

“The Pacific Research Platform: a Science-Driven Big-Data Freeway System.”

**NCSA Colloquium
National Center for Supercomputing Applications
University of Illinois, Urbana-Champaign
September 18, 2015**

**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
<http://lsmarr.calit2.net>**



Abstract

Research in data-intensive fields is increasingly multi-investigator and multi-institutional, depending on ever more rapid access to ultra-large heterogeneous and widely distributed datasets. The Pacific Research Platform (PRP) is a multi-institutional extensible deployment that establishes a science-driven high-capacity data-centric “freeway system.” The PRP spans all 10 campuses of the University of California, as well as the major California private research universities, four supercomputer centers, and several universities outside California. Fifteen multi-campus data-intensive application teams act as drivers of the PRP, providing feedback over the five years to the technical design staff. These application areas include particle physics, astronomy/astrophysics, earth sciences, biomedicine, and scalable multimedia, providing models for many other applications. The PRP partnership extends the NSF-funded campus Science DMZs to a regional model that allows high-speed data-intensive networking, facilitating researchers moving data between their labs and their collaborators’ sites, supercomputer centers or data repositories, and enabling that data to traverse multiple heterogeneous networks without performance degradation over campus, regional, national, and international distances



Vision: Creating a West Coast “Big Data Freeway” Connected by CENIC/Pacific Wave to Internet2 & GLIF

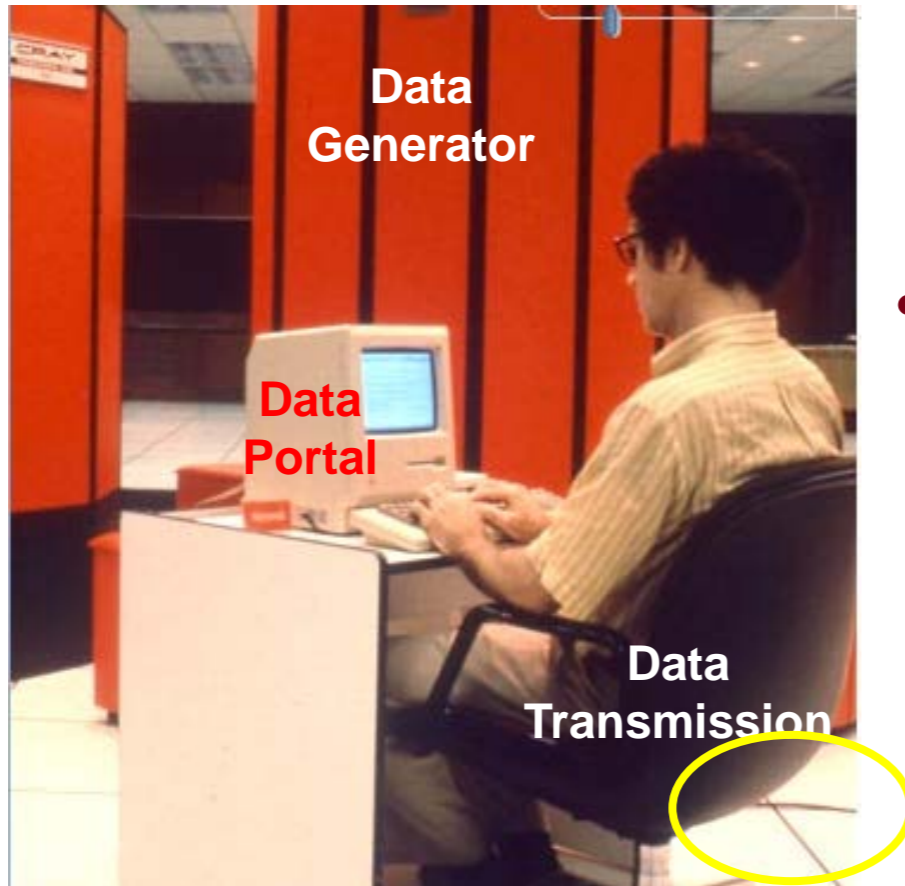
**Use Lightpaths to Connect
All Data Generators and Consumers,
Creating a “Big Data” Plane
Integrated With High Performance Global Networks**

*“The Bisection Bandwidth of a Cluster Interconnect,
but Deployed on a 10-Campus Scale.”*

This Vision Has Been Building for Over Two Decades



NCSA Telnet--“Hide the Cray” Paradigm That We Still Use Today



**John Kogut Simulating
Quantum Chromodynamics
He Uses a Mac—The Mac Uses the Cray**

- **NCSA Telnet -- Interactive Access**
 - From Macintosh or PC Computer
 - To Telnet Hosts on TCP/IP Networks
- **Allows for Simultaneous Connections**
 - To Numerous Computers on The Net
 - Standard File Transfer Server (FTP)
 - Lets You Transfer Files to and from Remote Machines and Other Users

Interactive Supercomputing Collaboratory Prototype: Using Analog Communications to Prototype the Fiber Optic Future

“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

SIGGRAPH 1989



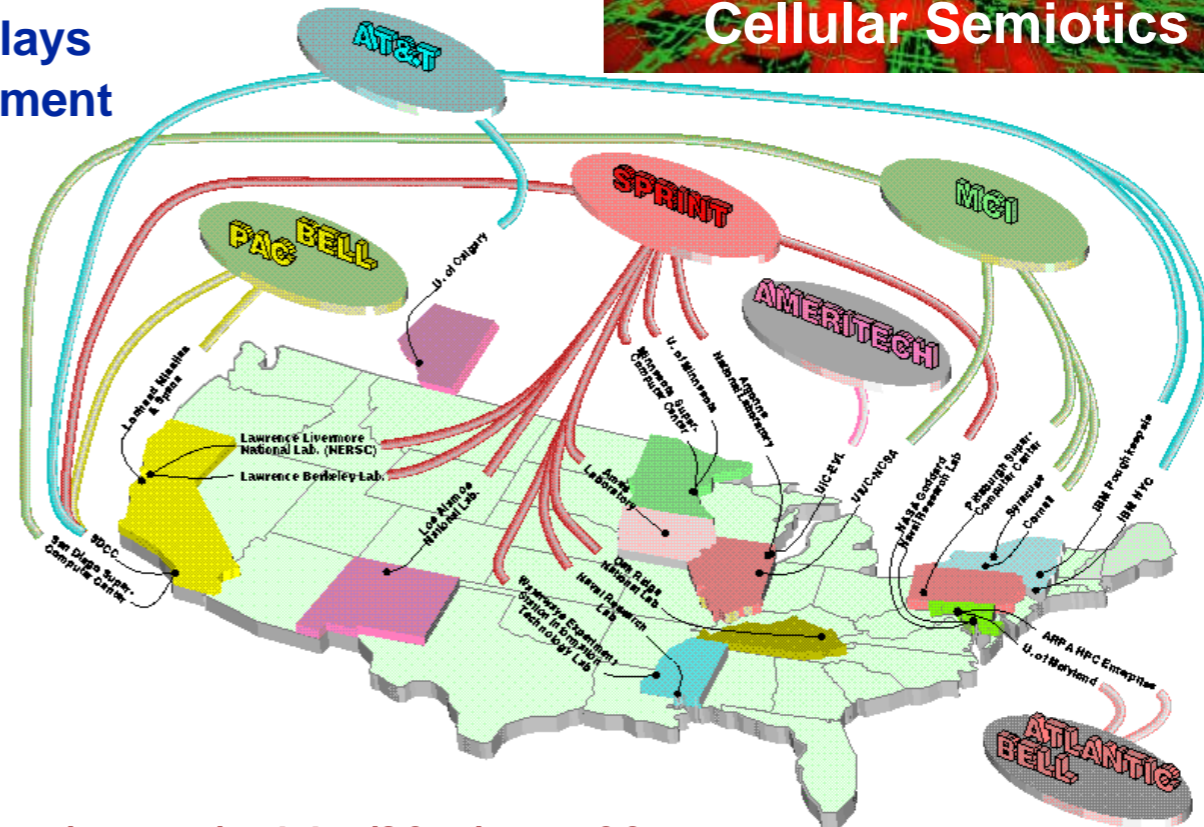

“We’re using satellite technology... to demo what it might be like to have high-speed fiber-optic links between advanced computers in two different geographic locations.”

