

# Global Lambda Integrated Facility (GLIF)

CANS 2004, November 30, 2004

**Maxine D. Brown**

**Associate Director, Electronic Visualization Laboratory**

**Co-Principal Investigator, STAR TAP/StarLight**

**Co-Principal Investigator, Euro-Link/TransLight**

**University of Illinois at Chicago**



**STARLIGHT<sup>SM</sup> Euro-Link<sup>SM</sup> TRANSLIGHT<sup>SM</sup>**

# What is GLIF?

- **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**
- **GLIF is under the leadership of SURFnet and University of Amsterdam in The Netherlands.**
- **[www.glif.is](http://www.glif.is)**



# Why GLIF?

## Motivations, Starting in 2001

- **Scientific:** All science is global.
- **Political:** A *neutral forum* in which to collaborate with colleagues worldwide to build a production quality *Global LambdaGrid* in support of e-science experiments.
- **Economic:** As the cost of transoceanic bandwidth continues to become more affordable, National Research Networks have additional capacity they are willing to make available for use by application scientists, computer scientists and engineers.
- **Technical:** Need to interconnect and interoperate production quality infrastructure for scientific experiments.




# What is the LambdaGrid?

- **Today's Grids** enable scientists to schedule computer resources and remote instrumentation over today's "best effort" networks.
- **LambdaGrids** enable scientists to *also* schedule bandwidth. Wave Division Multiplexing (WDM) technology divides white light into individual wavelengths (or "lambdas") on optical fiber, creating parallel networks.
- LambdaGrids provide *deterministic networks* with known and knowable characteristics.
  - Guaranteed Bandwidth (data movement)
  - Guaranteed Latency (collaboration, visualization, data analysis)
  - Guaranteed Scheduling (remote instruments)



