Shrinking the Planet—How Dedicated Optical Networks are Transforming Computational Science and Collaboration

Invited Presentation to the
Advisory Committee on Cyberinfrastructure
National Science Foundation
Arlington, VA
December 16, 2008

Dr. Larry Smarr
Director, California Institute for Telecommunications and Information Technology
Harry E. Gruber Professor,
Dept. of Computer Science and Engineering
Jacobs School of Engineering, UCSD

OptIPuter

it²

The Three Underpinnings of Cyberinfrastructure
The More Things Change, the More They Stay the Same!

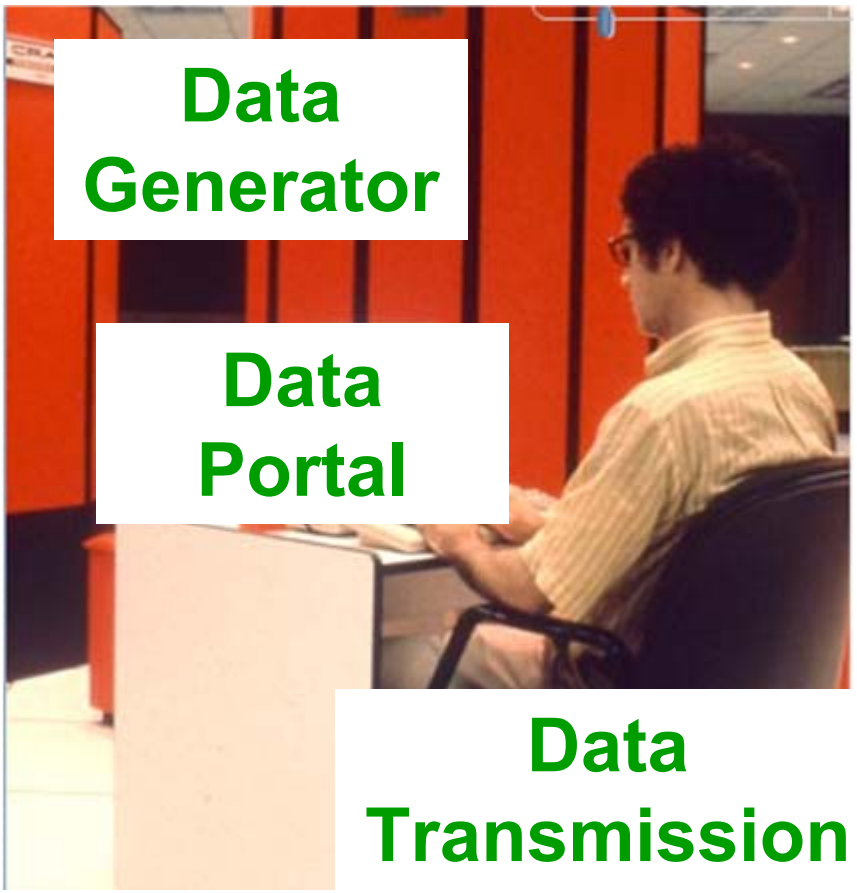

John Kogut Simulating Quantum Chromodynamics

- He Uses NCSA Telnet on the Mac
- The Mac Communicates with Cray
- The Simulation Runs on the Remote Cray

NCSA

NSF

OptIPuter

it²

Source: Larry Smarr 1985

The 20 Year Pursuit of a Dream: Building a CI for Collaborative Computational Science

"What we really have to do is eliminate distance between individuals who want to interact with other people and with other computers."
— Larry Smarr, Director, NCSA

- **Televisualization:**
  - **Telepresence Collaboration**
  - **Remote Interactive Visual Supercomputing**
  - **Multi-disciplinary Scientific Visualization**

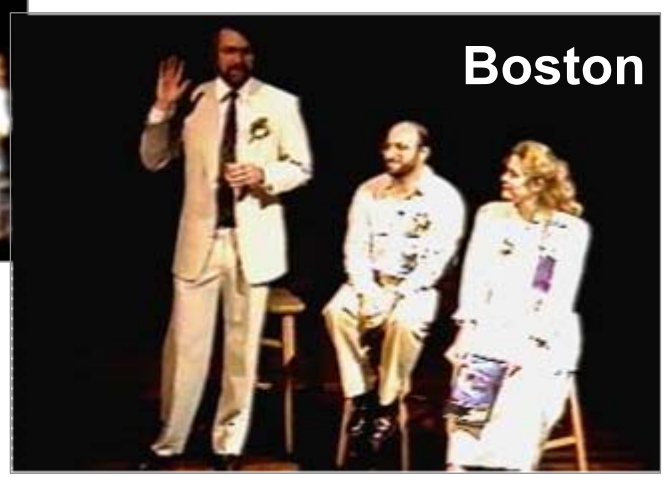

"We're using satellite technology…to demo what It might be like to have high-speed fiber-optic links between advanced computers in two different geographic locations."

— **Al Gore, Senator**
Chair, US Senate Subcommittee on Science, Technology and Space

ATT & Sun

SIGGRAPH 1989

OptIPuter

NCSA

it²

Data Transmission:
From Shared Internet to Dedicated Lightpaths

OptIPuter

it²

Dedicated Optical Fiber Channels Makes High Performance Cyberinfrastructure Possible

WAVELENGTH DIVISION MULTIPLEXING (WDM)

- EXPLOITS
  - ENORMOUS BANDWIDTH OF SILICA FIBER
  - HIGH-GAIN WIDEBAND OPTICAL AMPLIFIERS

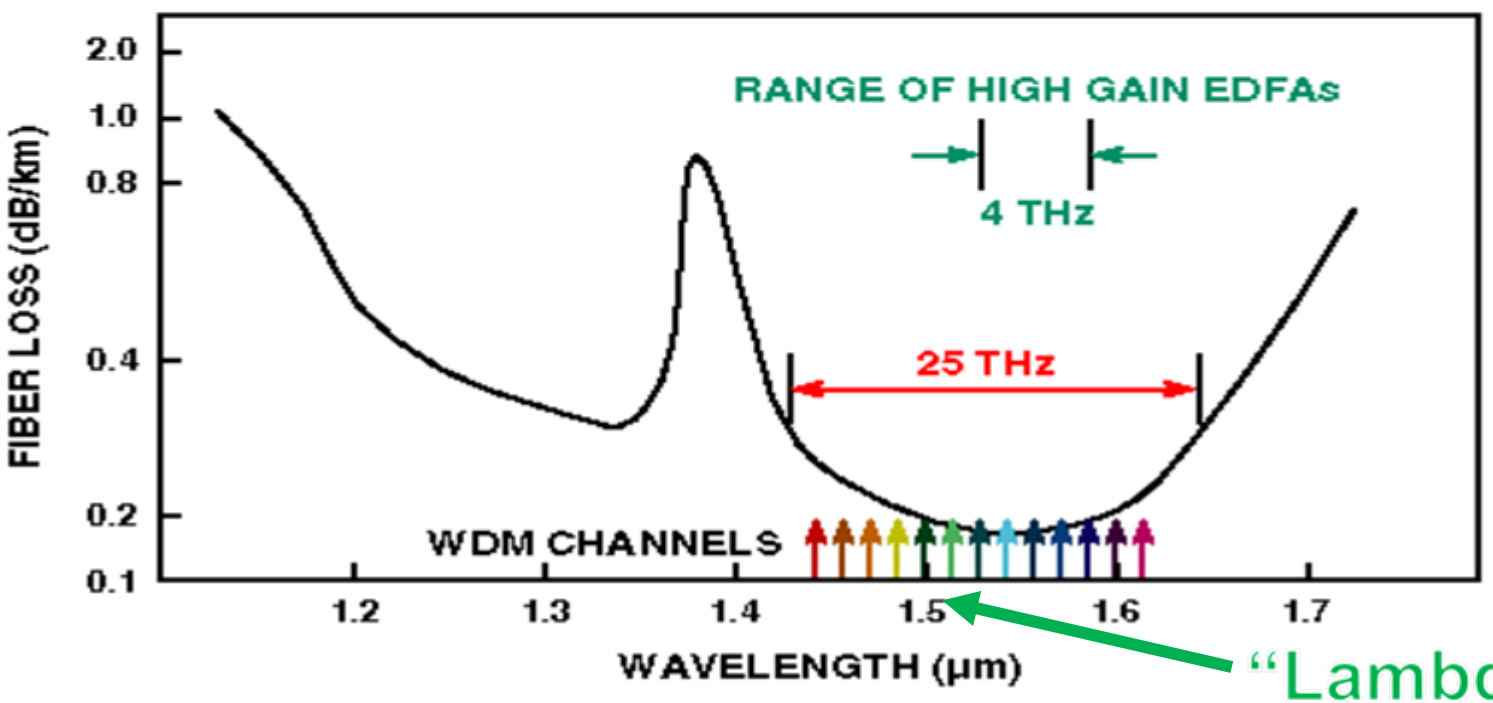

WDM Enables 10Gbps Shared Internet on One Lambda and a Personal 10Gbps Lambda on the Same Fiber!

Dedicated 10Gbps Lightpaths Tie Together State and Regional Fiber Infrastructure

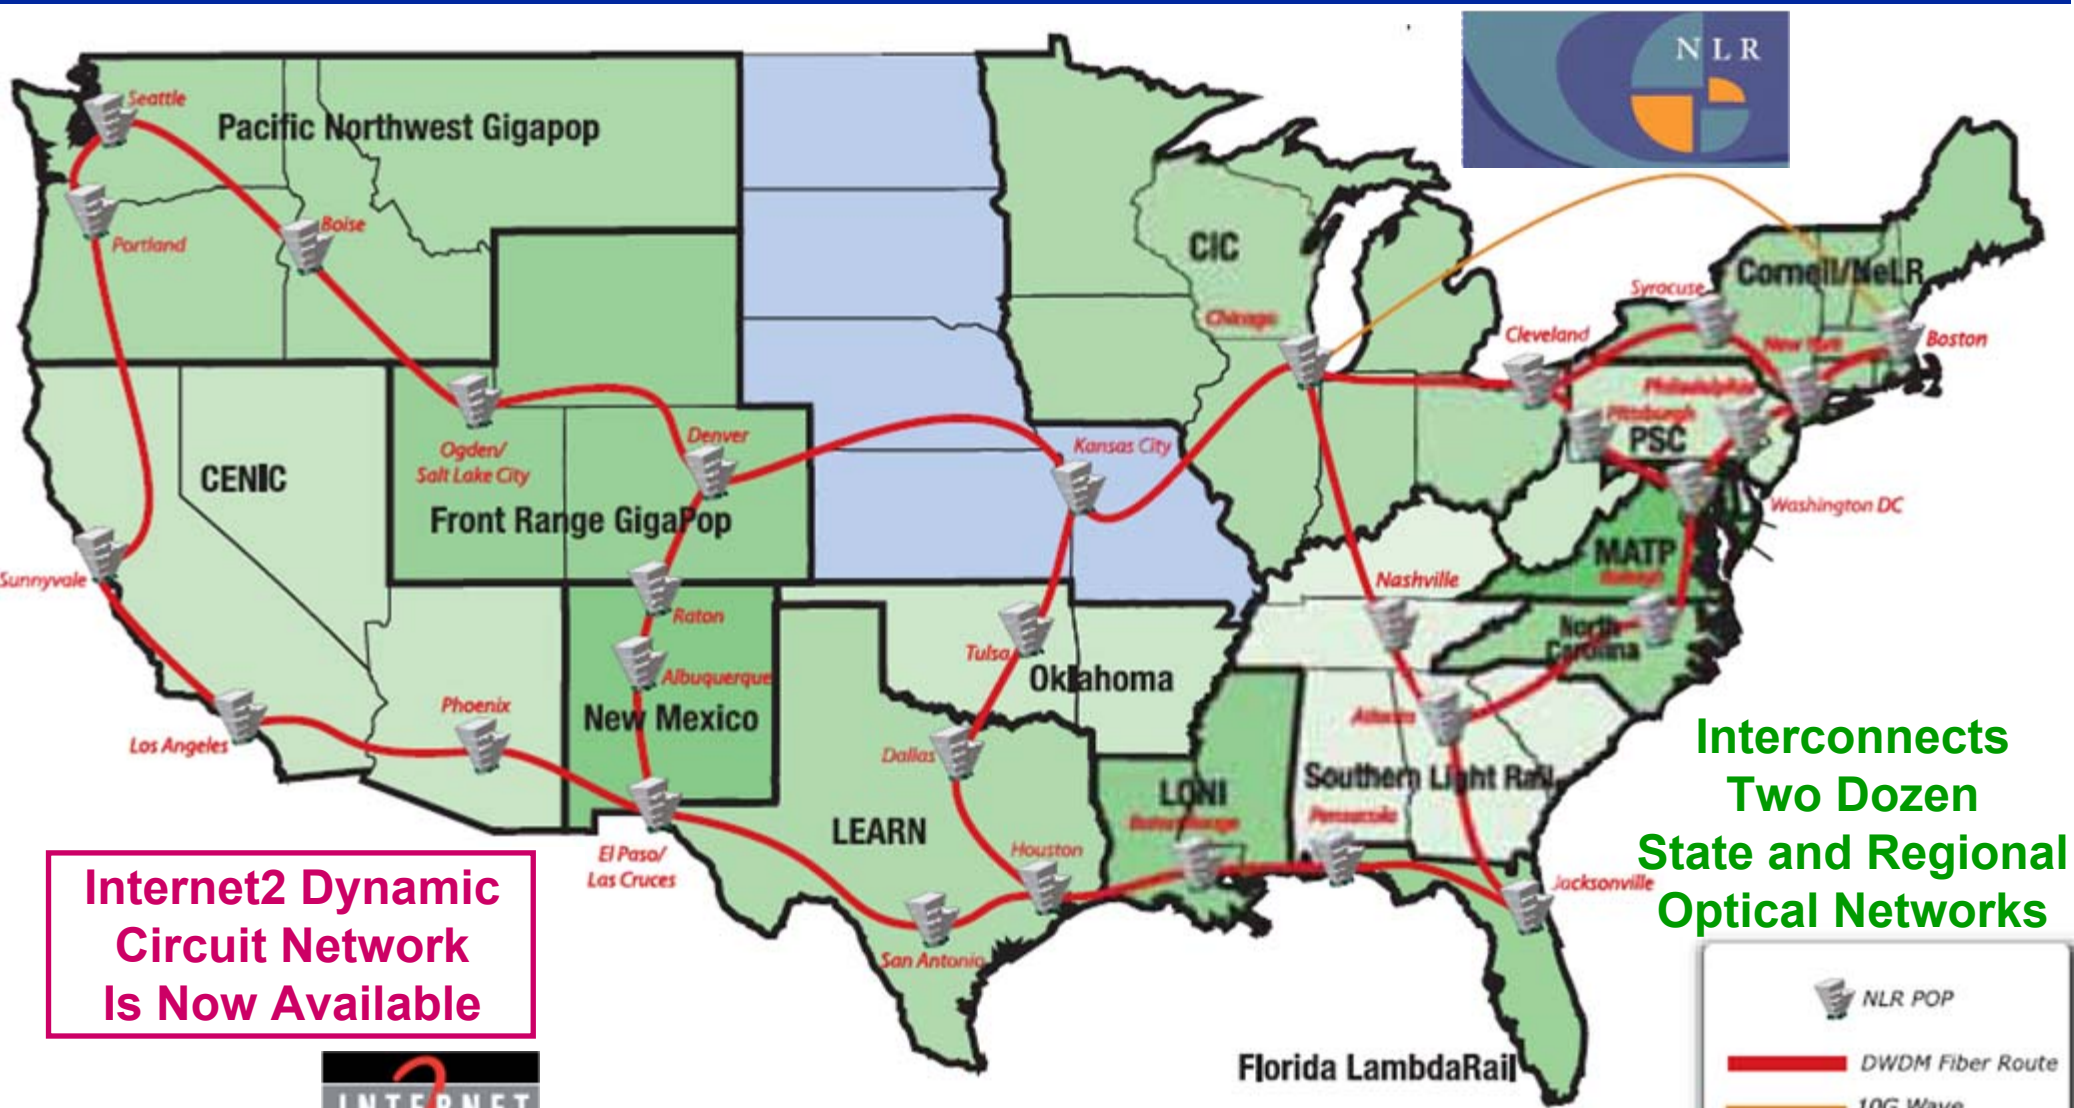

Discovering New Applications and Services Enabled by 1-10 Gbps Lambdas

Maxine Brown, Tom DeFanti, Co-Chairs

iGrid 2005

www.igrid2005.org

THE GLOBAL LAMBDA INTEGRATED FACILITY

September 26-30, 2005

Calit2 @ University of California, San Diego

California Institute for Telecommunications and Information Technology

the UC Discovery Grant

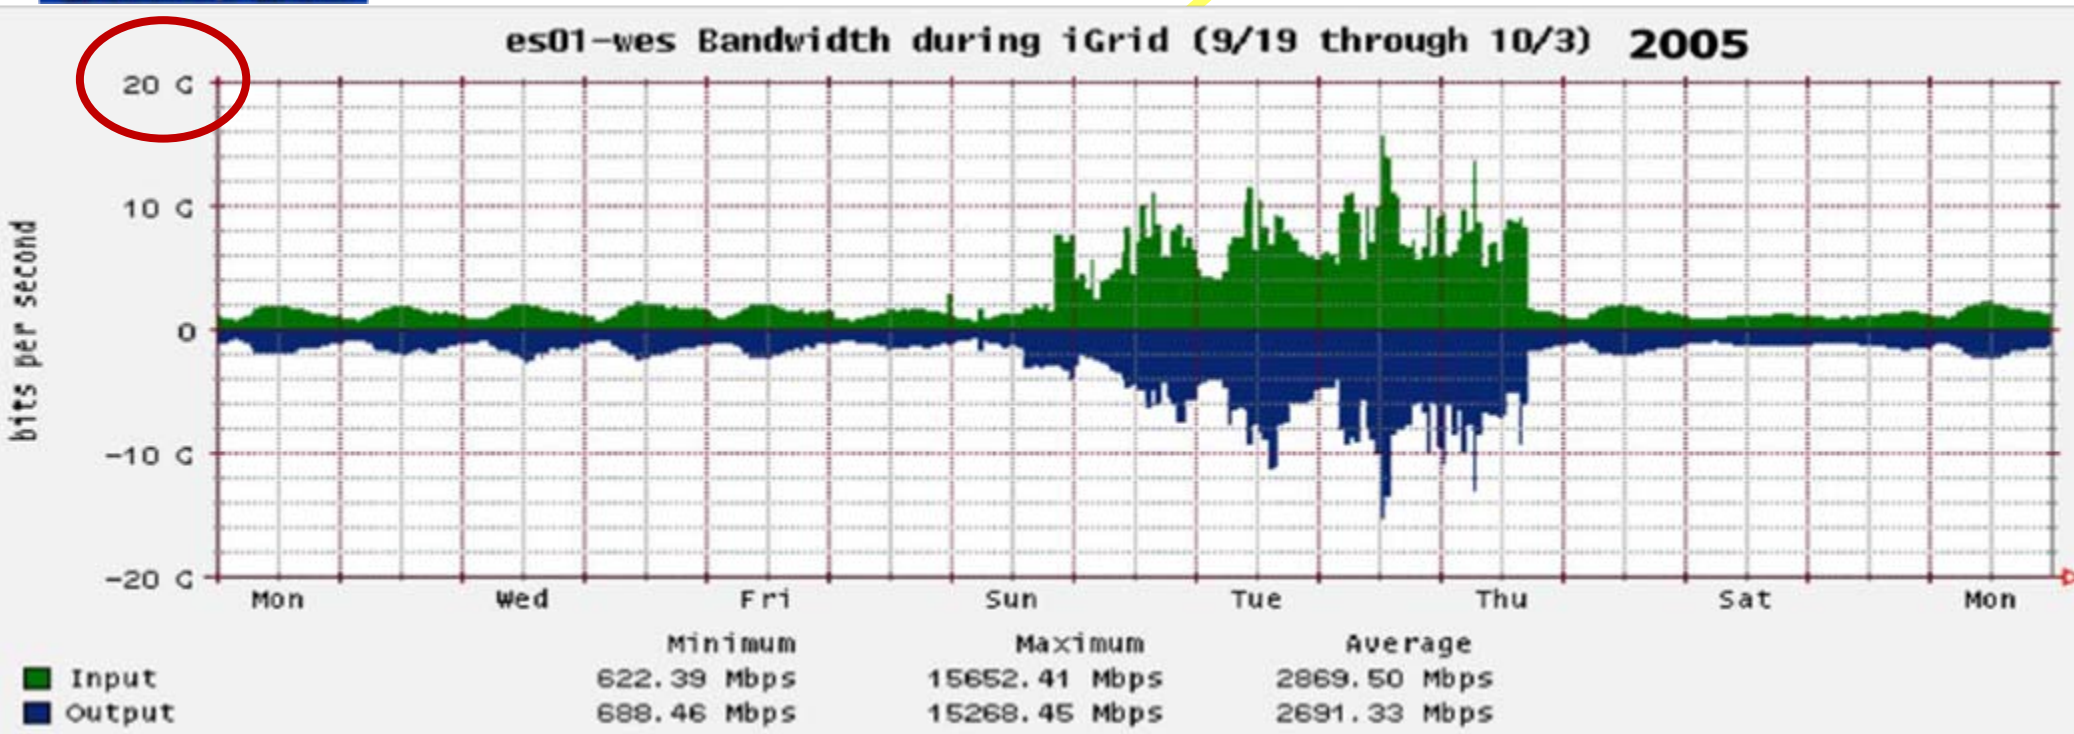

Major NSF Instruments Require Enormous Bandwidth

"The VLA facility is now able to generate 700 Gigabits/s of astronomical data and the Extended VLA will reach 3.2 Terabits/sec by 2009."

--Dr. Steven Durand, National Radio Astronomy Observatory, e-VLBI Workshop, MIT Haystack Observatory., Sep 2006.

ALMA Has a Requirement for a 120 Gbps Data Rate per Telescope

OptIPuter

it²

Next Great Planetary Instrument: The Square Kilometer Array Requires Dedicated Fiber

www.skatelescope.org

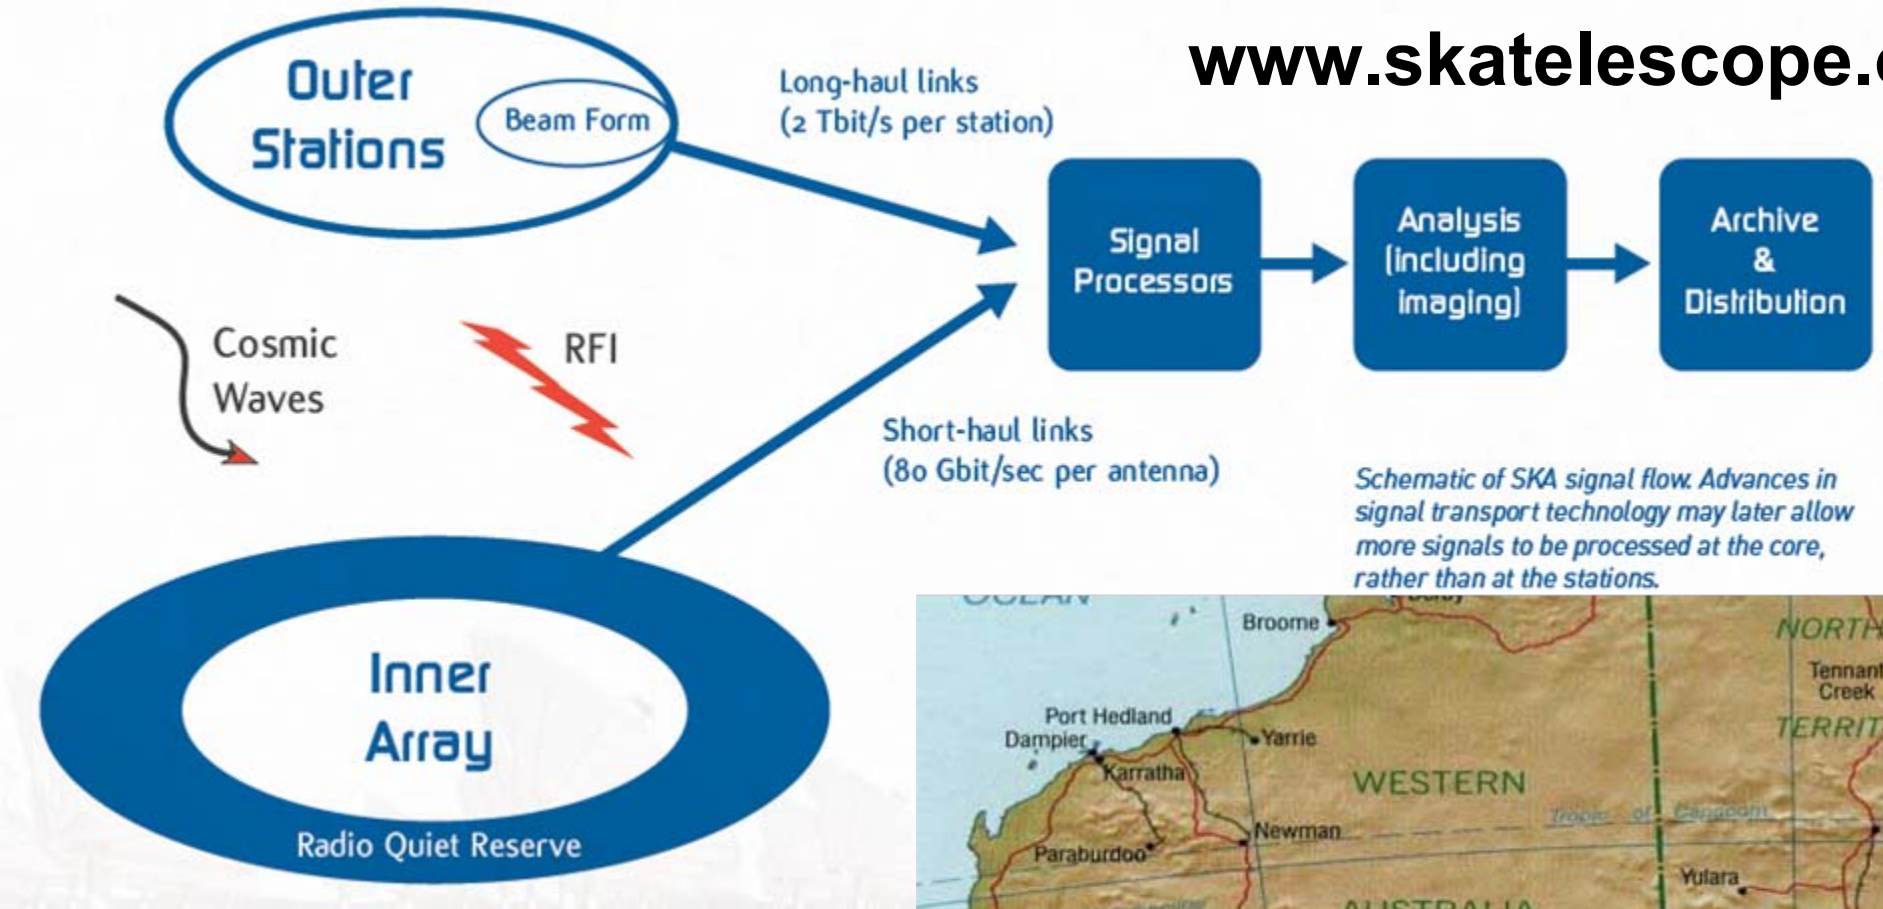

Schematic of SKA signal flow. Advances in signal transport technology may later allow more signals to be processed at the core, rather than at the stations.