— Al Gore, Senator

Chair, US Senate Subcommittee on Science, Technology and Space

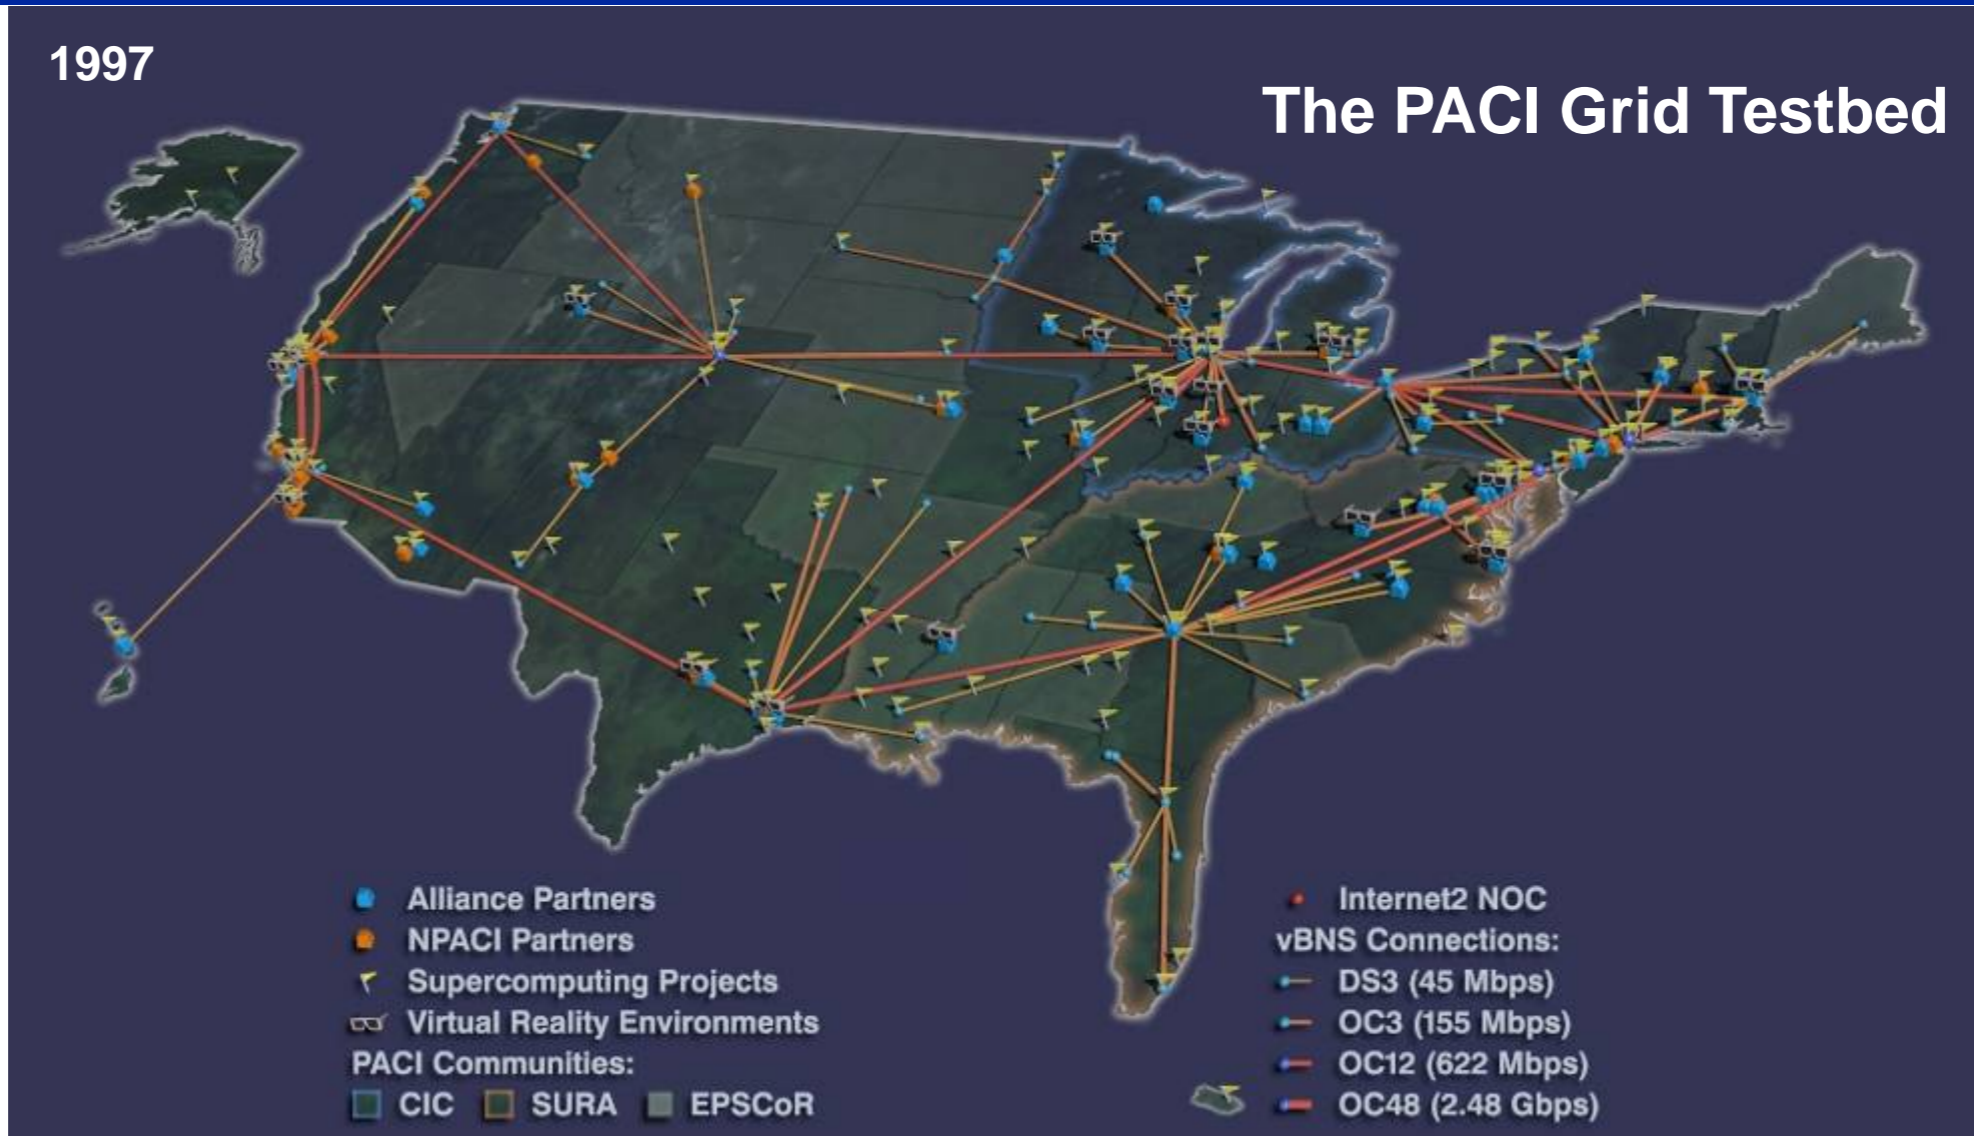
I-WAY: Information Wide Area Year Supercomputing '95

- The First National 155 Mbps Research Network
 - 65 Science Projects
 - Into the San Diego Convention Center
- I-Way Featured:
 - Networked Visualization Application Demonstrations
 - Large-Scale Immersive Displays
 - I-Soft Programming Environment



<http://archive.ncsa.uiuc.edu/General/Training/SC95/GII.HPCC.html>

PACI is Prototyping America's 21st Century Information Infrastructure



Chesapeake Bay Simulation Collaboratory: vBNS Linked CAVE, ImmersaDesk, Power Wall, and Workstation

Alliance Project: Collaborative Video Production via Tele-Immersion and Virtual Director

Alliance Application Technologies
Environmental Hydrology Team



4 MPixel PowerWall



Alliance 1997



Donna Cox, Robert Patterson, Stuart Levy, NCSA Virtual Director Team
Glenn Wheless, Old Dominion Univ.