# The Next International Optical Network According to GLIF



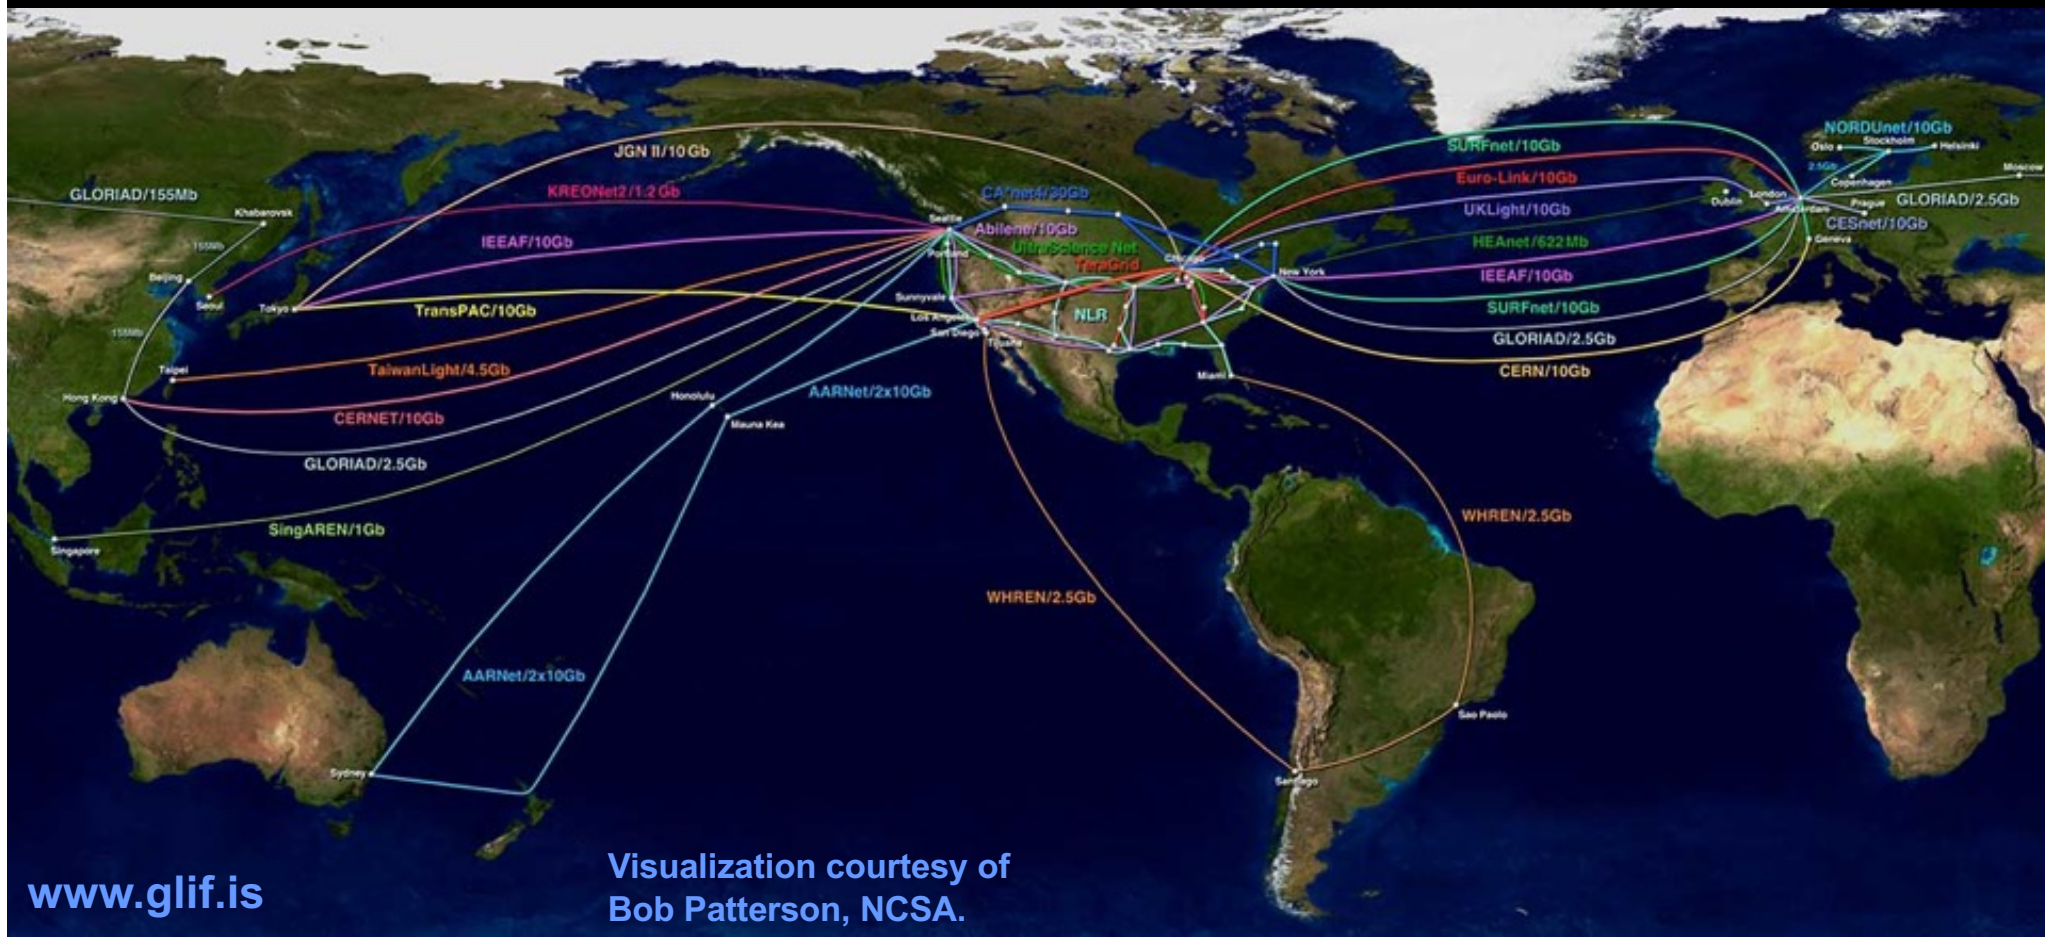
Source: Bill St. Arnaud

ST<sup>★</sup>RLIGHT<sup>SM</sup> Euro-Link<sup>SM</sup> TRANSLIGHT<sup>SM</sup>



# Global Lambda Integrated Facility World Map – December 2004

Predicted international Research & Education Network bandwidth, to be made available for scheduled application and middleware research experiments by December 2004.

