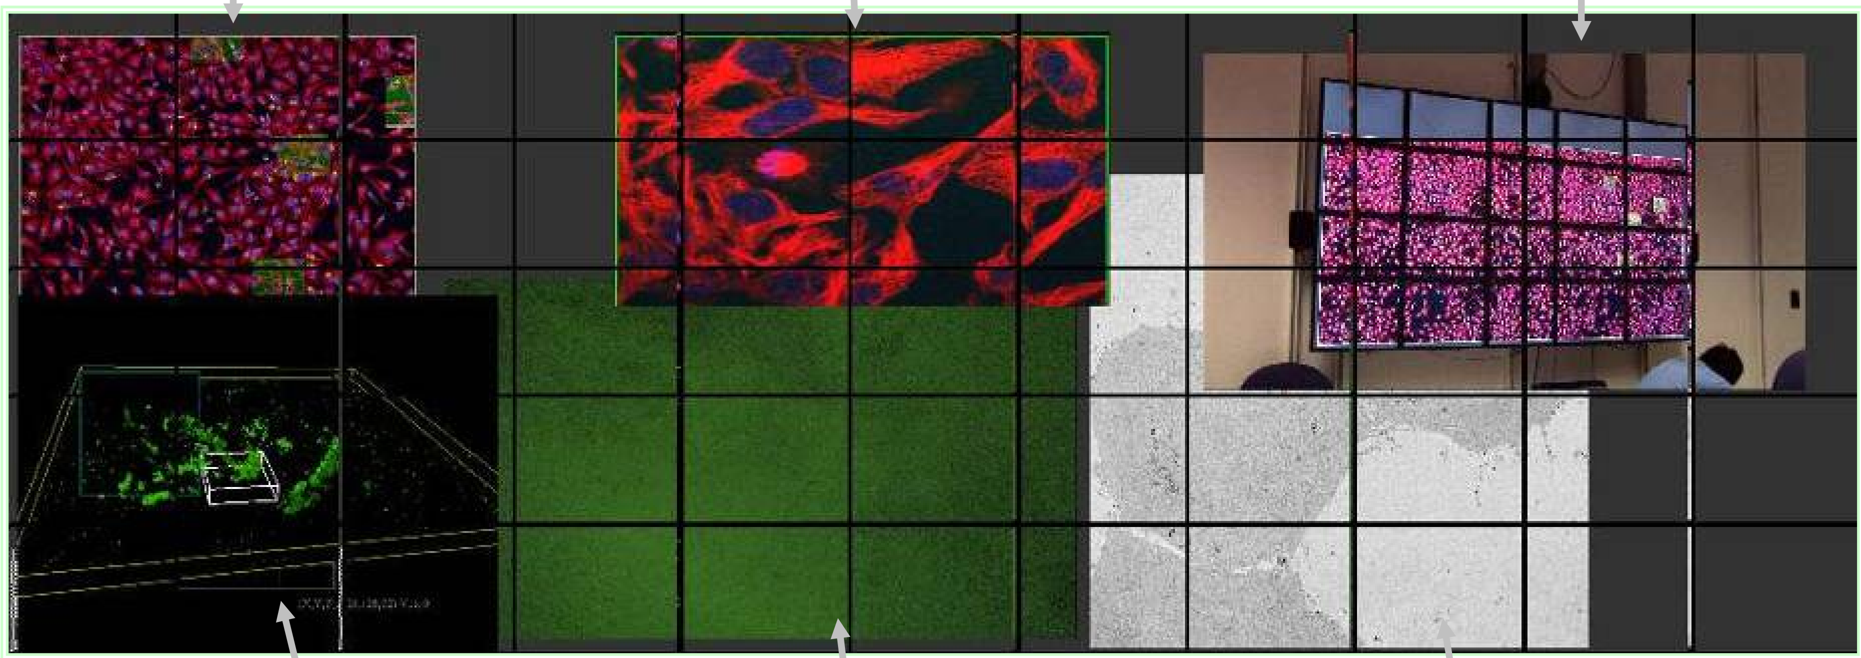

Volume rendering with Vol-a-Tile in Chicago

HDTV video stream from UHVEM in Osaka, Japan.

4K x 4K Digital images from NCMIR IVEM

Multi-scale Correlated Microscopy Experiment

Source: Steven T. Peltier

Active investigation of a biological specimen during UHVEM using multiple microscopies, data sources, and collaboration technologies

Collaboration Technologies and Remote Microscope Control

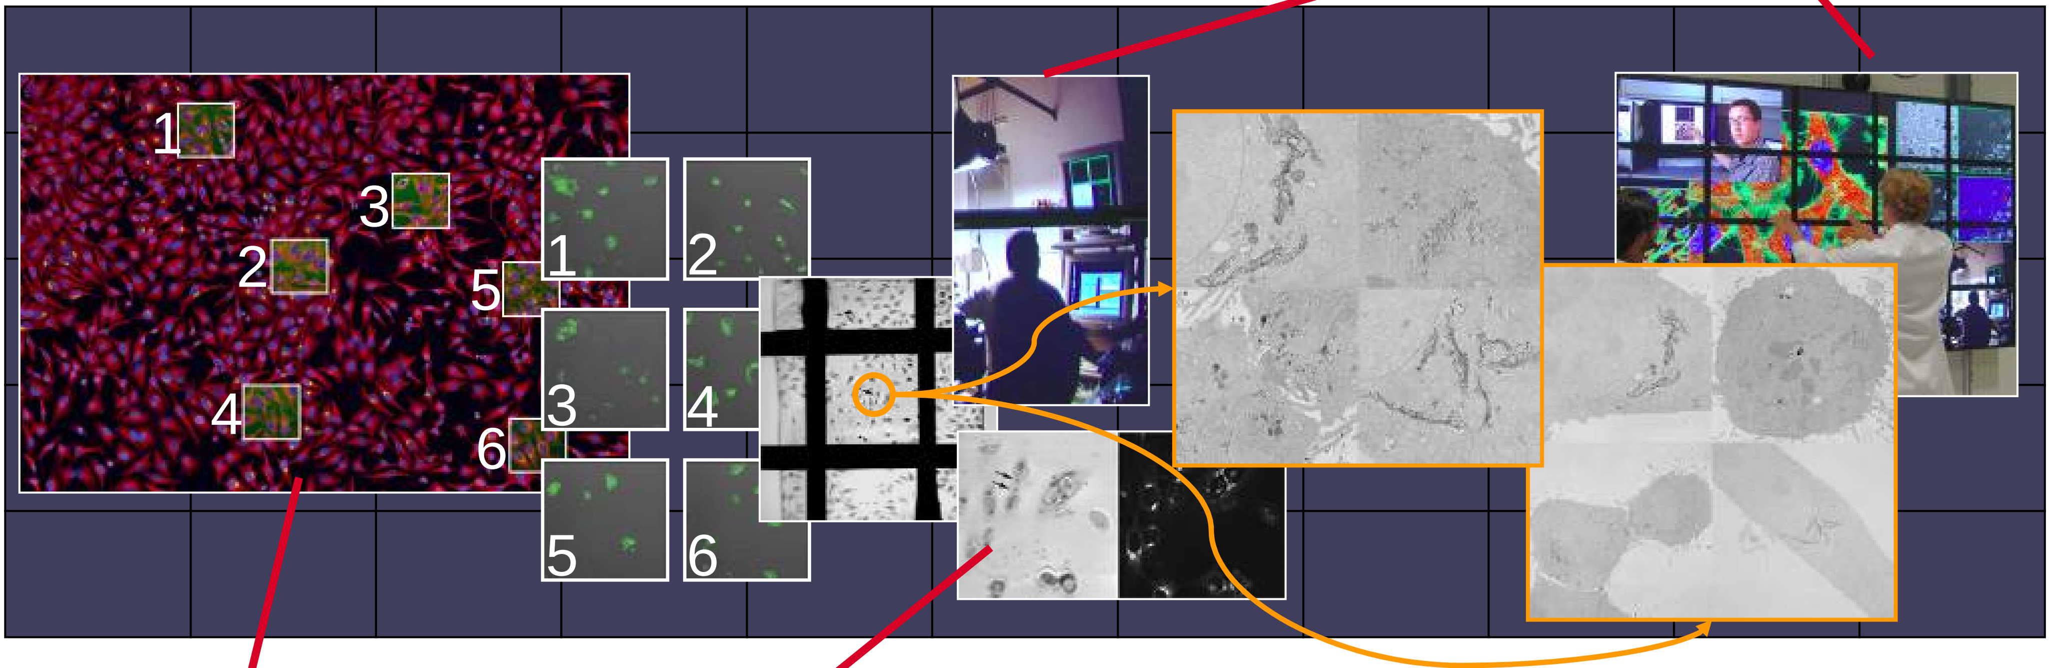

Light Microscopy Montage

Regions of Interest Time Lapse Movies

UHVEM HDTV Osaka, Japan

TRANSLIGHT℠ glif it² iGrid 2005 evl

iGrid 2005 Lambda Control Services Transform Batch Process to Real-Time Global e-VLBI

- Real-Time VLBI (Very Long Baseline Inferometry) Radio Telescope Data Correlation
- Radio Telescopes Collecting Data are Located Around the World
- Optical Connections Dynamically Managed Using the DRAGON Control Plane and Internet2 HOPI Network
- Achieved 512Mbps Transfers from USA and Sweden to MIT
- Results Streamed to iGrid2005 in San Diego
- Will be expanded to Japan, Australia, other European locations

Photonic Networks for Genomics

PI: Larry Smarr

TRANSLIGHT℠ glif iGrid 2005 evl

Marine Genome Sequencing Project
Measuring the Genetic Diversity of Ocean Microbes

CAMERA will include All Sorcerer II Metagenomic Data

TRANSLIGHT℠ glif it² J. Craig Venter INSTITUTE evl

Calit2 and the Venter Institute Combine Telepresence with Remote Interactive Analysis