Illinois is Positioned to Seize National Optical Networking Leadership with I-WIRE Infrastructure Investment

- **State-Funded Infrastructure**

- **Application Driven**

- **High Definition Streaming Media**

- Telepresence and Media

- **Computational Grids**

- Cloud Computing

- **Data Grids**

- Search & Information Analysis

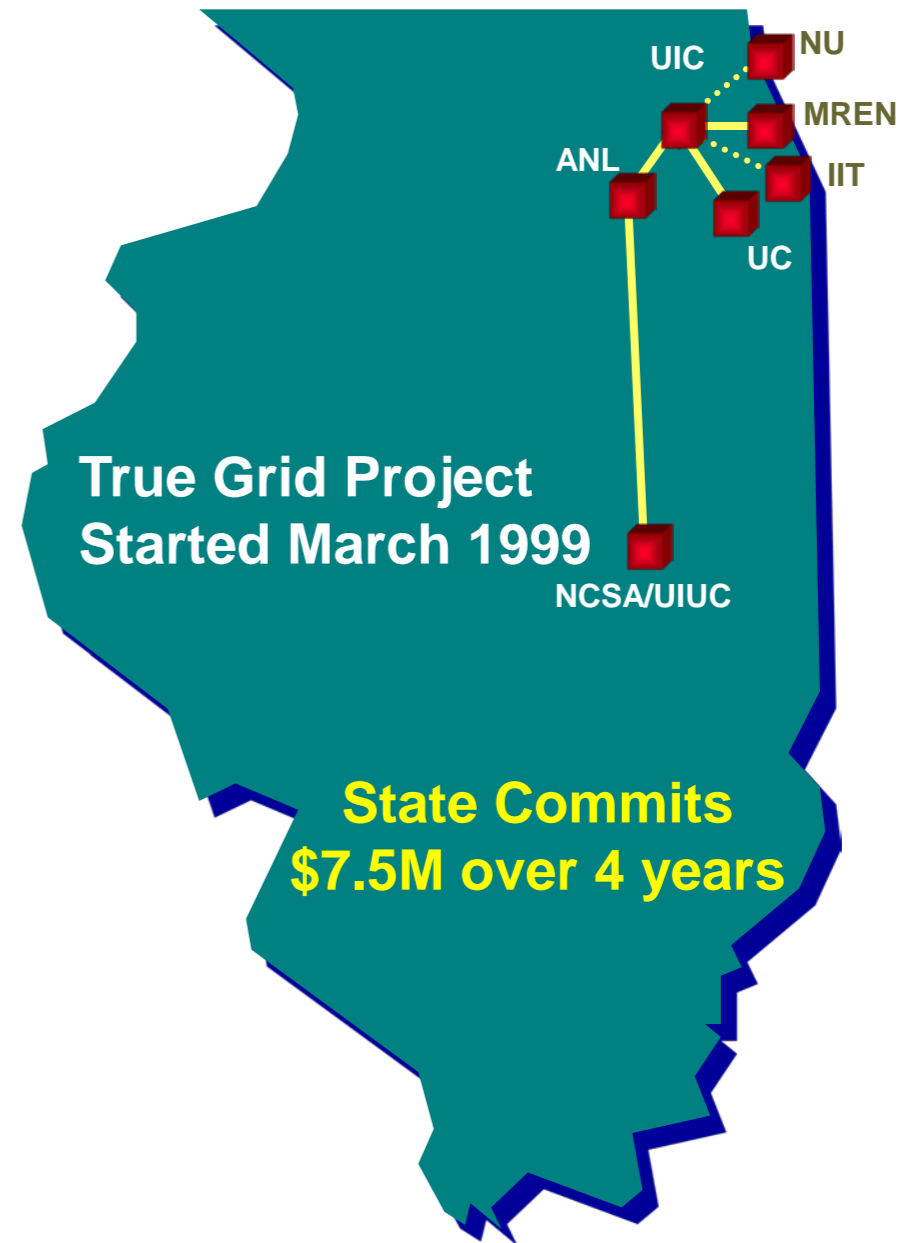
- **Emerging Tech Proving Ground**

- **Optical Switching**

- **Dense Wave Division Multiplexing**

- **Advanced Middleware Infrastructure**

- **Wireless Extensions**



Two New Calit2 Buildings Provide New Laboratories for “Living in the Future”



UC Irvine

- “Convergence” Laboratory Facilities
 - Nanotech, BioMEMS, Chips, Radio, Photonics
 - Virtual Reality, Digital Cinema, HDTV, Gaming
- Over 1000 Researchers in Two Buildings
 - **Linked via Dedicated Optical Networks**



www.calit2.net

Preparing for a World in Which
Distance is Eliminated...



Linking the Calit2 Auditoriums at UCSD and UCI With HD Streams




Photo by Erik Jepsen, UC San Diego





NSF's OptIPuter Project: Using Supernetworks to Meet the Needs of Data-Intensive Researchers

OptIPortal—
Termination
Device
for the
OptIPuter
Global
Backplane



2003-2009
\$13,500,000

In August 2003, Jason Leigh and his students used RBUDP to blast data from NCSA to SDSC over the TeraGrid DTFnet, achieving 18Gbps file transfer out of the available 20Gbps



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



High Resolution Uncompressed HD Streams Require Multi-Gigabit/s Lambdas

U. Washington




Prof. Smarr

Telepresence Using Uncompressed 1.5 Gbps
HDTV Streaming Over IP on Fiber Optics--
75x Home Cable "HDTV" Bandwidth!

JGN II Workshop
Osaka, Japan
Jan 2005

Osaka



Prof. Aoyama



"I can see every hair on your head!"—Prof. Aoyama

Source: U Washington Research Channel



First Trans-Pacific Super High Definition Telepresence Meeting Using Digital Cinema 4k Streams



4k = 4000x2000 Pixels = 4xHD

100 Times
the Resolution
of YouTube!

Streaming 4k
with JPEG 2000
Compression
1/2 gigabit/sec

Keio University
President Anzai

UCSD
Chancellor Fox

September 26-30, 2005

21 Countries Driving 50 Demonstrations
1 or 10Gbps to Calit2@UCSD Building Sept 2005



Lays
Technical
Basis for
Global
Digital
Cinema

Sony
NTT
SGI



Globally 10Gbp Optically Connected Digital Cinema Collaboratory




**Streaming 4K Live From NCSA Servers to Calit2@UCSD Auditorium
Content: Donna Cox, Robert Patterson, NCSA**



Project StarGate Goals: Combining Supercomputers and Supernetworks

- Create an “End-to-End”
10Gbps Workflow
- Explore Use of OptIPortals as
Petascale Supercomputer
“Scalable Workstations”
- Exploit Dynamic 10Gbps
Circuits on ESnet
- Connect Hardware Resources
at ORNL, ANL, SDSC
- Show that Data Need Not be
Trapped by the Network
“Event Horizon”

OptIPortal@SDSC



Rick Wagner

Mike Norman

Source: Michael Norman, SDSC, UCSD

- ANL * Calit2 * LBNL * NICS * ORNL * SDSC



Using Supernetworks to Couple End User to Remote Supercomputers and Visualization Servers

Source: Mike Norman,
Rick Wagner, SDSC



Demoed
SC09

**Real-Time Interactive
Volume Rendering Streamed
from ANL to SDSC**

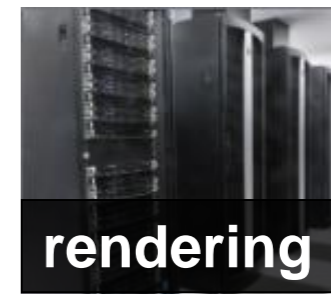
SDSC



Calit2/SDSC OptIPortal1
20 30" (2560 x 1600 pixel) LCD panels
10 NVIDIA Quadro FX 4600 graphics
cards > 80 megapixels
10 Gb/s network throughout

Argonne NL

DOE Eureka
100 Dual Quad Core Xeon Servers
200 NVIDIA Quadro FX GPUs in 50
Quadro Plex S4 1U enclosures
3.2 TB RAM



ESnet
10 Gb/s fiber optic network

**NICS
ORNL**


NSF TeraGrid Kraken
Cray XT5
8,256 Compute Nodes
99,072 Compute Cores
129 TB RAM



***ANL * Calit2 * LBNL * NICS * ORNL * SDSC**
www.calit2.net/newsroom/release.php?id=1624



Integrated "OptPlatform" Cyberinfrastructure System: A 10Gbps Lightpath Cloud



End User
OptIPortal



Instruments



HD/4k Telepresence



HD/4k Video Cams



Campus
Optical
Switch

10G
Lightpath



National LambdaRail




HPC

TeraGrid



Data Repositories & Clusters



HD/4k Video Images

LS 2009
Slide


So Why Don't We Have a National Big Data Cyberinfrastructure?