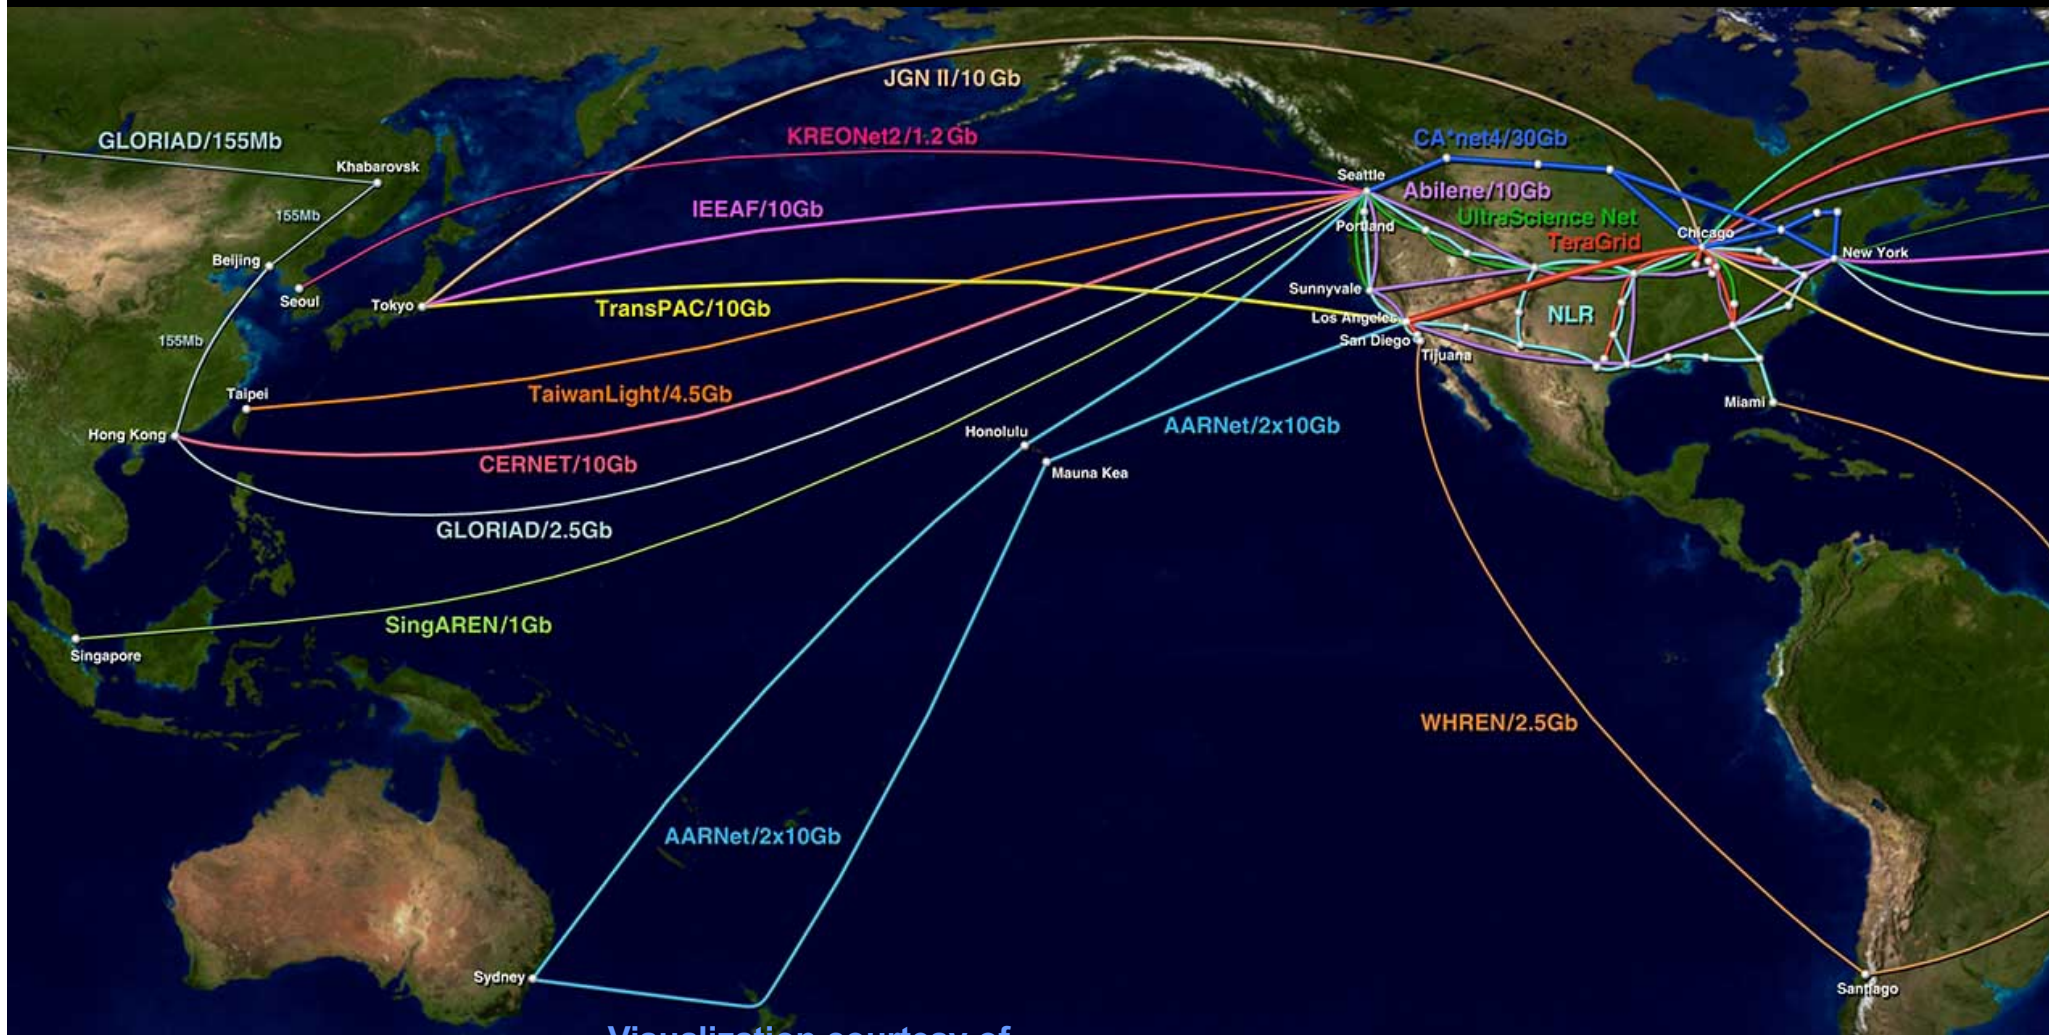


[www.glif.is](http://www.glif.is)



# Global Lambda Integrated Facility

## Predicted Bandwidth for Scheduled