Live Demonstration of 21st Century National-Scale Team Science

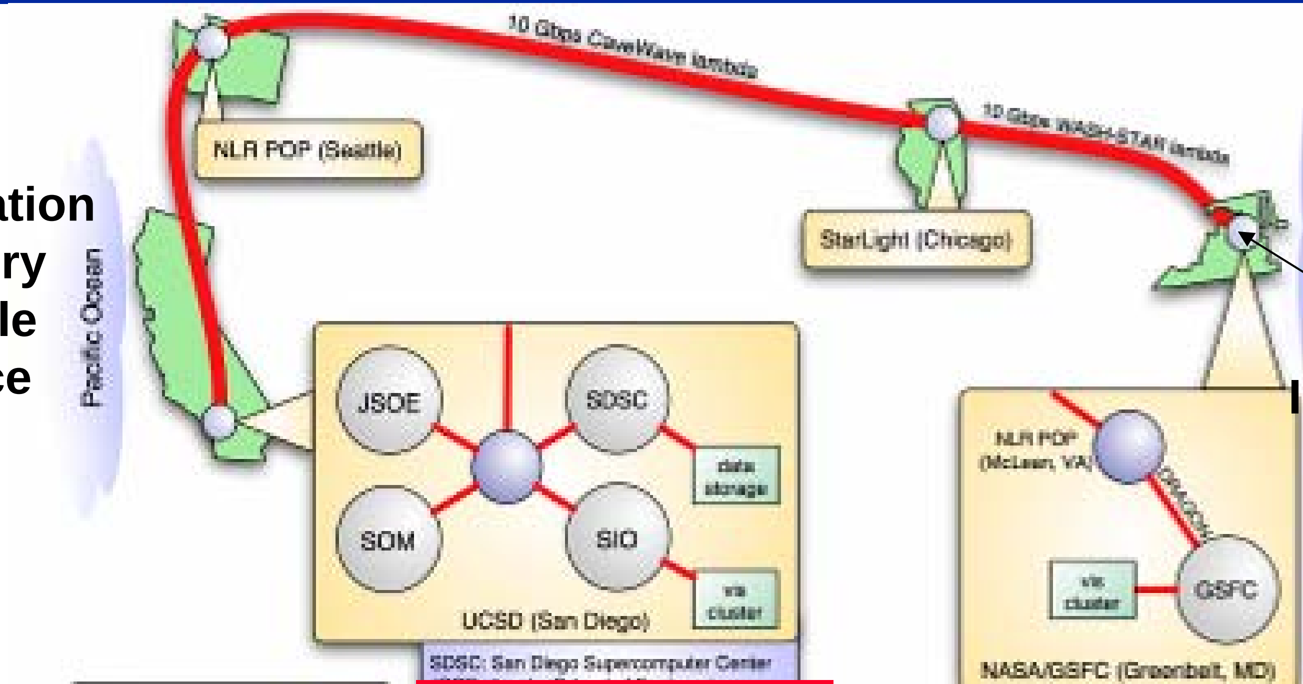

25 Miles

Venter Institute

OptIPuter Visualized Data

HDTV Over Lambda

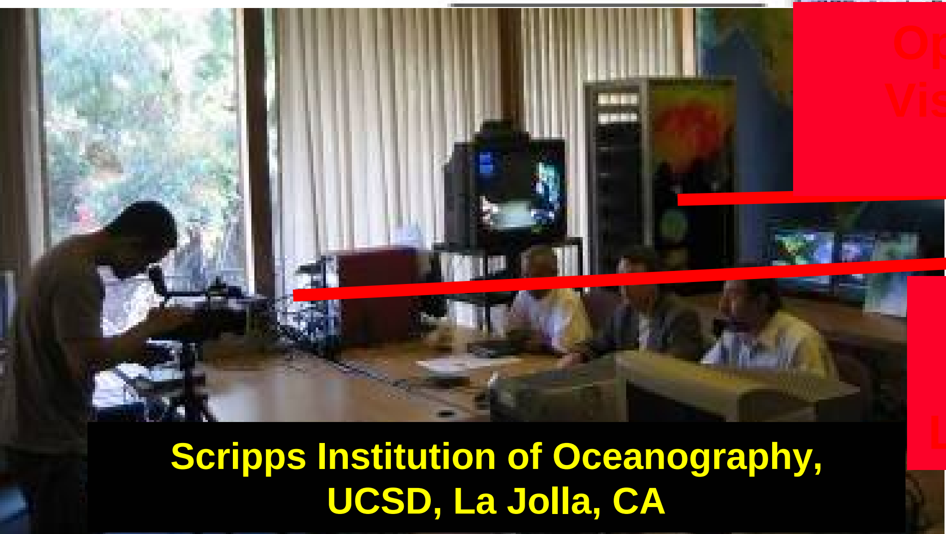

Scripps Institution of Oceanography, UCSD, La Jolla, CA

Goddard Space Flight Center, Maryland

CAMERA Metagenomics Server
Calit2's Direct Access Core Architecture

- Sargasso Sea Data
- Sorcerer II Expedition (GOS)
- JGI Community Sequencing Project
- Moore Marine Microbial Project
- NASA Goddard Satellite Data
- Community Microbial Metagenomics Data

Source: Phil Papadopoulos, SDSC, Calit2

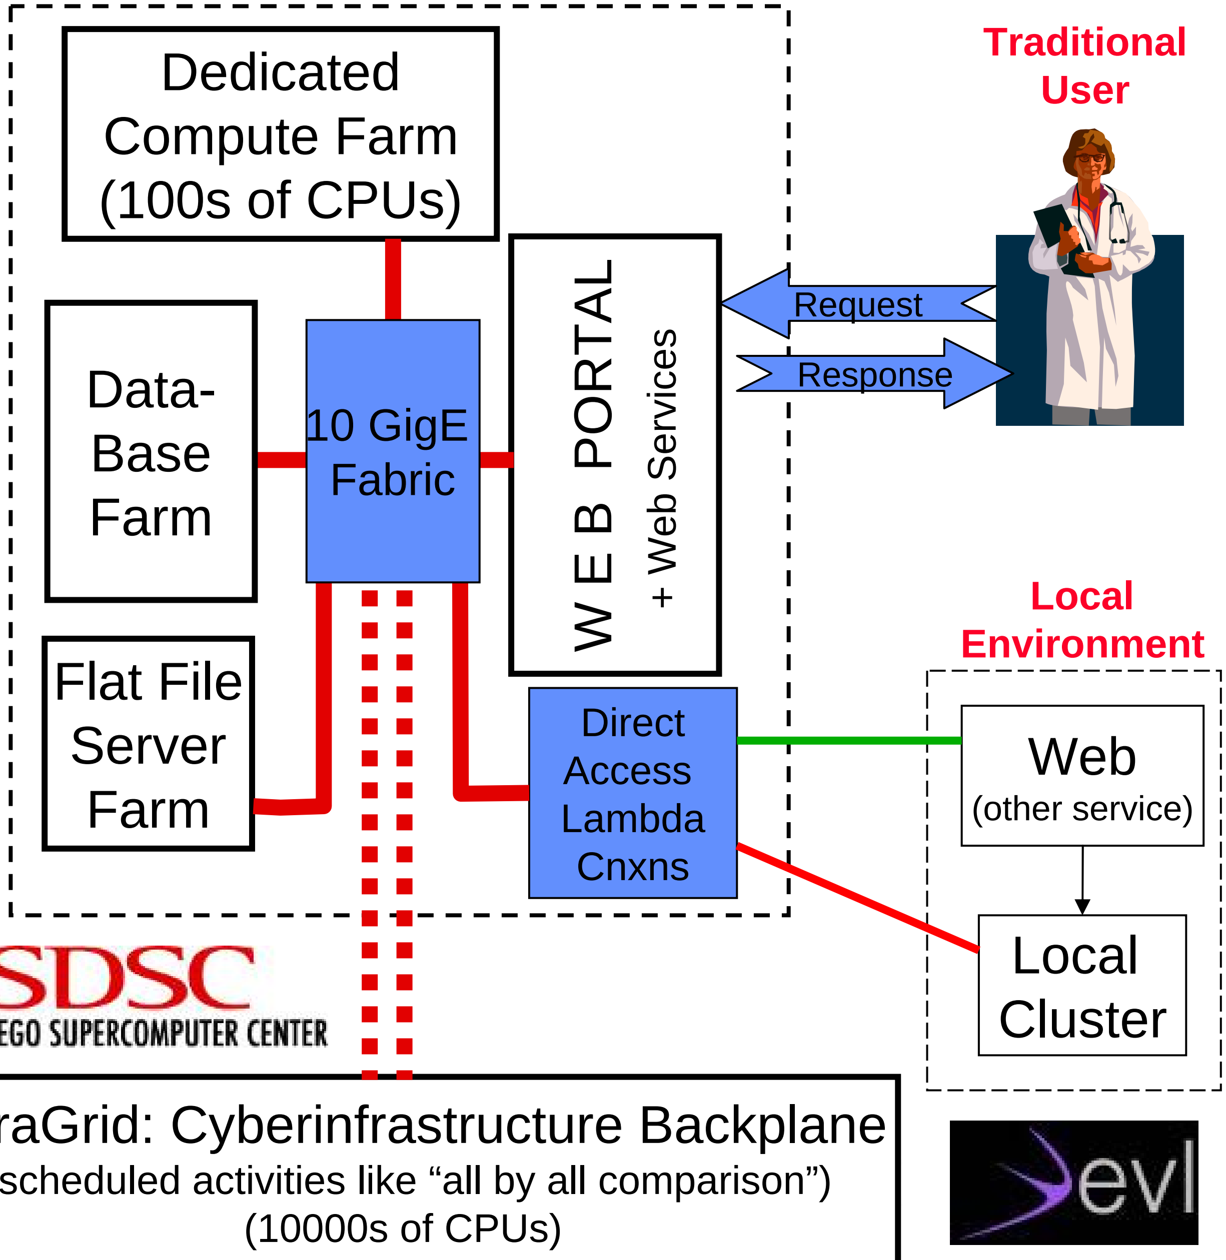

J. Craig Venter INSTITUTE

SDSC SAN DIEGO SUPERCOMPUTER CENTER

TRANSLIGHT℠

evl

Video over IP Experiments

- DV = 25Mbps as an I-frame codec with relatively low latency. WIDE has demoed this repeatedly, see www.sfc.wide.ad.jp/DVTS/
- HDV prosumer HD camcorders using either 18 or 25Mbps MPEG2 Long GOP. High latency if using native codec. However, its possible to use just the camera and do encoding externally to implement different bit rate (higher or lower) and different latency (lower or higher)
- WIDE did demos of uncompressed SD DTV at iGrid 2000 @ 270 Mbps over IPv6 from Osaka to Yokohama
- UW did multi-point HD teleconference over IP uncompressed at 1.5 Gbps at iGrid 2005 and SC05 http://www.researchchannel.org/news/press/sco5_demo.asp
- CalViz installed at Calit2 January 2006 uses HDV with MPEG2 at 25 Mbps for remote presentations at conferences
- NTT's iVISTO system capable of multi-stream HD over IP uncompressed at 1.5 Gbps with extremely low latency
- At iGrid 2005, demo by Keio, NTT Labs and UCSD in USA sent 4K over IP using JPEG 2000 at 400 Mbps, with back-channel of HDTV using MPEG2 I-frame at 50 mbps.
- Next challenge is bi-directional 4K and multi-point HD with low-latency compression.