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



DOE Esnet's Science DMZ: A Scalable Network Design Model for Optimizing Science Data Transfers

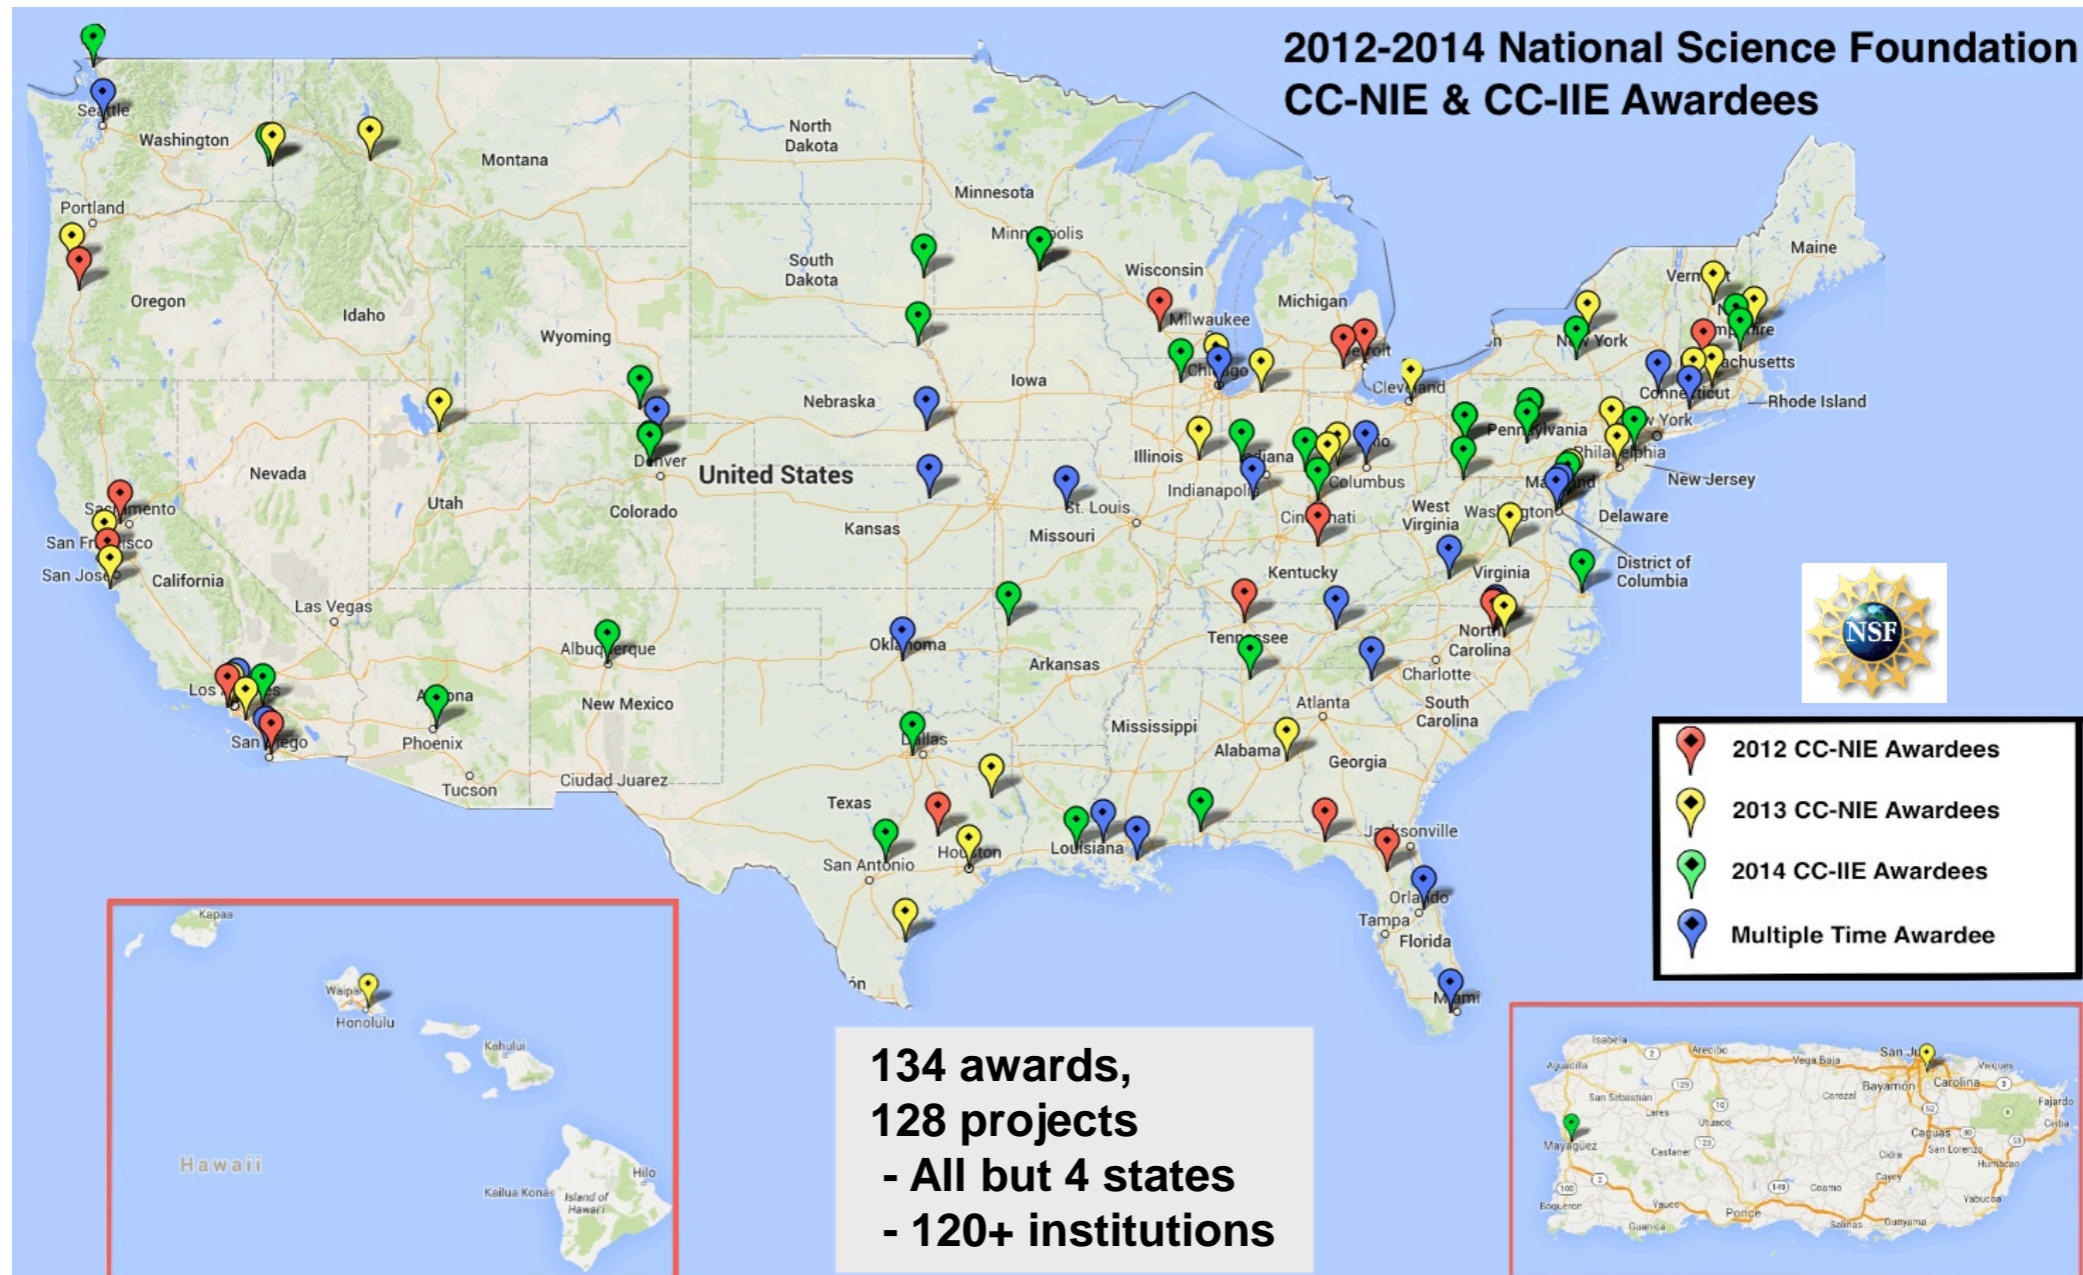
- **A Science DMZ integrates 4 key concepts into a unified whole:**
 - A network architecture designed for high-performance applications, with the science network distinct from the general-purpose network
 - The use of dedicated systems for data transfer
 - Performance measurement and network testing systems that are regularly used to characterize and troubleshoot the network
 - Security policies and enforcement mechanisms that are tailored for high performance science environments