Experiments, December 2004



[www.glif.is](http://www.glif.is)

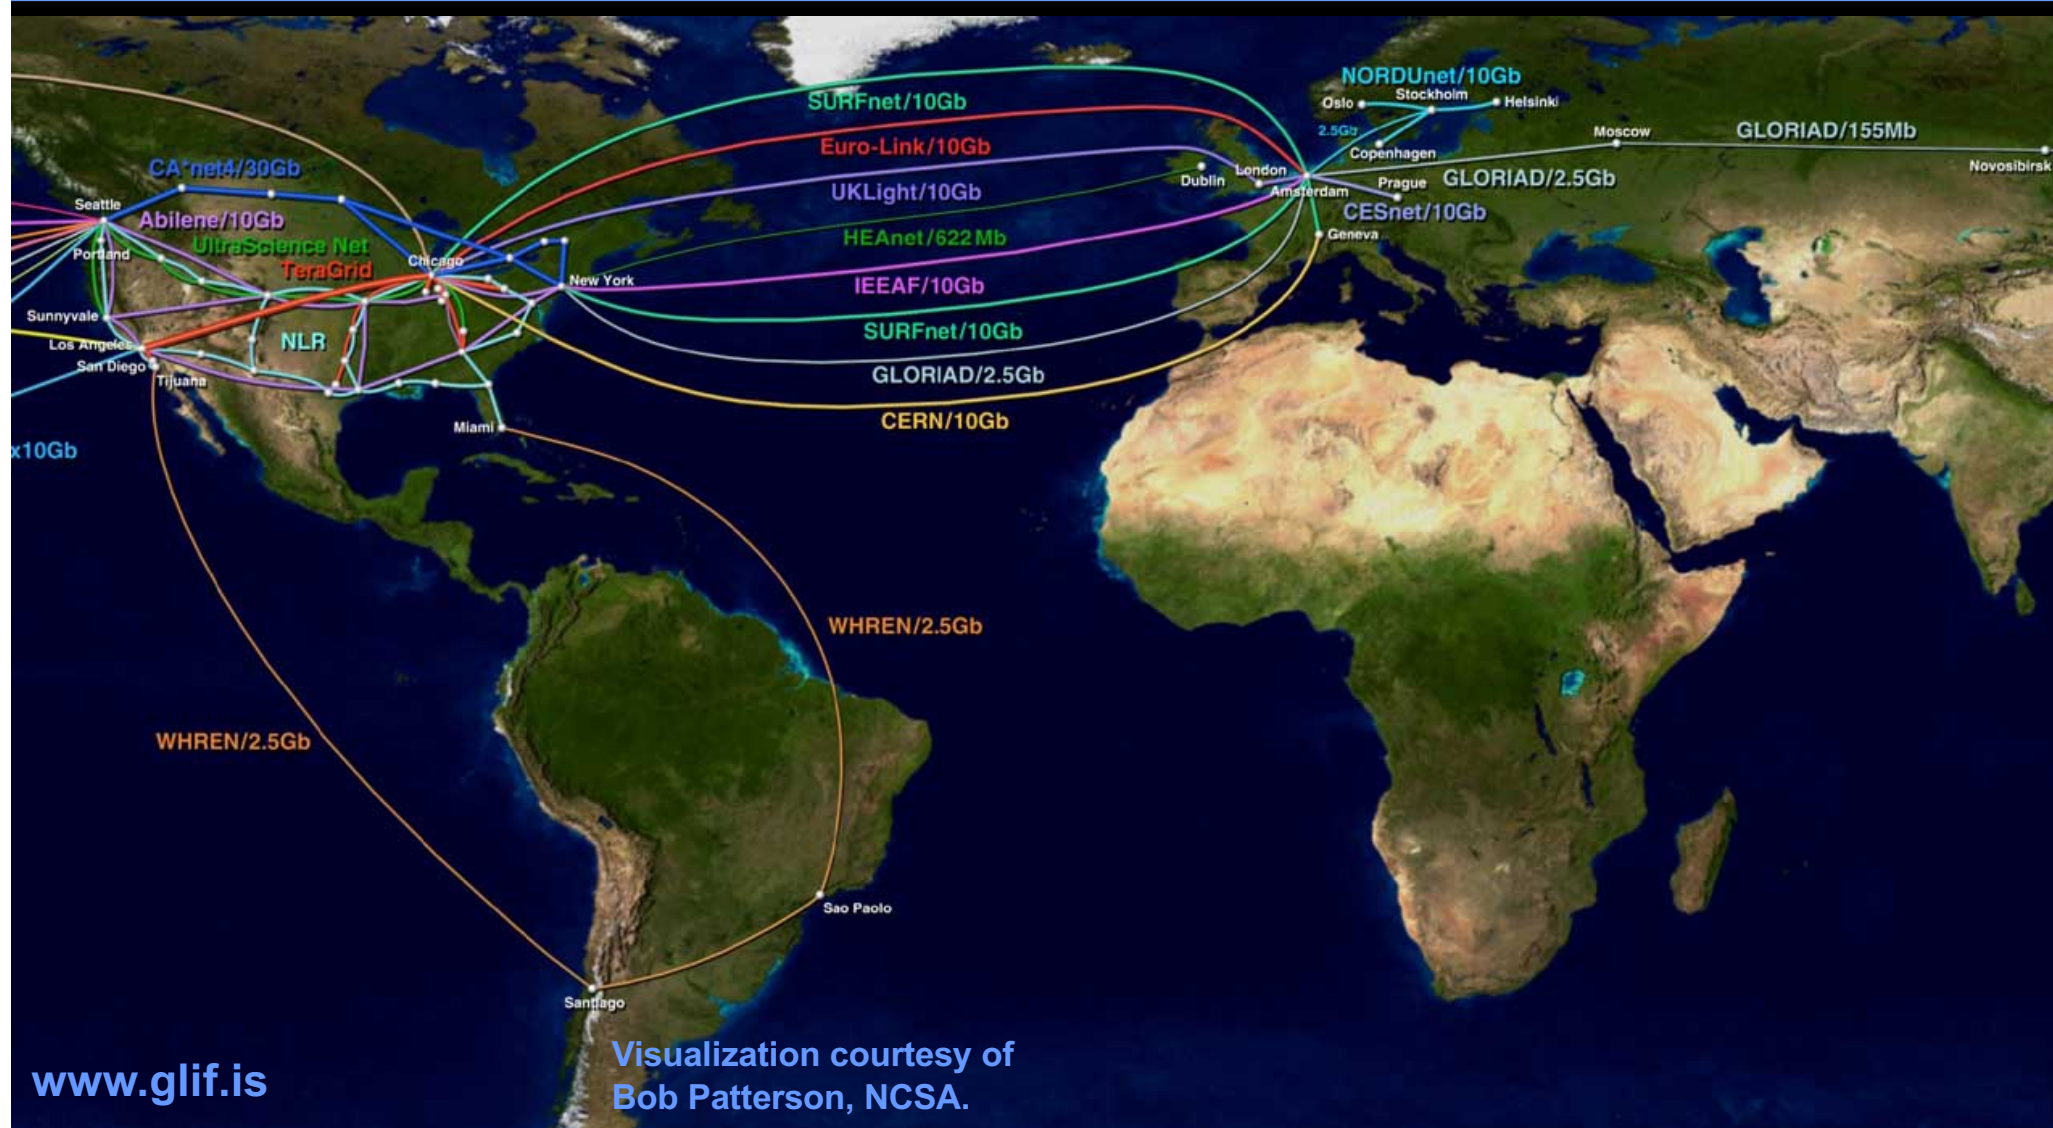
Visualization courtesy of  
Bob Patterson, NCSA.