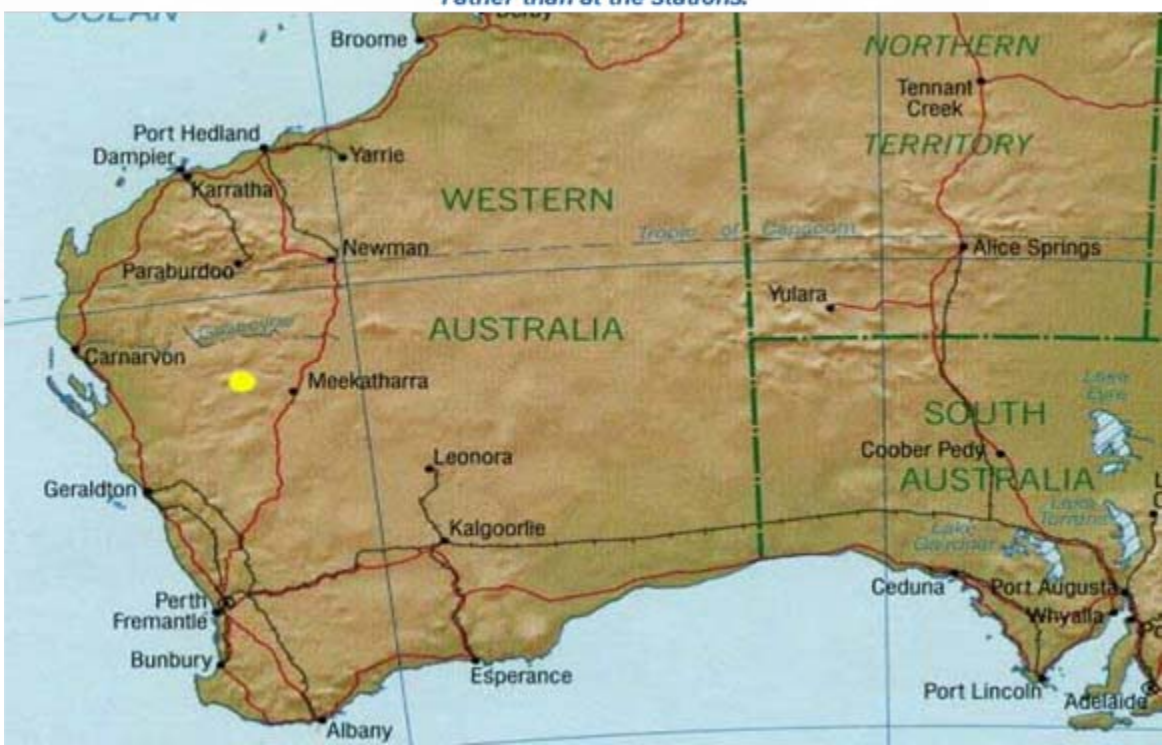

Transfers Of 1 TByte Images World-wide Will Be Needed Every Minute!

OptIPuter

Data Portals:
From User Analysis on PCs to OptIPortals

NSF: "Distribution of Balanced Computational Capacity at the Local Nodes is as Important as the Provision of Maximum Capability of the Principal Node"

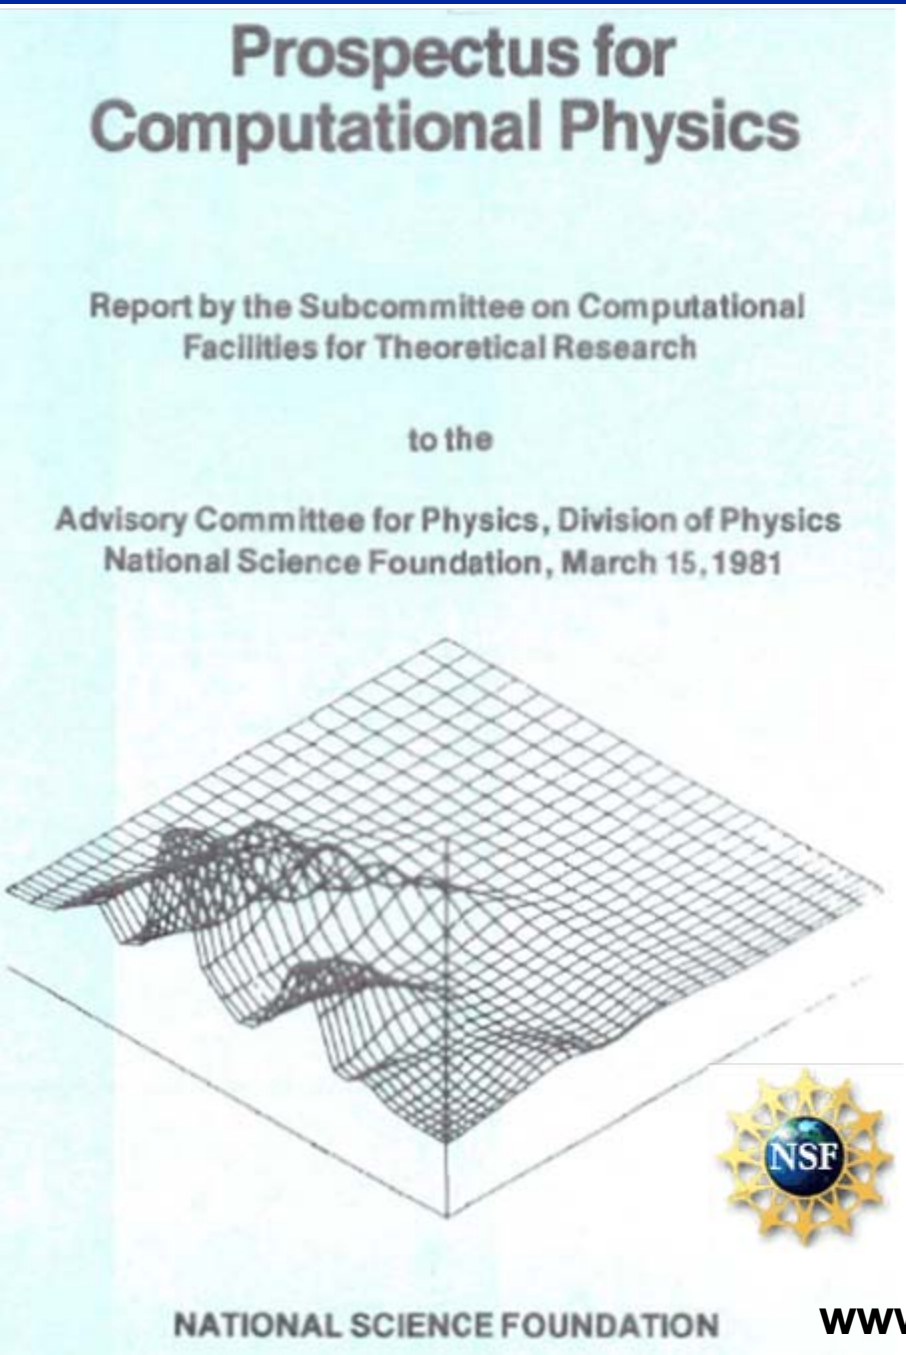

- **The Local Nodes Make Possible Modes of Operation and Scientific Investigations Not Possible via Remote Acess to the Central Node**
- **A Local Node Also Provides for Much Higher I/O Rates to the User than can be Provided with Remote Access to a Central Facility**
- **Any Commitment to Increased Computational Resources Must be Accompanied by an Expansion of Graphics Facilities to Help Manage the Increased Data Flow**

www.nr.com/whp/NSFCompPhys1981.pdf

Challenge—How to Bring Scalable Visualization Capability to the Supercomputer End User?

NCSA 4 MPixel
NSF Alliance PowerWall

LLNL 20 Mpixel Wall

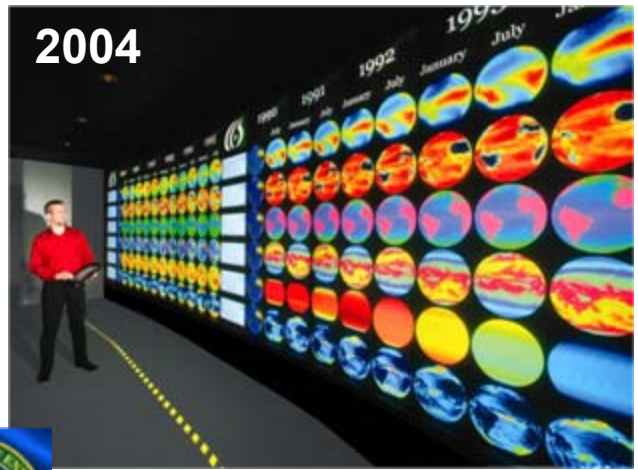

ORNL 35Mpixel EVEREST

EVL 100 Mpixel LambdaVision
NSF MRI

Calit2@UCI 200 Mpixel HiPerWall
NSF MRI

TACC 307 Mpixel Stallion
NSF TeraGrid

OptIPuter

NSF

A Decade of NSF Investment
Two Orders of Magnitude Growth!

it²

The OptIPuter Creates an OptIPlanet Collaboratory: Enabling Data-Intensive e-Research

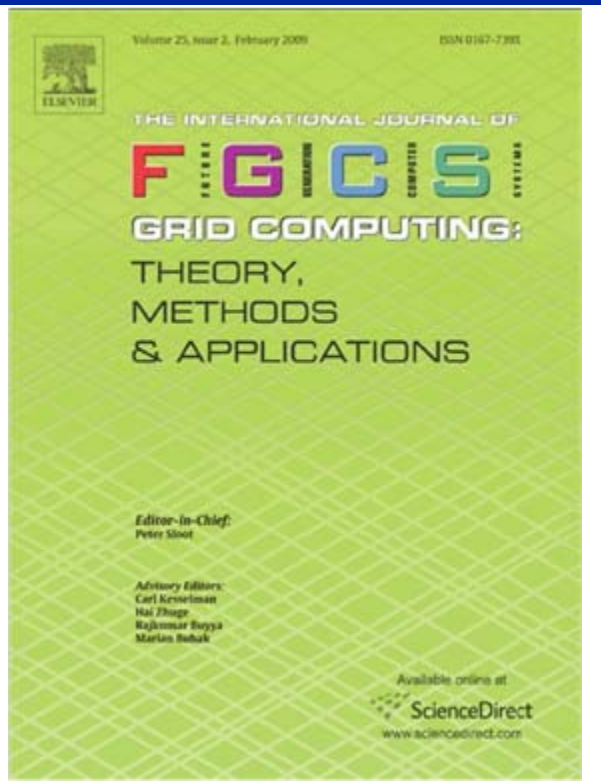

"OptIPlanet: The OptIPuter Global Collaboratory" –
Special Section of Future Generations Computer Systems, Volume 25, Issue 2, February 2009

evl electronic visualization laboratory

NSF

ROCKS

SAGE

OptIPuter

Calit2 (UCSD, UCI), SDSC, and UIC Leads—Larry Smarr PI
Univ. Partners: NCSA, USC, SDSU, NW, TA&M, UvA, SARA, KISTI, AIST
Industry: IBM, Sun, Telcordia, Chiaro, Calient, Glimmerglass, Lucent

it²

My OptIPortal™ – Affordable Termination Device for the OptIPuter Global Backplane

- 20 Dual CPU Nodes, 20 24” Monitors, ~$50,000
- 1/4 Teraflop, 5 Terabyte Storage, 45 Mega Pixels--Nice PC!
- Scalable Adaptive Graphics Environment (SAGE) Jason Leigh, EVL-UIC

evl electronic visualization laboratory

NSF

OptIPuter

Powered by Rocks ROCKS

Source: Phil Papadopoulos SDSC, Calit2

it²

Rocks Cluster Toolkit is the OptIPortal Foundation and its Software Deployment System

- **Enables Scientists in *Their* Lab to Easily Reap the Benefits of Cluster Computing**
- **Rocks is Extensible through Rolls to Build Reproducible Specialized Computing Facilities**
  - **OptIPortal Windowing and Graphics**
  - **Virtual Machines for Cluster Extension**
- **Large User Community**
  - **Volunteer Registry of Installed Clusters > 1100 Clusters, 100K CPUS 600+TF**
  - **Active User Support List of 1800+ Users**
  - **Used at Universities Throughout the US.**
  - **Impacts NSF, NIH, DoE Funded Research**
- **Supports RedHat Enterprise Linux et al (CentOS, Scientific Linux)**
  - **Alpha Support for Solaris**

Funded NSF SDCI Program award #OCI-0721623 Sept 2007

Use of OptIPortal to Interactively View Multi-Scale Biomedical Imaging

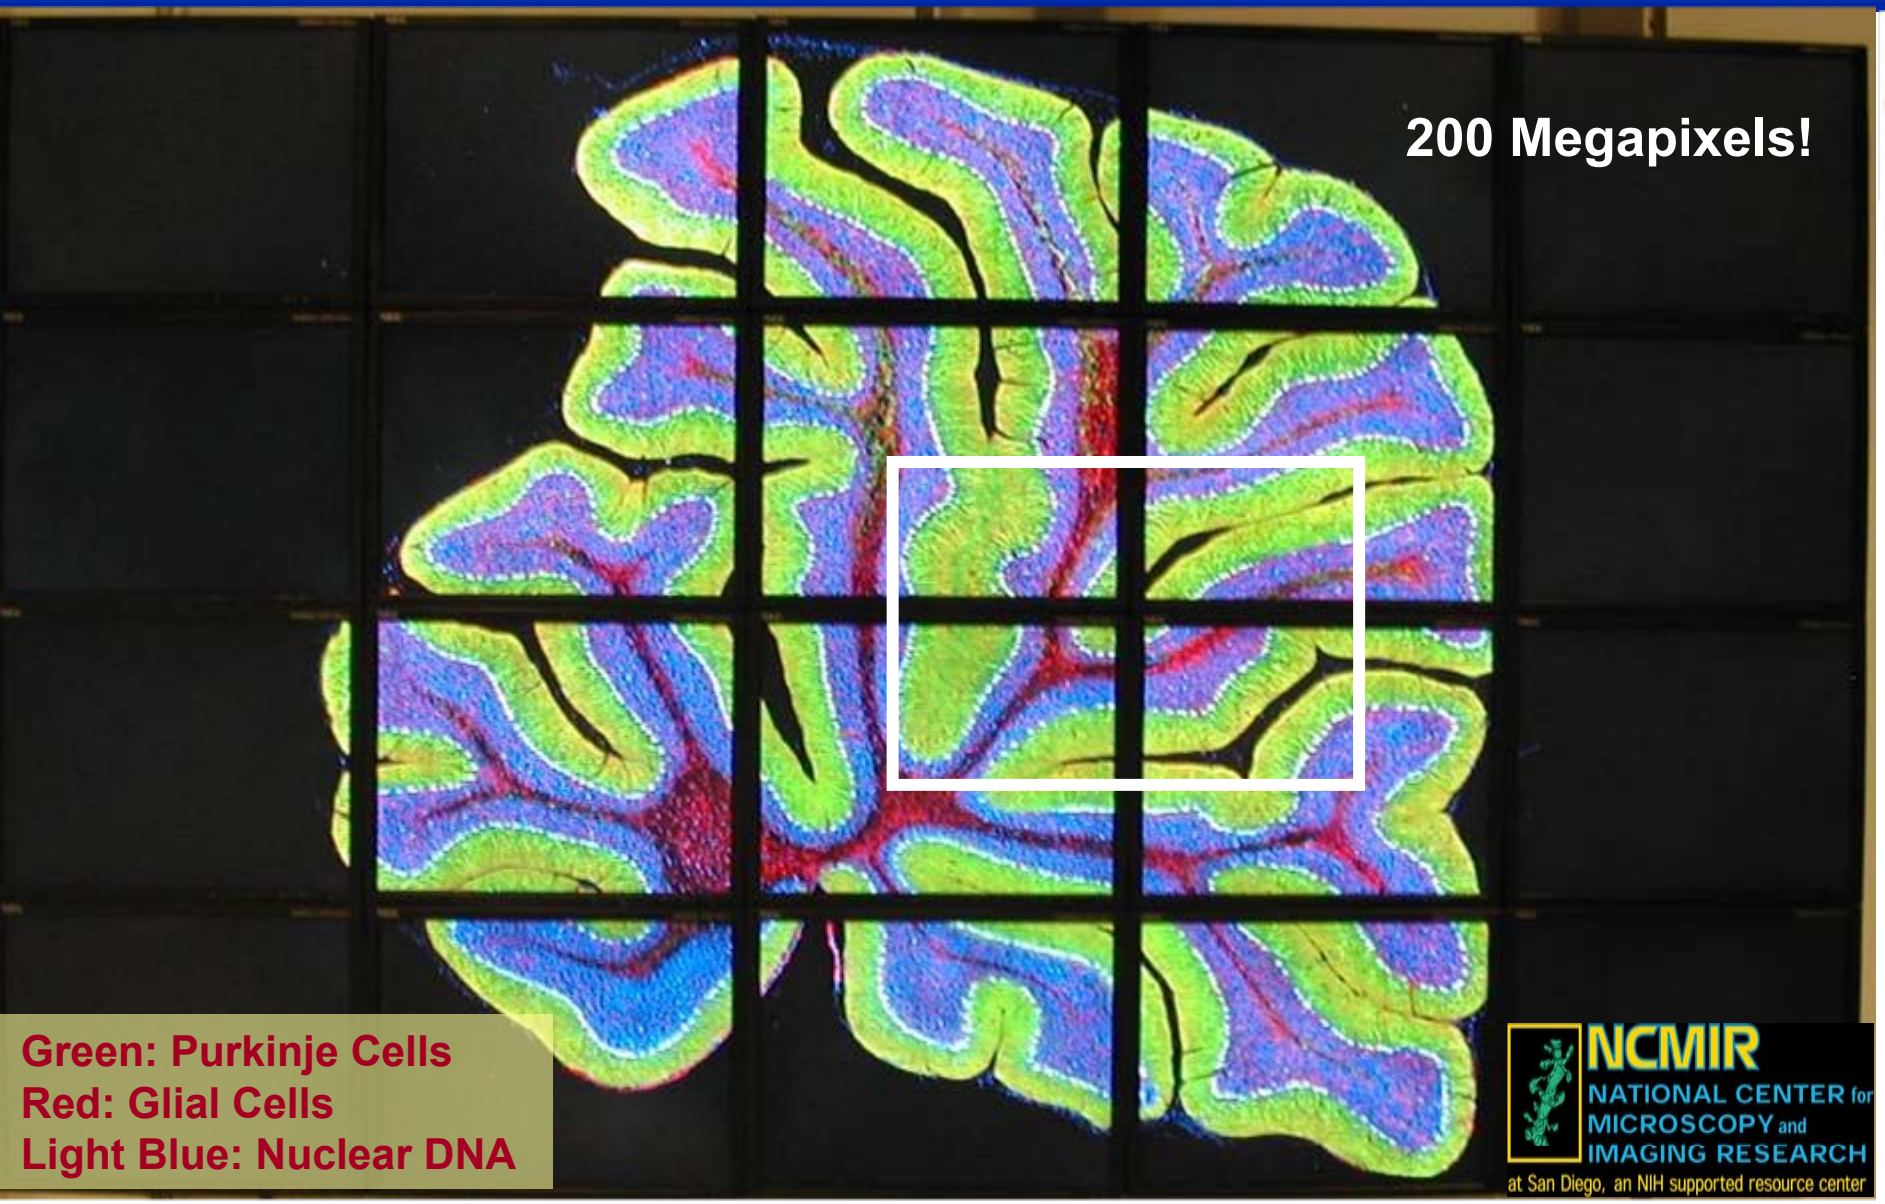

Source: Mark Ellisman, David Lee, Jason Leigh

Two-Photon Laser Confocal Microscope Montage of 40x36=1440 Images in 3 Channels of a Mid-Sagittal Section of Rat Cerebellum Acquired Over an 8-hour Period