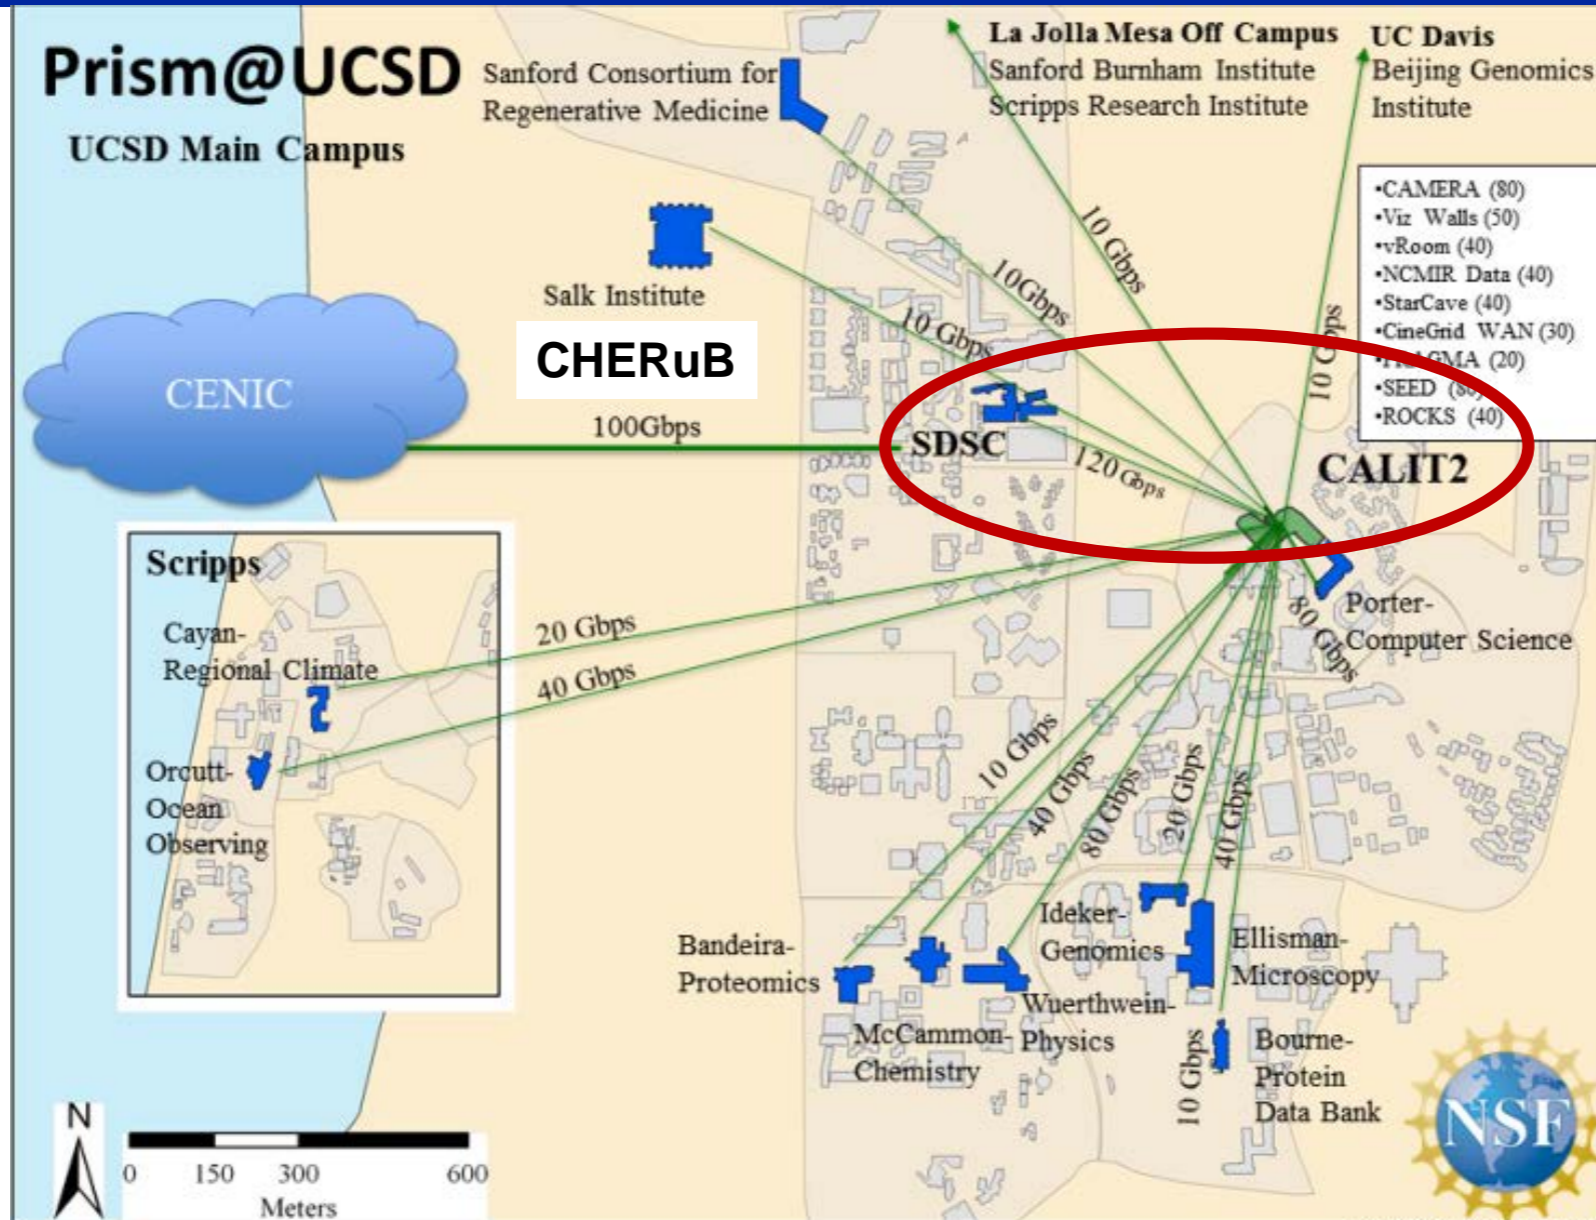
<http://fasterdata.es.net/science-dmz/>



The National Science Foundation Has Funded Over 100 Campuses to Build Local Data Freeways