STARLIGHT<sup>SM</sup> Euro-Link<sup>SM</sup> TRANSLIGHT<sup>SM</sup>

# Global Lambda Integrated Facility

## Predicted Bandwidth for Scheduled Experiments, December 2004





# Steps Leading Up to GLIF

- **Creation** of international Open Exchanges
  - StarLight (2001)
  - NetherLight (2001)
- **Cooperation** among institutions, organizations, consortia and country National Research Networks who voluntarily share optical networking resources and expertise for the advancement of scientific collaboration and discovery
  - LambdaGrid meetings in 2001, 2002, 2003, 2004
- **Collaborations** among discipline scientists, computer scientists, network engineers
  - iGrid 1998, iGrid 2000, iGrid 2002
  - SC conferences
  - Etc.



# Early Creation of International Open Exchanges (examples)




STARLIGHT<sup>SM</sup> Euro-Link<sup>SM</sup> TRANSLIGHT<sup>SM</sup>

# STAR TAP

## Science Technology And Research Transit Access Point

NSF funded the development of STAR TAP in 1997 to provide a persistent infrastructure for the long-term interconnection and interoperability of advanced international networking, in support of applications, performance measuring, and technology evaluations. By 2000, STAR TAP became a model for Next-Generation Internet eXchanges (NGIXs).



## Open Exchange “By Researchers For Researchers”

StarLight is a 1GE and 10GE switch/router facility for high-performance access to participating networks and also offers true optical switching for wavelengths.