TRANSLIGHT℠ glif it² iGrid 2005 evl

CalViz--25Mb/s HDV Streaming Internationally

Studio on 4th Floor of Calit2@UCSD Building
Two Talks to Australia in March 2006

Source: Harry Ammons

Calit2—UCSD Digital Cinema Theater

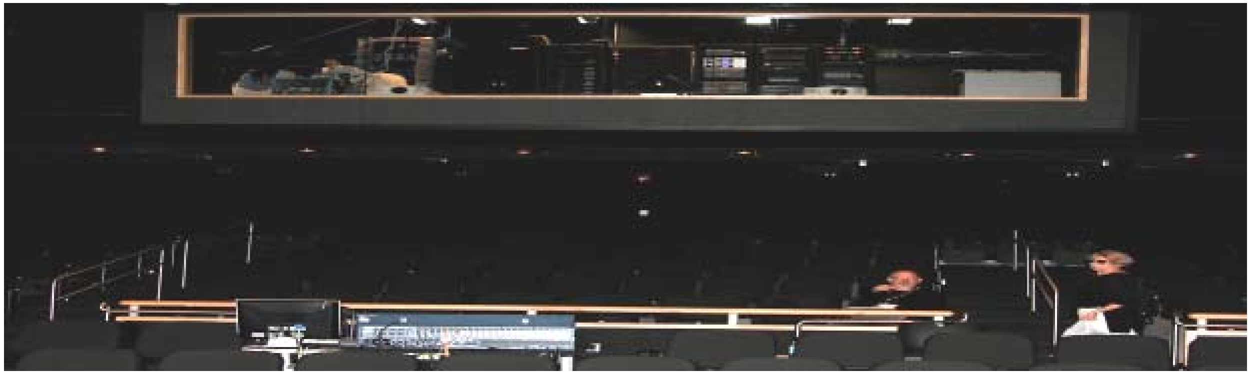

200 Seats, 8.2 Sound, Sony SRX-R110, SGI Prism w/21TB, 10GE to Computers/Data

CineGrid International Real-time Streaming 4K Digital Cinema at iGrid 2005

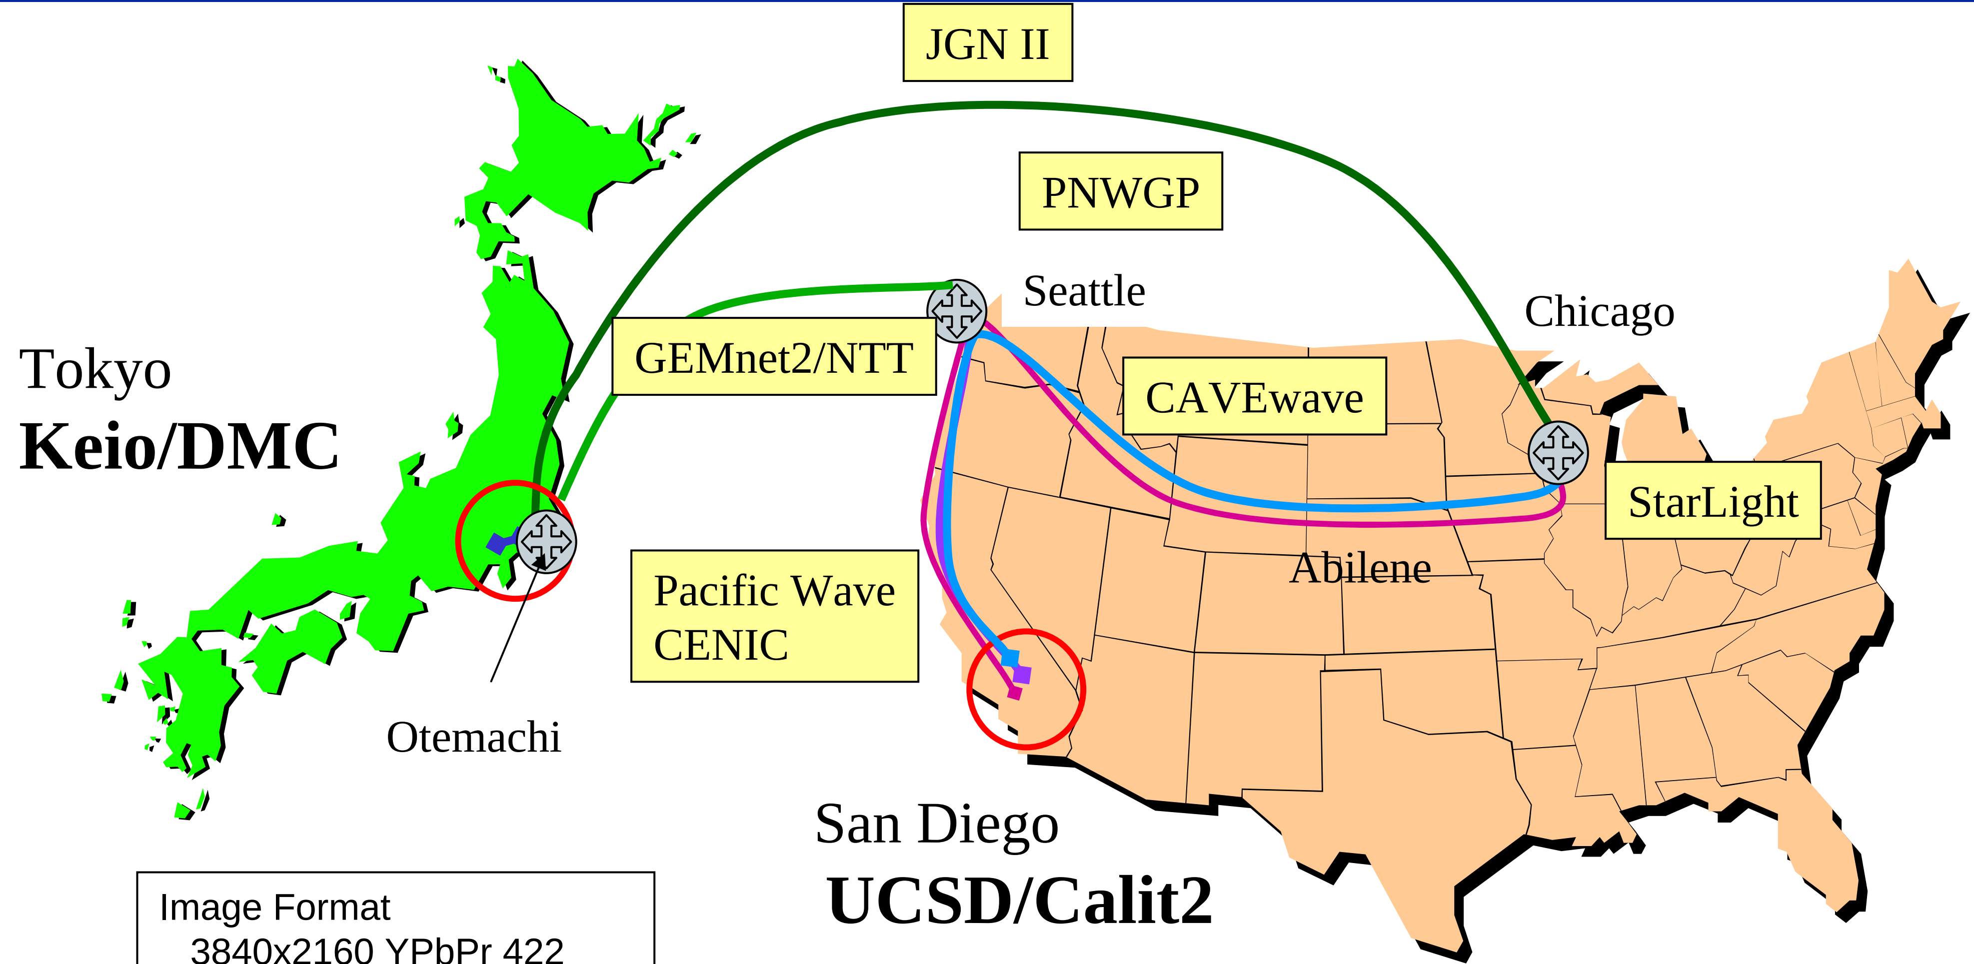

Image Format
 3840x2160 YPbPr 422
 24 or 29.97 frame/sec
Audio Format
 2ch or 5.1ch .WAV
 24 bit/48 KHz

iGrid 2005 International Real-Time Streaming 4K Digital Cinema ~500Mb/s

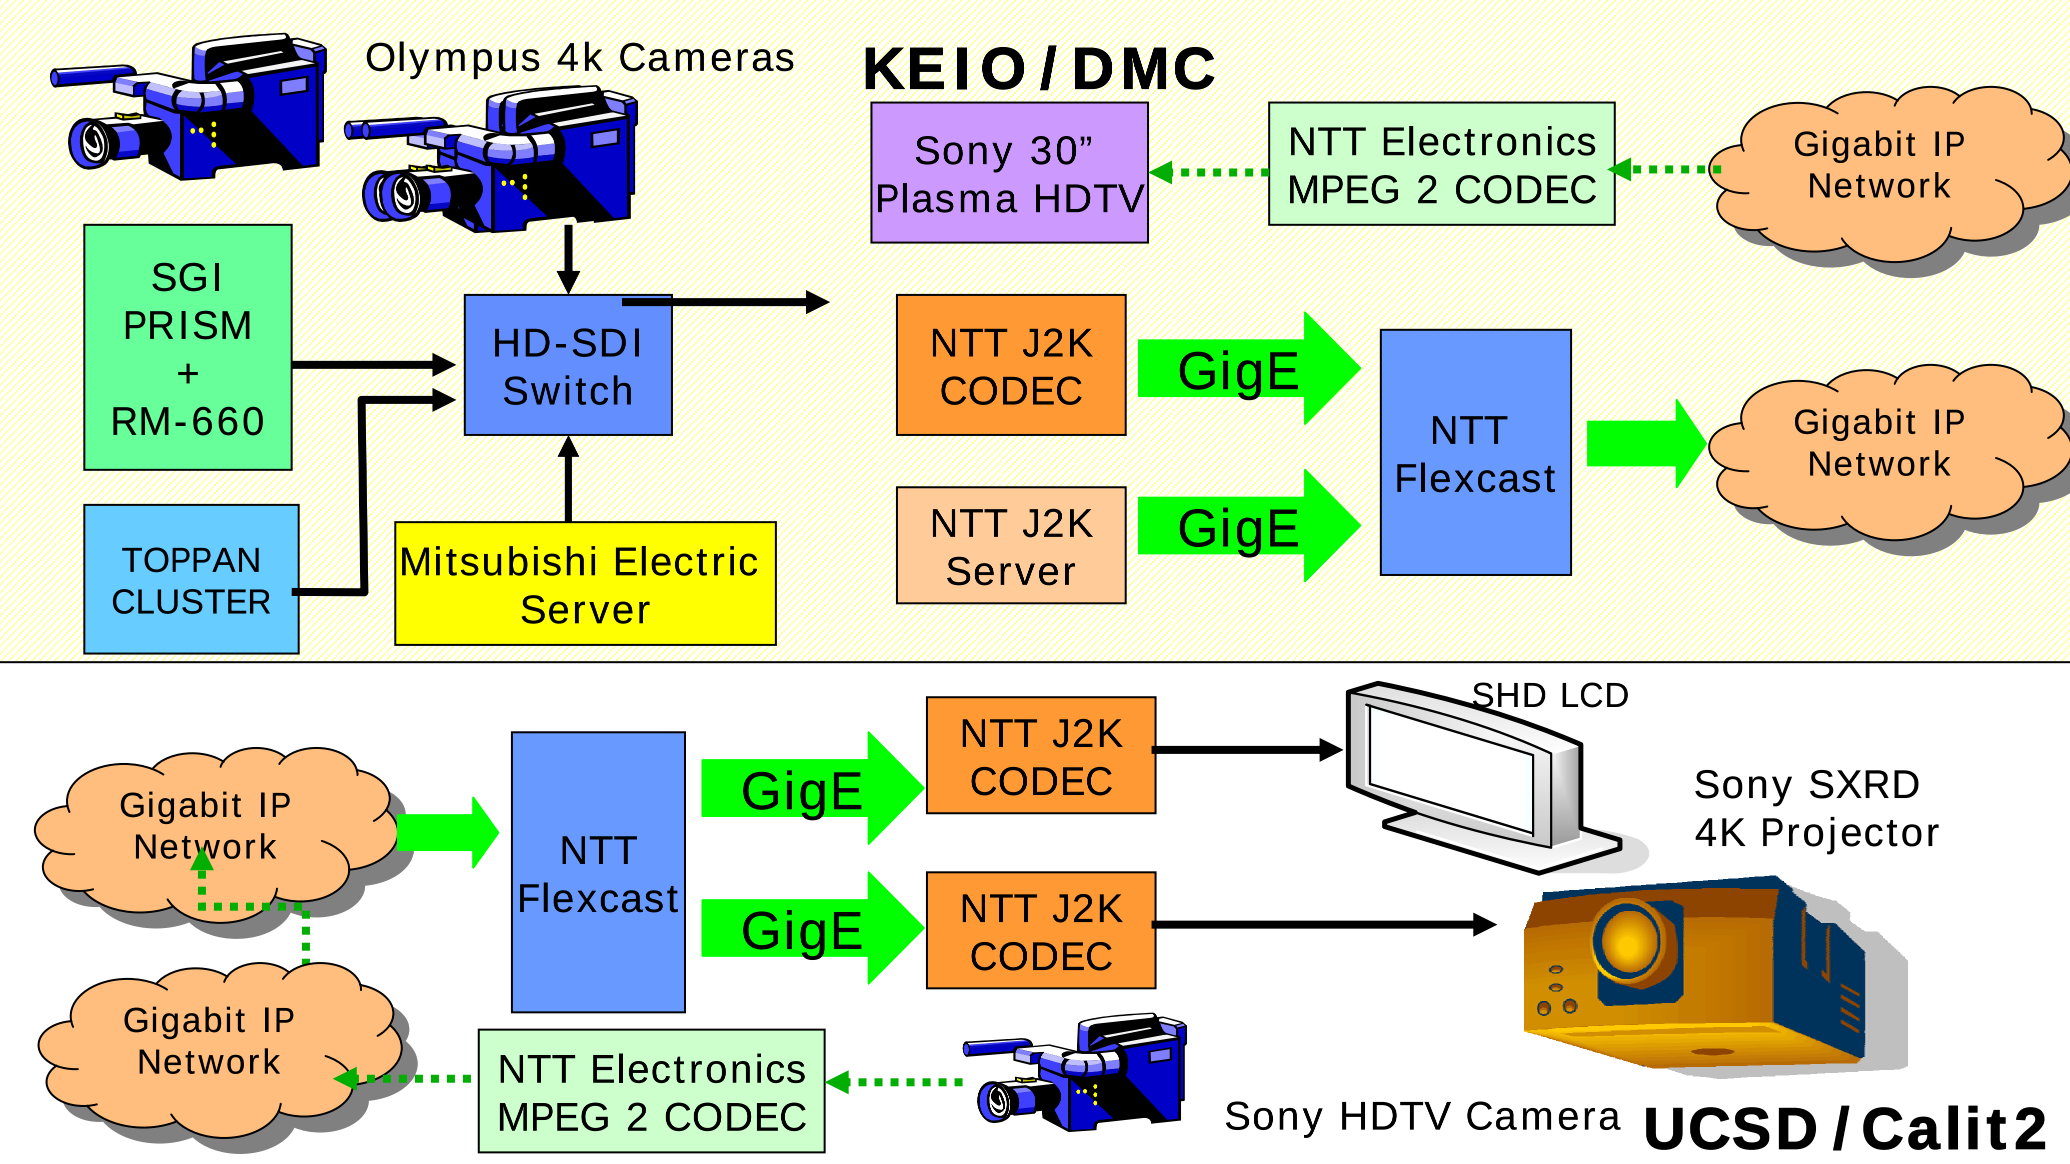

4K Telepresence over IP at iGrid 2005 Lays Technical Basis for Global Digital Cinema