Creating a “Big Data” Plane on Campus: NSF CC-NIE Funded Prism@UCSD and CHerUB



Prism@UCSD, Phil Papadopoulos, SDSC, Calit2, PI
CHerUB, Mike Norman, SDSC PI

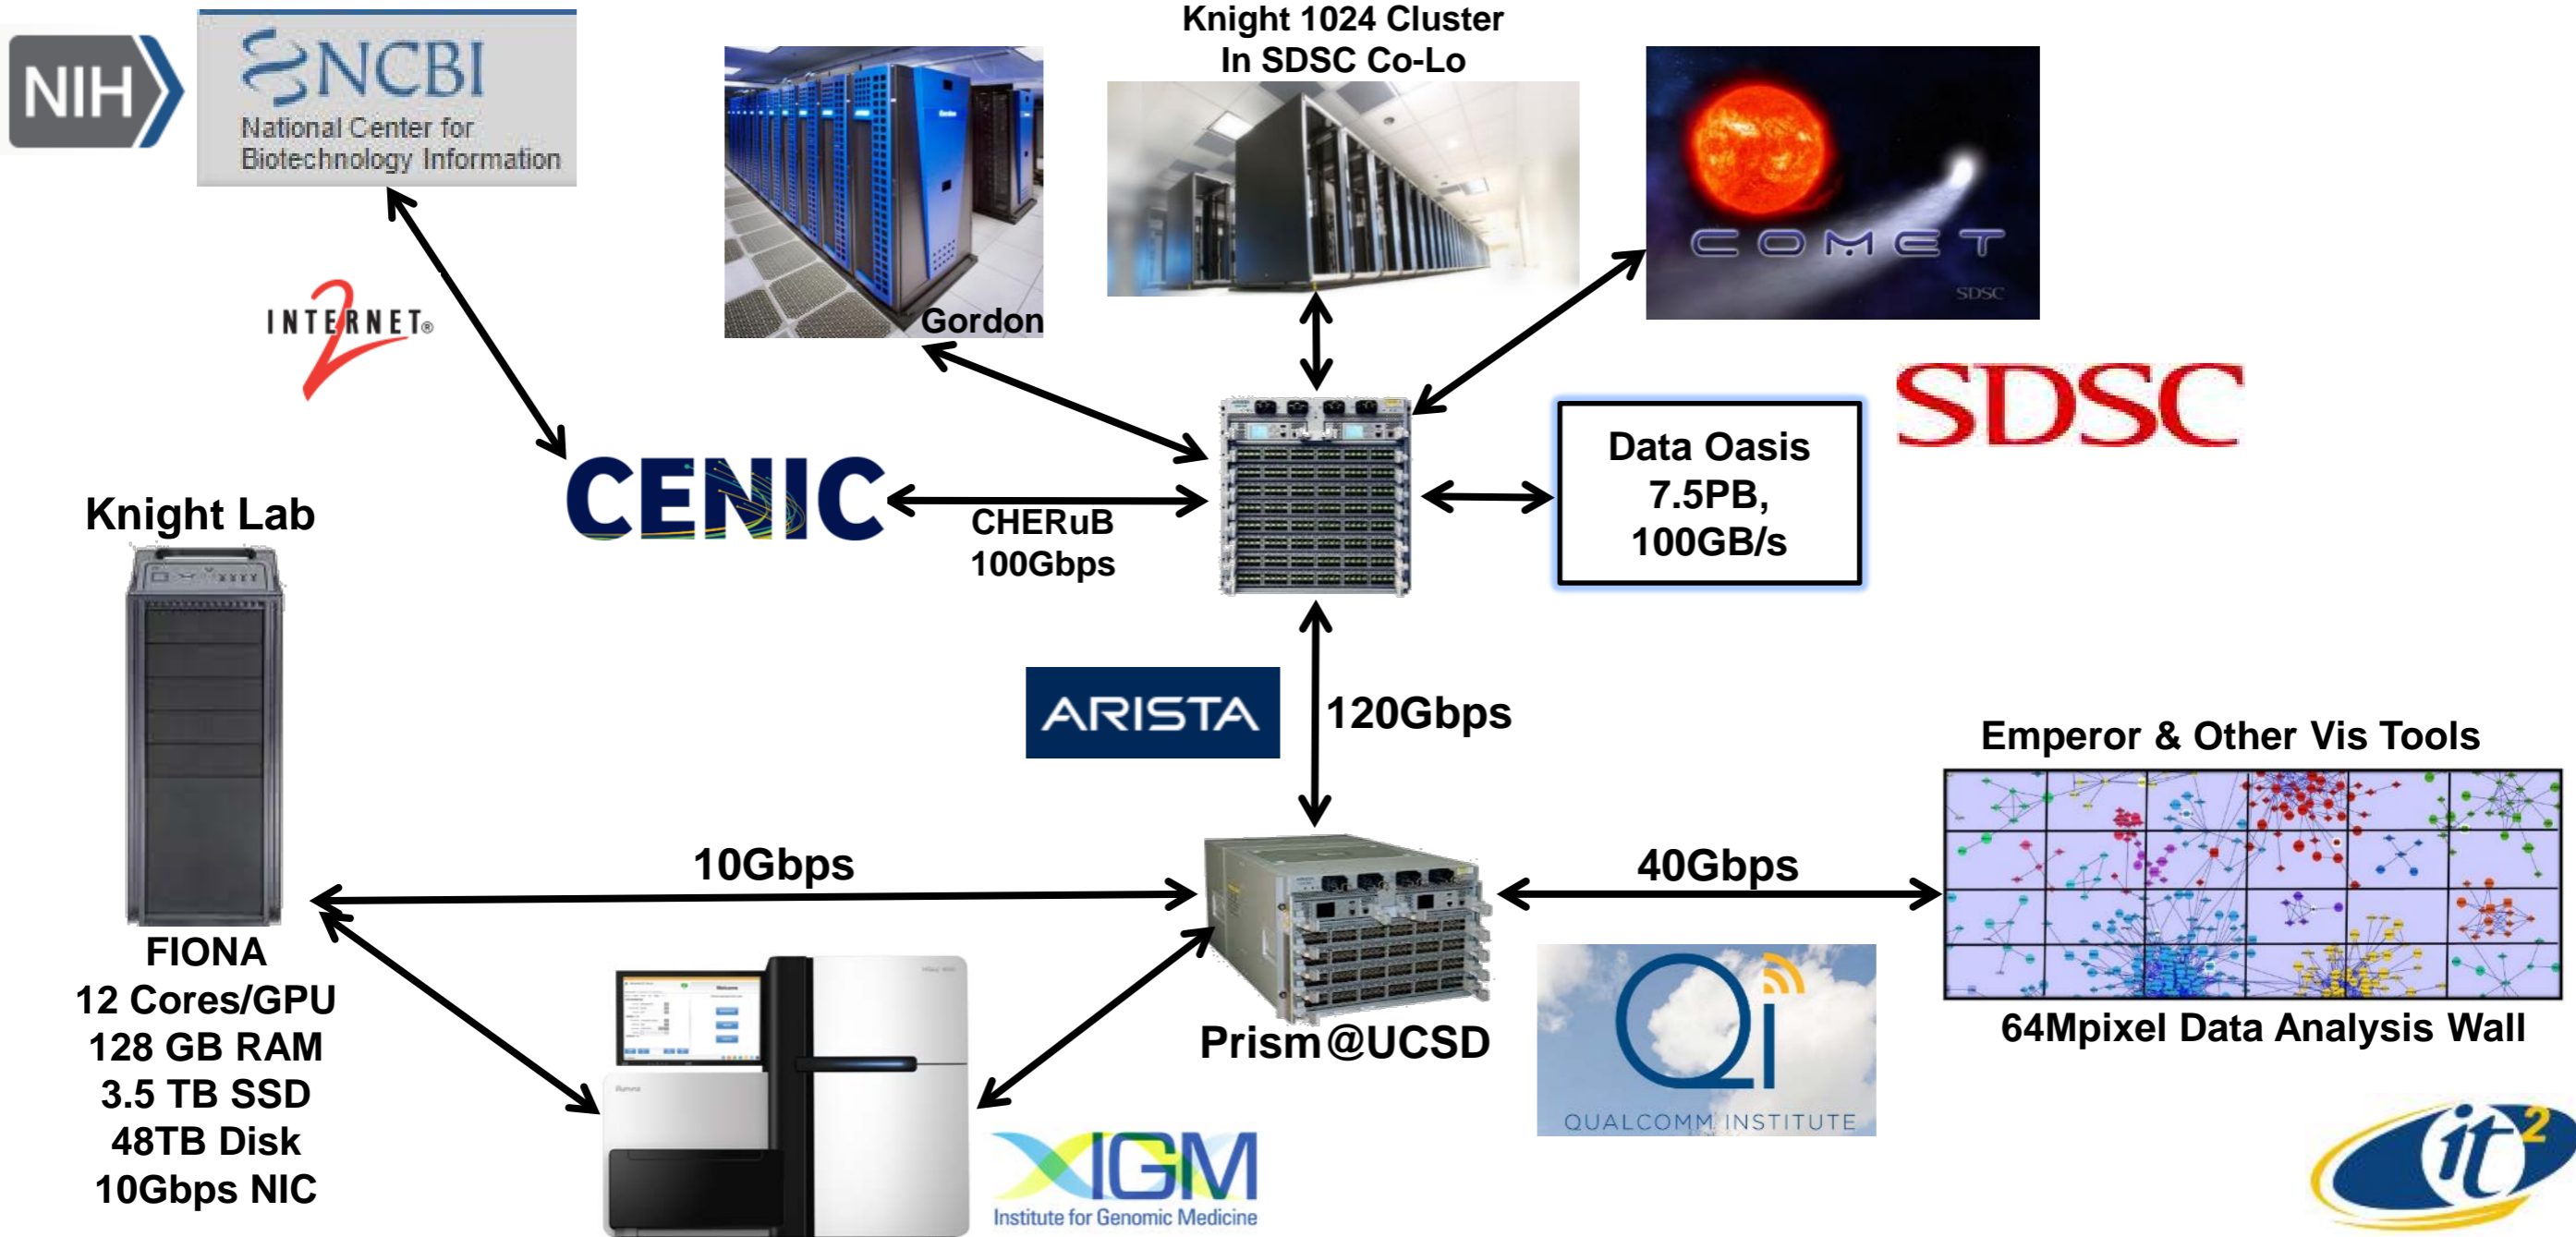


Science DMZ Data Transfer Nodes - Optical Network Termination Devices: Inexpensive PCs Optimized for Big Data


- **FIONA – Flash I/O Network Appliance**
 - Combination of Desktop and Server Building Blocks
 - US\$5K - US\$7K
 - Desktop Flash up to 16TB
 - RAID Drives up to 48TB
 - 10GbE/40GbE Adapter
 - Tested speed 40Gbs
 - Developed Under UCSD CC-NIE Prism Award by UCSD's
 - **Phil Papadopoulos**
 - **Tom DeFanti**
 - **Joe Keefe**



Integrated Digital Cyberinfrastructure Supporting Knight Lab



Interactively Exploring Microscope Images of Brains: 40Gbps From NCMIR to Calit2 64Mpixel Wall



NCMIR

National Center for Microscopy and Imaging Research



Why Now?