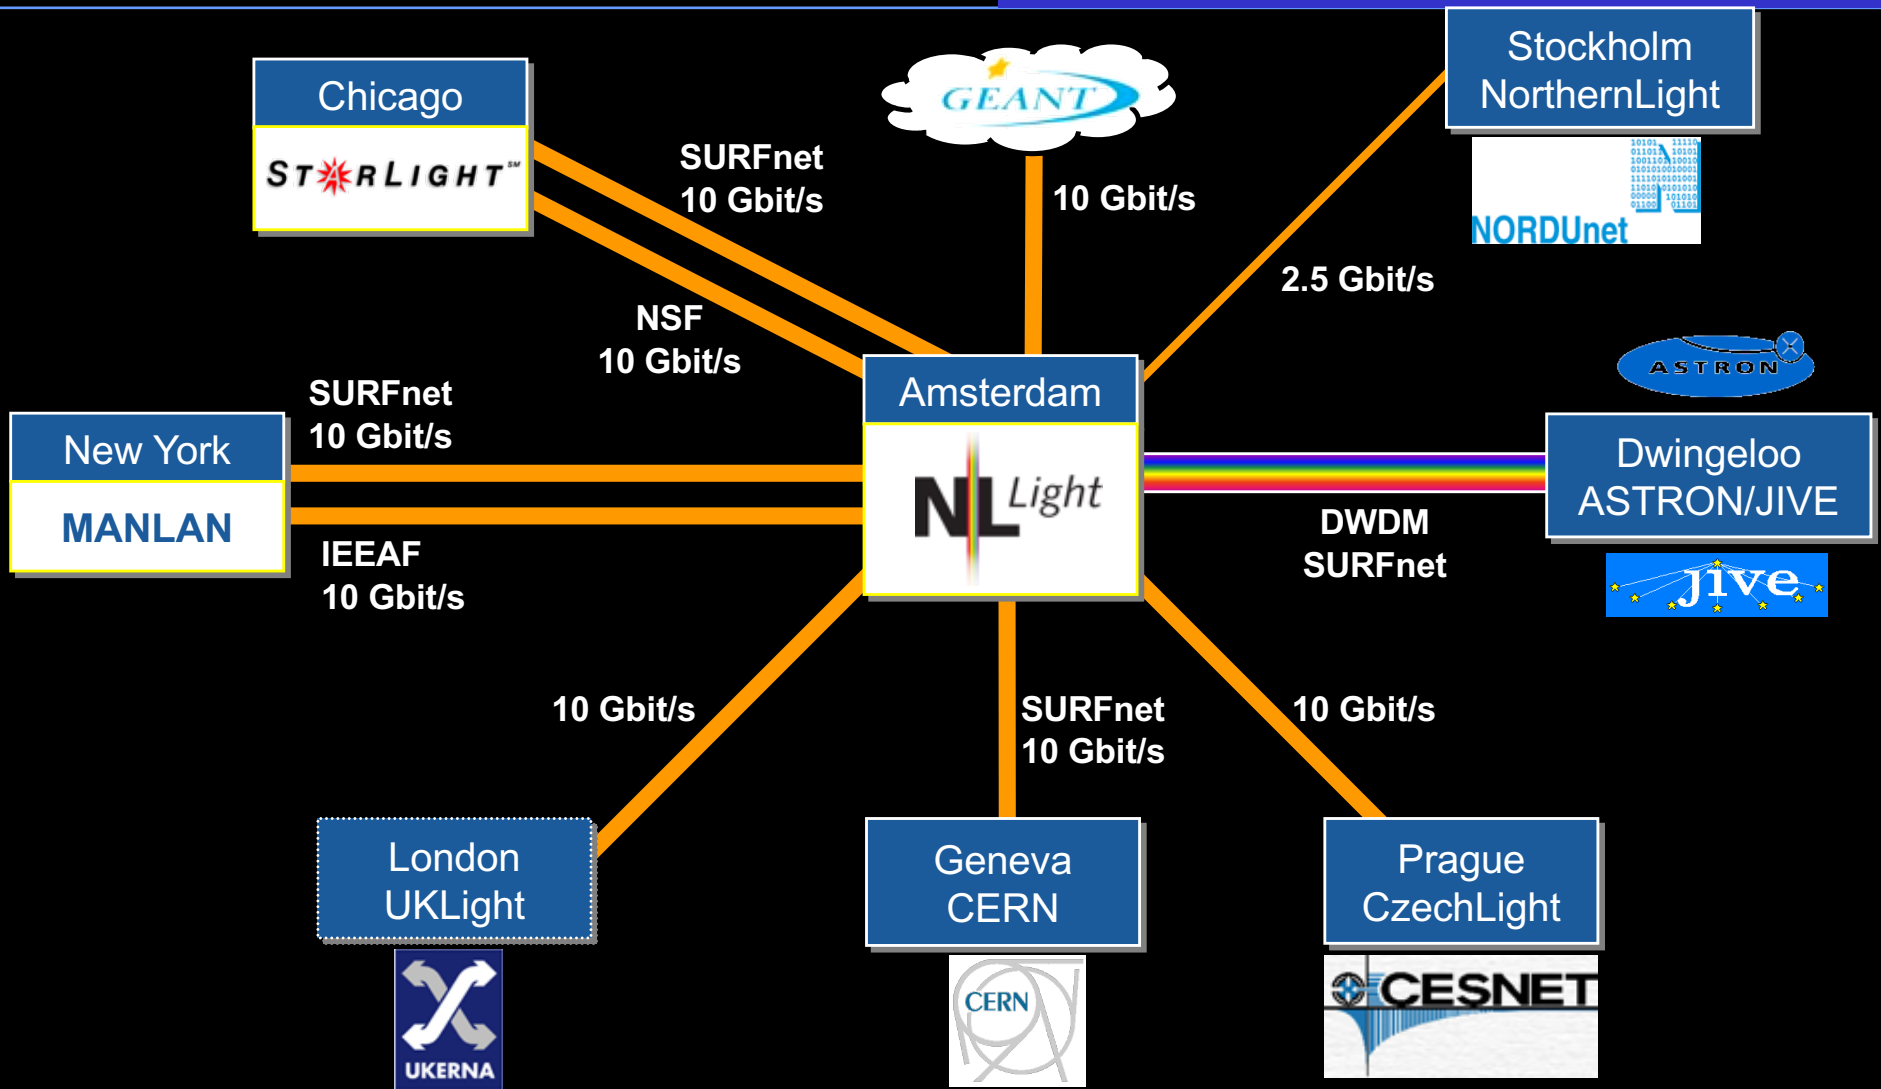


View from StarLight

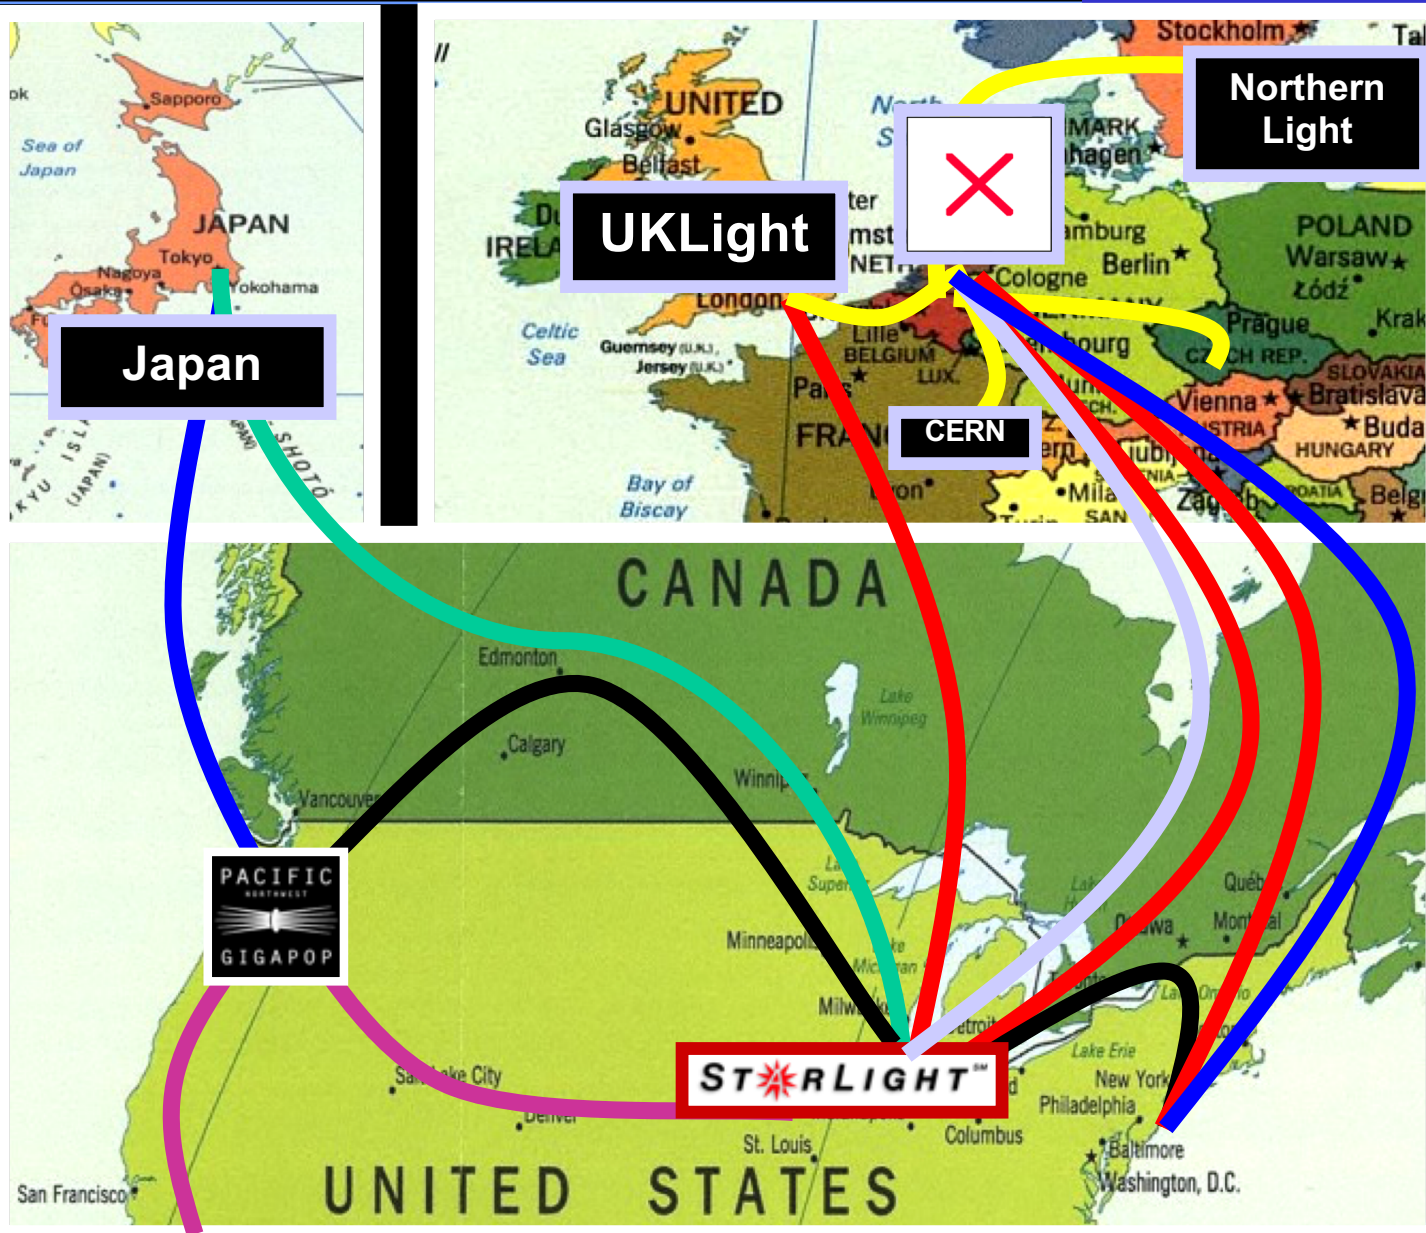


Abbott Hall, Northwestern University's Chicago downtown campus

# NetherLight Open Exchange



# Actual TransLight \*Lambdas Today: TransLight Governance Ends in 2004; Supports GLIF



**European lambdas to US (red)**  
 -10Gb Amsterdam—Chicago  
 -10Gb London—Chicago  
 -10Gb Amsterdam—NYC

**Canadian lambdas to US (white)**  
 -30Gb Chicago-Canada-NYC  
 -30Gb Chicago-Canada-Seattle

**US sublambdas to Europe (grey)**  
 -6Gb Chicago—Amsterdam

**Japan JGN II lambda to US (cyan)**  
 -10Gb Chicago—Tokyo

**European lambdas (yellow)**  
 -10Gb Amsterdam—CERN  
 -2.5Gb Prague—Amsterdam  
 -2.5Gb Stockholm—Amsterdam  
 -10Gb London—Amsterdam

**IEEAF lambdas (blue)**  
 -10Gb NYC—Amsterdam  
 -10Gb Seattle—Tokyo

**CAVEWave/PacificWave (purple)**  
 -10Gb Chicago—Seattle  
 -10Gb Seattle—LA—San Diego

# Cooperation to Share Optical Networking Resources and Expertise



**STARLIGHT**<sup>SM</sup> **Euro-Link**<sup>SM</sup> **TRANSLIGHT**<sup>SM</sup>

# GLIF History

- **Invitation-only annual meetings to discuss optical networking and the Global LambdaGrid.**
  - 2001 in Amsterdam, hosted by the Trans-European Research and Education Networking Association (TERENA, Europe)
  - 2002 in Amsterdam, hosted by the Amsterdam Science and Technology Centre



- 2003 in Reykjavik, Iceland, hosted by NORDUnet
- 2004 in Nottingham, UK, hosted by UKERNA