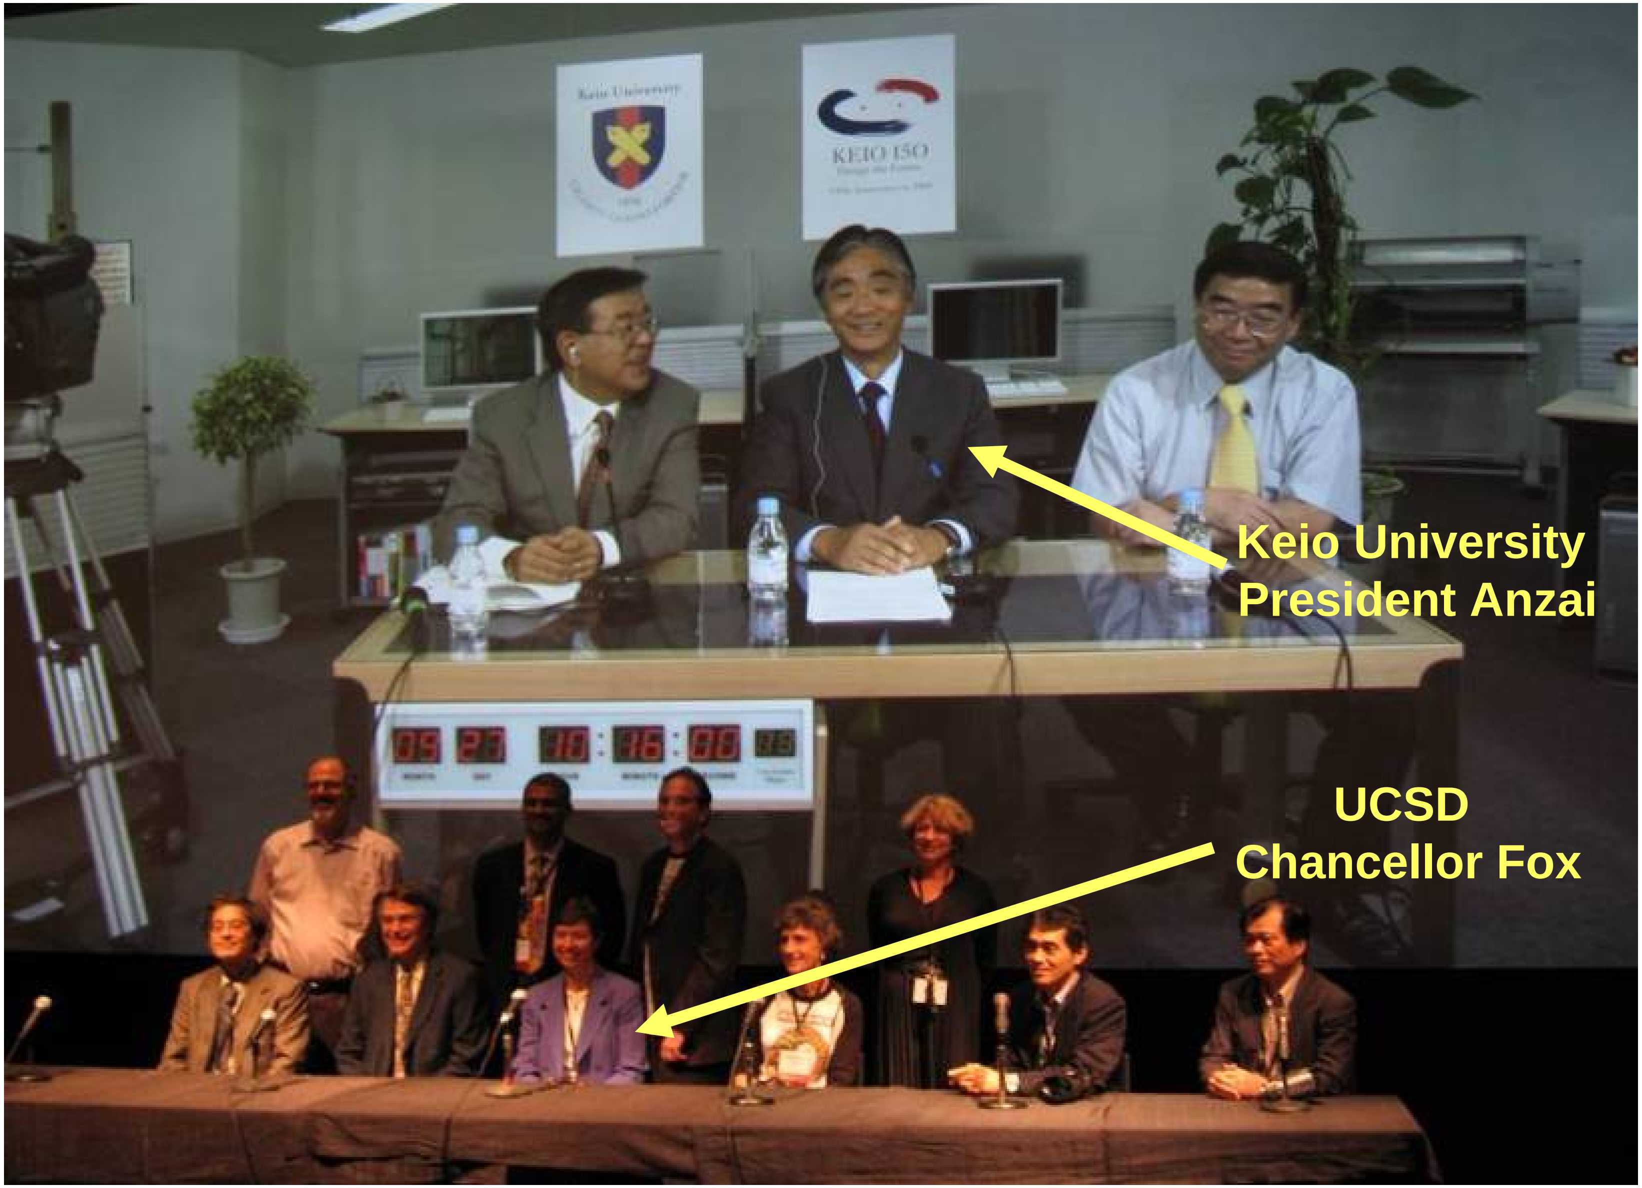

Calit2 is Partnering with CENIC to Connect Digital Media Researchers Into CineGrid

Partnering with SFSU's Institute for Next Generation Internet

CineGrid will Link UCSD/Calit2 and USC School of Cinema TV with Keio Research Institute for Digital Media and Content

In addition, 1Gb and 10Gb Connections to:

- Seattle, Asia, Australia, New Zealand
- Chicago, Europe, Russia, China
- Tijuana