Scalable Displays Allow Both Global Content and Fine Detail

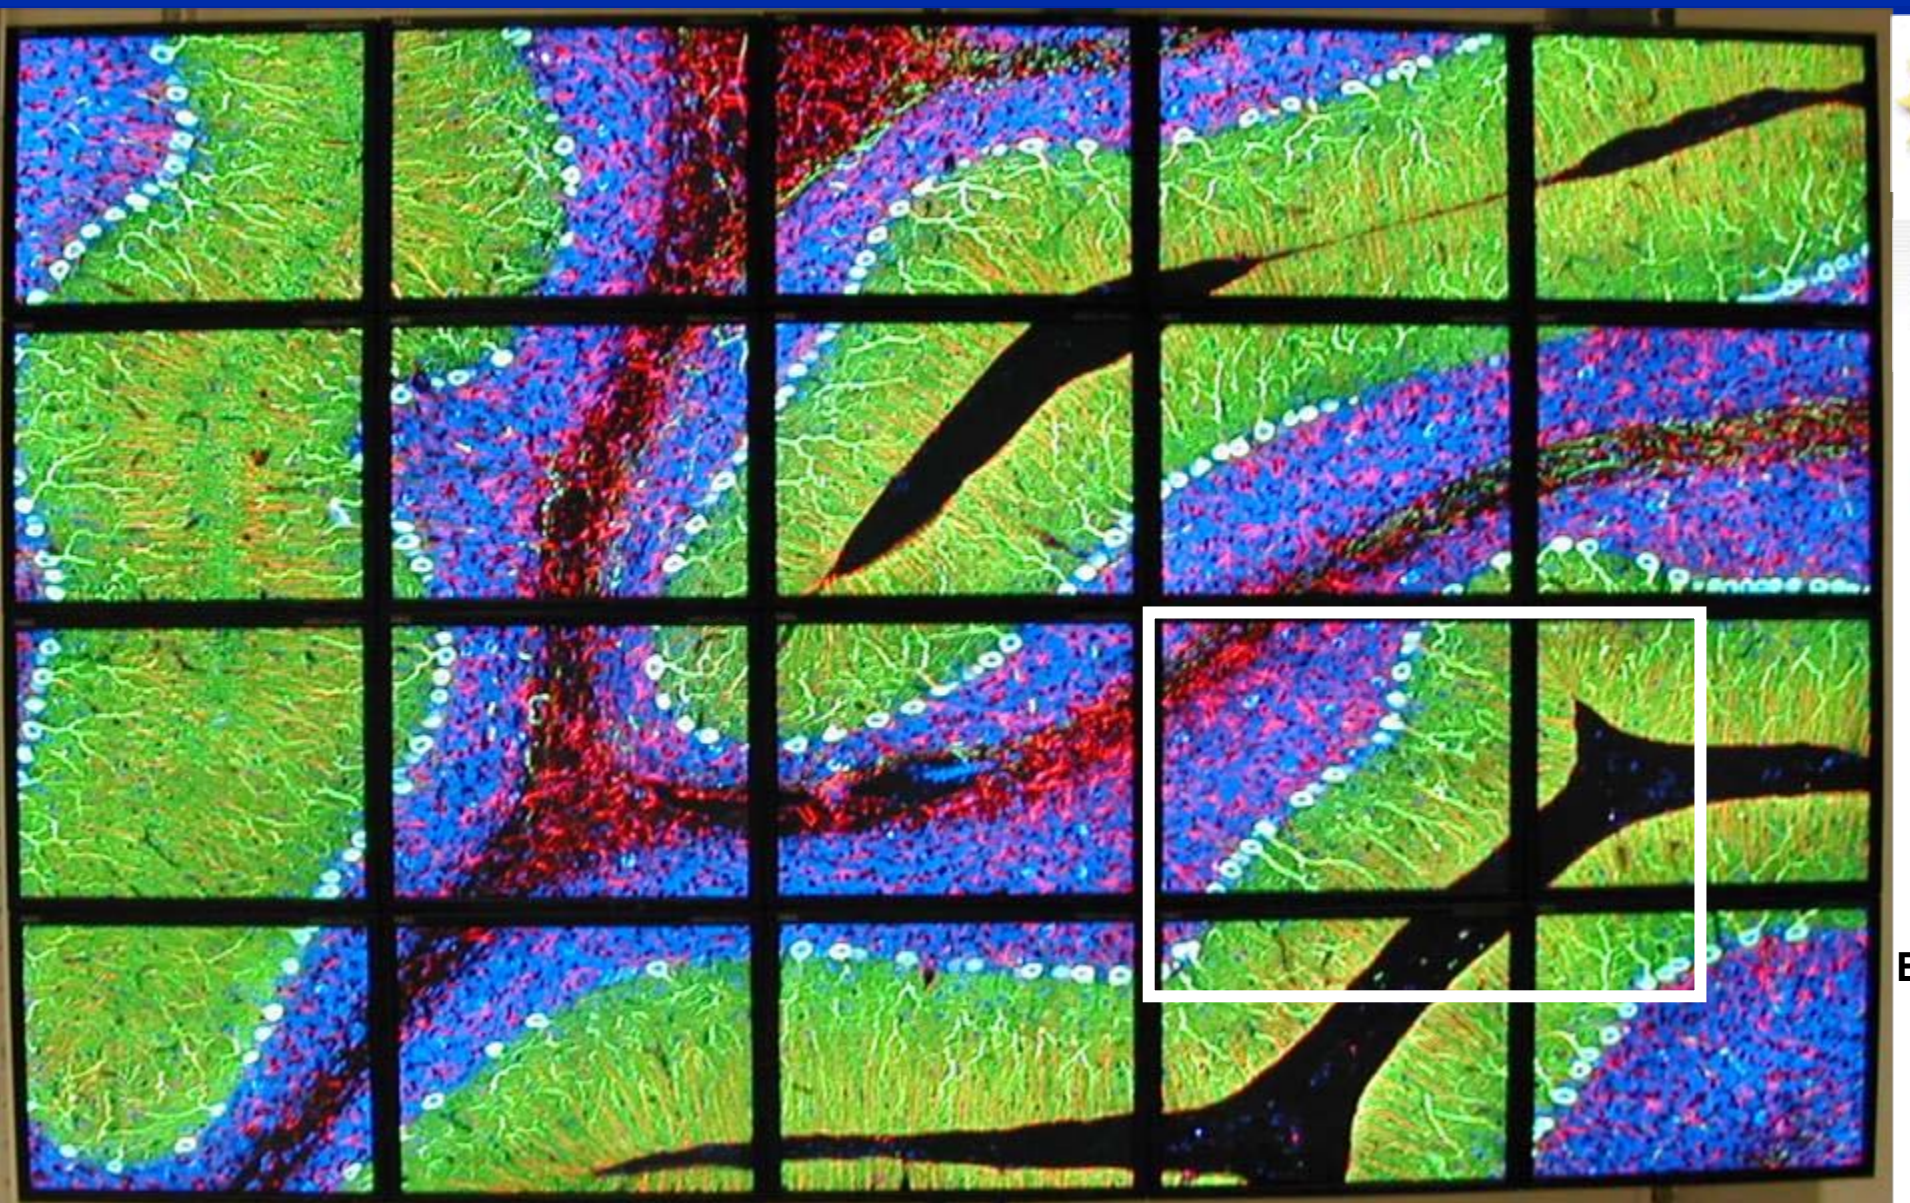

Source: Mark Ellisman, David Lee, Jason Leigh

OptIPuter

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

it²

Allows for Interactive Zooming from Cerebellum to Individual Neurons

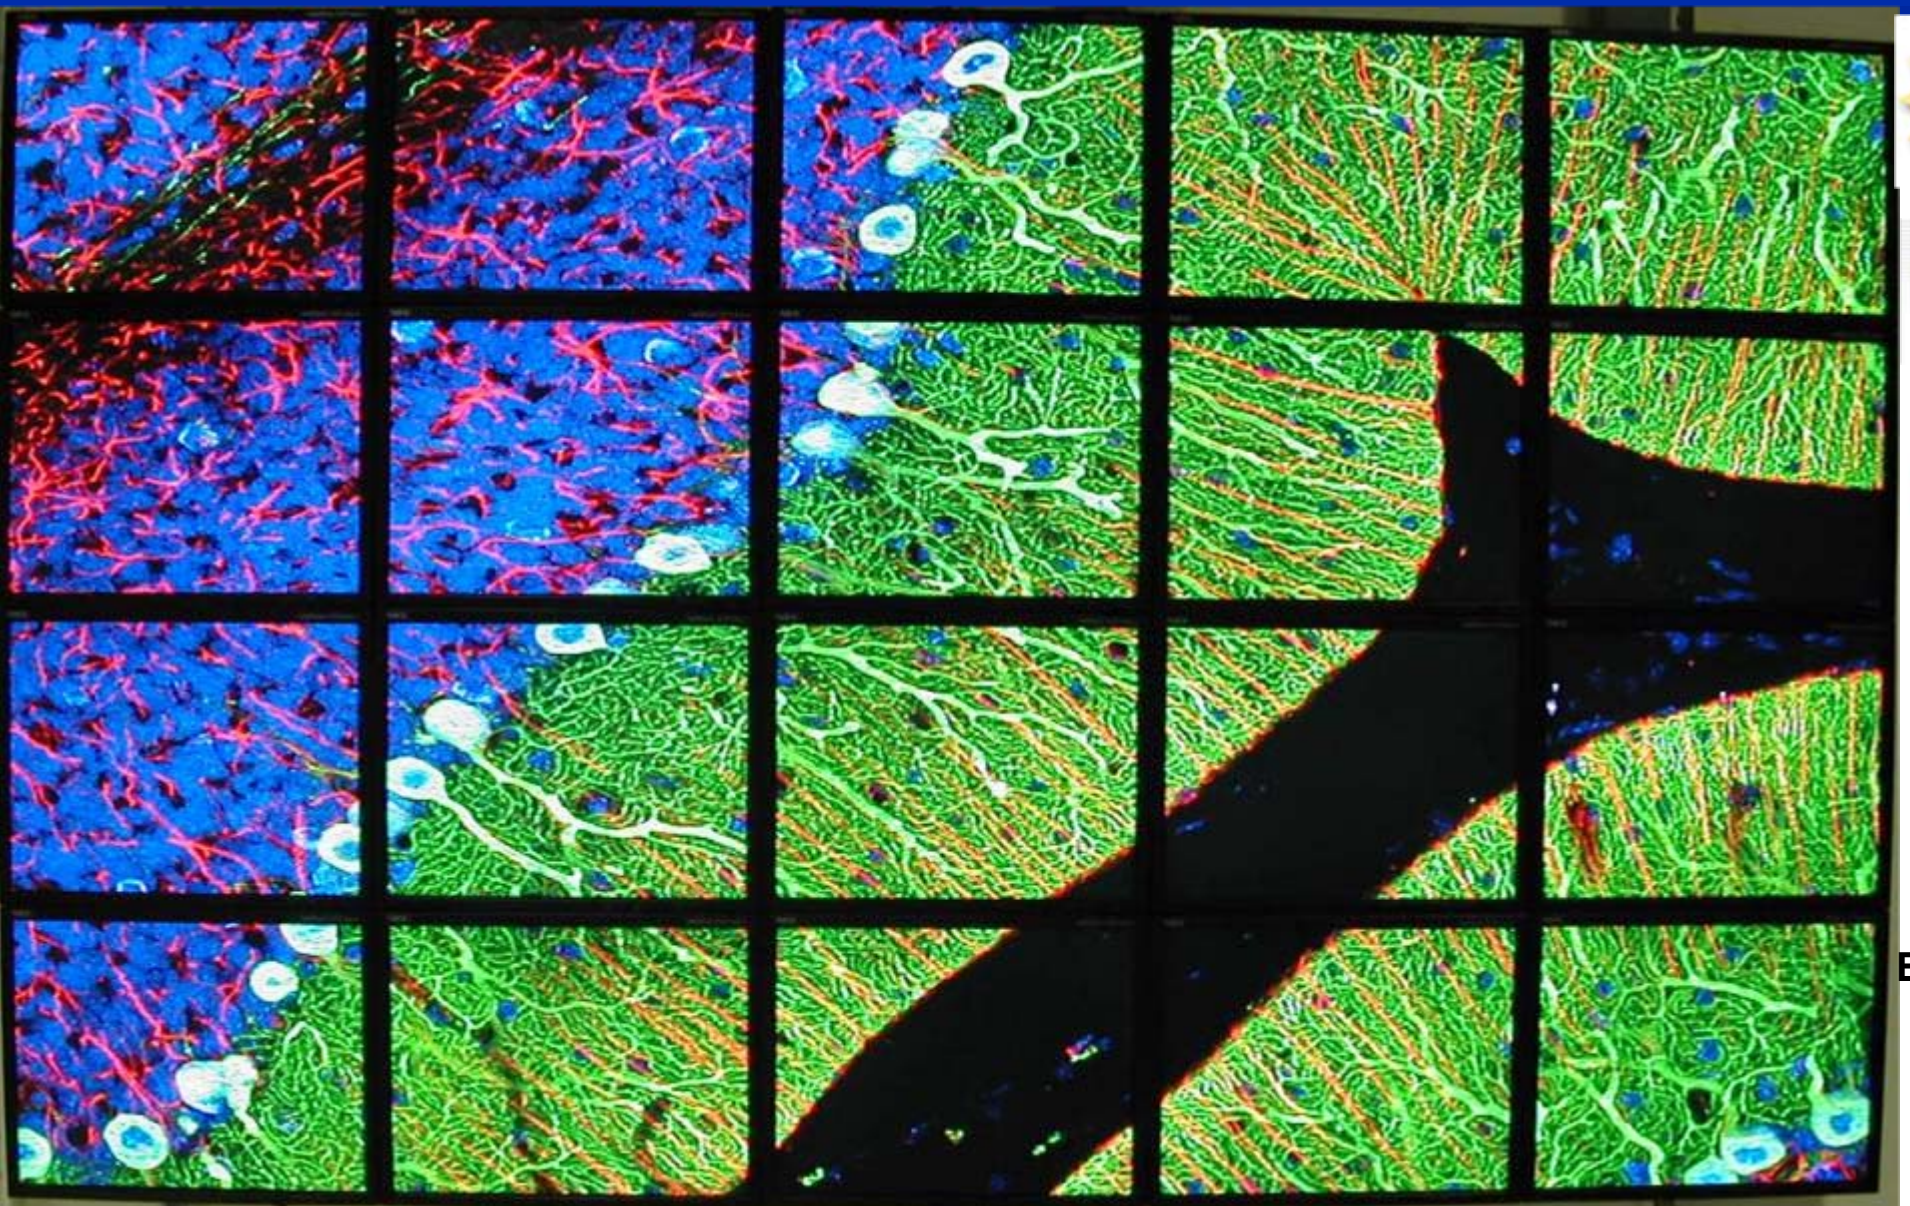

Source: Mark Ellisman, David Lee, Jason Leigh

OptIPuter

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

it²

OptIPortals Scale to 1/3 Billion Pixels Enabling Viewing of Very Large Images or Many Simultaneous Images

Spitzer Space Telescope (Infrared)

NASA Earth Satellite Images Bushfires October 2007 San Diego

Source: Falko Kuester, Calit2@UCSD

Students Learn Case Studies in the Context of Diverse Medical Evidence

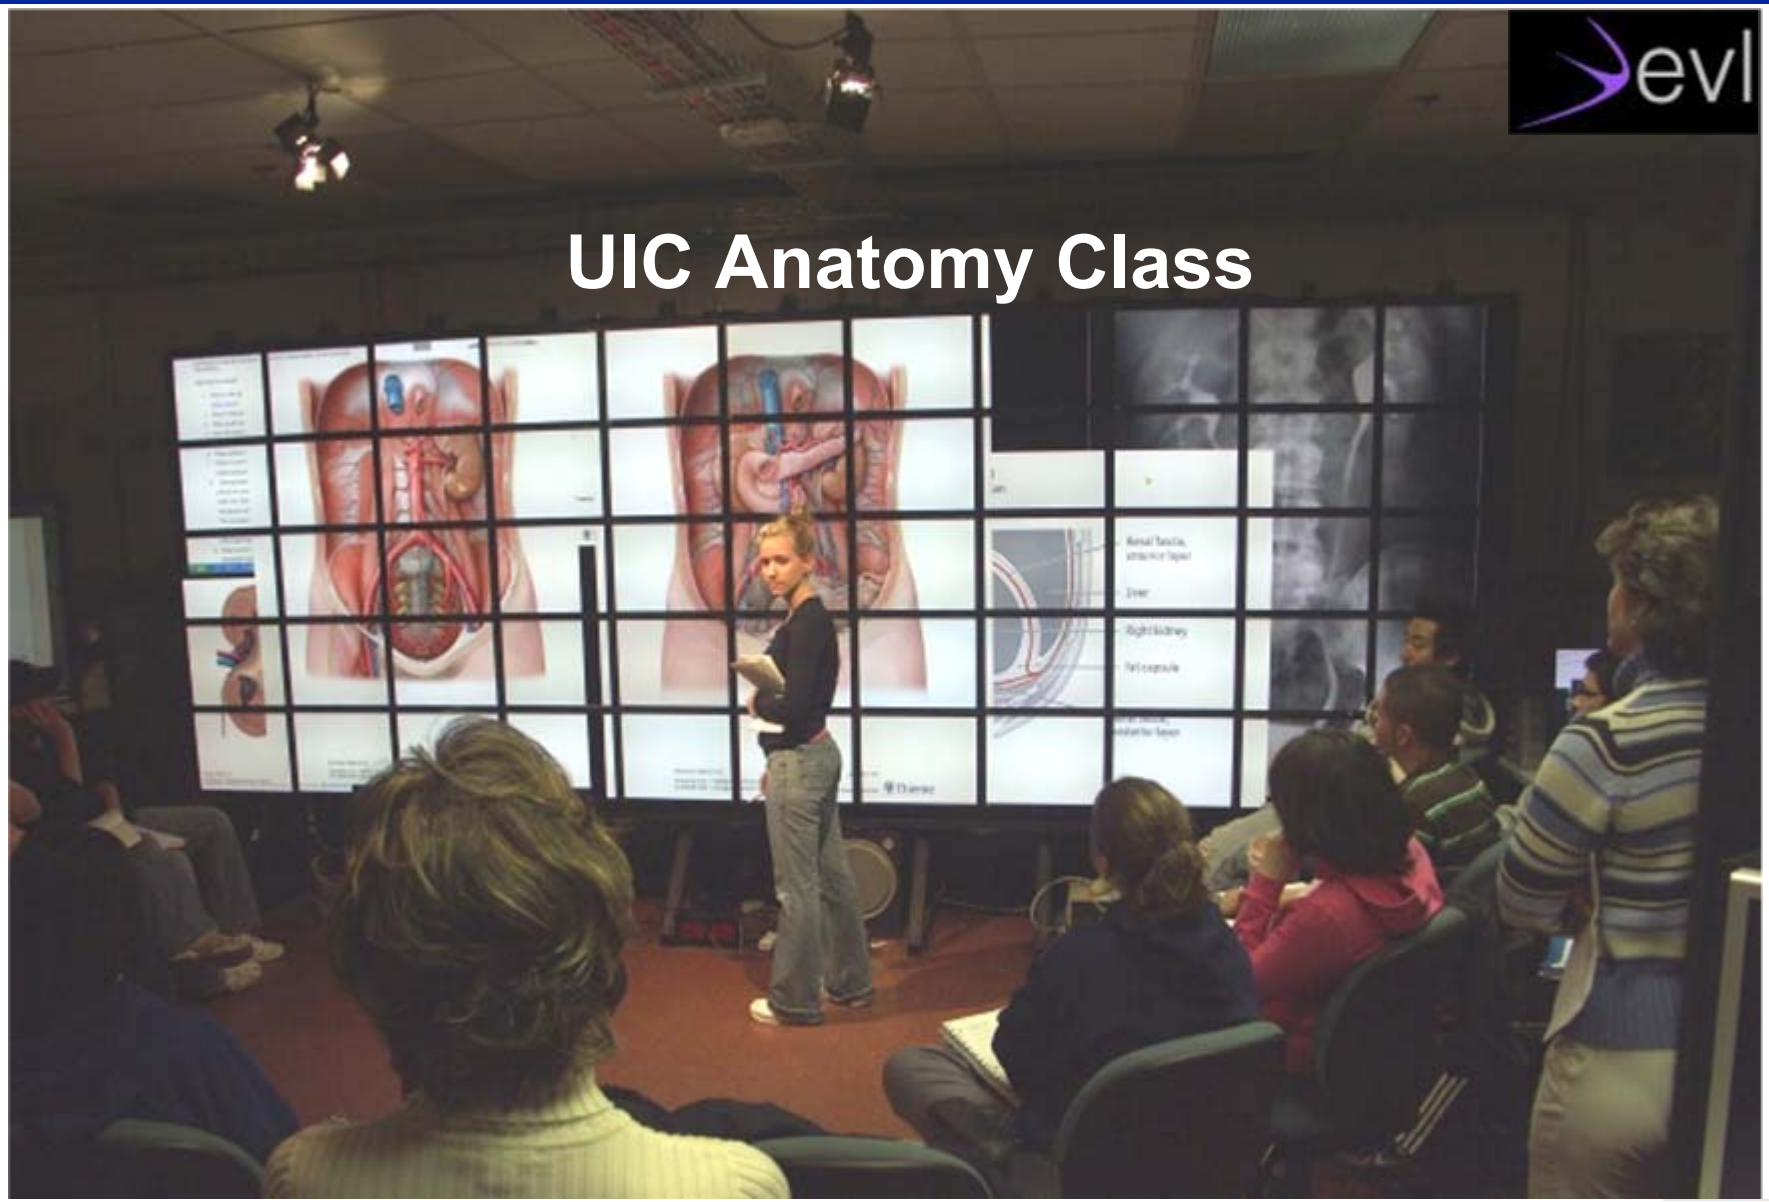

On-Line Resources
Help You Build Your Own OptIPortal

www.optiputer.net

http://wiki.optiputer.net/optiportal

www.evl.uic.edu/cavern/sage

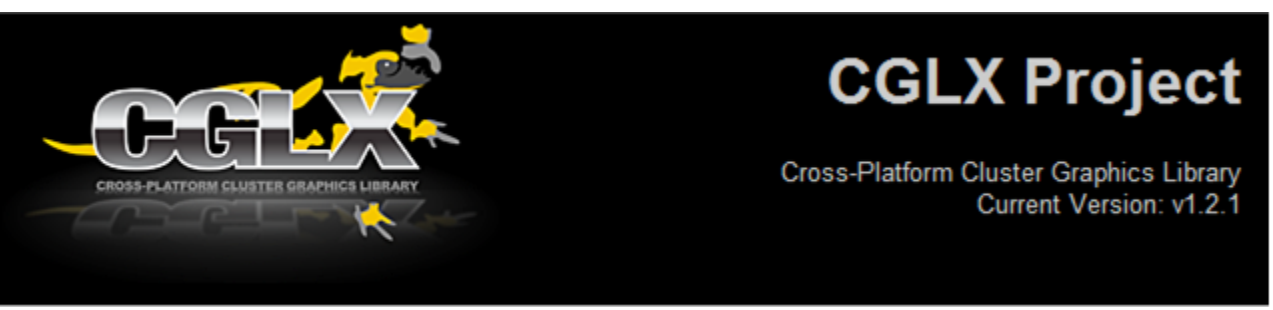

http://vis.ucsd.edu/~cglx/

OptIPuter

it²

Calit2 3D Immersive StarCAVE OptIPortal: Enables Exploration of High Resolution Simulations

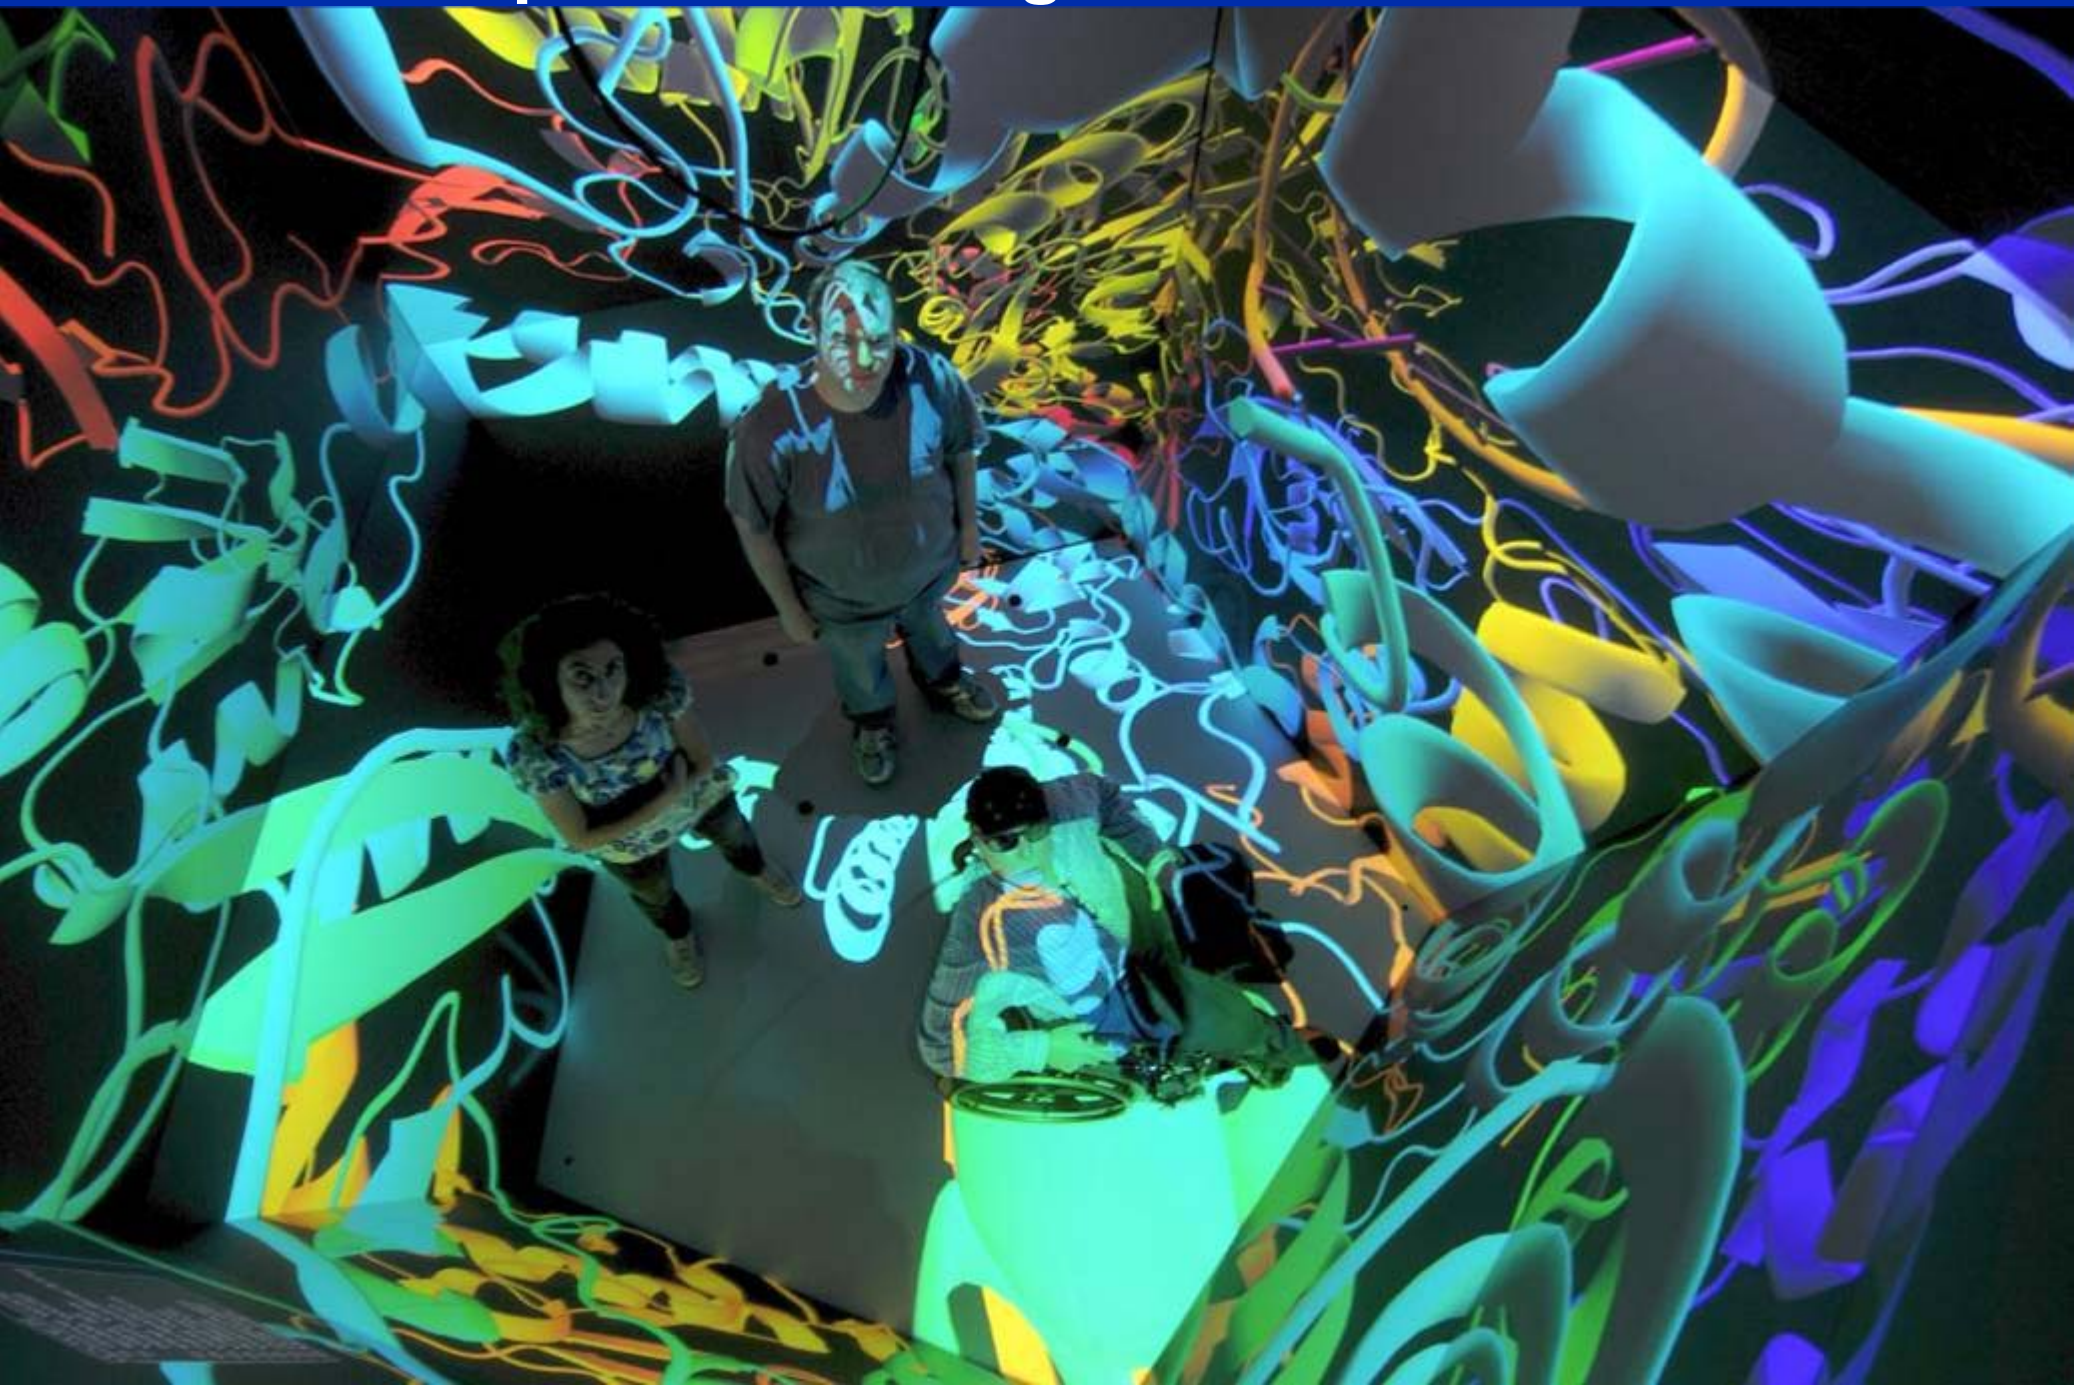

Cyberinfrastructure Integration: Integration of Data Generators, Transmission, and Portals