Federating the Six UC CC-NIE Grants

UC IT Leadership Council
Oakland, CA
May 19, 2014

- **2011 ACCI Strategic Recommendation to the NSF #3:**
 - NSF should create a new program funding high-speed (currently 10 Gbps) connections from campuses to the nearest landing point for a national network backbone. The design of these connections must include support for dynamic network provisioning services and must be engineered to support rapid movement of large scientific data sets."
 - - pg. 6, NSF Advisory Committee for Cyberinfrastructure Task Force on Campus Bridging, Final Report, March 2011
 - www.nsf.gov/od/oci/taskforces/TaskForceReport_CampusBridging.pdf
 - Led to Office of Cyberinfrastructure RFP March 1, 2012
- **NSF's Campus Cyberinfrastructure – Network Infrastructure & Engineering (CC-NIE) Program**
 - 85 Grants Awarded So Far (NSF Summit Last Week)
 - 6 Are in UC

UC Must Move Rapidly or Lose a Ten-Year Advantage!



CENIC is Rapidly Moving to Connect at 100 Gbps Across the State and Nation

DOE

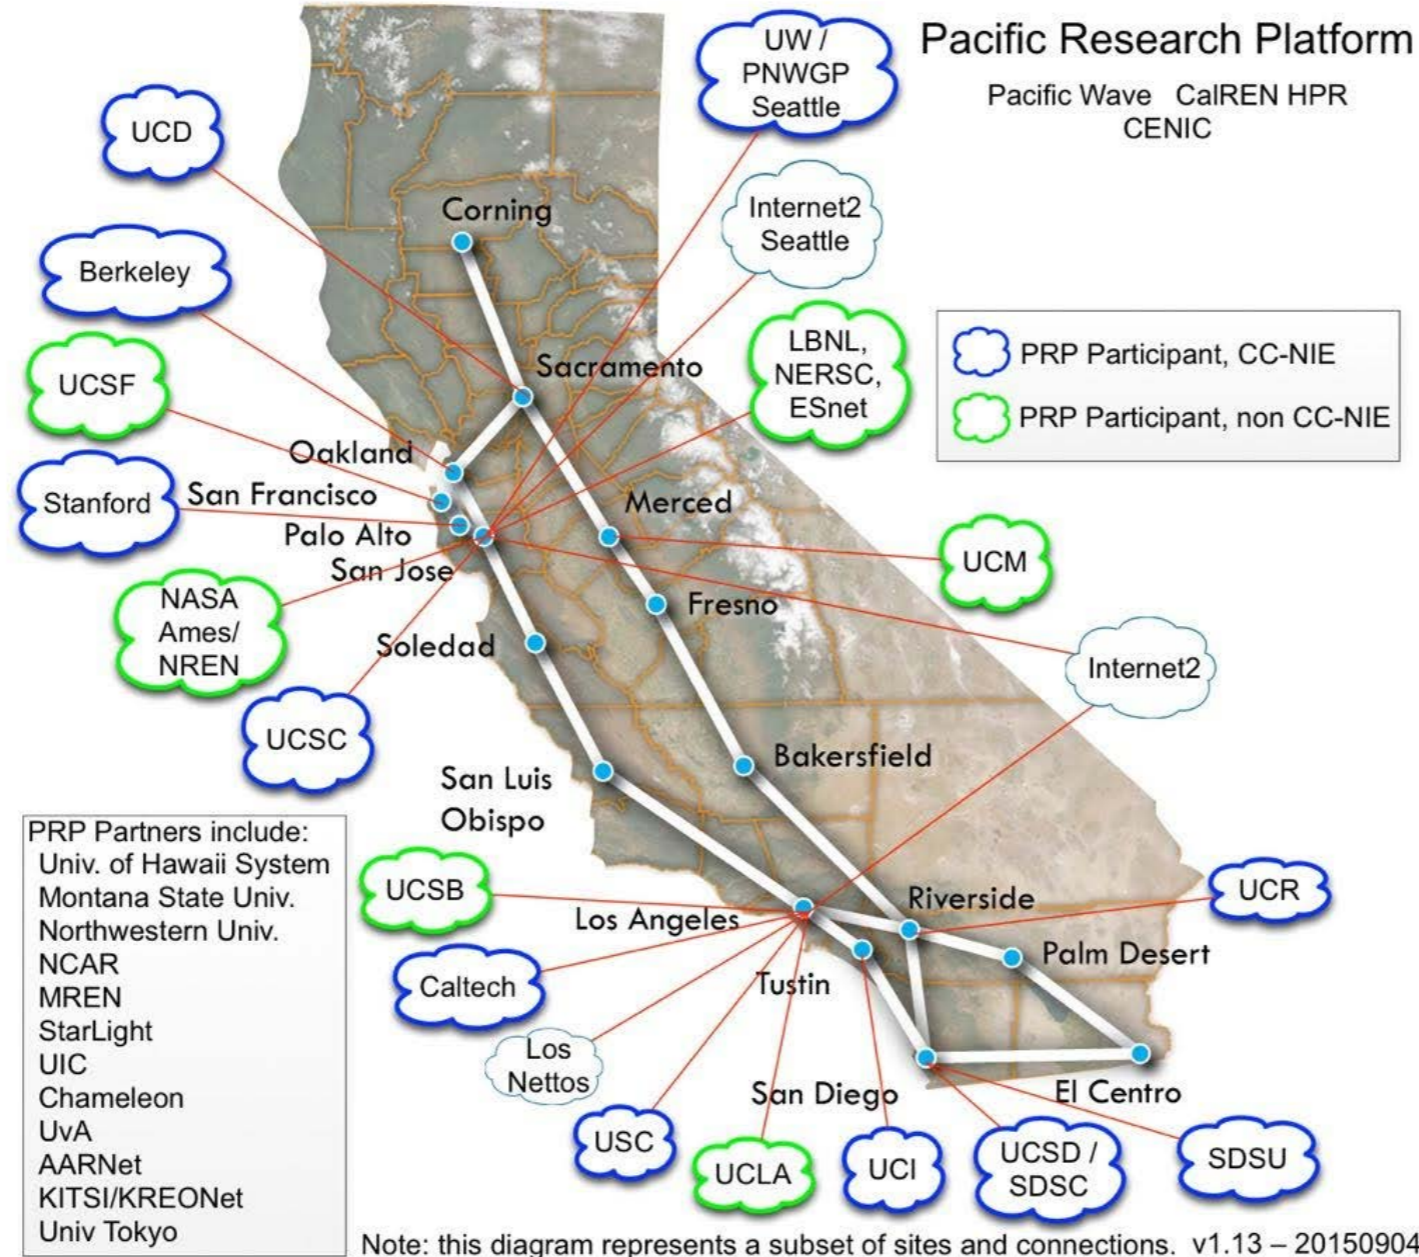


Internet2



The Pacific Wave Platform

Creates an End-to-End Regional Science Big Data Freeway



Source:
John Hess, CENIC



Ten Week Sprint to Demonstrate the West Coast Big Data Freeway System: PRPv0

FIONA DTNs Now Deployed to All UC Campuses And Most PRP Sites

As of 3/9/15, the Pacific Research Platform (PRPv0) as a facility, logs rather good performance:


From	To	Measured Bandwidth	Data Transfer Utility
San Diego State Univ.	UC Los Angeles	5Gb/s out of 10	GridFTP
UC Riverside	UC Los Angeles	9Gb/s out of 10	GridFTP
UC Berkeley	UC San Diego	9.6Gb/s out of 10	GridFTP
UC Davis	UC San Diego	9.6Gb/s out of 10	GridFTP
UC Irvine	UC Los Angeles	9.6Gb/s out of 10	GridFTP
UC Santa Cruz	UC San Diego	9.6Gb/s out of 10	FDT
Stanford	UC San Diego	12Gb/s out of 40	FDT
Univ. of Washington	UC San Diego	12Gb/s out of 40	FDT
UC Los Angeles	UC San Diego	36Gb/s out of 40	FDT
Caltech	UC San Diego	36Gb/s out of 40	FDT

Table I.2.1: Bandwidth of flash disk-to-flash disk file transfers shown between several sites for the existing experimental facility “PRPv0.”