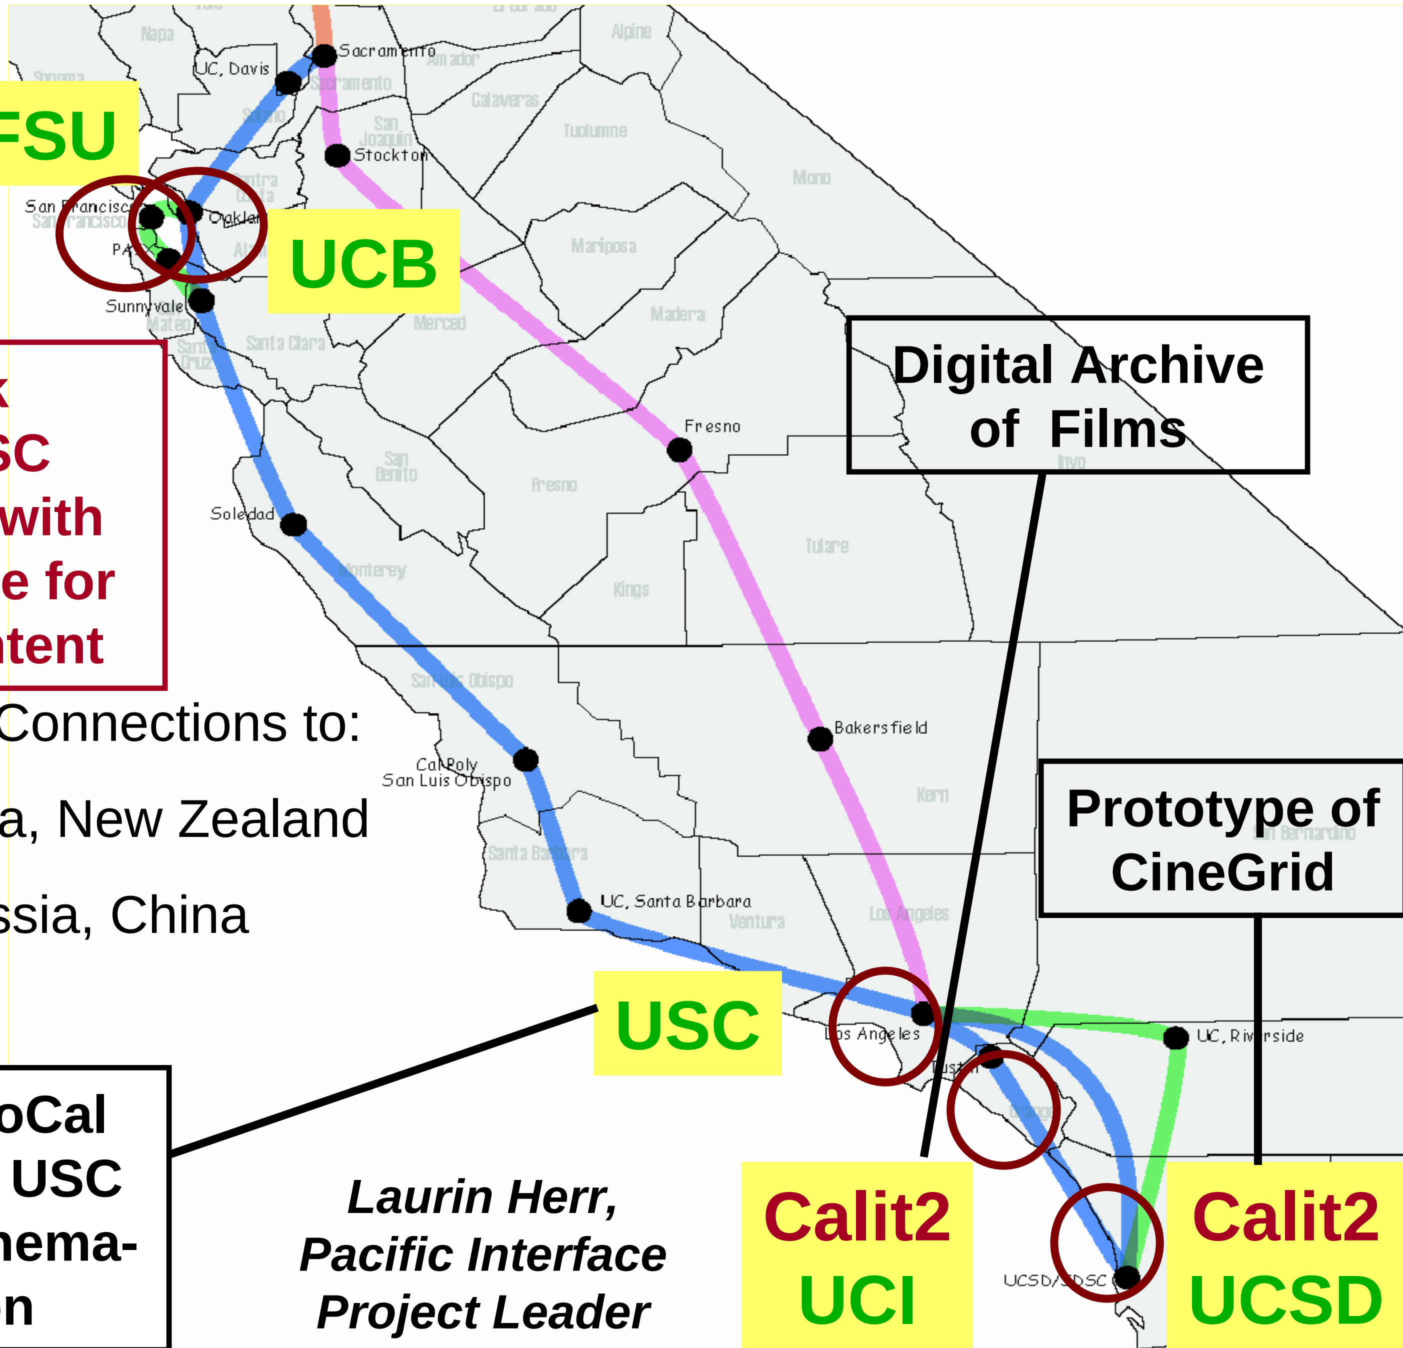

Extended SoCal OptIPuter to USC School of Cinema-Television

Laurin Herr, Pacific Interface Project Leader

GLIF = Global Lambda Integrated Facility
www.glif.is

- **A worldwide laboratory for application and middleware development**
- **Networks of interconnected optical wavelengths (also known as lambda grids).**
- **Takes advantage of the cost and capacity advantages offered by optical multiplexing**
- **Supports powerful distributed systems that utilize processing power, storage, and instrumentation at various sites around the globe.**
- **Aim is to encourage the shared used of resources by eliminating a lack of network capacity as the traditional performance bottleneck**

TRANSLIGHT[SM] glif it[2] iGrid 2005 evl

GLIF—the Global Lambda Integrated Facility

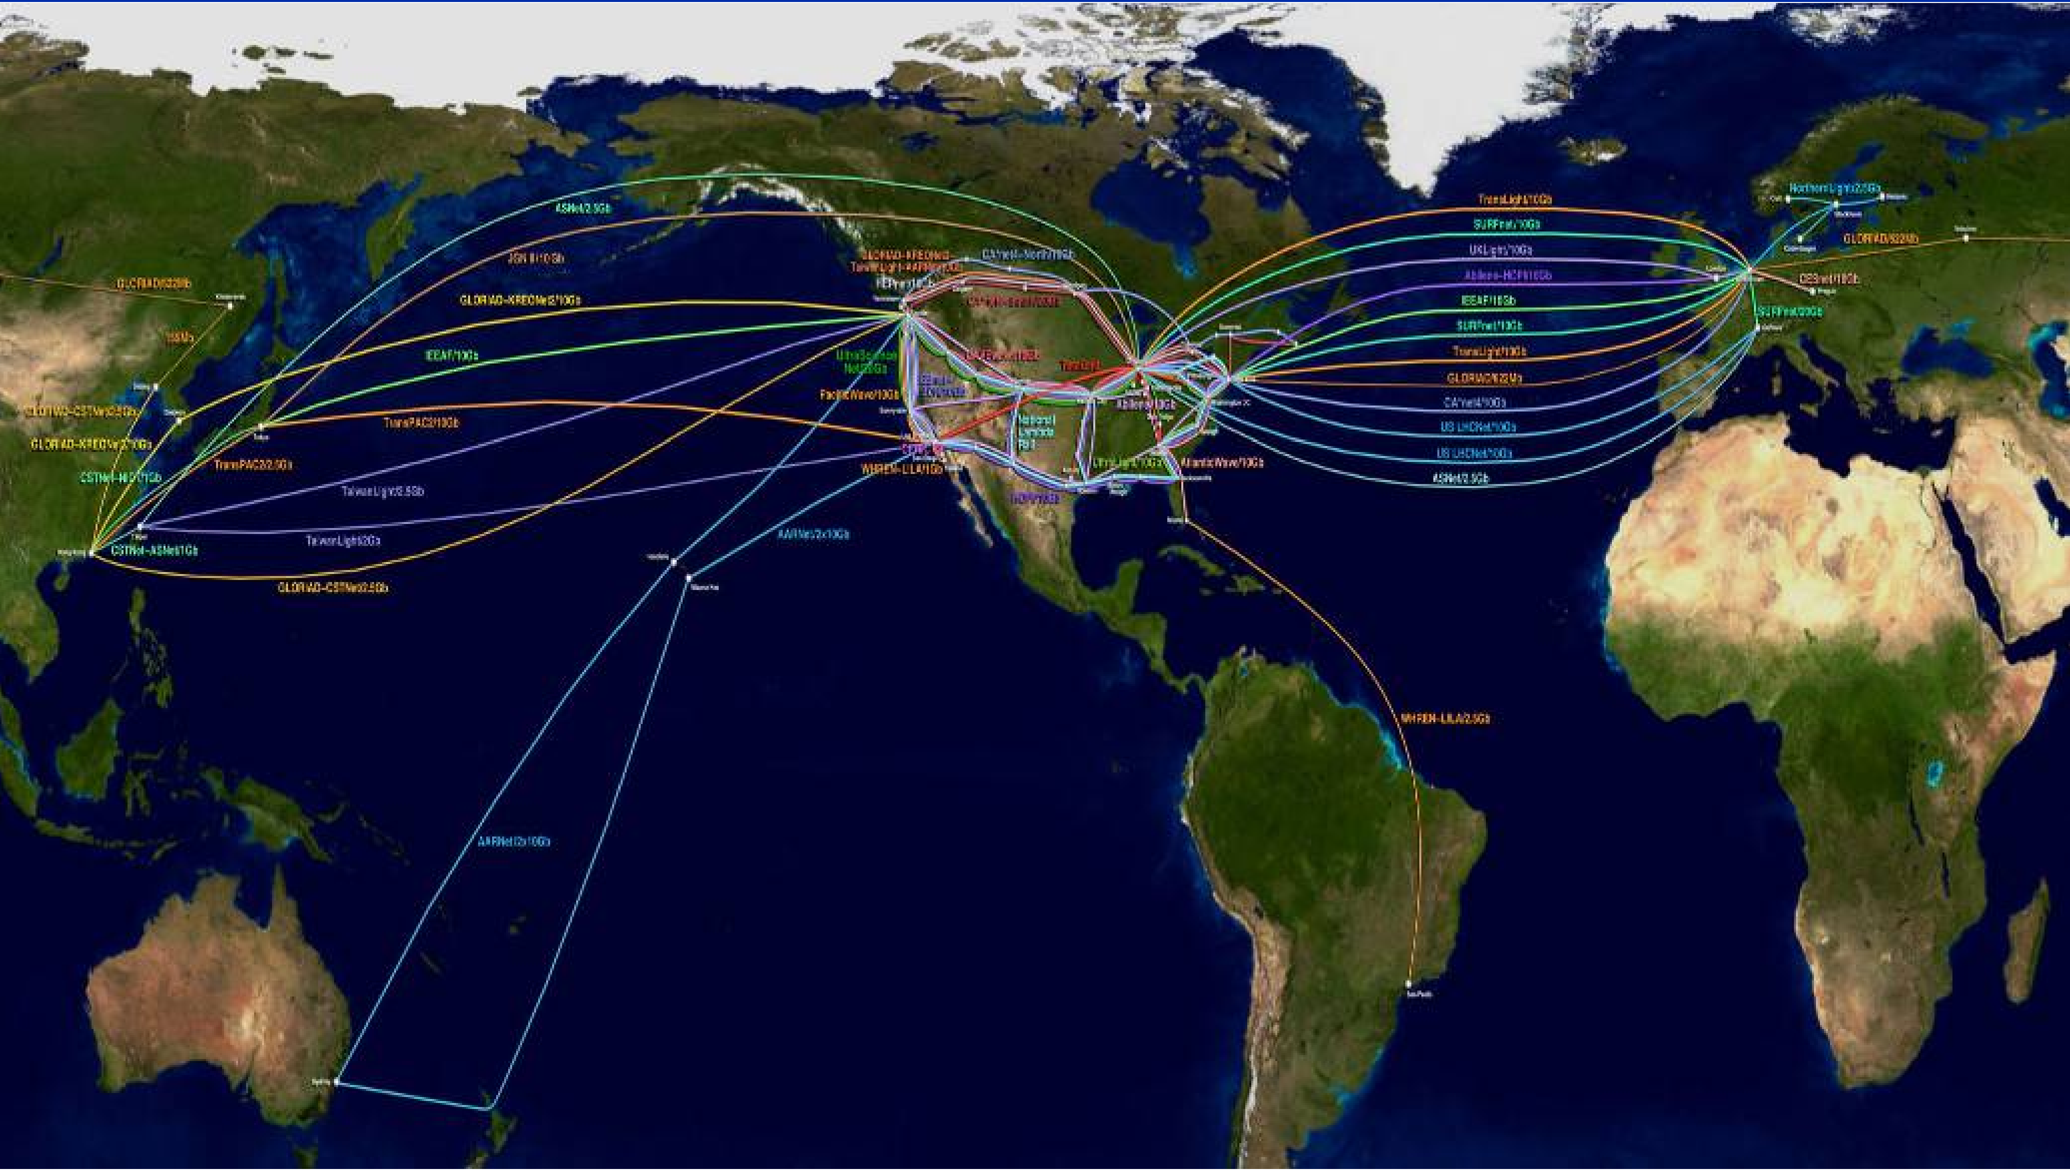

GLIF Uses Lambdas

- Lambdas are dedicated high-capacity circuits over optical wavelengths
- A *lightpath* is a communications channel (virtual circuit) established over lambdas, that connects two end-points in the network.
- Lightpaths can take-up some or all of the capacity of individual GLIF lambdas, or indeed can be concatenated across several lambdas.
- Lightpaths can be established using different protocol mechanisms, depending on the application.
  - Layer 1
  - Layer 2
  - Layer 3
  - Many in GLIF community are finding advantage to implement a lightpath as a 1 or 10 Gigabit Ethernet, so the virtual circuit acts as a virtual local area network, or VLAN.
- GLIF relies on a number of lambdas contributed by the GLIF participants who own or lease them

GLIF Participants

- The GLIF participants are organizations that
  - share the vision of optical interconnection of different facilities
  - voluntarily contribute network resources (equipment and/or lambdas)
  - and/or actively participate in activities in furtherance of these goals
- Seamless end-to-end connections require a high degree of interoperability between different transmission, interface and service implementations, and also require harmonization of contracting and fault management processes
- The GLIF Technical and Control Plane Working Groups are technical forums for addressing these operational issues
- The network resources that make-up GLIF are provided by independent network operators who collaborate to provide end-to-end lightpaths across their respective optical domains
- GLIF does not provide any network services itself, so research users need to approach an appropriate GLIF network resource provider to obtain lightpath services
- GLIF participants meet at least once per year
  - 2003 - Reykjavik, Iceland
  - 2004 - Nottingham, UK
  - 2005 - San Diego, US
  - 2006 - Tokyo, Japan

TRANSLIGHT[SM] glif it² iGrid 2005 evl

GOLE = Global Open Lambda Exchange

- **GLIF is interconnected through a series of exchange points known as GOLEs (pronounced "goals"). GOLE is short for "Global Open Lambda Exchange"**
- **GOLEs are usually operated by GLIF participants, and are comprised of equipment that is capable of terminating lambdas and performing lightpath switching.**
- **At GOLEs, different lambdas can be connected together, and end-to-end lightpaths established over them.**
- **Normally GOLEs must interconnect at least two autonomous optical domains in order to be designated as such.**