OptIPuter

it²

The Calit2 OptIPortals at UCSD and UCI Are Now a Gbit/s HD Collaboratory

UCSD cluster: 15 x Quad core Dell XPS with Dual nVIDIA 5600s
UCI cluster: 25 x Dual Core Apple G5

OptIPuter Software Architecture--a Service-Oriented Architecture Integrating Lambdas Into the Grid

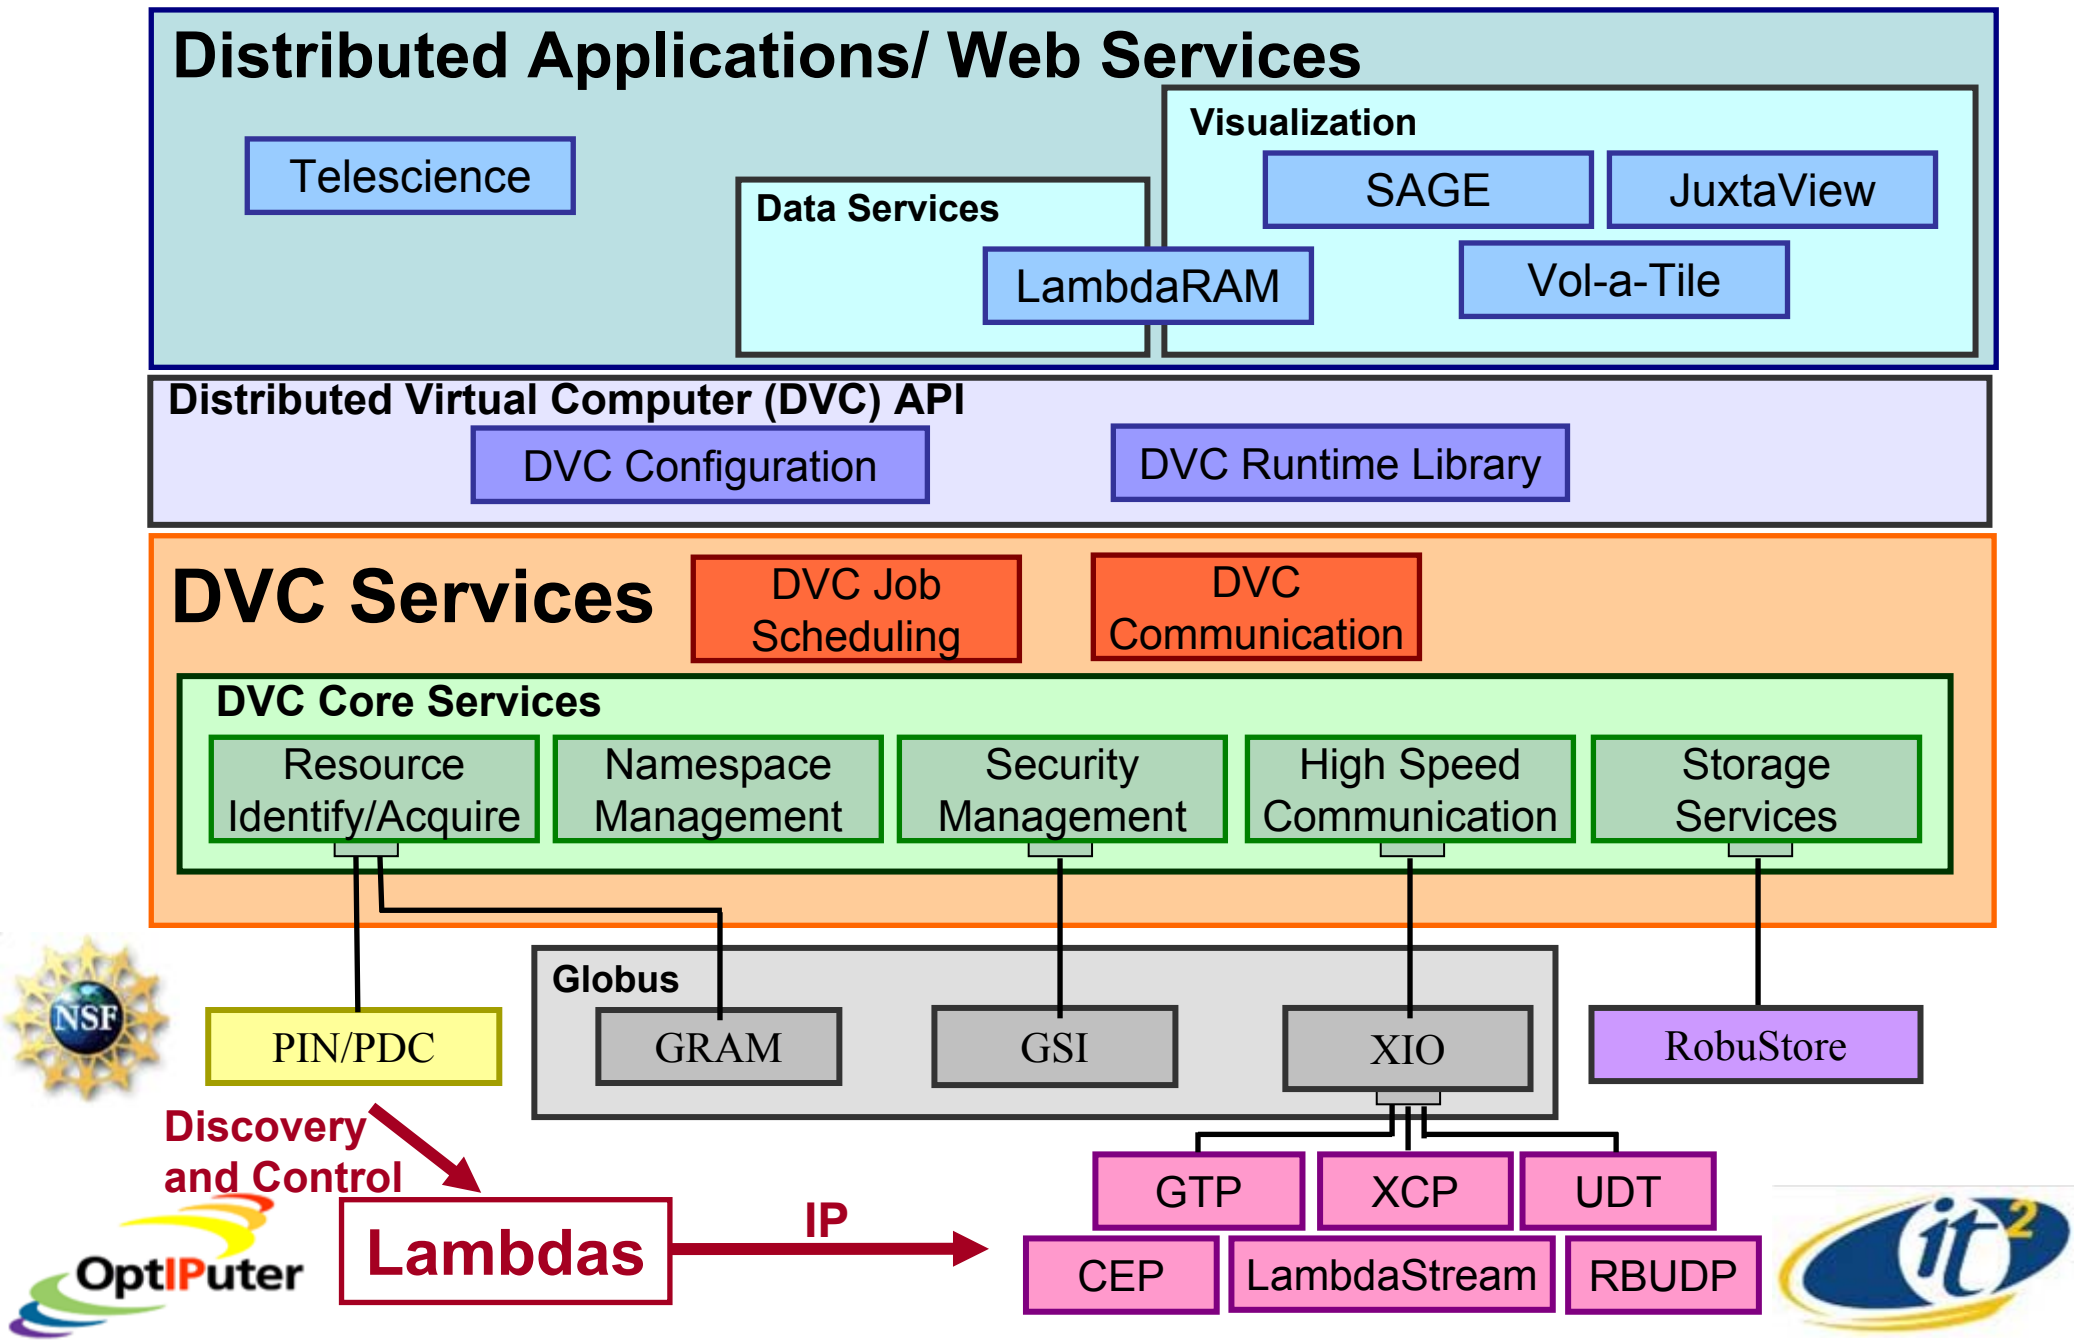

Calit2 is Creating a Global CI to Support Marine Microbial Metagenomics

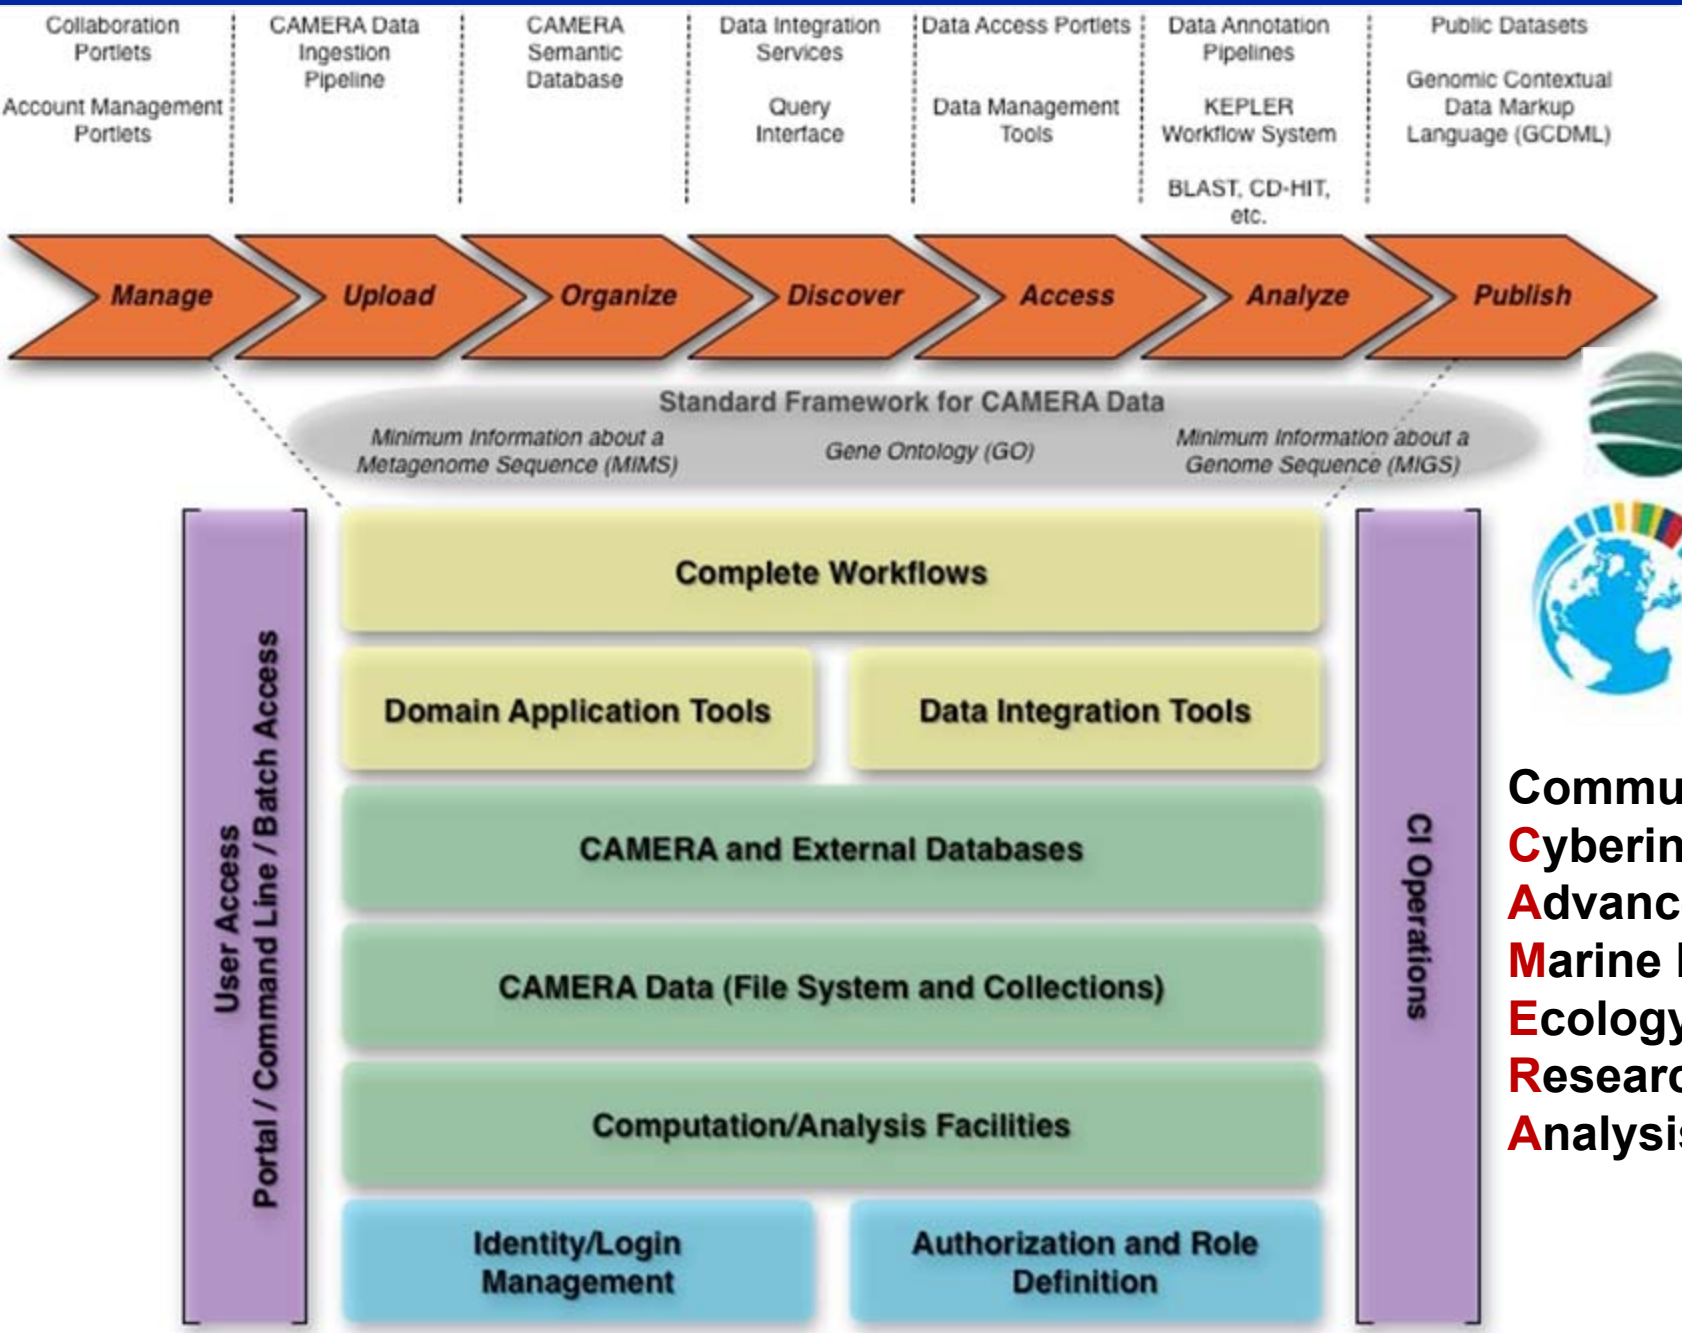

Source: CAMERA CTO Mark Ellisman

Gordon and Betty MOORE FOUNDATION

camera
Marine Microbial Ecology

Community Cyberinfrastructure for Advanced Marine Microbial Ecology Research and Analysis

it²

Calit2 Microbial Metagenomics Cluster- Next Generation Optically Linked Science Data Server

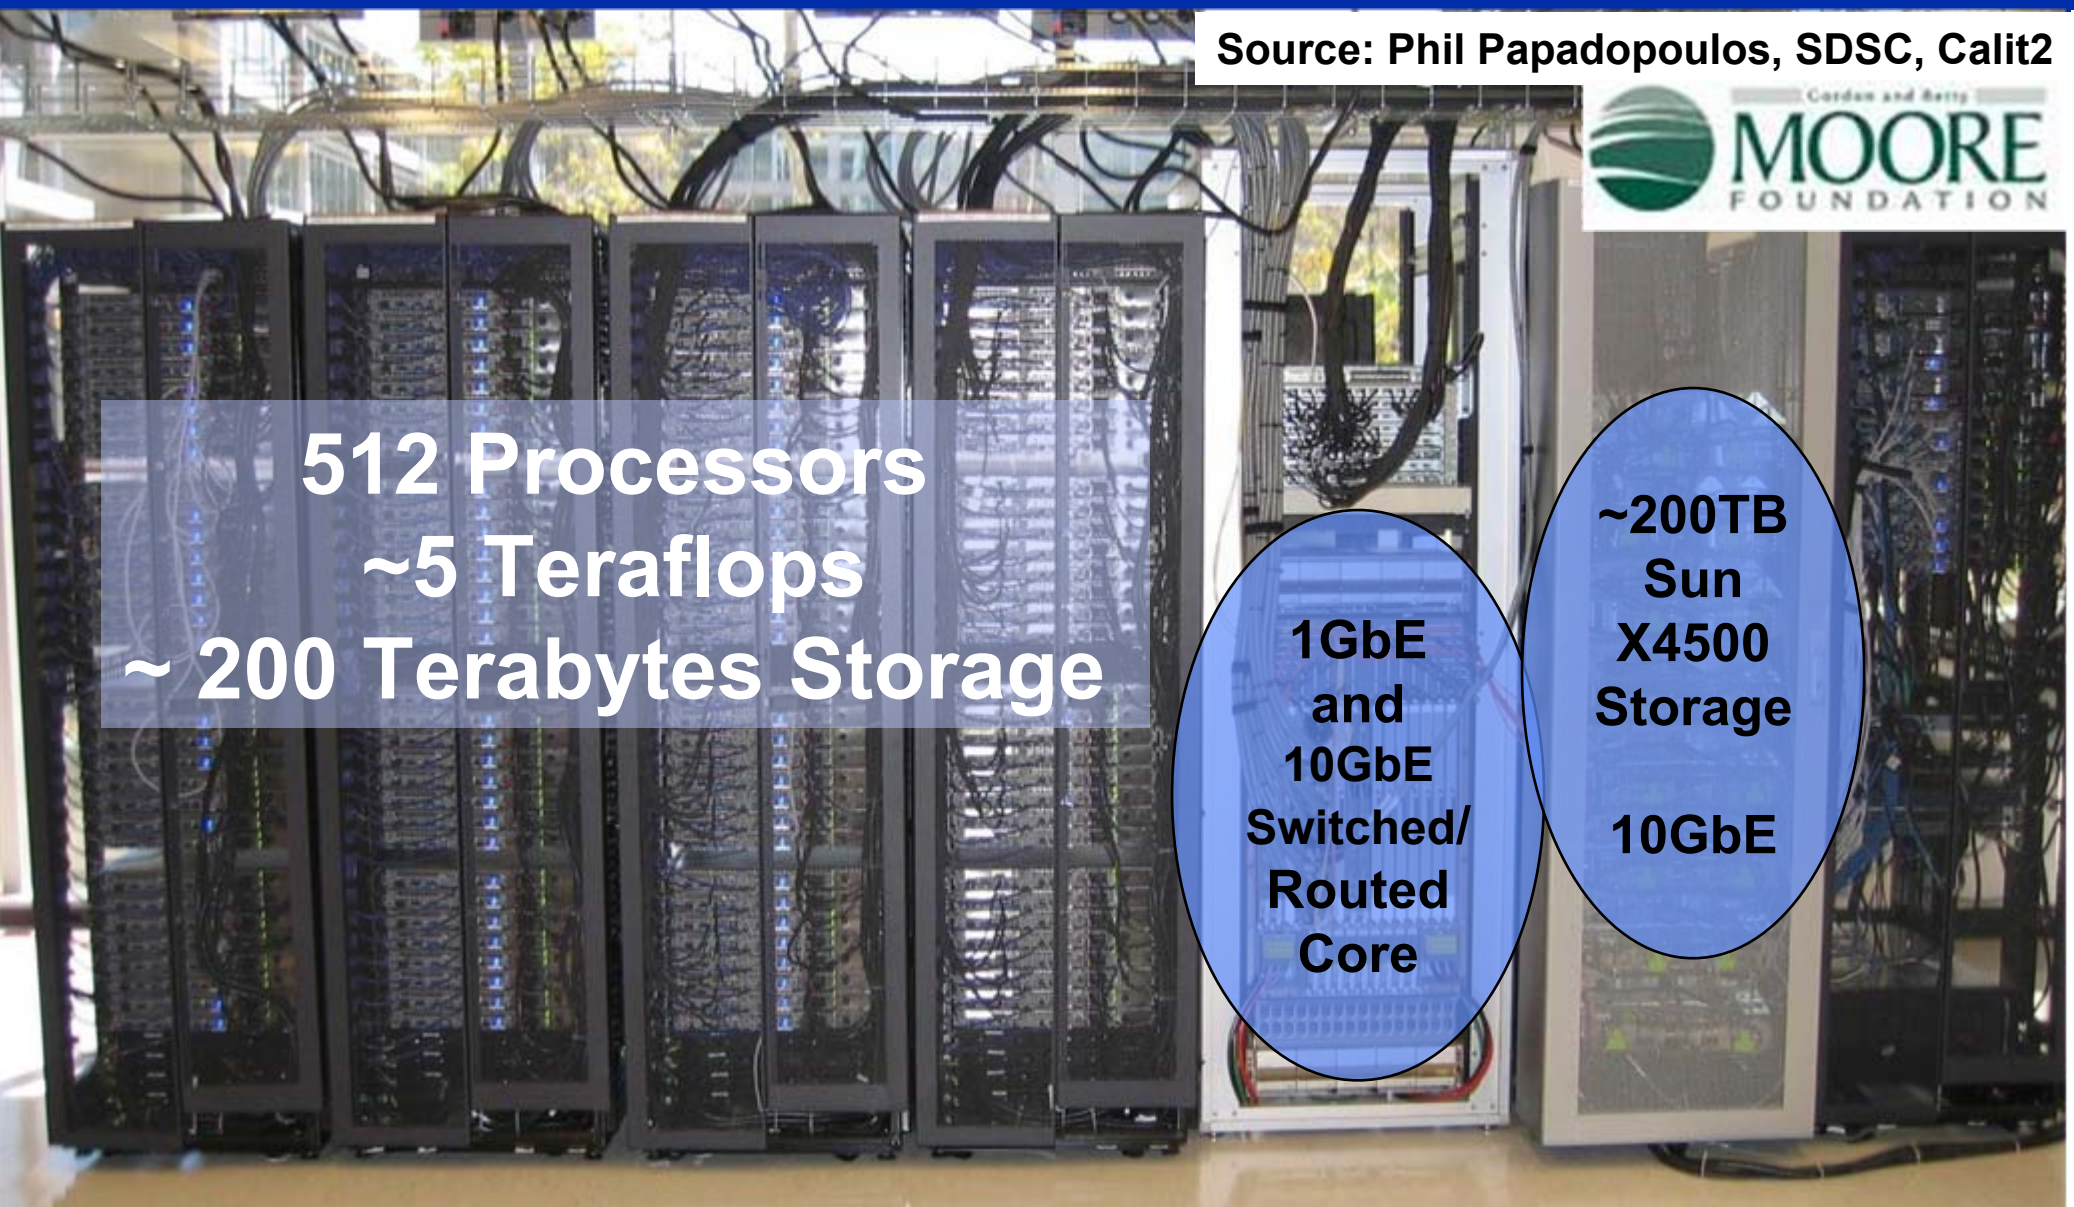

OptIPuter SDSC ROCKS camera Marine Microbial Ecology J. Craig Venter INSTITUTE it²

CAMERA's
Global Microbial Metagenomics CyberCommunity

Nearly 2500 Registered Users
From 55 Countries

OptIPuter Persistent Infrastructure Enables Calit2 and U Washington CAMERA Collaboratory