# GLIF 2004: 60 World Leaders in Advanced Networking and the Scientists Who Need It



GLIF 2004

Photo courtesy of Steve Wallace



STARLIGHT<sup>SM</sup> Euro-Link<sup>SM</sup> TRANSLIGHT<sup>SM</sup>

# GLIF 2004 Who's Who

- **Australia's Research and Education Network (AARNet)**
- **CANARIE (Canada)**
- **CERN**
- **CESNET (Czech Republic)**
- **Chinese Academy of Science**
- **DANTE/GÉANT (Europe)**
- **European Commission**
- **HEAnet (Ireland)**
- **Japanese Gigabit Network 2 (JGN-II)**
- **Korea Institute of Science and Technology Information (KISTI)/KREONet2**
- **National Center for High Performance Computing (NCHC, Taiwan)**
- **National Institute of Advanced Industrial Science and Technology (AIST, Japan)**



# GLIF 2004 Who's Who

- **NORDUnet (Nordic countries)**
- **SURFnet/NetherLight (The Netherlands)**
- **Trans-European Research and Education Networking Association (TERENA, Europe)**
- **UK Joint Information Systems Committee (JISC)**
- **UKERNA/UKLight (United Kingdom)**
- **WIDE (Japan)**
- **USA**
  - National LambdaRail, Internet2, DoE ESnet, TeraGrid, Illinois' I-WIRE initiative, California's CENIC network, NSF StarLight, NSF High Performance International Internet Services awardees (Euro-Link, TransPAC, GLORIAD and AMPATH), major GigaPoPs (PNWGP and Pacific Wave, MREN, MAX) the Internet Educational Equal Access Foundation (IEEAF), and major universities and government laboratories.



# GLIF Working Groups

- **Governance:** To create an open, neutral community for anyone who wants to contribute resources and/or services (bandwidth, software, application drivers), to build the Global LambdaGrid
- **Engineering:** To define the types of links and the minimum/maximum configurations of Optical Exchange facilities in order to assure the interoperability and interconnectivity of participating networks
- **Applications:** To enable the super-users providing the application drivers; to find new e-science drivers; and, to move scientific experiments into production usage as they mature, and to document these advancements
- **Control Plane and Grid Integration Middleware** (proposed): To agree on the interfaces and protocols for lambda provisioning and management



# Collaborations Among Discipline Scientists, Computer Scientists, Network Engineers



# iGrid 1998 at SC'98

November 7-13, 1998, Orlando, Florida, USA

- **10 countries:** Australia, Canada, Germany, Japan, Netherlands, Russia, Singapore, Switzerland, 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



[www.startup.net/igrd98](http://www.startup.net/igrd98)




STARLIGHT<sup>SM</sup> Euro-Link<sup>SM</sup> TRANSLIGHT<sup>SM</sup>

# iGrid 2000 at INET 2000

## July 18-21, 2000, Yokohama, Japan

- **14 regions:** 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
- **100Mb transpacific bandwidth carefully managed**



[www.startap.net/igrd2000](http://www.startap.net/igrd2000)







STARLIGHT<sup>SM</sup> Euro-Link<sup>SM</sup> TRANSLIGHT<sup>SM</sup>

# iGrid 2002

## Application Demonstrations

- **28 demonstrations from 16 countries:** Australia, Canada, CERN, 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, tele-science
- **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!)**



[www.startap.net/igrd2002](http://www.startap.net/igrd2002)

STARLIGHT<sup>SM</sup> Euro-Link<sup>SM</sup> TRANSLIGHT<sup>SM</sup>




# Electronic Very-Long Baseline Interferometry (eVLBI)

## 22 September 2004 – Creating a Global Virtual Telescope

**VLBI** is a technique used by radio astronomers to image the sky in detail. Arrays of telescopes across countries/continents record data onto tape and then ship the tapes to a central processing facility (correlator) for analysis. The resulting image has a resolution equal to that of a telescope as large as the maximum antenna separation.

**E-VLBI** links telescopes electronically, enabling real-time analysis.

- A 20-hour long observation of the star IRC+10420, heading toward a cataclysmic “supernova” explosion, involved telescopes in the UK, Sweden, the Netherlands, Poland and Puerto Rico. The maximum separation of the antennas was 8200 km, giving a resolution of 20 milliarcseconds (mas); about 5 times better than the Hubble Space Telescope. The antenna at Arecibo, in Puerto Rico, increased the sensitivity of the telescope array by a factor of 10.
- Each telescope was connected to its country’s NREN and the data routed at 32 Mbps per telescope through GÉANT to SURFnet. The data was then sent to the Joint Institute for VLBI in Europe (JIVE) in the Netherlands (6Gbps link from NetherLight), the central processing facility for the European VLBI Network (EVN). 9 Terabits of data were correlated and the resulting image returned to the participating sites.



[www.evlbi.org](http://www.evlbi.org)

STARLIGHT<sup>SM</sup> Euro-Link<sup>SM</sup> TRANSLIGHT<sup>SM</sup>

# LOFAR

## Lambdas as part of research instruments

LOFAR is a modern radio telescope that uses simple omni-directional antennas instead dish antennas. (Basic radio telescope technology has not changed since the 1960's. A telescope 100x larger than existing instruments is therefore unaffordable.) A more sensitive telescope will see stars, galaxies, black holes and other objects that are farther away.

LOFAR is an IT-telescope – electronic signals from the antennas are digitized, transported to a central digital processor, and combined in software to emulate a conventional antenna. The Dutch have funded phase 1; 15,000 antennas and maximum baselines of 100 km will be built. Data transport requirements are in the range of many Tbps.

- Many data collection points
- Processing in Groningen
- Large datasets distributed to many destinations in The Netherlands and abroad



[www.lofar.org](http://www.lofar.org)



# SC 2004

## Data Reservoir Project – University of Tokyo

- University of Tokyo
- Fujitsu Computer Technologies LTD
- WIDE
- NTT Communications
- SARA
- University of Amsterdam



**Land Speed Record:** The world's longest 10Gbps circuit ever recorded – 31,248 km circuit between Pittsburgh and CERN via Tokyo. This demonstration achieved a single stream TCP payload of 7.2 Gbps – for a new record of 224,985 terabit kilometer / second.