GOLEs and Lambdas
www.glif.is/resources/

- CANARIE-StarLight - Chicago
- CANARIE-PNWGP - Seattle
- CERN - Geneva
- KRLight - Seoul
- MAN LAN - New York
- NetherLight - Amsterdam
- NorthernLight - Stockholm
- Pacific Northwest GigaPoP - Seattle
- StarLight - Chicago
- T-LEX - Tokyo
- UKLight - London
- UltraLight - Los Angeles

Linked GOLEs For GLIF - October 2005

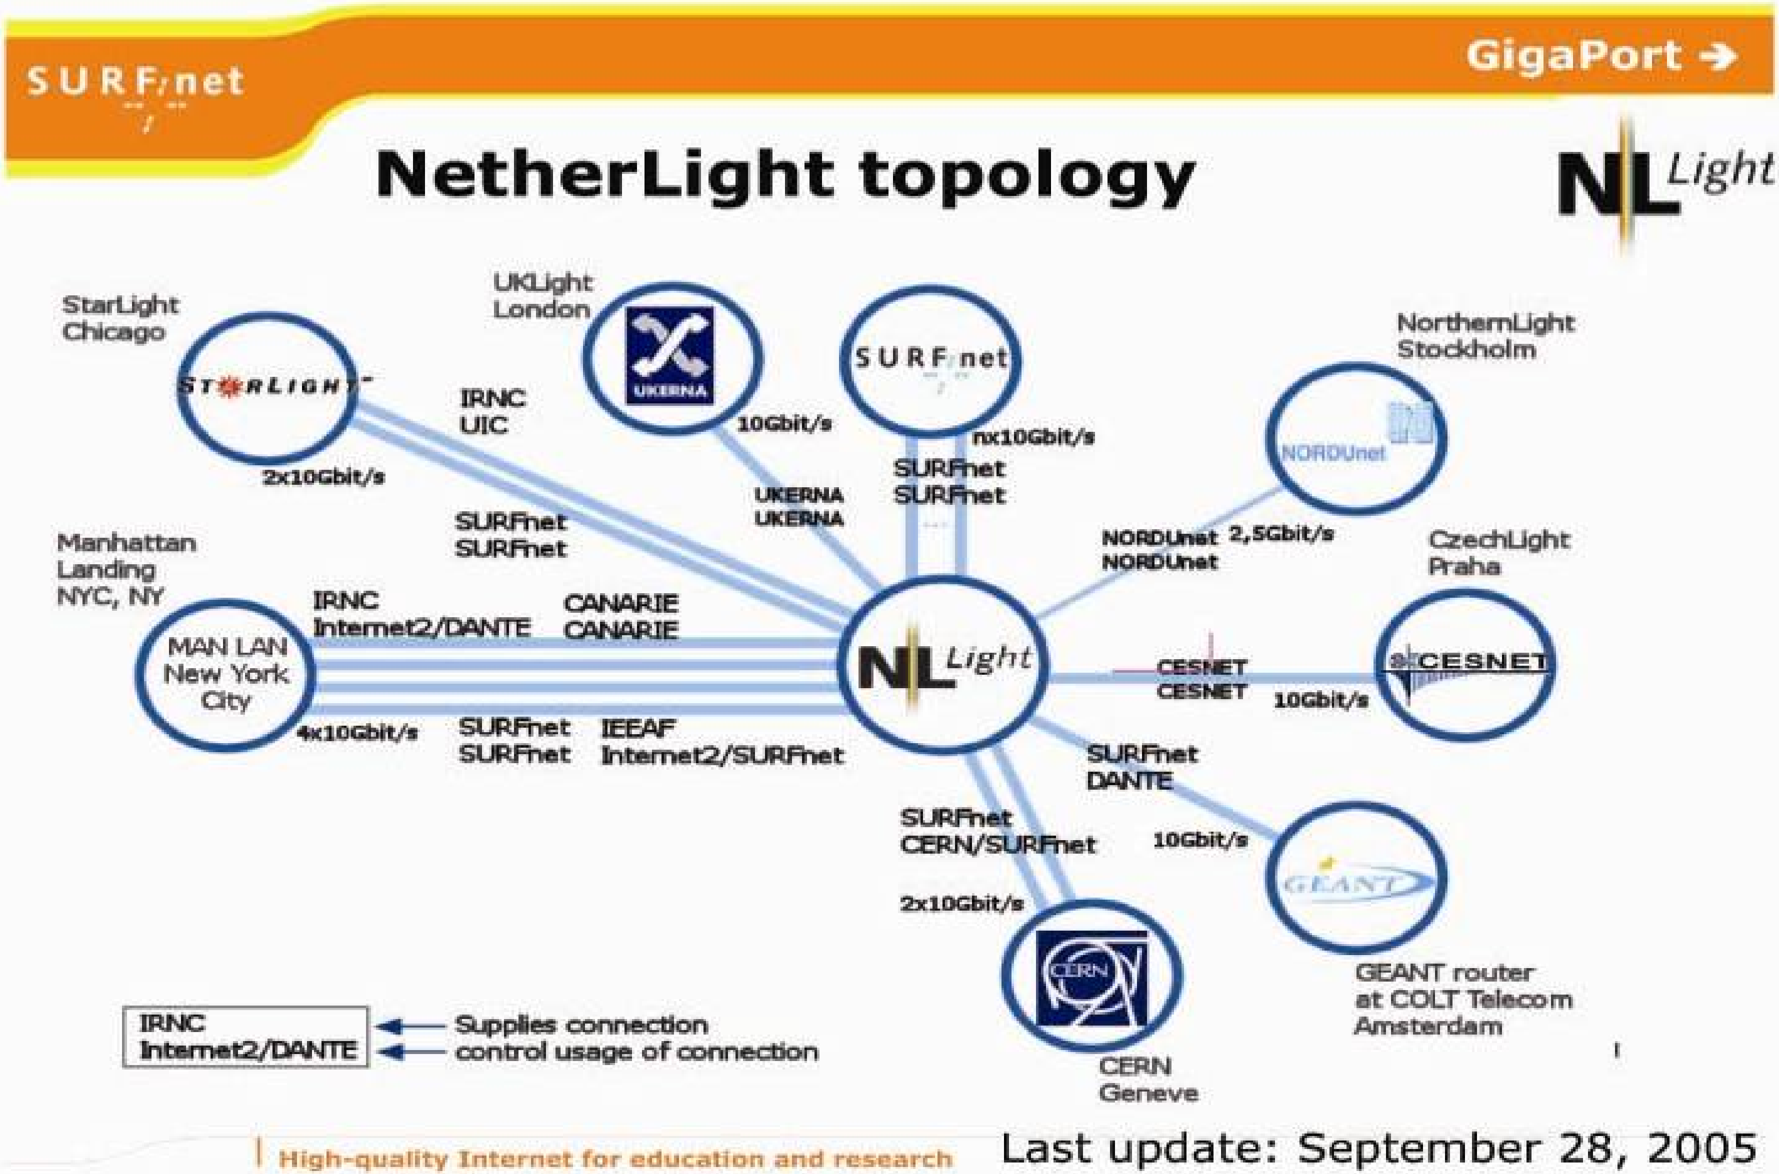

Marijke.Kaat
@SURFnet.nl
24 Oct 2005

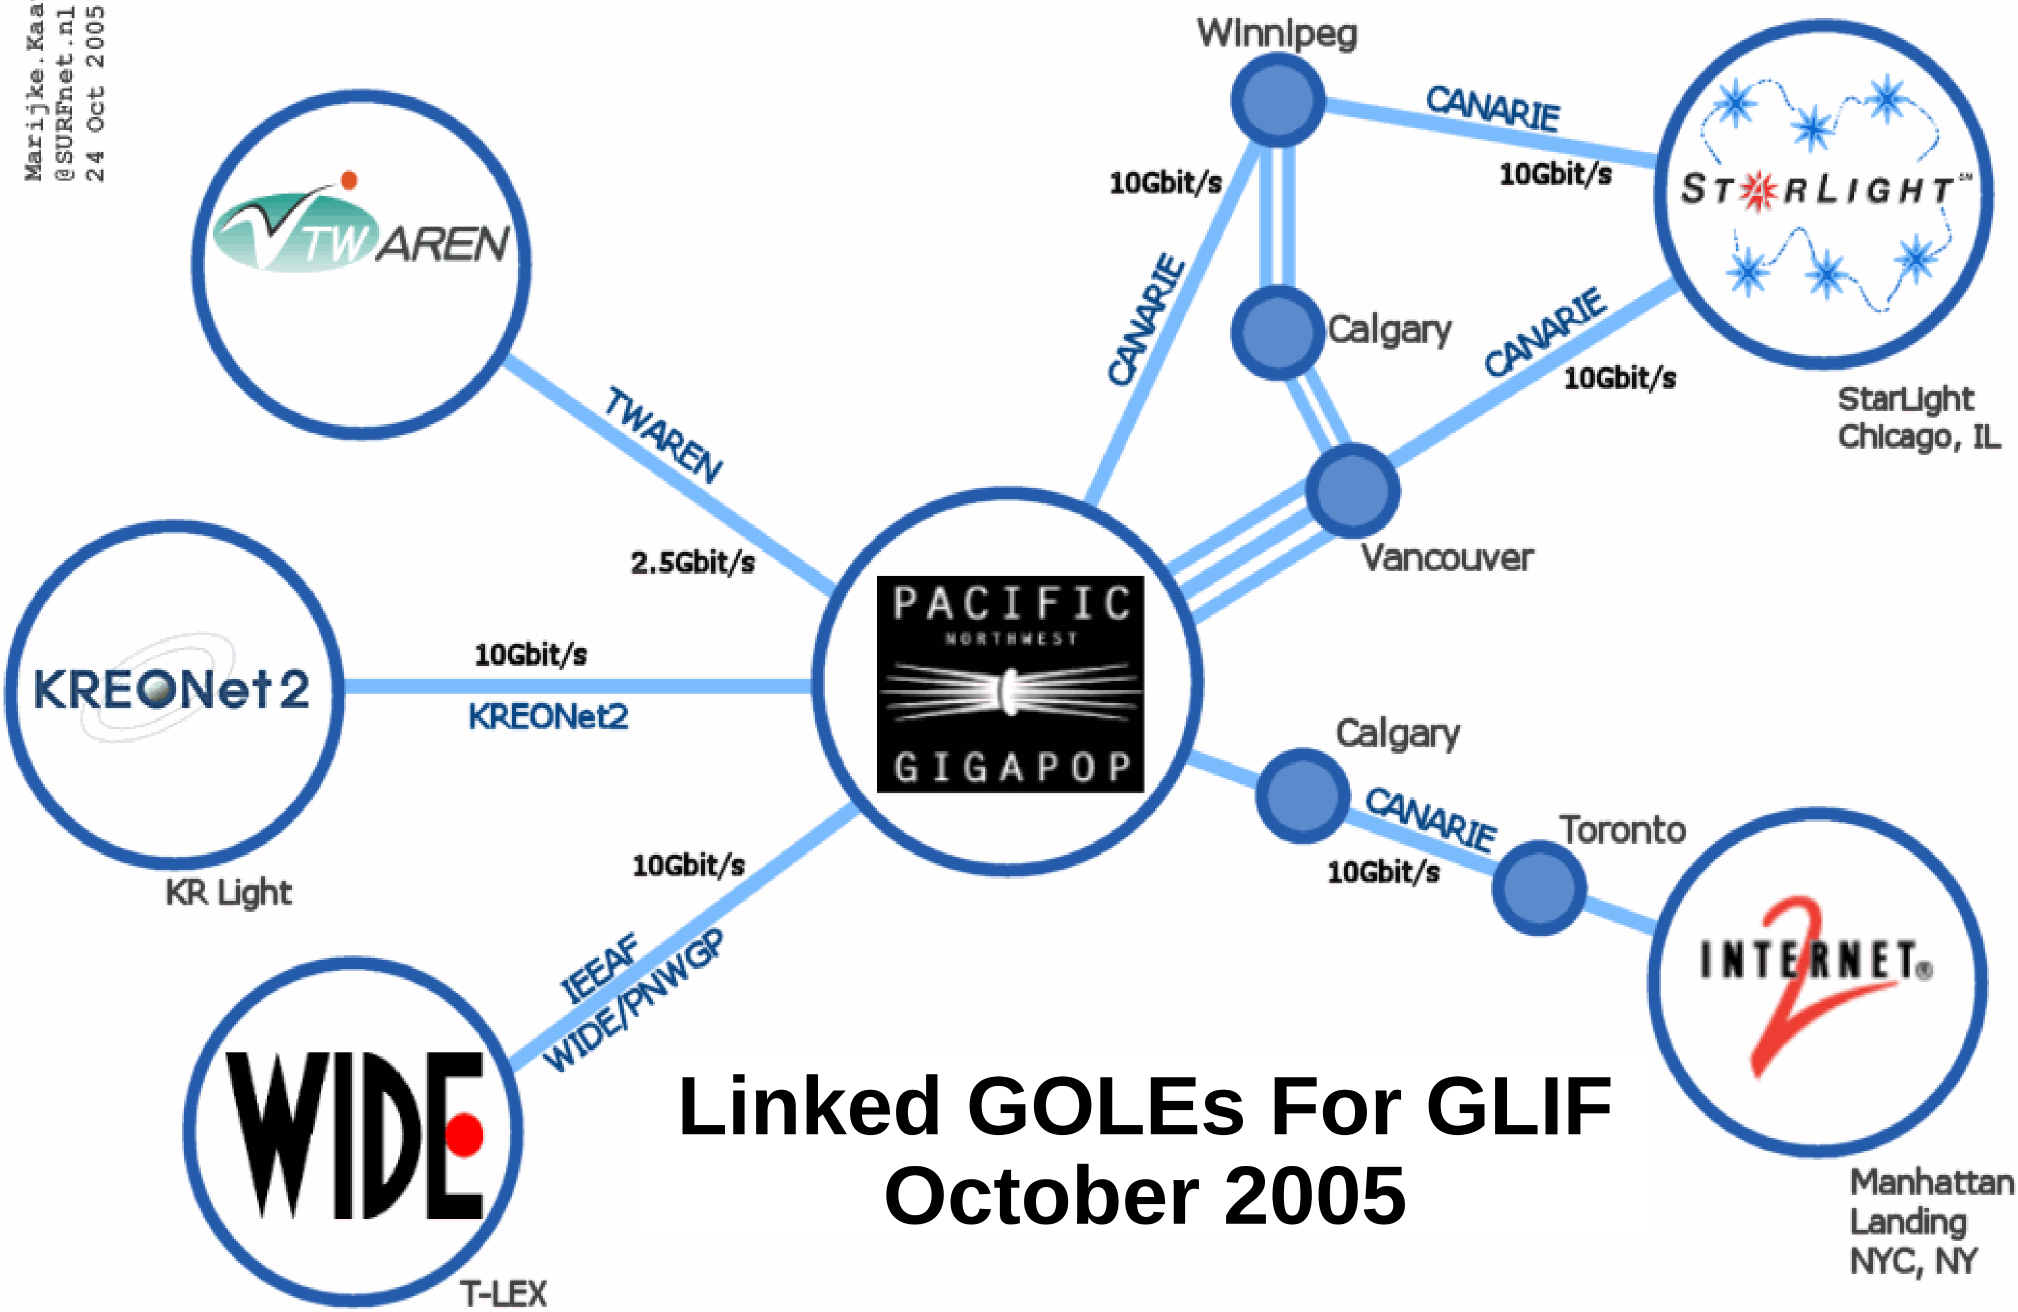

Linked GOLEs For GLIF
October 2005

TRANSLIGHT℠ glif it2 iGrid 2005 evl

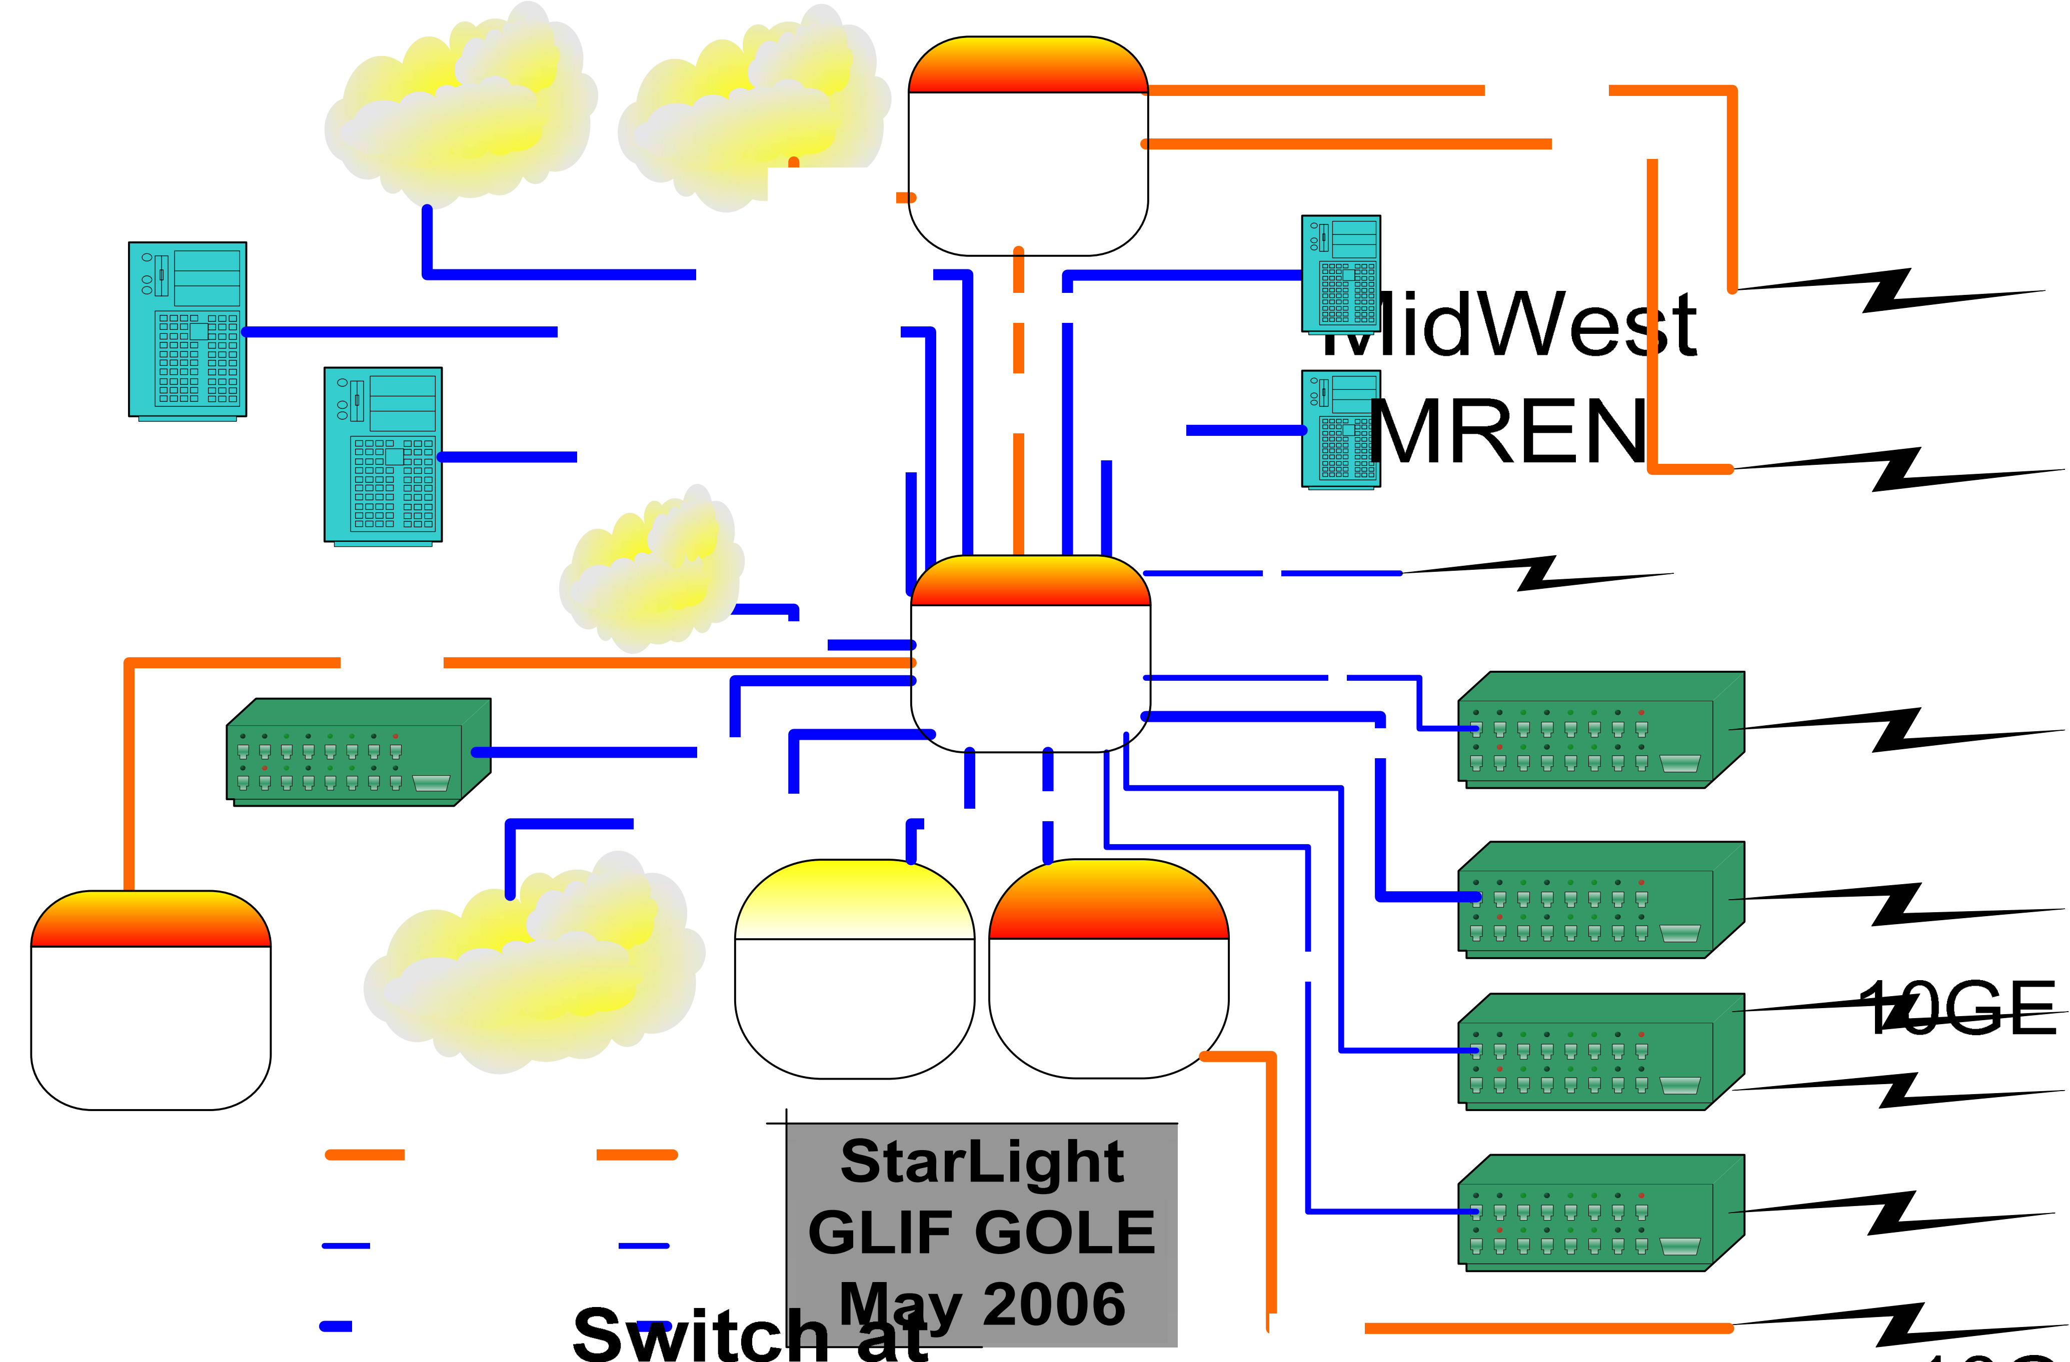

TRANSLIGHT℠

Linked GOLEs For GLIF

evl

Linked GOLEs For GLIF - October 2005

T-LEX in Tokyo, JAPAN

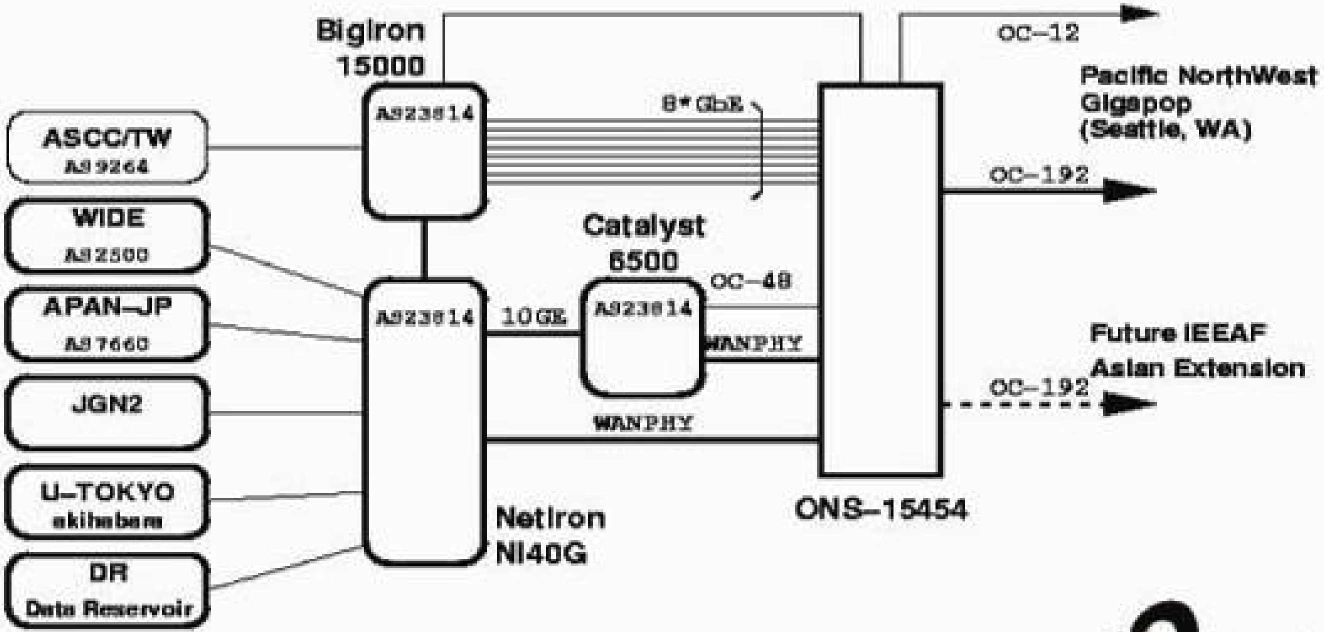

WIDE

TλEx
Tokyo Lambda EXchange

TRANSLIGHT[SM] glif it2 iGrid 2005 evl