OptIPuter

Ginger Armbrust's Diatoms: Micrographs, Chromosomes, Genetic Assembly

ROCKS

Gordon and Betty MOORE FOUNDATION

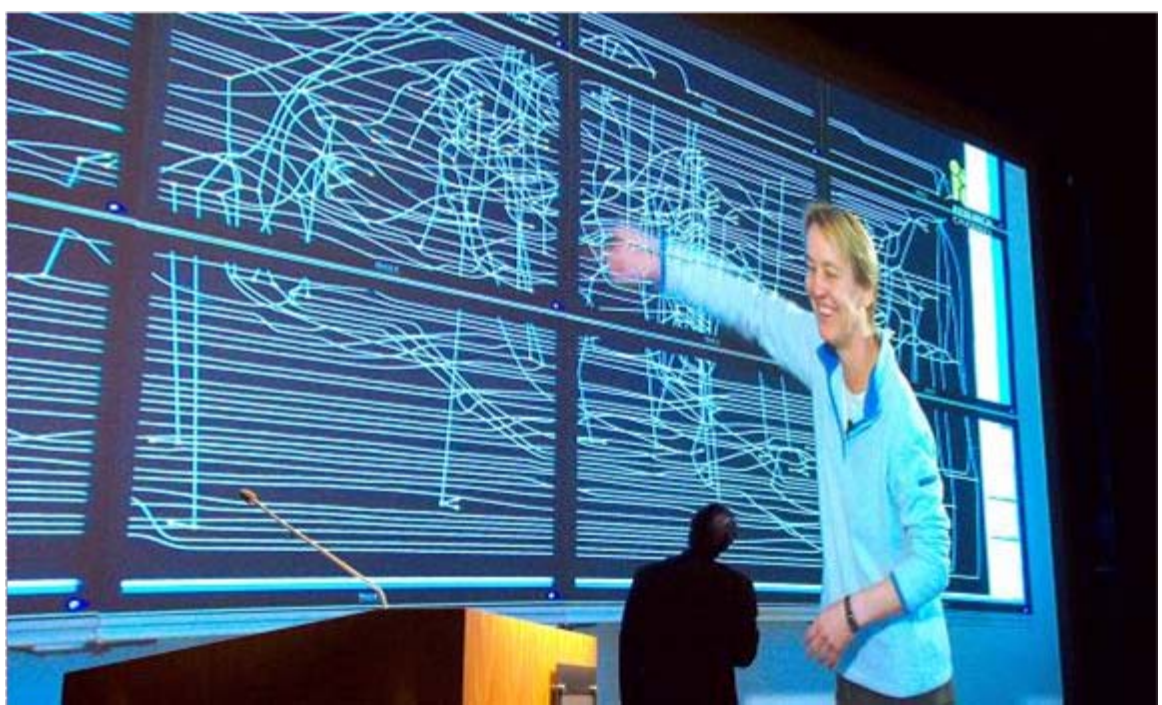

NSF

camera
Marine Microbial Ecology

OptIPuter

it²

Remote Control of Scientific Instruments: Live Session with JPL and Mars Rover from Calit2

OptIPuter

CGLX
CROSS-PLATFORM CLUSTER GRAPHICS LIBRARY

Powered by Rocks ROCKS

it²

Source: Falko Kuester, Calit2; Michael Sims, NASA

U Michigan Virtual Space Interaction Testbed (VISIT) Instrumenting OptIPortals for Social Science Research

- Using Cameras Embedded in the Seams of Tiled Displays and Computer Vision Techniques, we can Understand how People Interact with OptIPortals
  - Classify Attention, Expression, Gaze
  - Initial Implementation Based on Attention Interaction Design Toolkit (J. Lee, MIT)
- Close to Producing Usable Eye/Nose Tracking Data using OpenCV

Leading U.S. Researchers on the Social Aspects of Collaboration

Source: Erik Hofer, UMich, School of Information

Creating Environmental Observatories – Combining OptIPortals with DataTurbine

- **Streaming Data Middleware System Satisfies Common Critical Infrastructure Requirements Across NSF-Sponsored Observing Systems:**
  - **Framework for Integration of Heterogeneous Instruments with Reliable Real-Time Data Transport**
  - **Sensors /Sensor Streams Become First-Class Objects**
  - **Comprehensive Suite Of Services For Data Management, Routing, Synchronization, Monitoring, and Visualization**
- **All-Software Solution (Java)**
  - **Open Source Software- Apache 2.0 License**
- **Used In NSF, NASA, NOAA, DOE Projects**
  - **www.dataturbine.org**
  - **NSF Support from SDCI program (Sept 07)**

Source: Tony Fountain, Sameer Tilak

OPEN SOURCE DATA TURBINE INITIATIVE

Empowering the Scientific Community with Streaming Data Middleware

OptIPuter · Gordon and Betty MOORE FOUNDATION · CREON · GLEON global lake ecological observatory network · PRAGMA · it²

Streaming Underwater Video From Taiwan's Kenting Reef to Calit2's OptIPortal

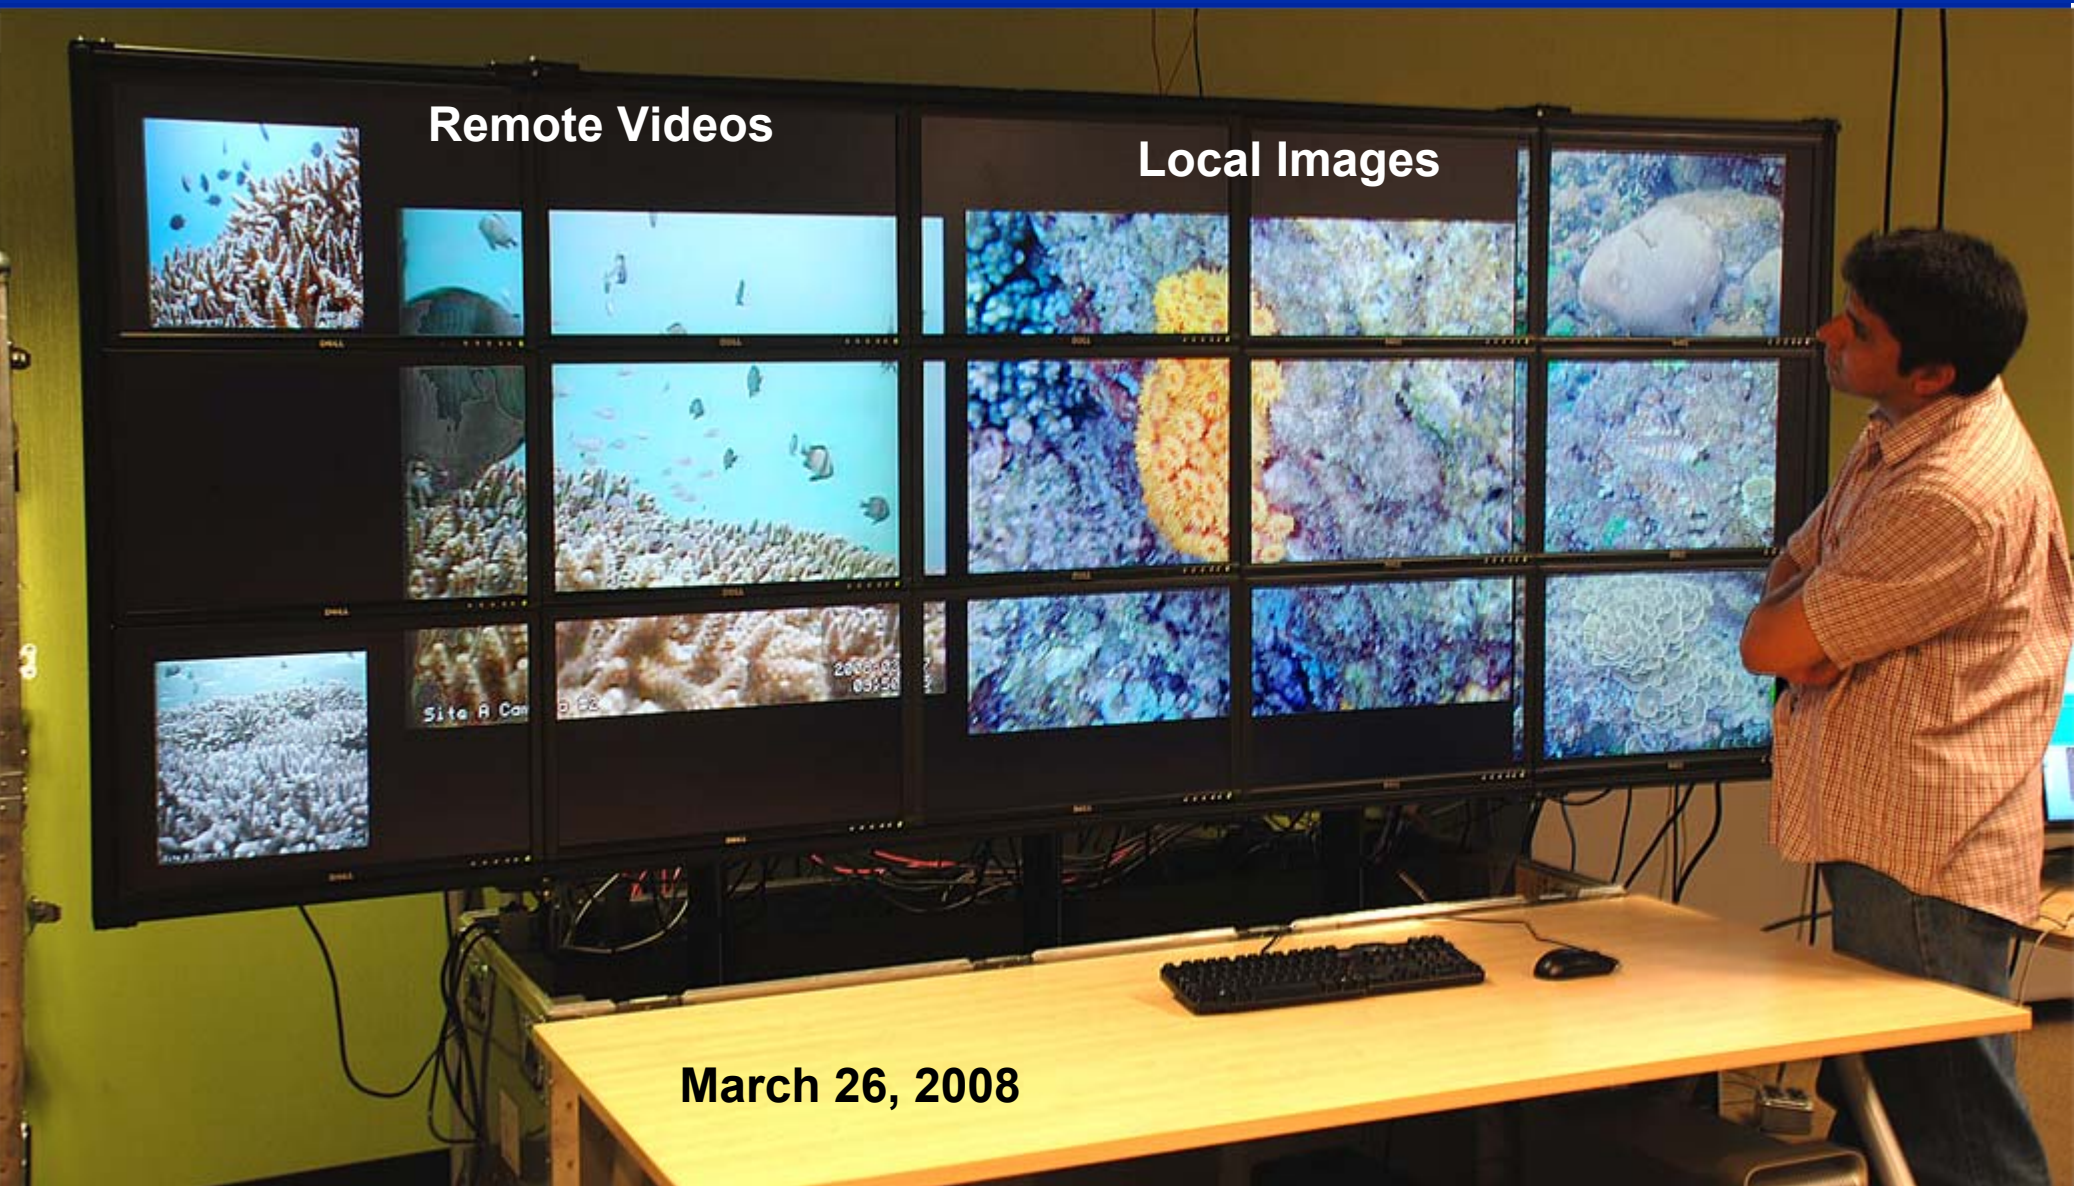

OPEN SOURCE DATA TURBINE INITIATIVE
Empowering the Scientific Community with Streaming Data Middleware

OptIPuter

UCSD: Rajvikram Singh, Sameer Tilak, Jurgen Schulze, Tony Fountain, Peter Arzberger
NCHC : Ebbe Strandell, Sun-In Lin, Yao-Tsung Wang, Fang-Pang Lin