This represents a collaboration among WIDE, APAN/JGN-II, NetherLight, SURFnet, StarLight, CANARIE, PNWGP, IEEAF, SCinet

<http://www.supercomputingonline.com/article.php?sid=7514>



# SC 2004

## Computing the Quantum Universe

- California Institute of Technology (Caltech)
- Stanford Linear Accelerator Center (SLAC)
- Fermi National Accelerator Laboratory
- CERN
- University of Florida
- Florida International University
- National LambdaRail and Abilene
- University of Manchester
- University College London and UKLight
- University of London
- Rio de Janeiro State University (Brazil)
- State University of Sao Paulo
- Academic Network of Sao Paulo (ANSP)
- Kyungpook National University (Korea)



**Bandwidth Challenge for Sustained Bandwidth:** Transferred data in and out of the SC site in Pittsburgh at a rate of 101.13 Gbps, roughly equivalent to transferring the contents of three DVDs in just one second.

Seven 10Gbps links were connected to the Caltech booth and three 10Gbps links to the SLAC/FermiLab booth. (Four dedicated NLR waves from the show floor to Los Angeles (2 waves), Chicago, and Jacksonville, as well as three 10Gbps connections to Abilene, the TeraGrid and ESnet. Also, 2.5Gbps from Pittsburgh to Miami (via Abilene) and to Brazil.

<http://supercomputing.fnal.gov>



**STARLIGHT<sup>SM</sup>** **Euro-Link<sup>SM</sup>** **TRANSLIGHT<sup>SM</sup>**

# SC 2004: The OptIPuter

## Scalable Adaptive Graphics Environment

- UIC Electronic Visualization Laboratory
- UCSD National Center for Microscopy and Imaging Research
- UCSD Scripps Institution of Oceanography
- USGS Earth Resources Observatio Systems Data Center
- University of Amsterdam
- National LambdaRail



**SAGE used optical networking to retrieve very-high-resolution 2D and 3D datasets from servers in San Diego and Amsterdam, render them on clusters in Chicago, and display them in Pittsburgh.**

[www.evl.uic.edu/cavern/sage](http://www.evl.uic.edu/cavern/sage), [www.optiputer.net](http://www.optiputer.net)



# SC 2004: The OptIPuter Trans-Pacific HDTV

- UCSD National Center for Microscopy and Imaging Research
- Biomedical Informatics Research Network
- Osaka University
- KDDI R&D Laboratories
- National LambdaRail



HDTV from the world's largest microscope in Japan is streamed live to SC 2004 and San Diego while being controlled by project scientists in San Diego. High-quality HDTV is essential for resolving useful information. Dedicated wavelengths offer lower latencies and control of network jitter, especially important in large streams of video data.

<http://ncmir.ucsd.edu>



STARLIGHT<sup>SM</sup>

Euro-Link<sup>SM</sup>

TRANSLIGHT<sup>SM</sup>



# SC 2004: The OptIPuter

## UDT Fairness and Friendliness

- UIC National Center for Data Mining
- Johns Hopkins University, Department of Physics and Astronomy
- SARA Computing and Networking Services, Amsterdam, The Netherlands
- National LambdaRail



UDP-based Data Transfer protocol (UDT) is a very fast and efficient transport protocol, used to transfer data over long distances. UDT's fairness and friendliness properties are visually displayed on a web interface.

<http://demos.dataspaceweb.net>



# SC 2004: Sloan Digital Sky Survey

- UIC National Center for Data Mining
- Johns Hopkins University, Department of Physics and Astronomy
- Tokyo
- SARA Computing and Networking Services, Amsterdam, The Netherlands



**Bandwidth Challenge 3rd Place:** 1.5 Terabytes of Sloan Digital Sky Survey (SDSS) data (release 3) was sent via network instead of shipping disks and computers. Using the UDT transport protocol, data was transferred to disk at 1.6 Gbps between two nodes at SC in Pittsburgh and two nodes in Tokyo (disk I/O bottleneck). Data was also moved memory to memory over a larger cluster at 16Gbps (out of 20Gbps) to nodes in Tokyo and Amsterdam. Next month, data will be sent from Chicago to Japan, Korea and Australia. Clusters are part of an NSF/Army sponsored Teraflow Testbed.

[www.sdss.org](http://www.sdss.org), <http://demos.dataspaceweb.net>



STARLIGHT<sup>SM</sup> Euro-Link<sup>SM</sup> TRANSLIGHT<sup>SM</sup>



# SC 2004: Multi-Gigabit Interactive Video Transmission

- AARNet (Australia's Academic and Research Network)
- ResearchChannel
- University of Washington/ Pacific Northwest GigaPoP
- University of Hawaii
- NLR



This is the first demonstration of high-definition uncompressed interactive video interaction across the Pacific, between Canberra, Australia, and SC 2004 in Pittsburgh, at 1.4Gbps in each direction. Future expansion of AARNet is planned to offer these and additional resources to other continents, and bring scientists and researchers together by exploiting new Internet technologies.

[www.aarnet.edu.au](http://www.aarnet.edu.au), [www.researchchannel.org](http://www.researchchannel.org)



STARLIGHT<sup>SM</sup> Euro-Link<sup>SM</sup> TRANSLIGHT<sup>SM</sup>

**iGrid 2005**  
**September 26-30, 2005 in San Diego**

# ***i*Grid 2005**

**THE GLOBAL LAMBDA INTEGRATED FACILITY**

**University of California, San Diego**  
**California Institute for Telecommunications and Information Technology**  
**[Cal-(IT)<sup>2</sup>]**



**STARLIGHT<sup>SM</sup> Euro-Link<sup>SM</sup> TRANSLIGHT<sup>SM</sup>**

# The Future of GLIF

- 2005 at UCSD, hosted by Cal-(IT)<sup>2</sup> in conjunction with iGrid 2005
- 2006 in Japan, hosted by the WIDE Project (Jun Murai) and JGN-II (Tomonori Aoyama)



# Thanks to NSF and Colleagues For These International Networking Opportunities

- **StarLight/Euro-Link/TransLight planning, research, collaborations, and outreach efforts are made possible, in major part, by funding from:**
  - National Science Foundation (NSF) awards SCI-9980480, SCI-9730202, CNS-9802090, CNS-9871058, SCI-0225642, and CNS-0115809
  - State of Illinois I-WIRE Program, and major UIC cost sharing
  - Northwestern University for providing space, power, fiber, engineering and management
  - Pacific Wave, StarLight, National LambdaRail, CENIC, PNWGP, CANARIE, SURFnet, UKERNA, and IEEAF for Lightpaths
- **DoE/Argonne National Laboratory for StarLight and I-WIRE network engineering and design**