Linked GOLEs For GLIF - October 2005

KRLight and GLORIAD-KR

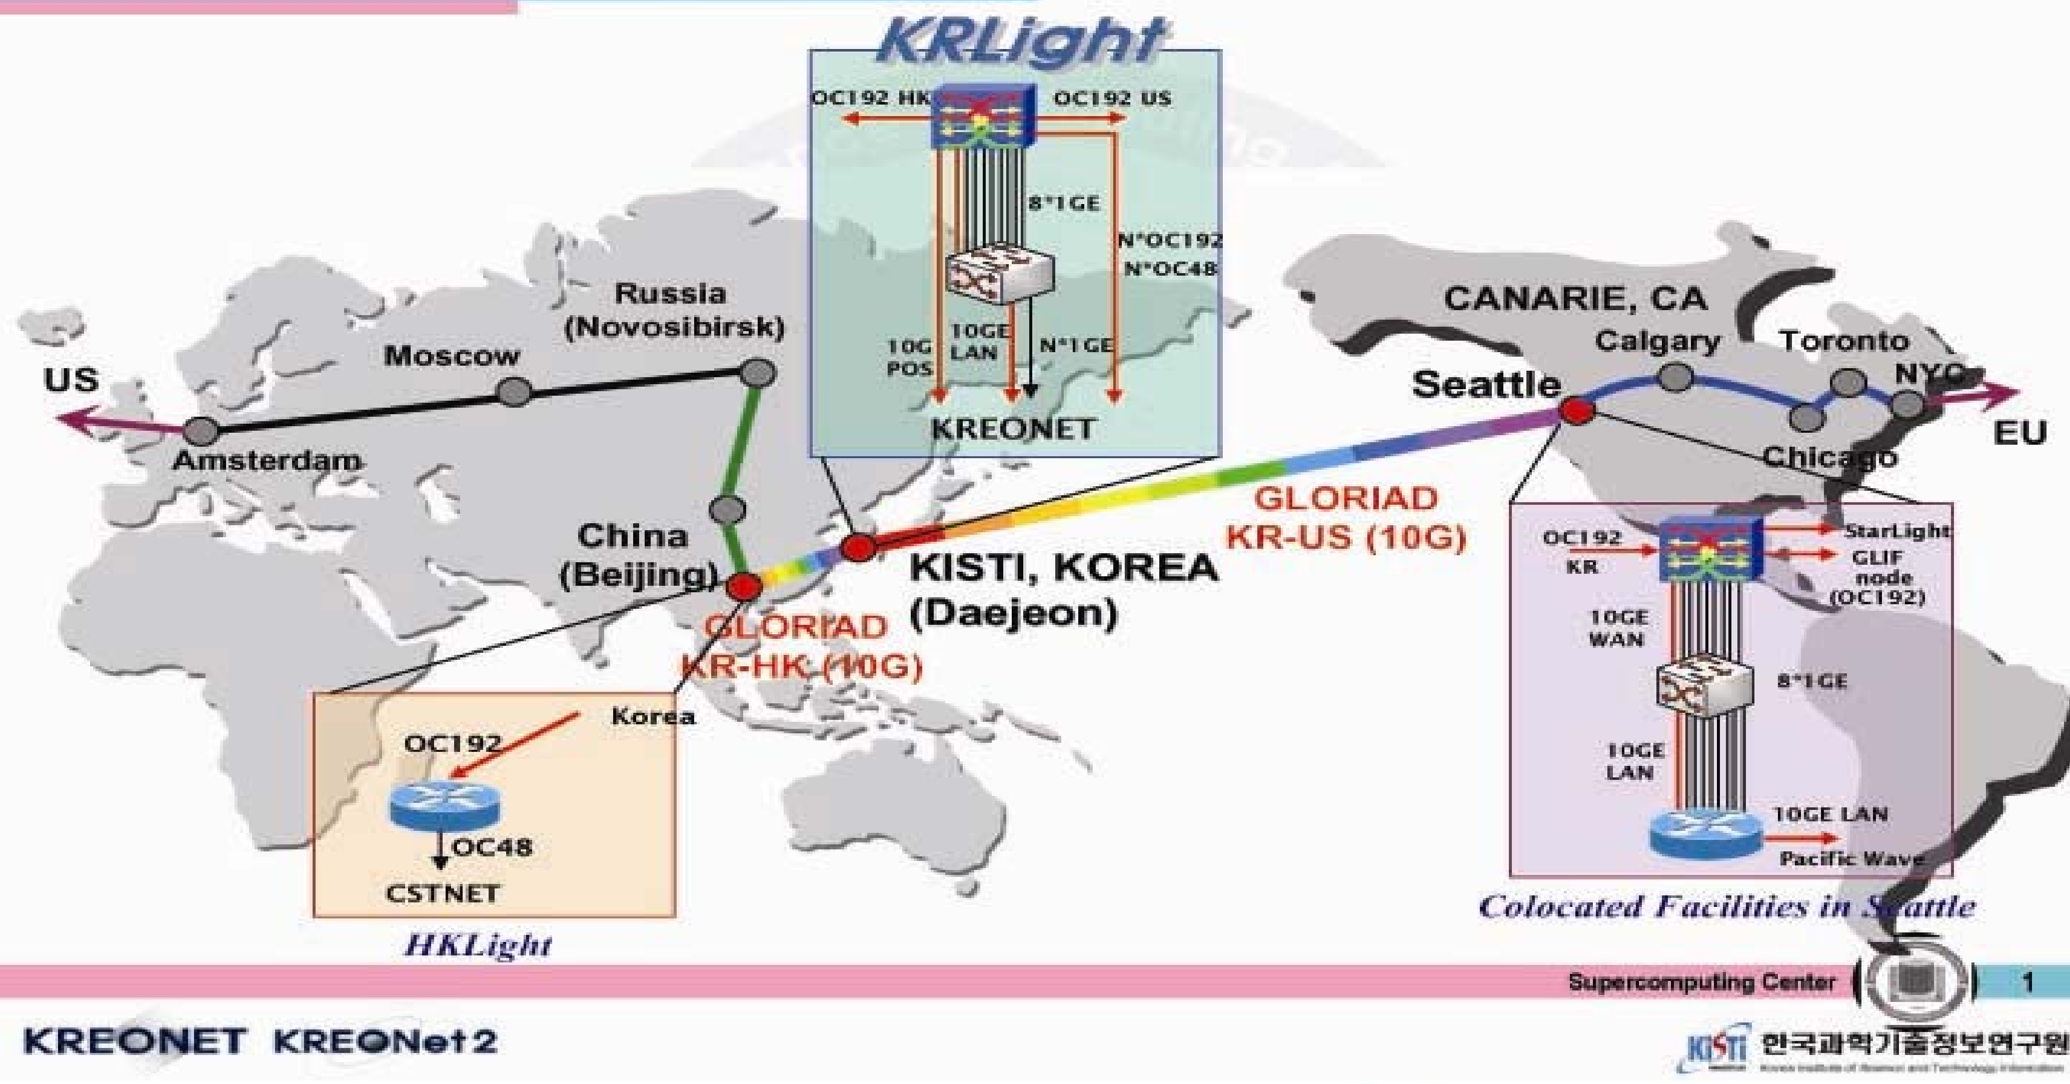

TRANSLIGHT℠ glif it² iGrid 2005 evl

Conclusion - GLIF and GOLE for 21st Century

- Applications need deterministic networks:
  - Known and knowable bandwidth
  - Known and knowable latency
  - Scheduling of entire 10G lighpaths when necessary
- iGrid2005 proved that the technologies for GLIF work (with great effort)
- GLIF partner activities are training the next generation of network engineers
- GLIF partners are building new GOLEs
- GLIF researchers are now implementing automation (e.g., UCLP)
- Scalability at every layer remains the challenge!

Special iGrid 2005 FGCS Issue

Coming Summer 2006!

Special iGrid 2005 issue

25 Refereed Papers!

Future Generation Computer Systems/ The International Journal of Grid Computing: Theory, Methods and Applications, Elsevier, B.V.

Guest Editors
Larry Smarr, Tom DeFanti,
Maxine Brown, Cees de Laat

Volume 19, Number 6, August 2003
Special Issue on iGrid 2002

Thank You!

- **Our planning, research, and education efforts are made possible, in major part, by funding from:**
  - US National Science Foundation (NSF) awards ANI-0225642, EIA-0115809, and SCI-0441094
  - State of Illinois I-WIRE Program, and major UIC cost sharing
  - State of California, UCSD Calit2
  - Many corporate friends and partners
  - Gordon and Betty Moore Foundation
- **Argonne National Laboratory and Northwestern University for StarLight and I-WIRE networking and management**
- **Laurin Herr and Maxine Brown for content and editing**

For More Information

www.glif.is

www.startap.net

www.evl.uic.edu

www.calit2.edu

www.igrid2005.org