it²

The Ocean Observatory Initiative (OOI) Cyberinfrastructure is Under Development

Source: John Orcutt, Matthew Arrott, SIO/Calit2

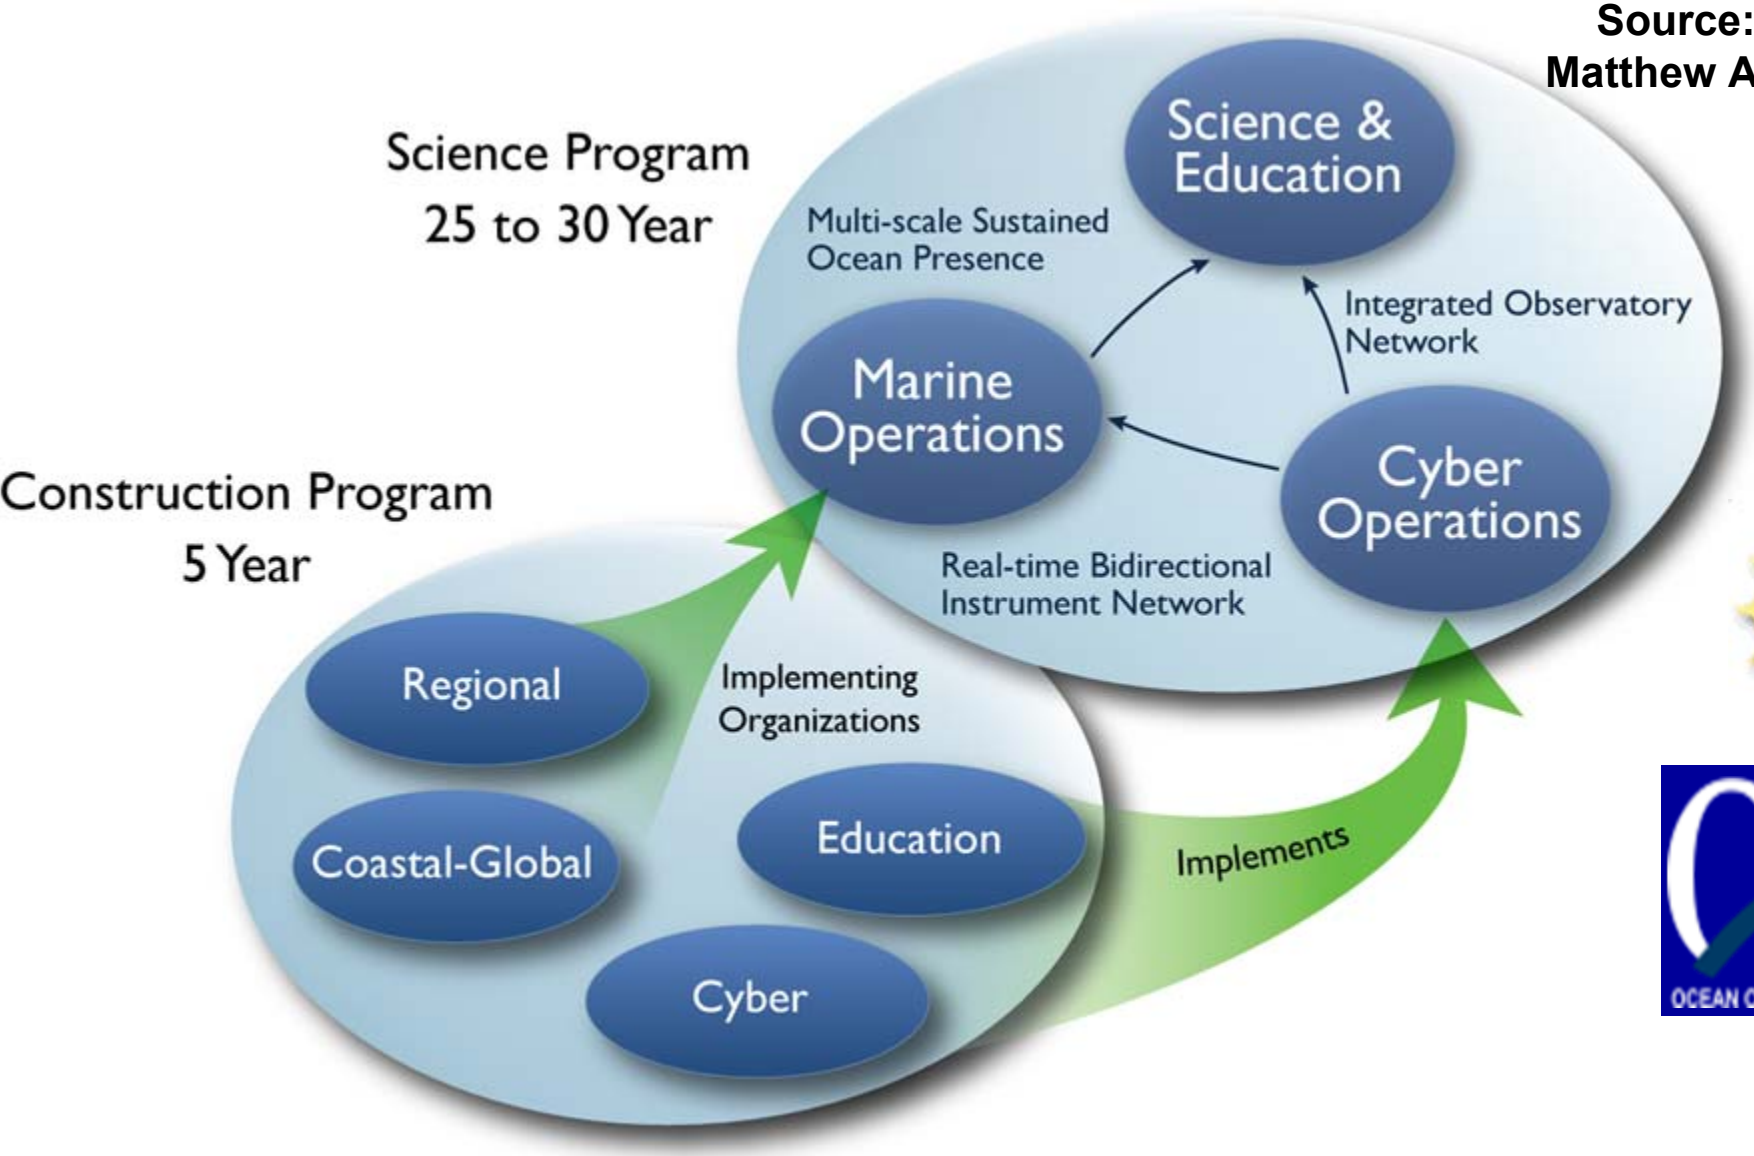

NSF

OCEAN OBSERVATORIES INITIATIVE

OptIPuter

OOI CI Executive Management Team: UCSD, WHOI, Rutgers

it2

OOI CI Prototype:
Tight Coupling Between Simulation and Experiment

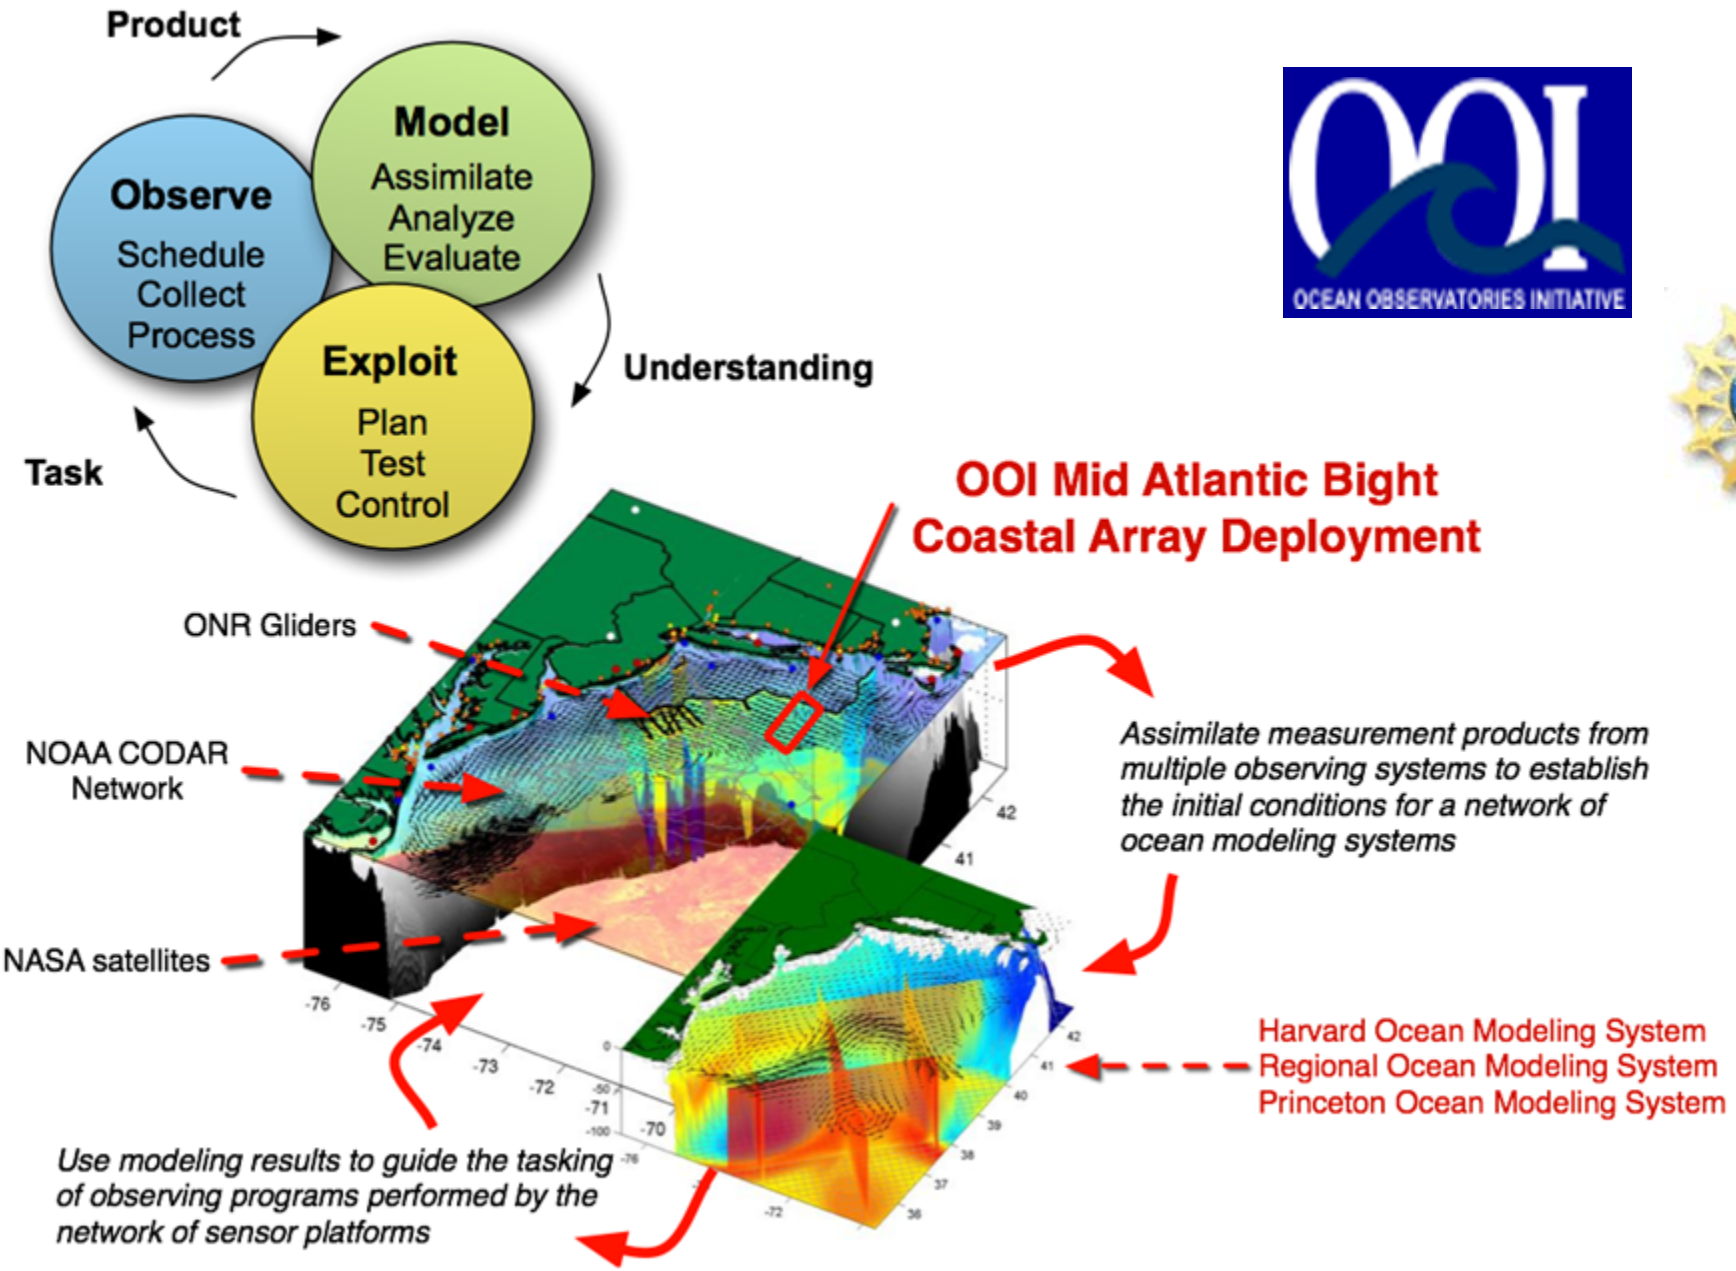

OOI CI
is Built on NLR/I2 Optical Infrastructure

OCEAN OBSERVATORIES INITIATIVE

Source: John Orcutt, Matthew Arrott, SIO/Calit2

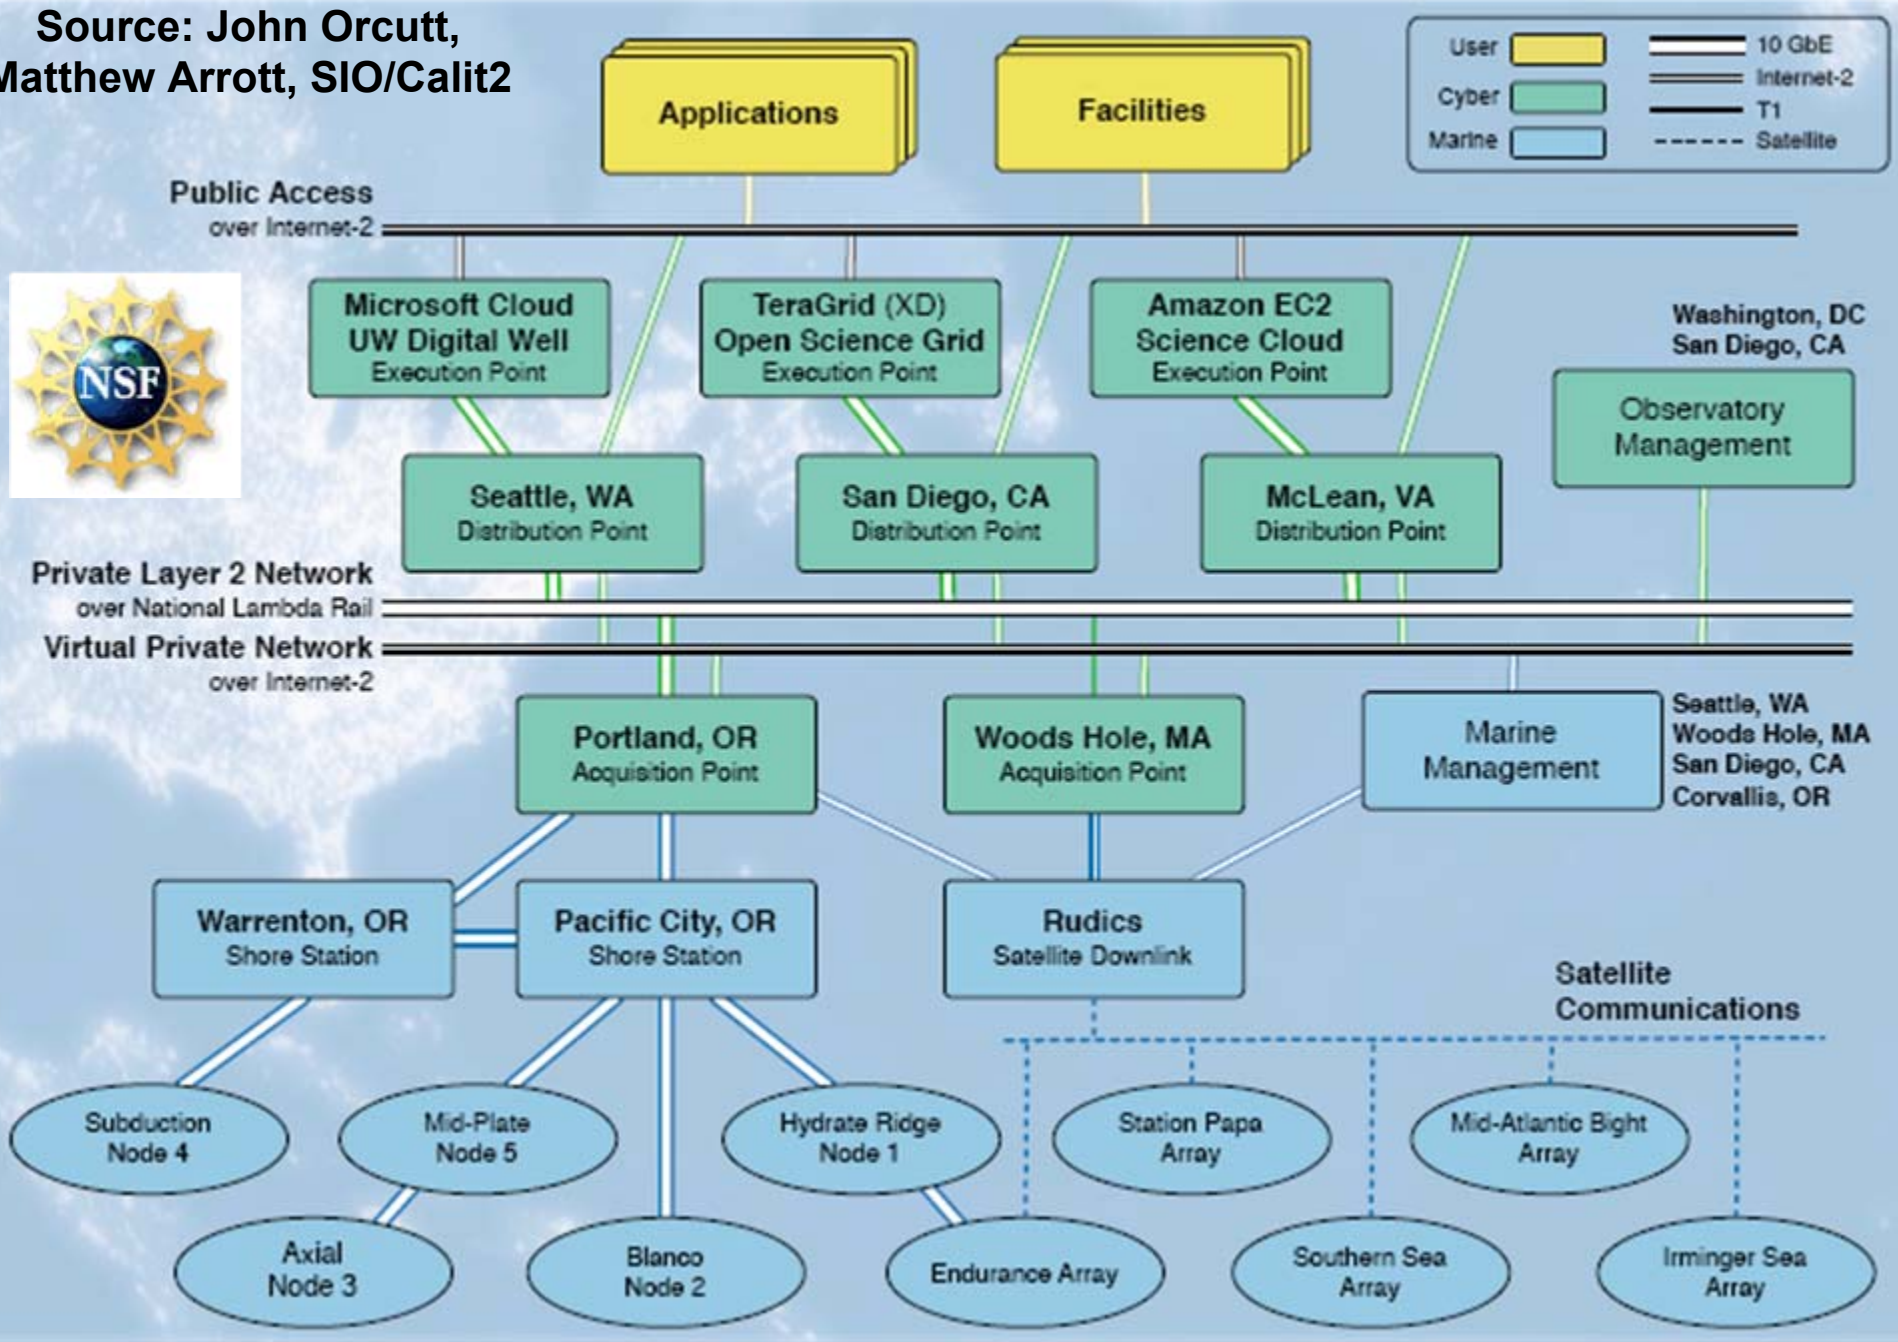

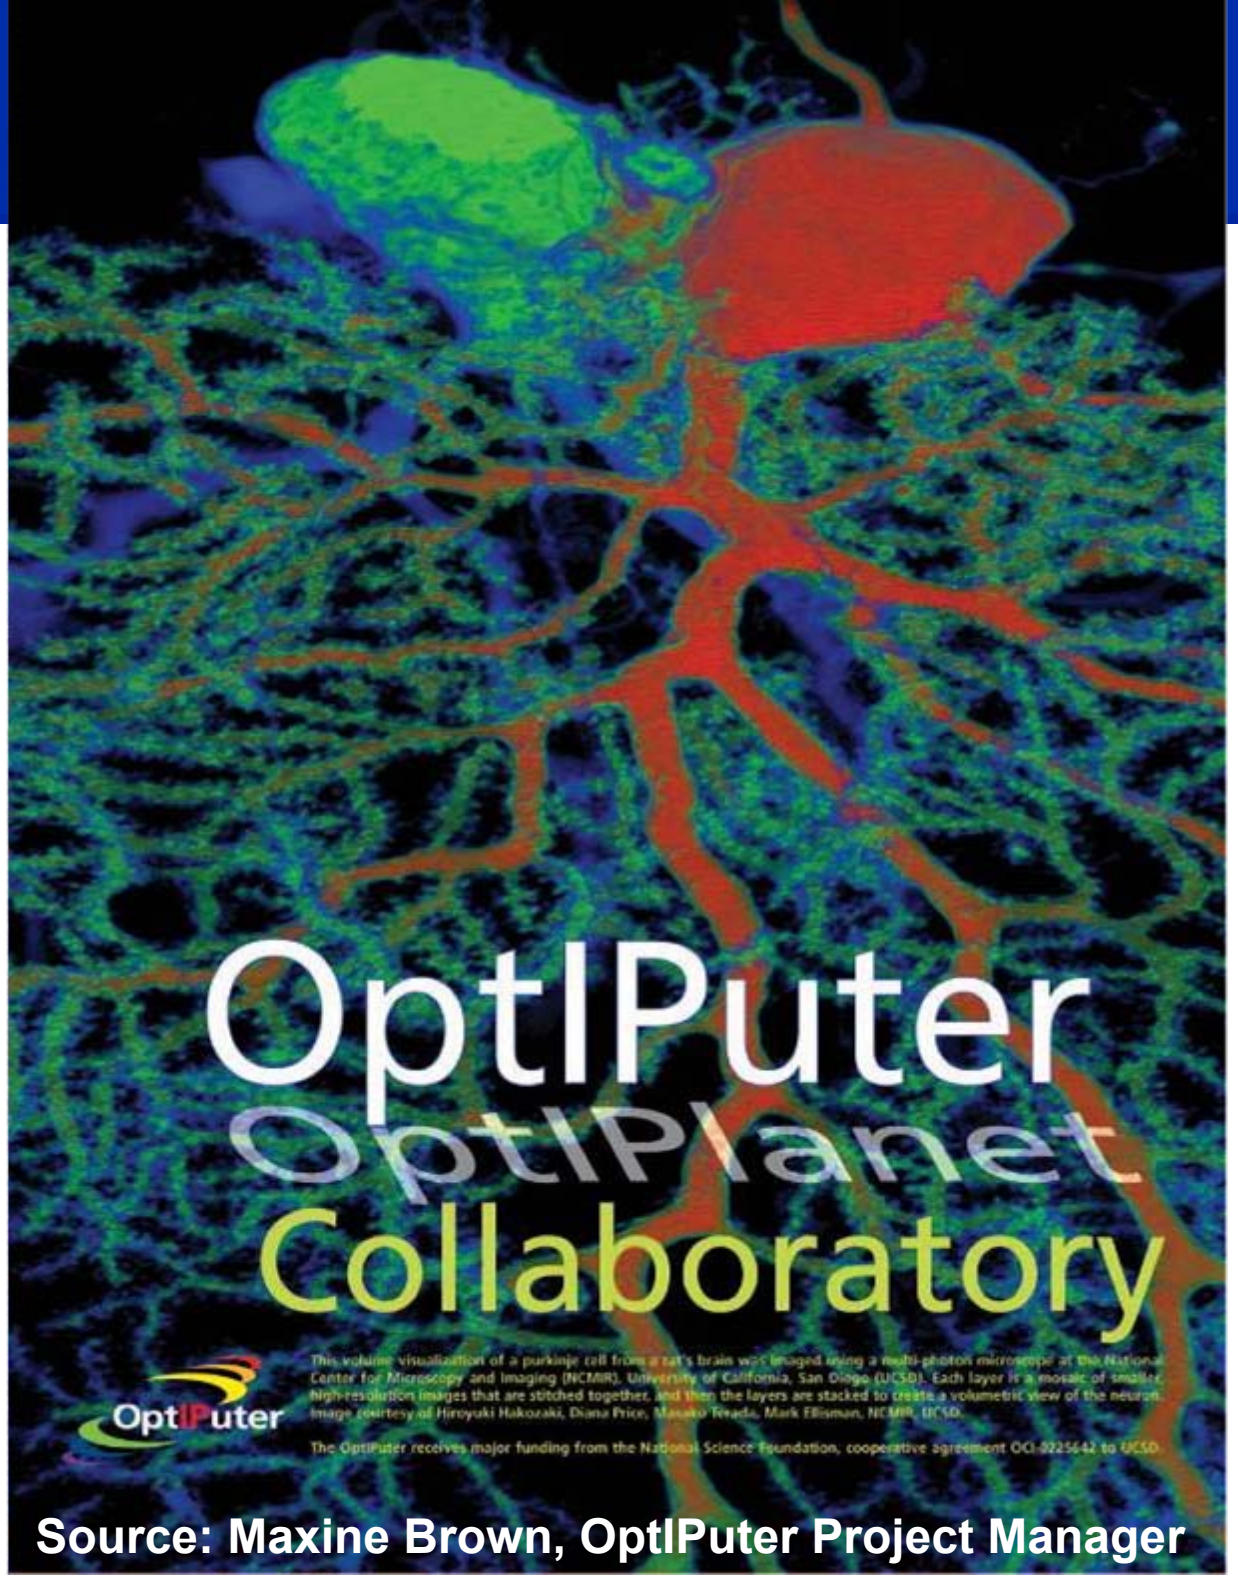

Source: Maxine Brown, OptIPuter Project Manager

Green Initiative:

Can Optical Fiber Replace Airline Travel for Continuing Collaborations ?

North American OptIPuter Team OptIPortals

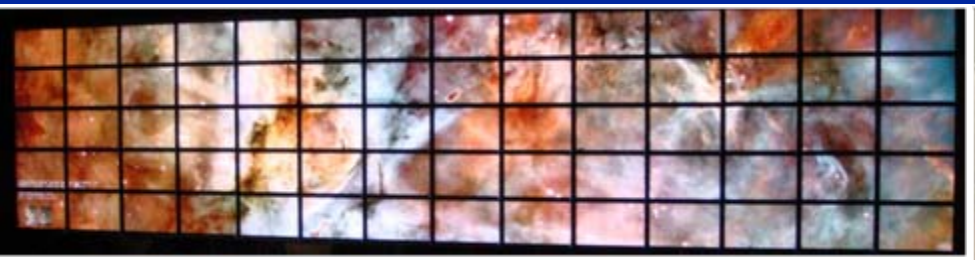

Calit2@UCSD

Calit2@UCSD

Calit2@UCSD

NCMIR-UCSD

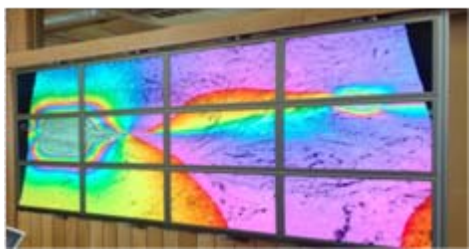

SIO-UCSD

Calit2@UCSD

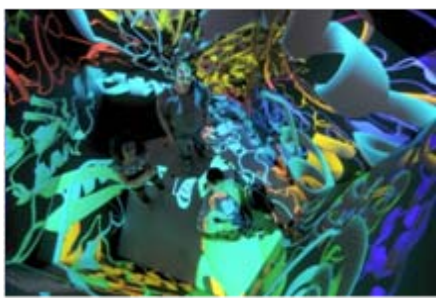

Calit2@UCSD

USGS

VISIT-U Michigan

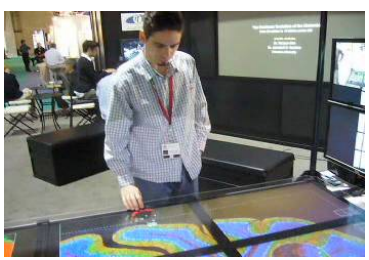

EVL@UIC

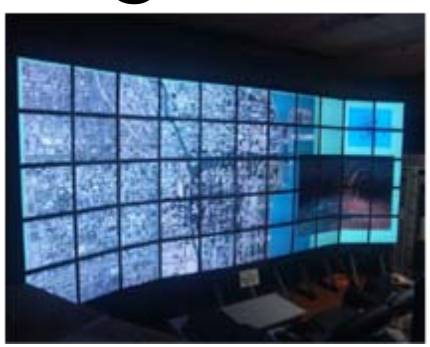

EVL@UIC

Calit2@UCI

iCAIR-
Northwestern U

Texas A&M

EVL@UIC

North American OptIPortals

U Washington

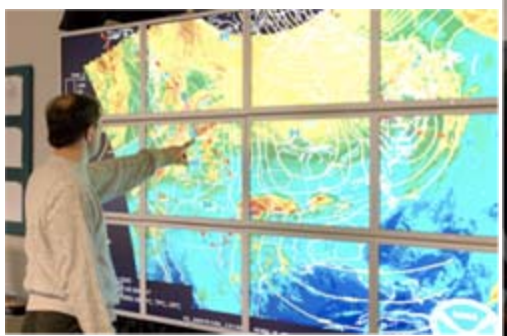

AOSS-U Michigan

Adler Planetarium - Chicago

TRCC, Chicago

TACC TeraGrid

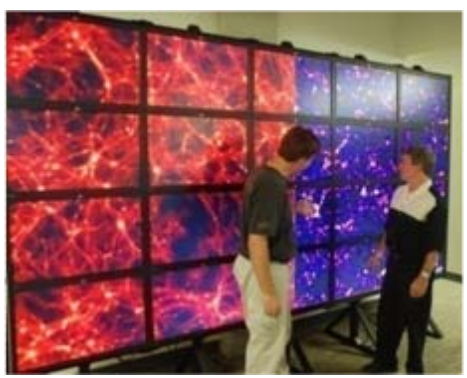

SDSC - UCSD

DMC-U Michigan

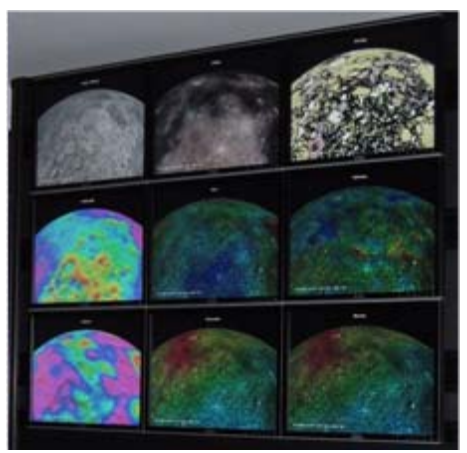

NASA Goddard

OptIPuter

North American OptIPortals

Sharp Corp.

Rincon

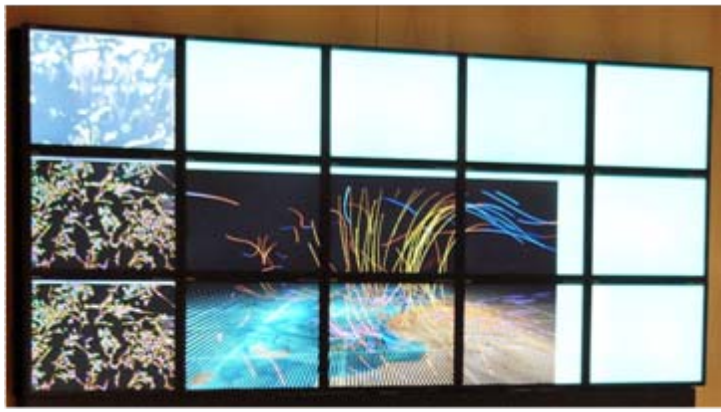

Nortel

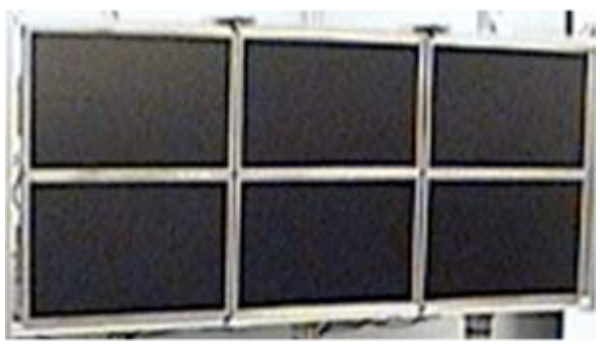

Florida International U

MIT

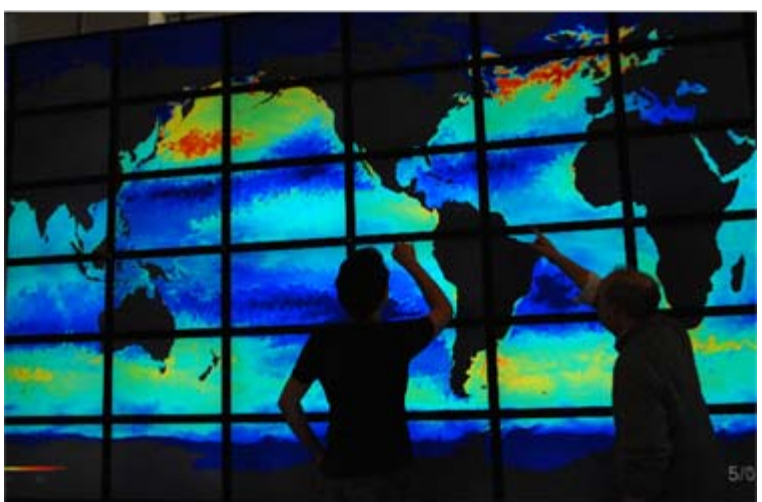

CICESE, Mexico

OptIPuter

it²

Launch of the 100 Megapixel OzIPortal Kicked Off a Rapid Build Out of Australian OptIPortals

ANU

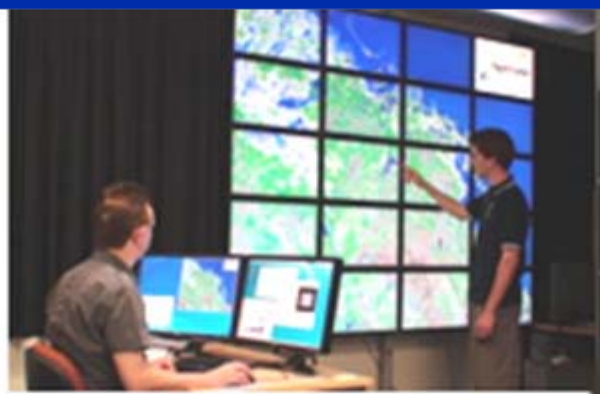

U Queensland

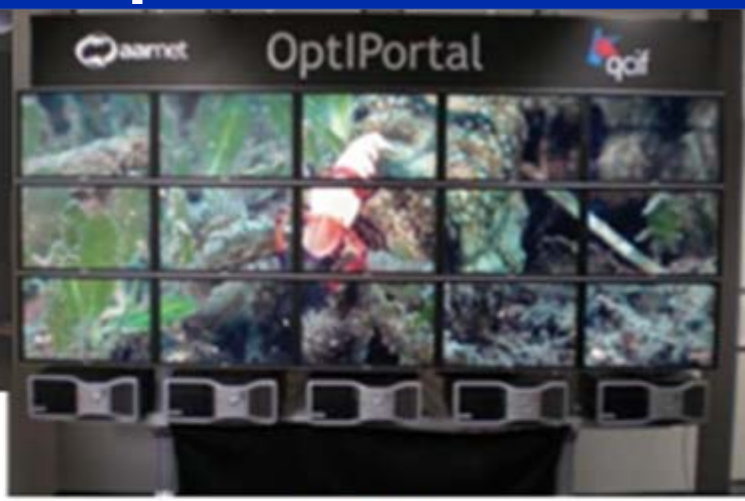

AARNet

Monash U

CSIRO

U Wellington, NZ

No Calit2 Person Physically Flew to Australia to Bring This Up!

Covise, Phil Weber, Jurgen Schulze, Calit2
CGLX, Kai-Uwe Doerr , Calit2
http://www.calit2.net/newsroom/release.php?id=1421

International OptIPortals Are Being Adopted Globally Connected by GLIF

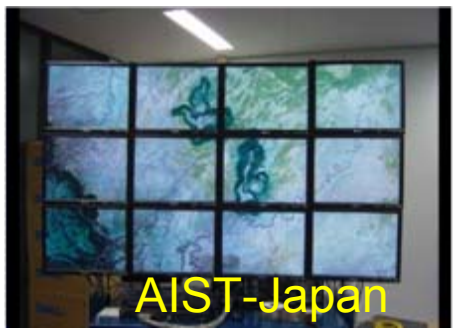

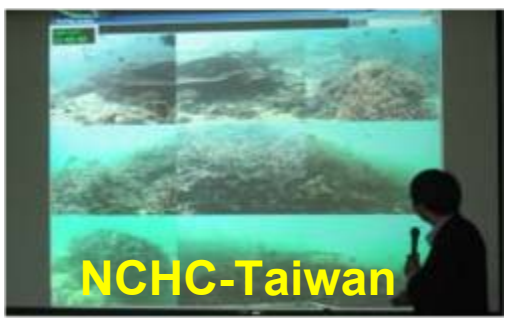

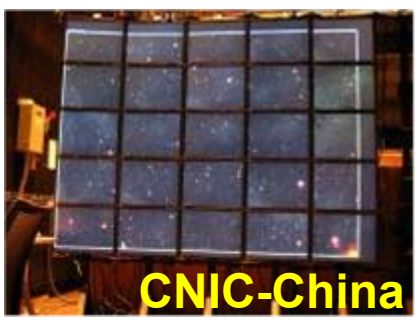

OptIPuter

it²

EVL's SAGE OptIPortal VisualCasting Multi-Site OptIPuter Collaboratory

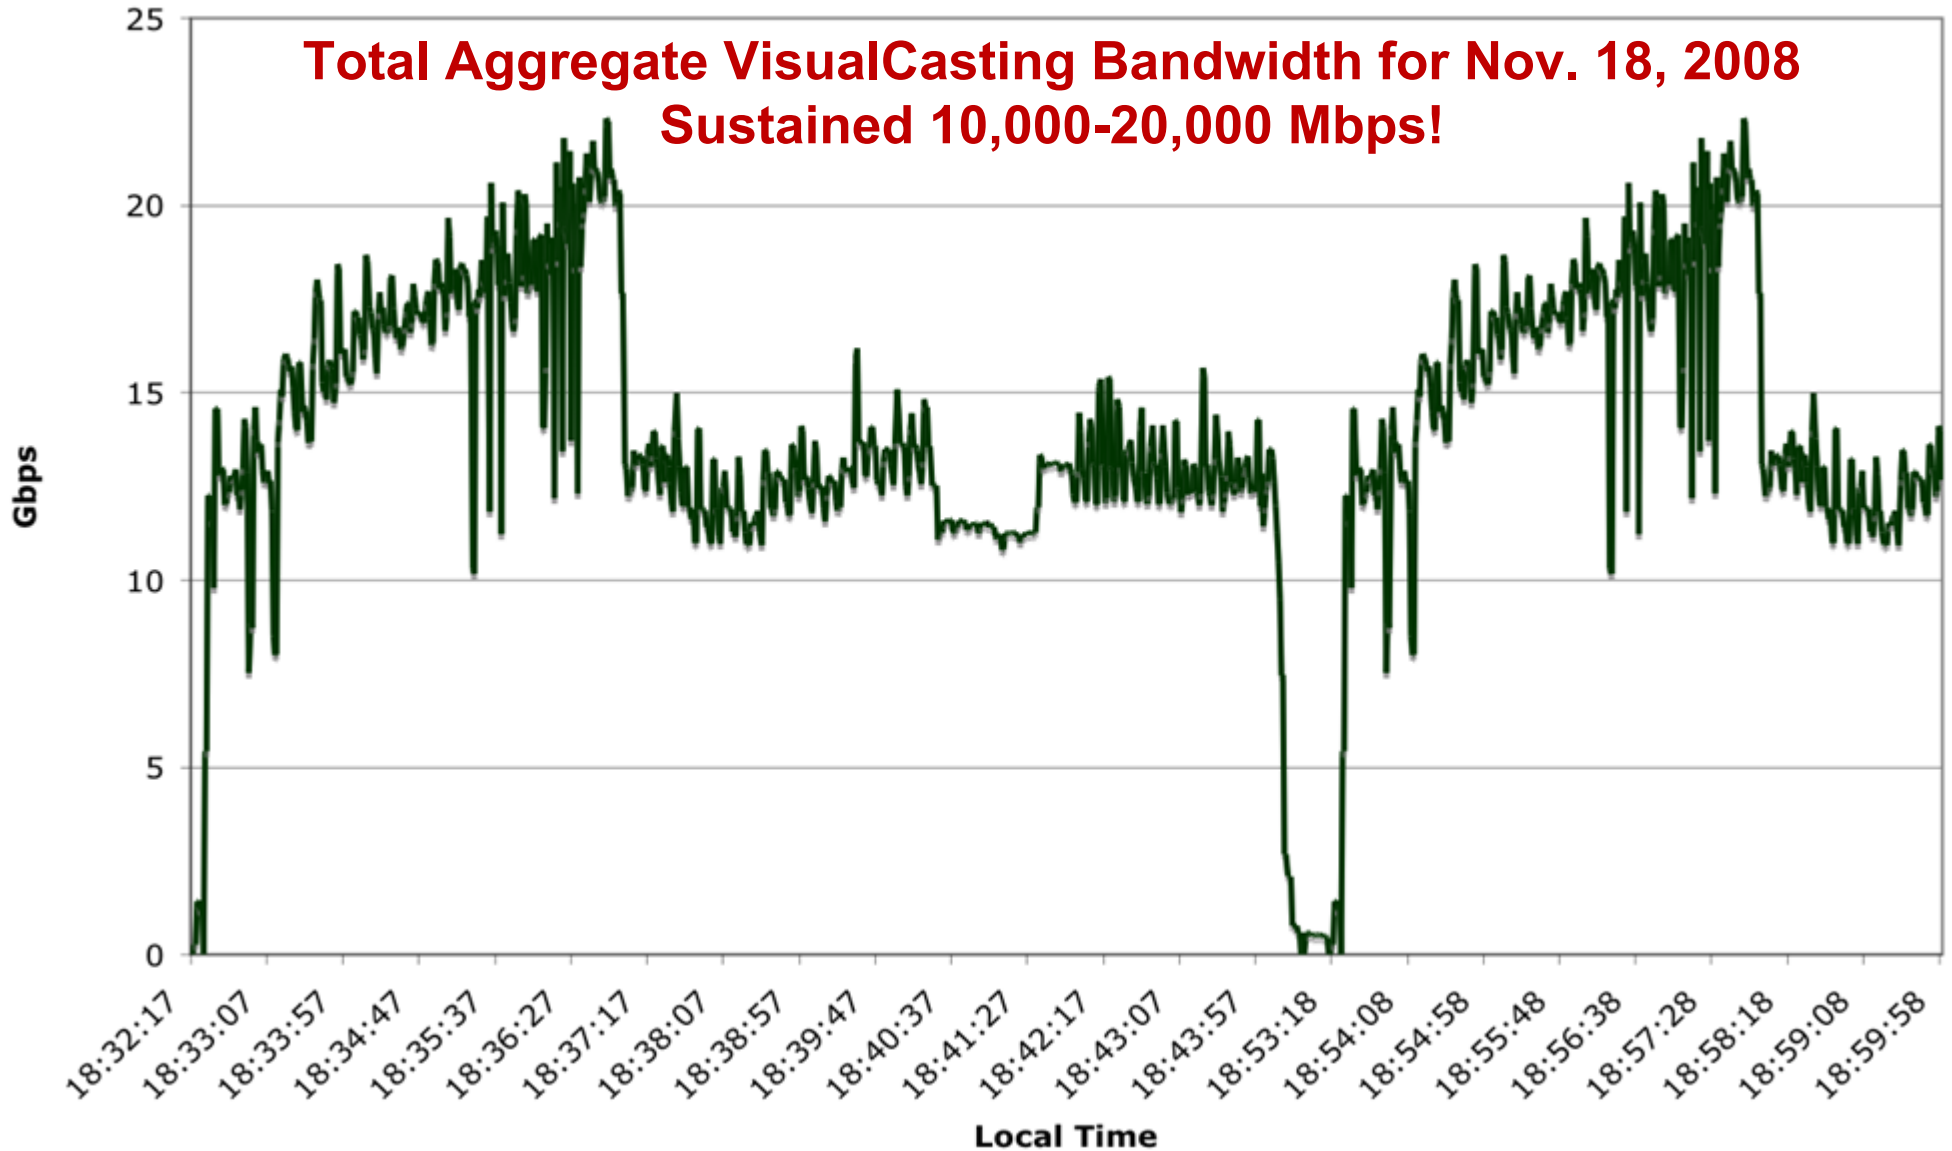

OptIPuter

Source: Jason Leigh, Luc Renambot, EVL, UI Chicago

it²

CI Bottleneck:
The Campus Last Mile

How Do You Get From Your Lab to the Regional Optical Networks?

"Research is being stalled by 'information overload,' Mr. Bement said, because data from digital instruments are piling up far faster than researchers can study. In particular, he said, campus networks need to be improved. High-speed data lines crossing the nation are the equivalent of six-lane superhighways, he said. But networks at colleges and universities are not so capable. "Those massive conduits are reduced to two-lane roads at most college and university campuses," he said. Improving cyberinfrastructure, he said, "will transform the capabilities of campus-based scientists."

-- Arden Bement, the director of the National Science Foundation

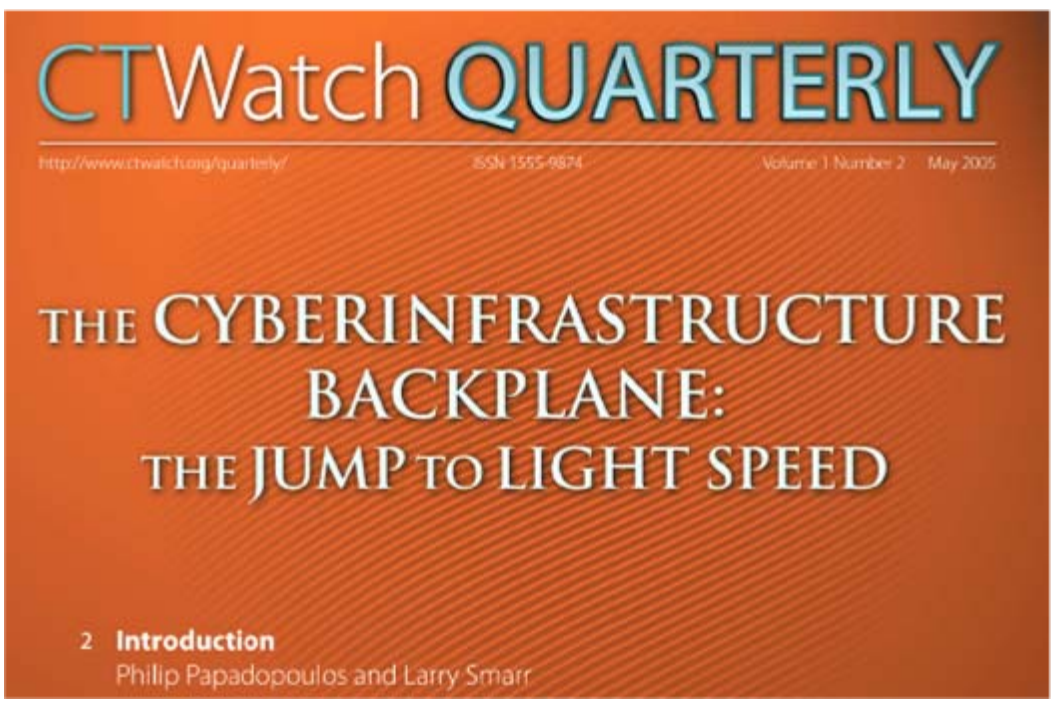

www.ctwatch.org

Creating a California Cyberinfrastructure of OptIPuter "On-Ramps" to NLR & TeraGrid Resources

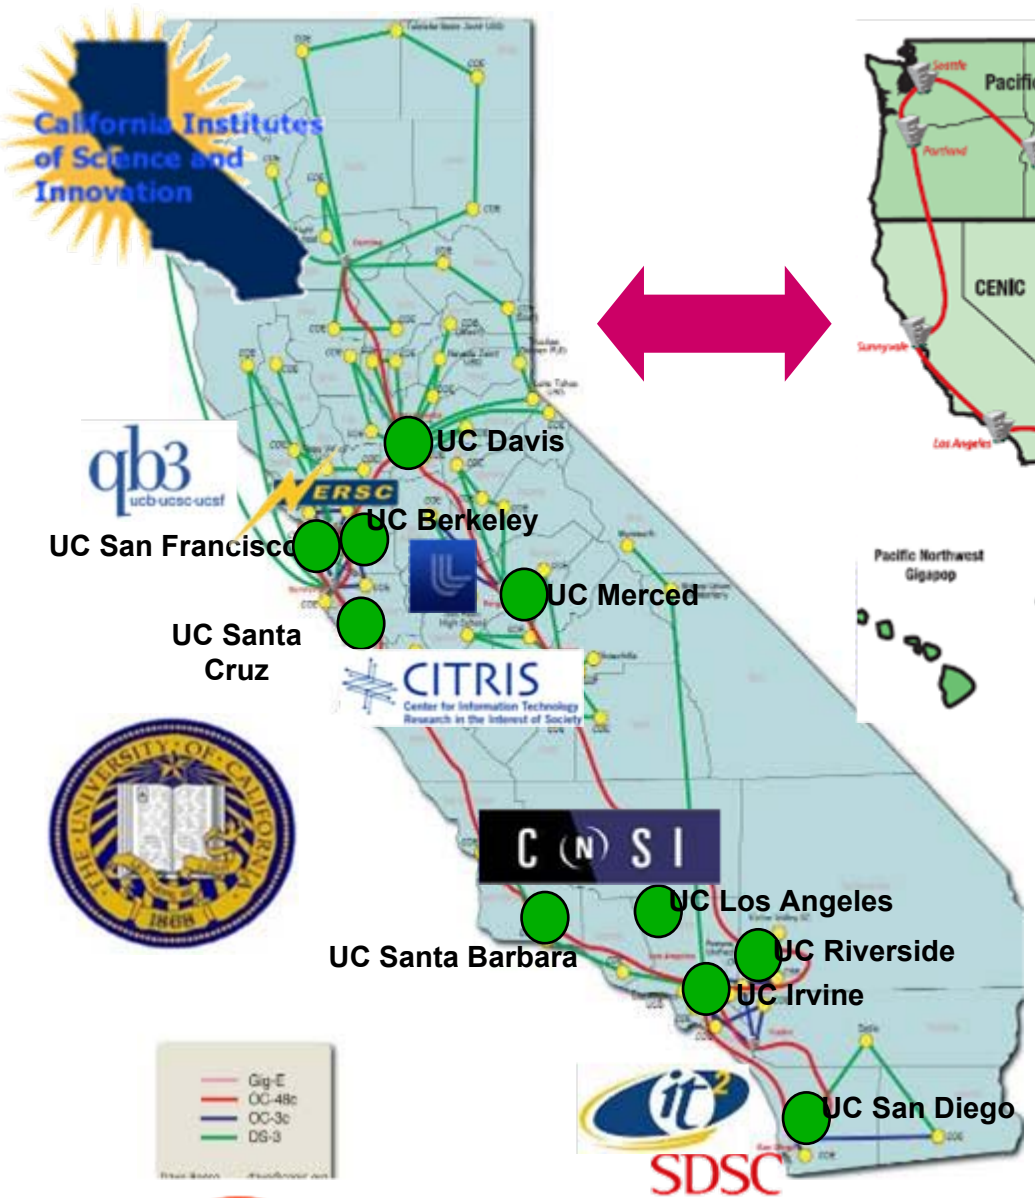

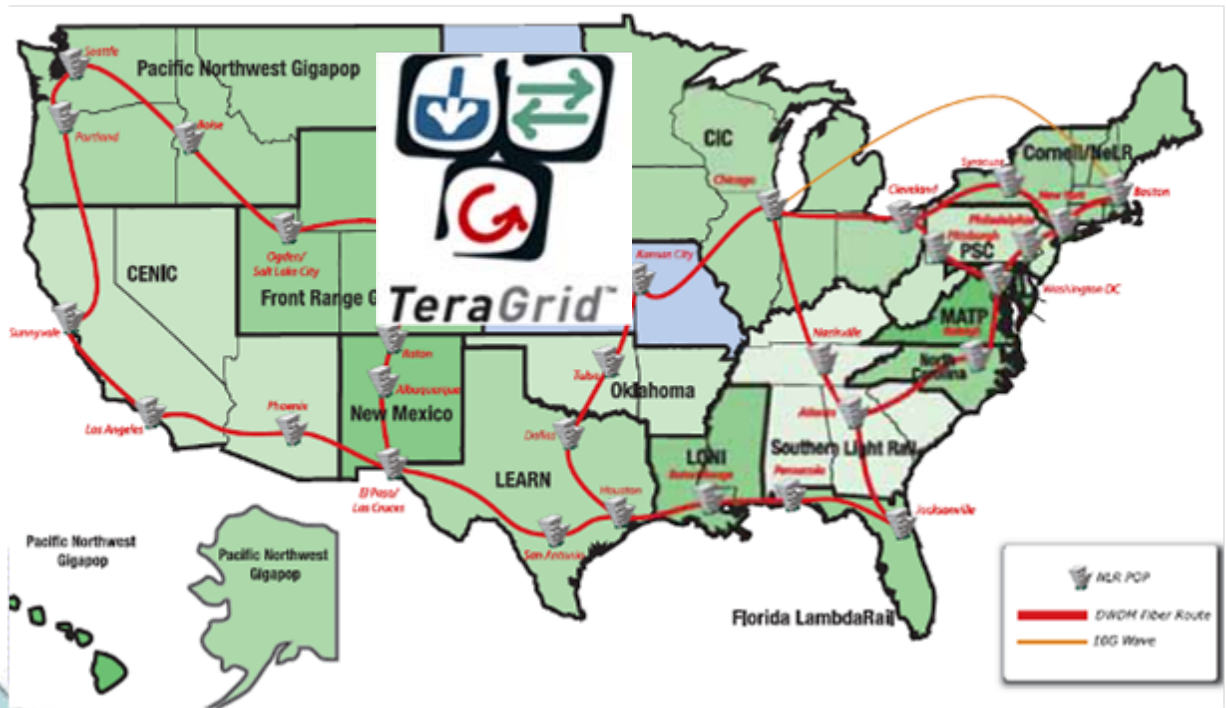

Creating a Critical Mass of OptIPuter End Users on a Secure LambdaGrid

CENIC Workshop at Calit2 Sept 15-16, 2008

OptIPuter

it²

CENIC's New "Hybrid Network" - Traditional Routed IP and the New Switched Ethernet and Optical Services

~ $14M Invested in Upgrade

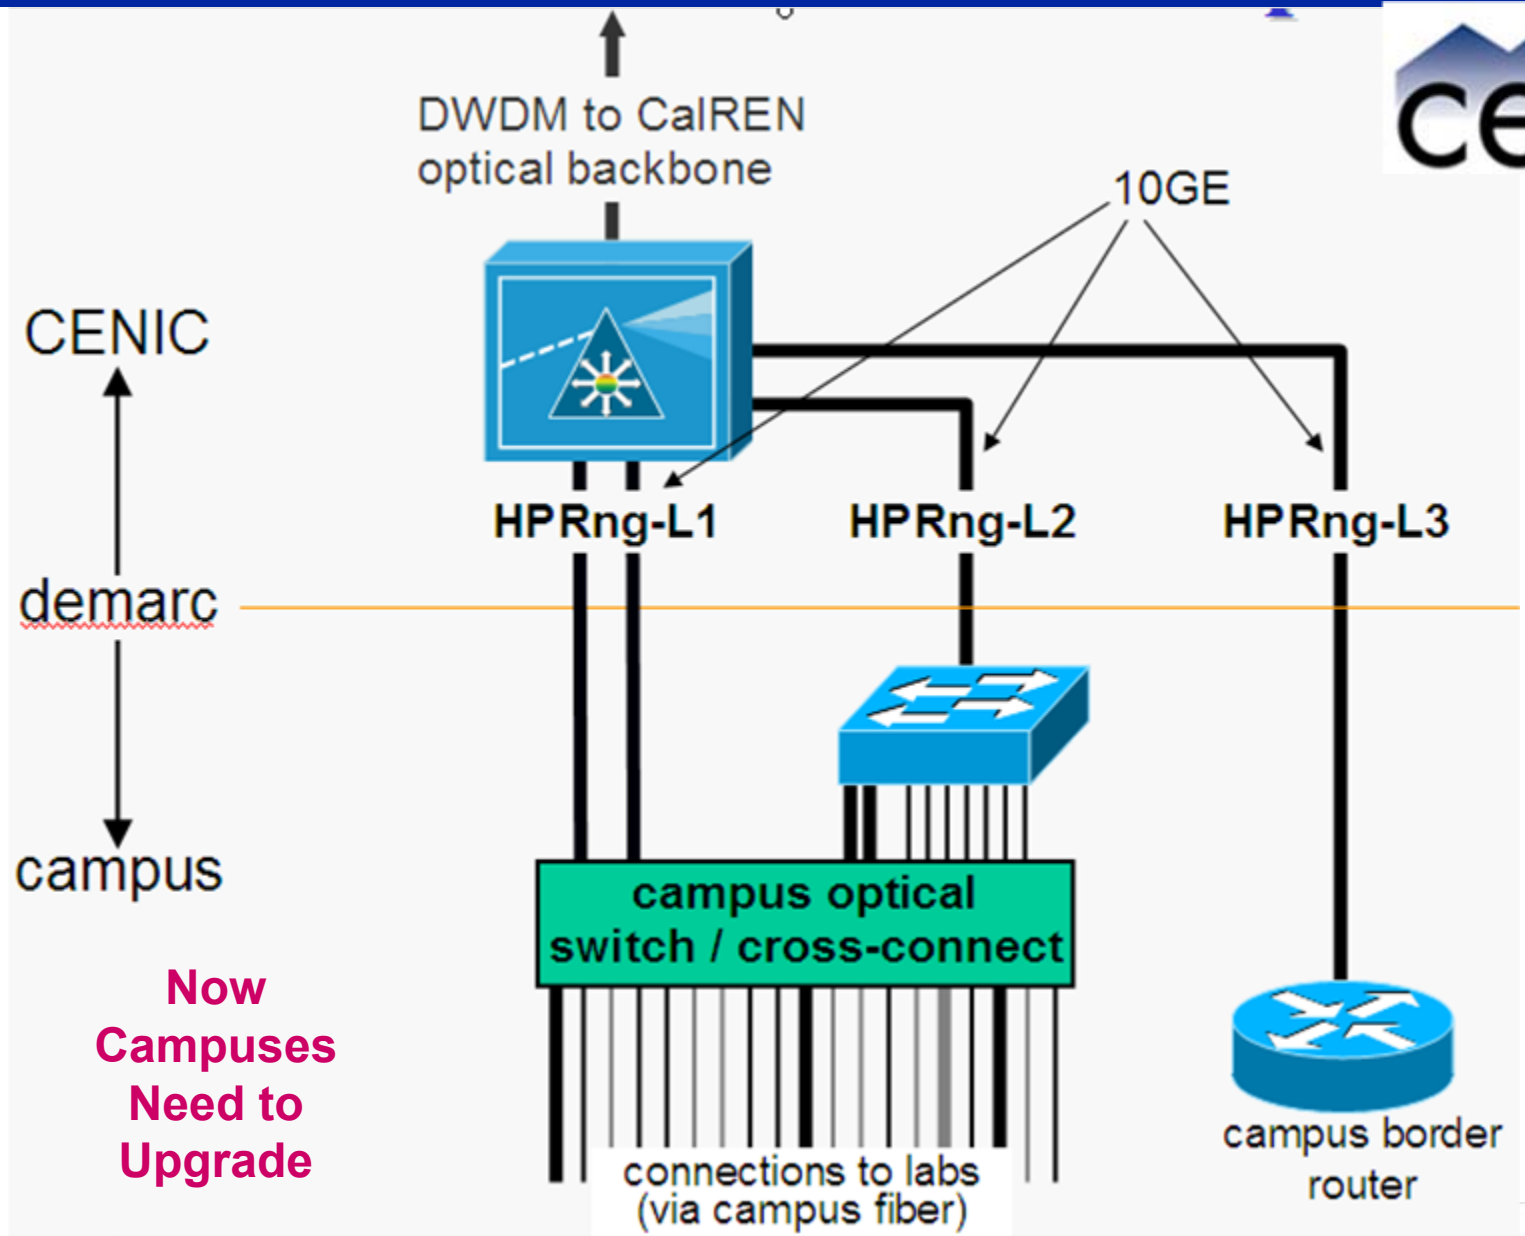

Now Campuses Need to Upgrade

Source: Jim Dolgonas, CENIC

The "Golden Spike" UCSD Experimental Optical Core: Ready to Couple Users to CENIC L1, L2, L3 Services

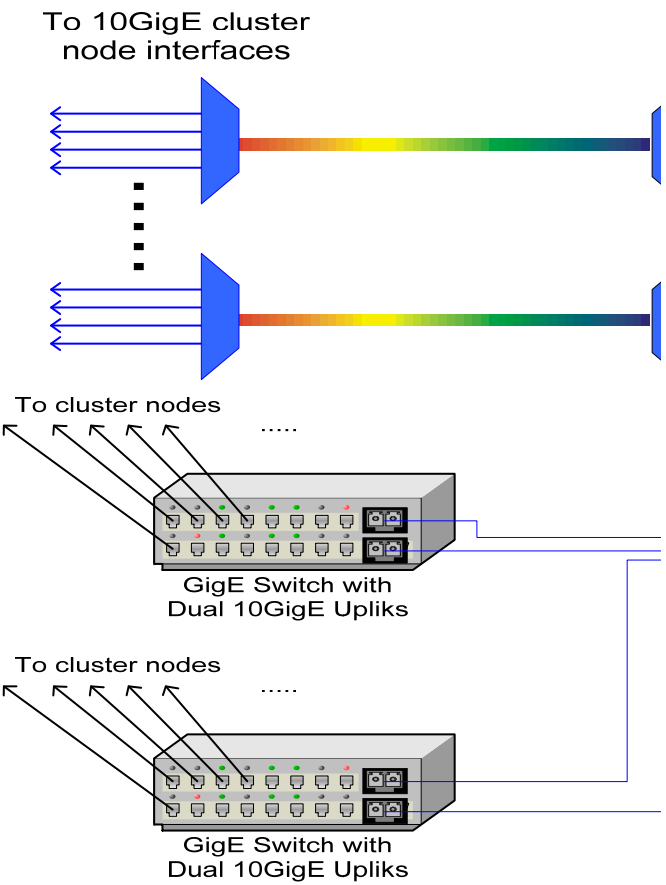

Goals by 2008:

>= 60 endpoints at 10 GigE

>= 30 Packet switched

>= 30 Switched wavelengths

>= 400 Connected endpoints

Approximately 0.5 Tbps Arrive at the "Optical" Center of Hybrid Campus Switch

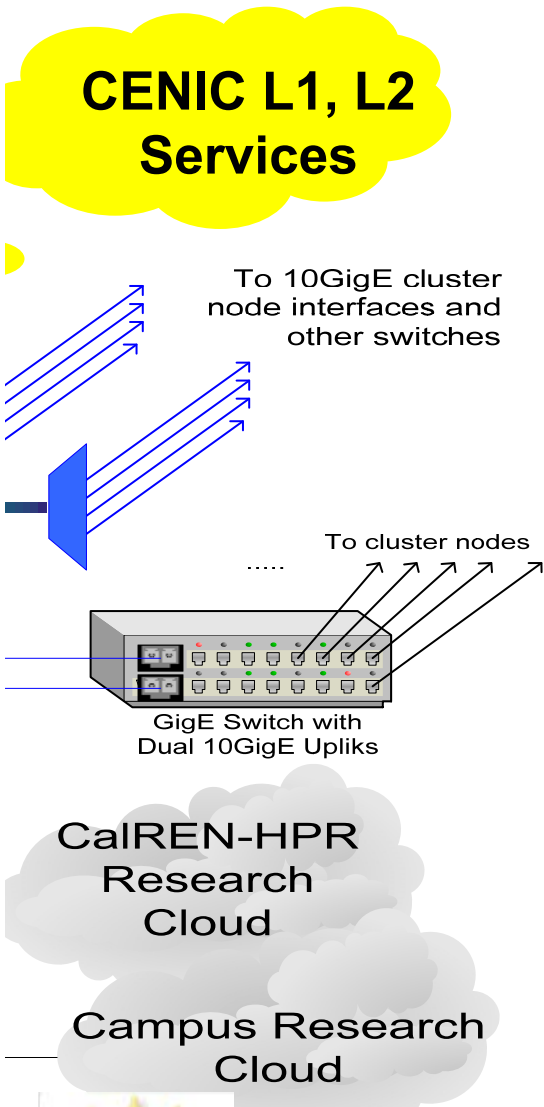

Funded by NSF MRI Grant

Source: Phil Papadopoulos, SDSC/Calit2 (Quartzite MRI PI, OptIPuter co-PI)

OCI Early Goal:
Provide End-to-End CI for Petascale End Users

SDSC
SAN DIEGO SUPERCOMPUTER CENTER

Two 64K Images From a Cosmological Simulation of Galaxy Cluster Formation

ROCKS

CGLX

OptIPuter

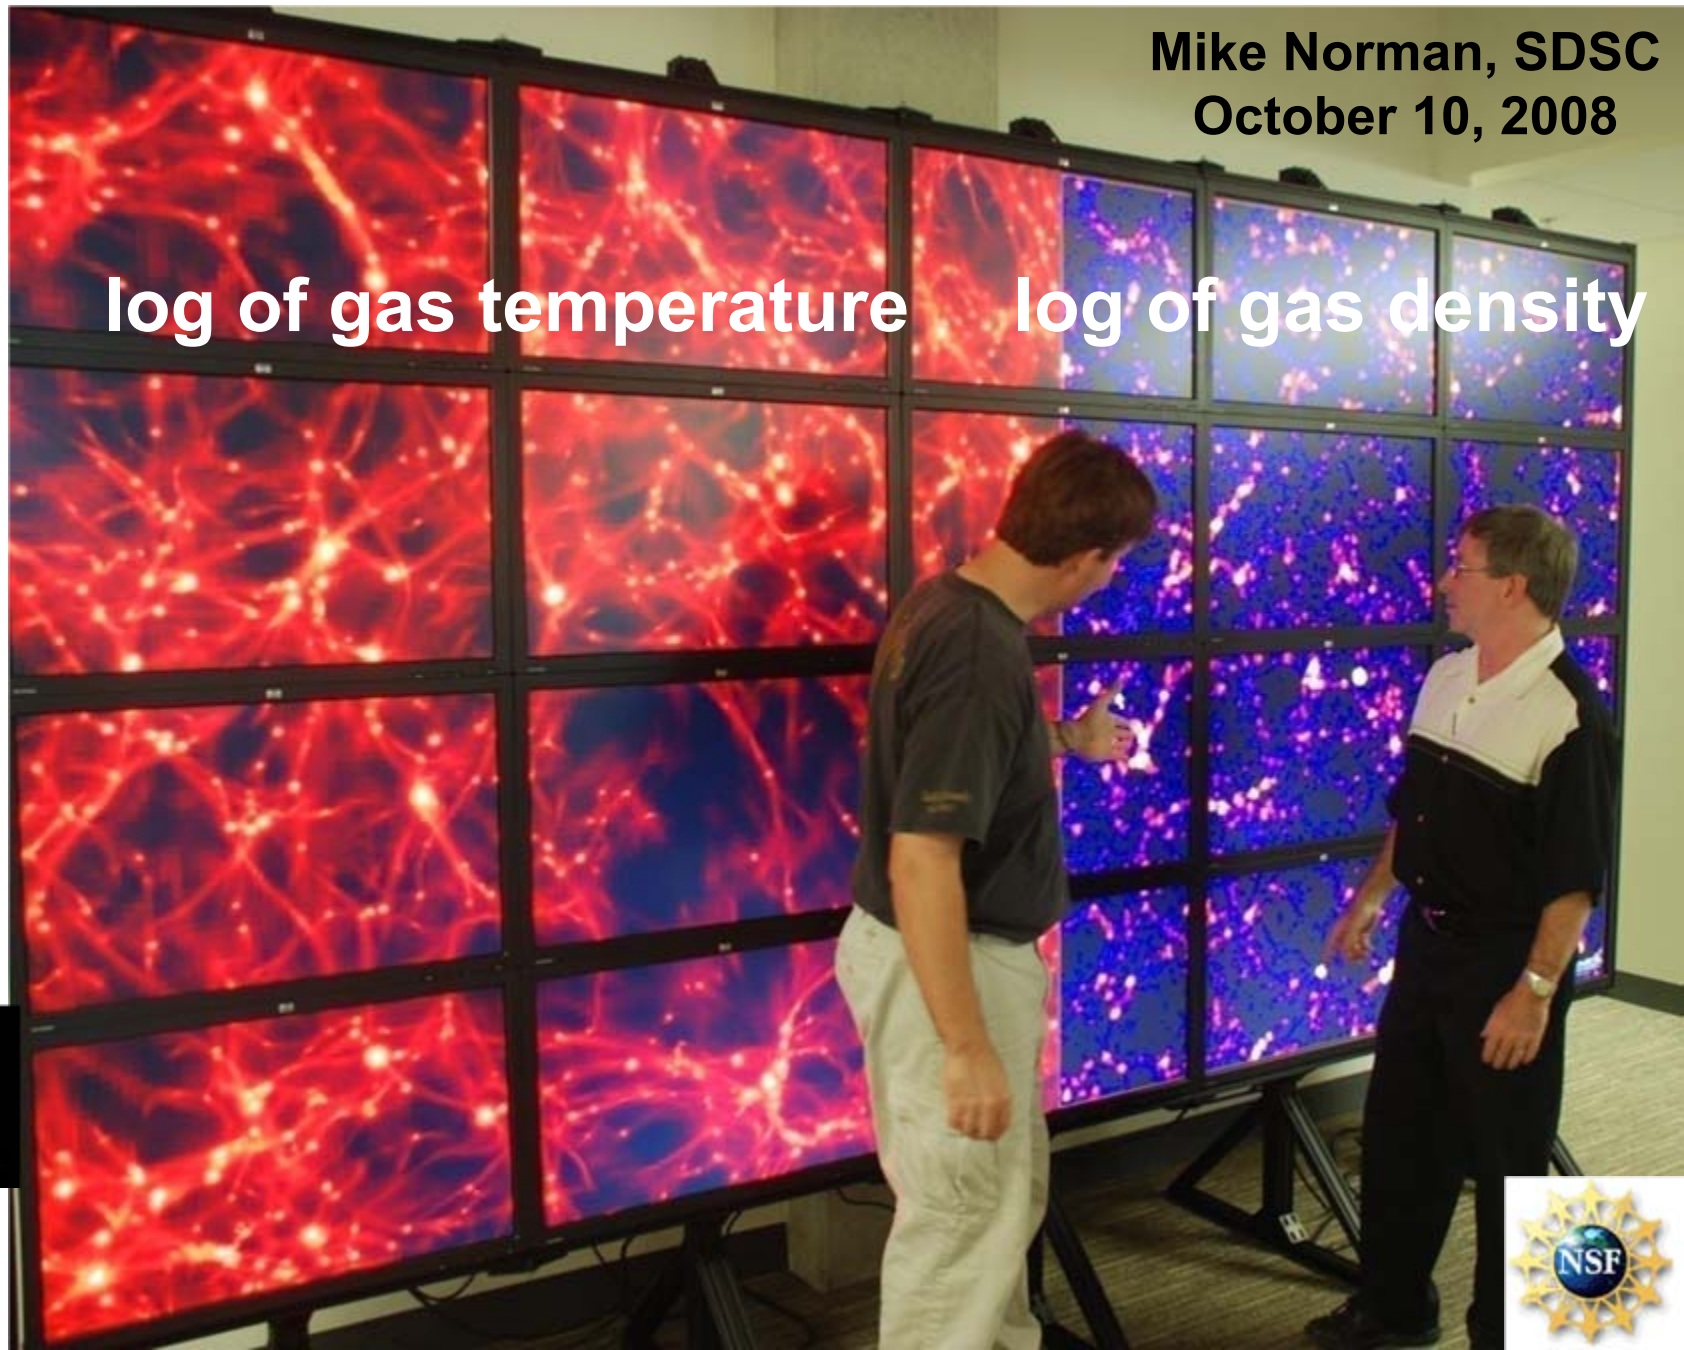

NSF

SDSC OptIPortal Uses UCSD Research Network to Get to TeraGrid with 10Gbps Clear Channel

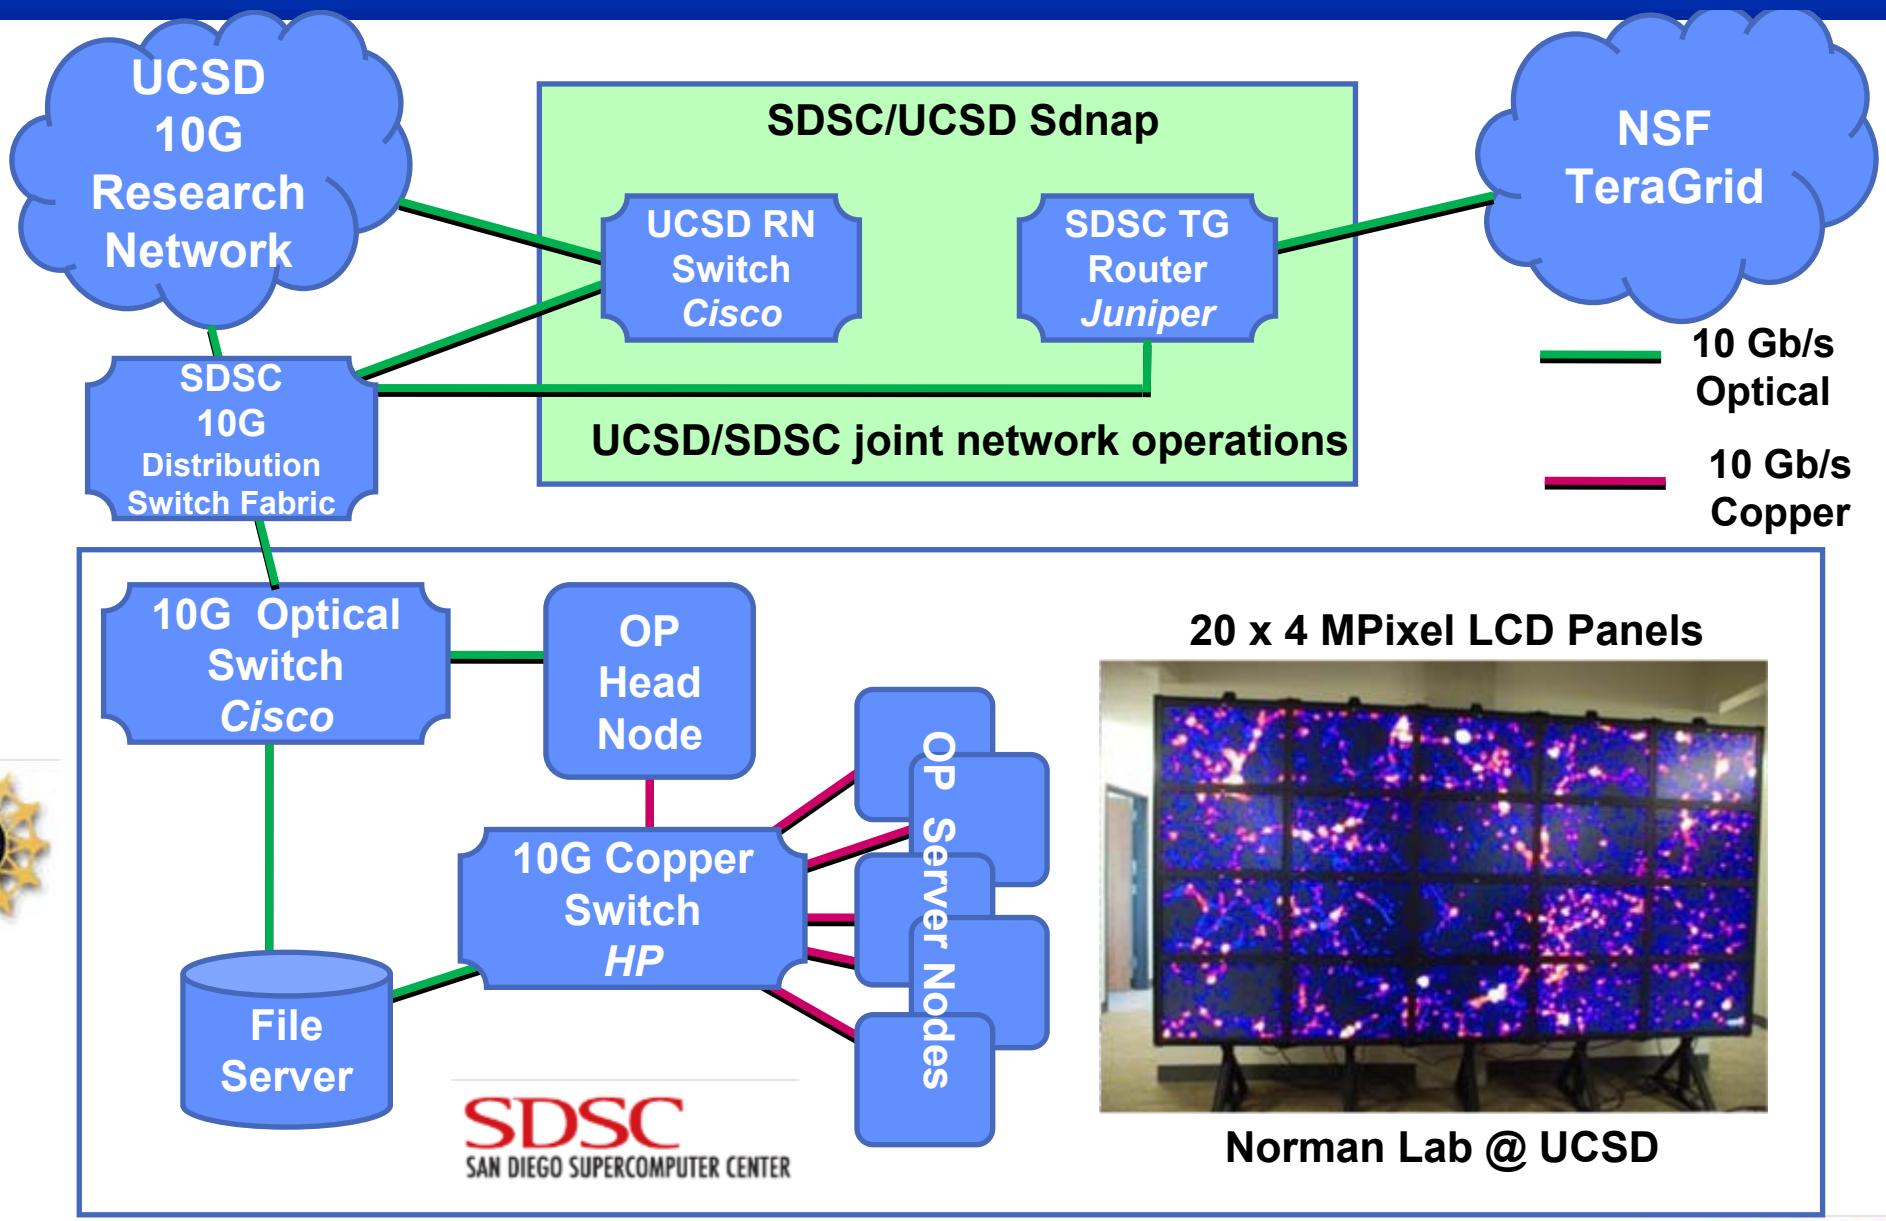

Source: Mike Norman, Tom Hutton, Rick Wagner, SDSC

Fiber Optic Networks Create Campus-Scale Data Utilities

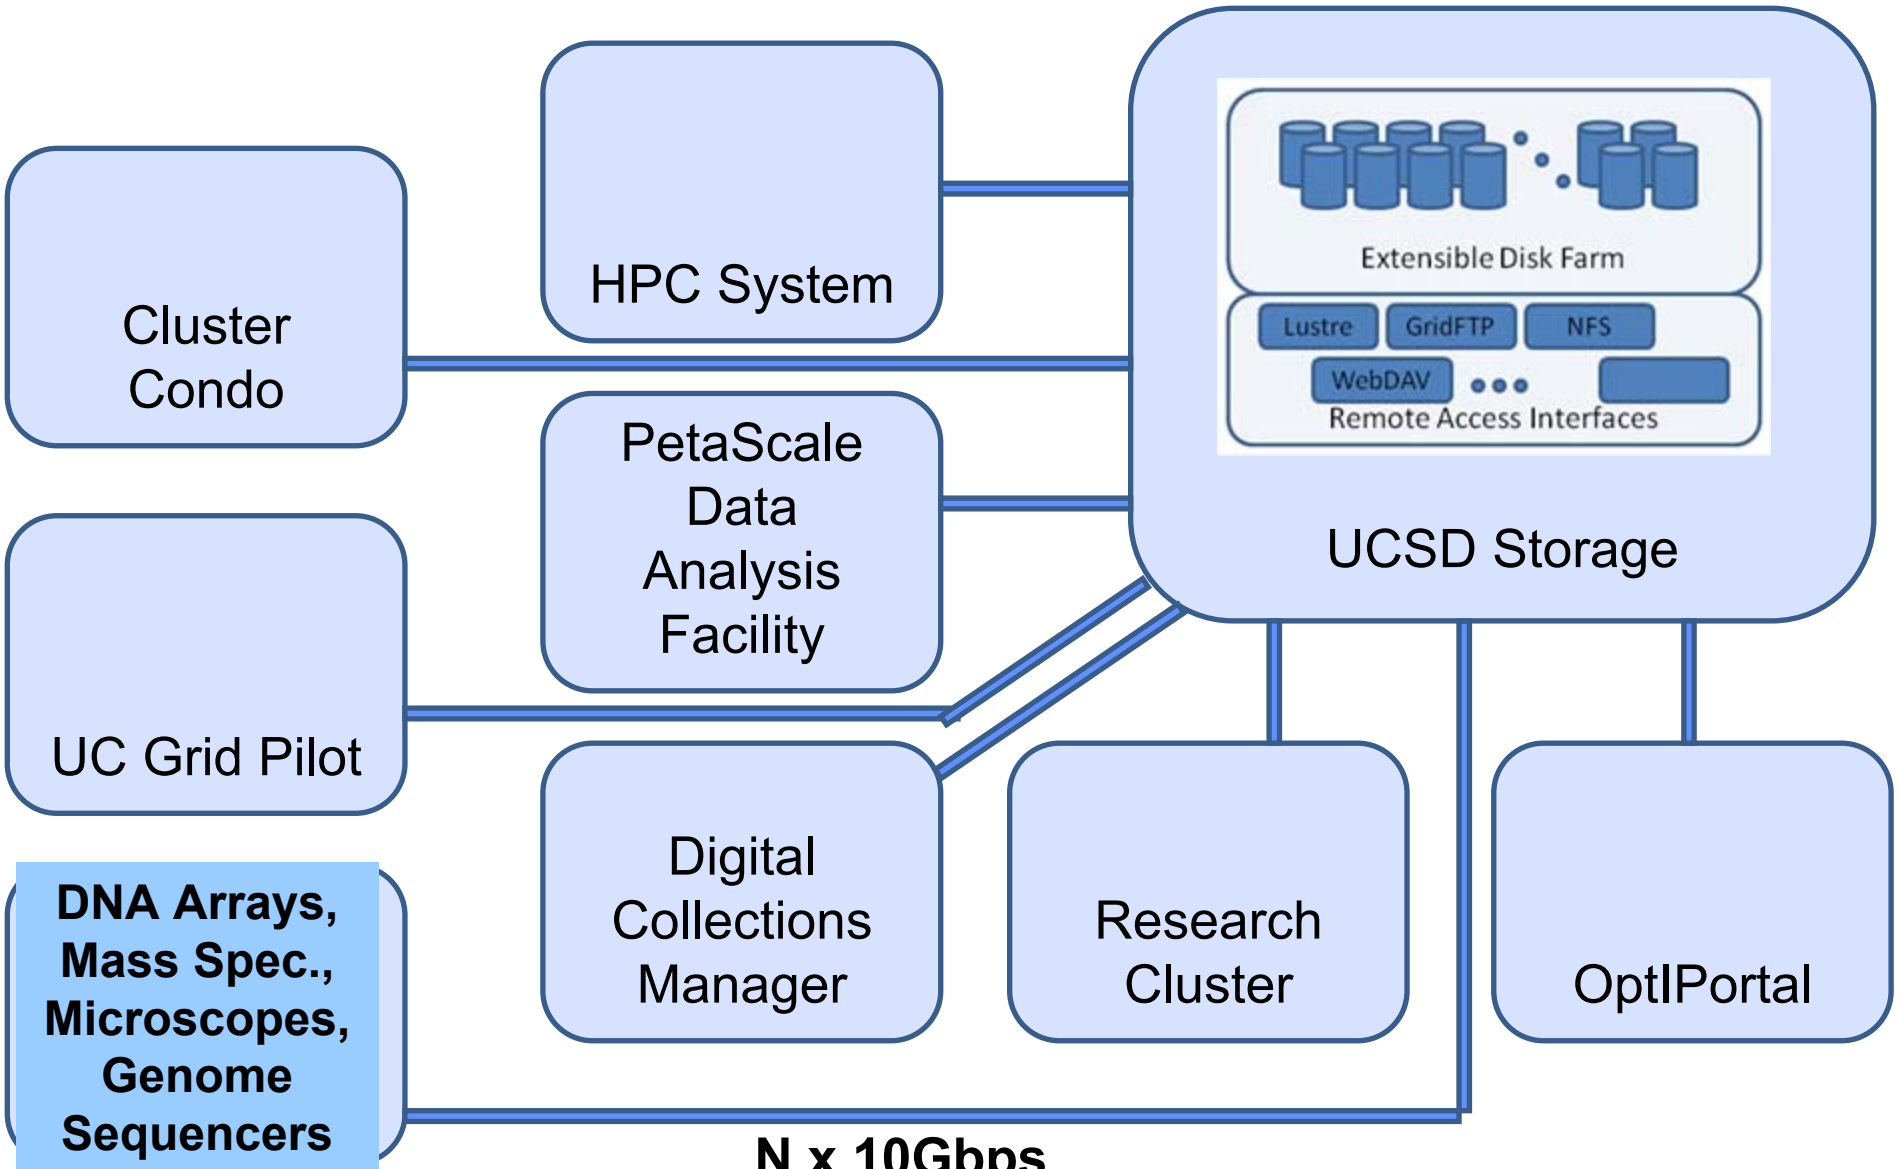

Source:Phil Papadopoulos, SDSC/Calit2

Threat to CI Deployment—Research Needed on How to Deploy a Green CI

UCSD Structural Engineering Dept. Conducted Tests May 2007

7 Racks plus Network

Takes up 2 Parking Spaces

Data Power Cooling

- **Measure and Control Energy Usage:**
  - Sun Has Shown up to 40% Reduction in Energy
  - Active Management of Disks, CPUs, etc.
  - Measures Temperature at 5 Spots in 8 Racks
  - Power Utilization in Each of the 8 Racks
  - Chilled Water Cooling Systems

NSF MRI GreenLight Project

UCSD (Calit2 & SOM) Bought Two Sun Boxes May 2008

Calit2 GreenLight MRI Project Enables Green IT Computer Science Research

- **Computer Architecture**
  - Rajesh Gupta/CSE
- **Software Architecture**
  - Amin Vahdat & Ingolf Kruger/ CSE
- **CineGrid Exchange**
  - Tom DeFanti/Calit2
- **Visualization**
  - Falko Kuster/Structural Engineering
- **Power and Thermal Management**
  - Tajana Rosing/CSE
- **Analyzing Power Consumption Data**
  - Jim Hollan/Cog Sci

http://greenlight.calit2.net

My OCI Recommendations

- **Build Out a National End-to-End Balanced HiPerf "Green" CI**
  - Lead with Petascale User Lambdas & Analysis OptIPortals
  - Provide CI Users Access to On-Demand WAN NLR/I2 L1/L2
  - Fund NLR/I2 Campus Gateway to Users 'vBNS Connections Program'
  - Add Green CI Review to NSF Transformational Evaluations
  - Find Ways to Co-Fund International CI Collaborations
- **Creating and Supporting Sustainable CI and Workforce**
  - Harvest from CISE Successful Research Grants
  - Use NSF SDCI-Like Programs to Sustain the CI Software
  - Use NSF MRIs to Prototype Campus CI Innovations
  - Grad, Postdoc, Career Awards for CI Development & Usage
- **With Other Directorates**
  - Prototype CI Using Global Distributed Data Sharing Communities
  - Encourage Common CI across MREFCs
  - Long Term Collaborative Computational Science for Complex Problems

OptIPuter

Source: Larry Smarr, Calit2

it²