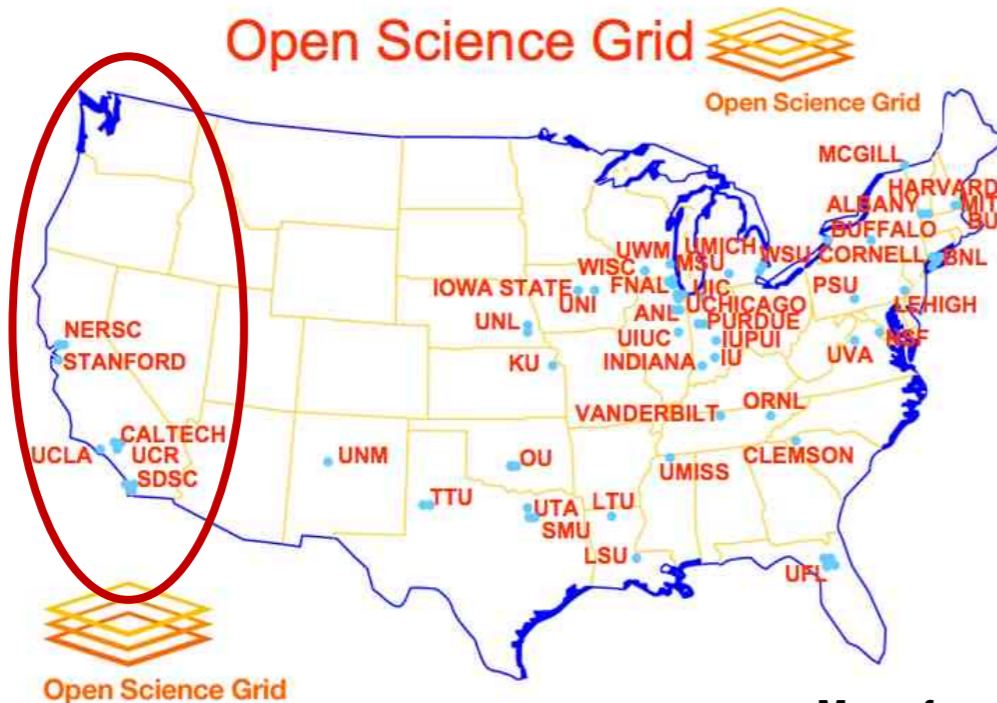

Pacific Research Platform Multi-Campus Science Driver Teams

- **Particle Physics**
- **Astronomy and Astrophysics**
 - Telescope Surveys
 - Galaxy Evolution
 - Gravitational Wave Astronomy
- **Biomedical**
 - Cancer Genomics Hub/Browser
 - Microbiome and Integrative 'Omics
 - Integrative Structural Biology
- **Earth Sciences**
 - Data Analysis and Simulation for Earthquakes and Natural Disasters
 - Climate Modeling: NCAR/UCAR
 - California/Nevada Regional Climate Data Analysis
 - CO2 Subsurface Modeling
- **Scalable Visualization, Virtual Reality, and Ultra-Resolution Video**

Particle Physics: Creating a 10-100 Gbps LambdaGrid to Support LHC Researchers



LHC Data
Generated by
CMS & ATLAS
Detectors
Analyzed
on OSG



Maps from www.uslhq.us



Two Automated Telescope Surveys Creating Huge Datasets Will Drive PRP



**Precursors to
LSST
And
NCSA**

INTERMEDIATE PALOMAR TRANSIENT FACTORY

**300 images per night.
100MB per raw image**

30GB per night

120GB per night

Dark Energy Spectroscopic Instrument

**250 images per night.
530MB per raw image**

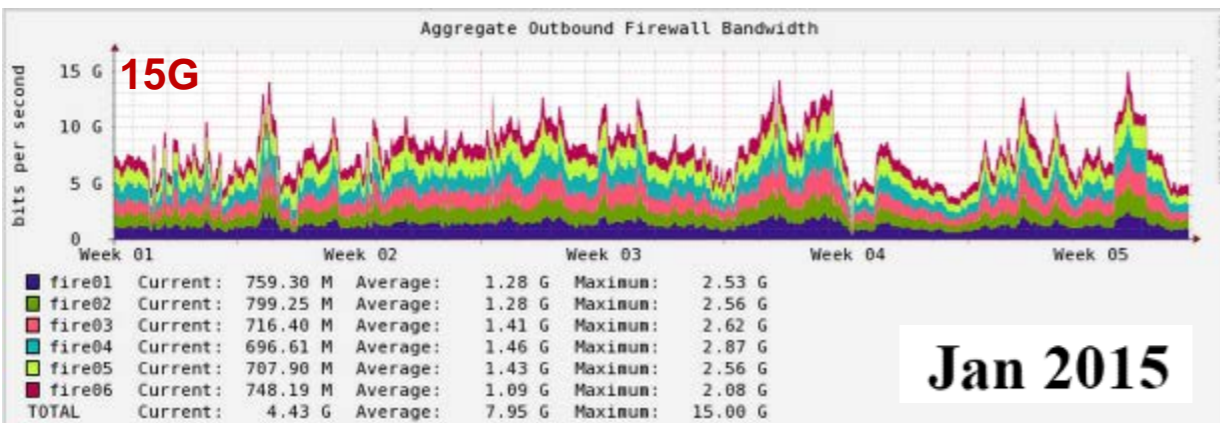


150 GB per night

800GB per night

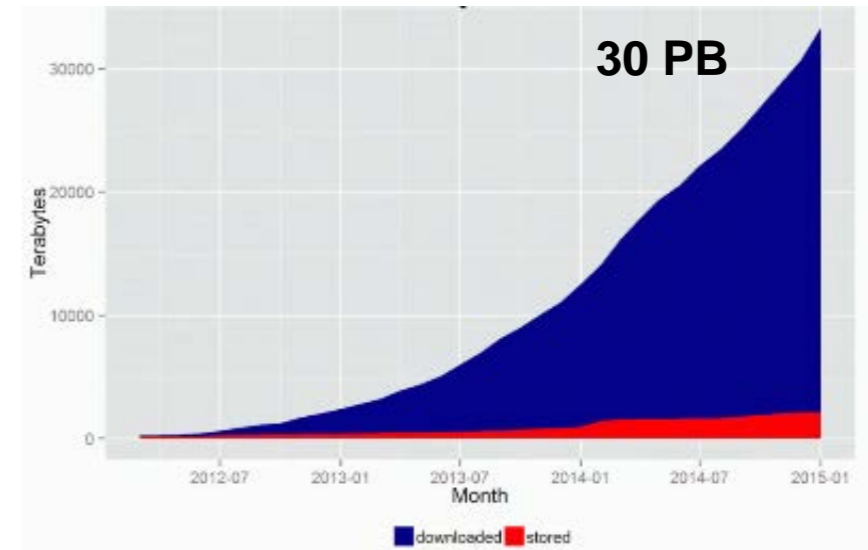
**When processed
at NERSC
Increased by 4x**



Cancer Genomics Hub (UCSC) is Housed in SDSC CoLo: Large Data Flows to End Users



**Cumulative TBs of CGH
Files Downloaded**



Data Source: David Haussler,
Brad Smith, UCSC



To Map Out the Dynamics of Autoimmune Microbiome Ecology Couples Next Generation Genome Sequencers to *Big Data* Supercomputers

Source: Weizhong Li, UCSD

Our Team Used 25 CPU-years
to Compute
Comparative Gut Microbiomes
Starting From
2.7 Trillion DNA Bases
of My Samples
and Healthy and IBD Controls


SDSC Gordon Data Supercomputer



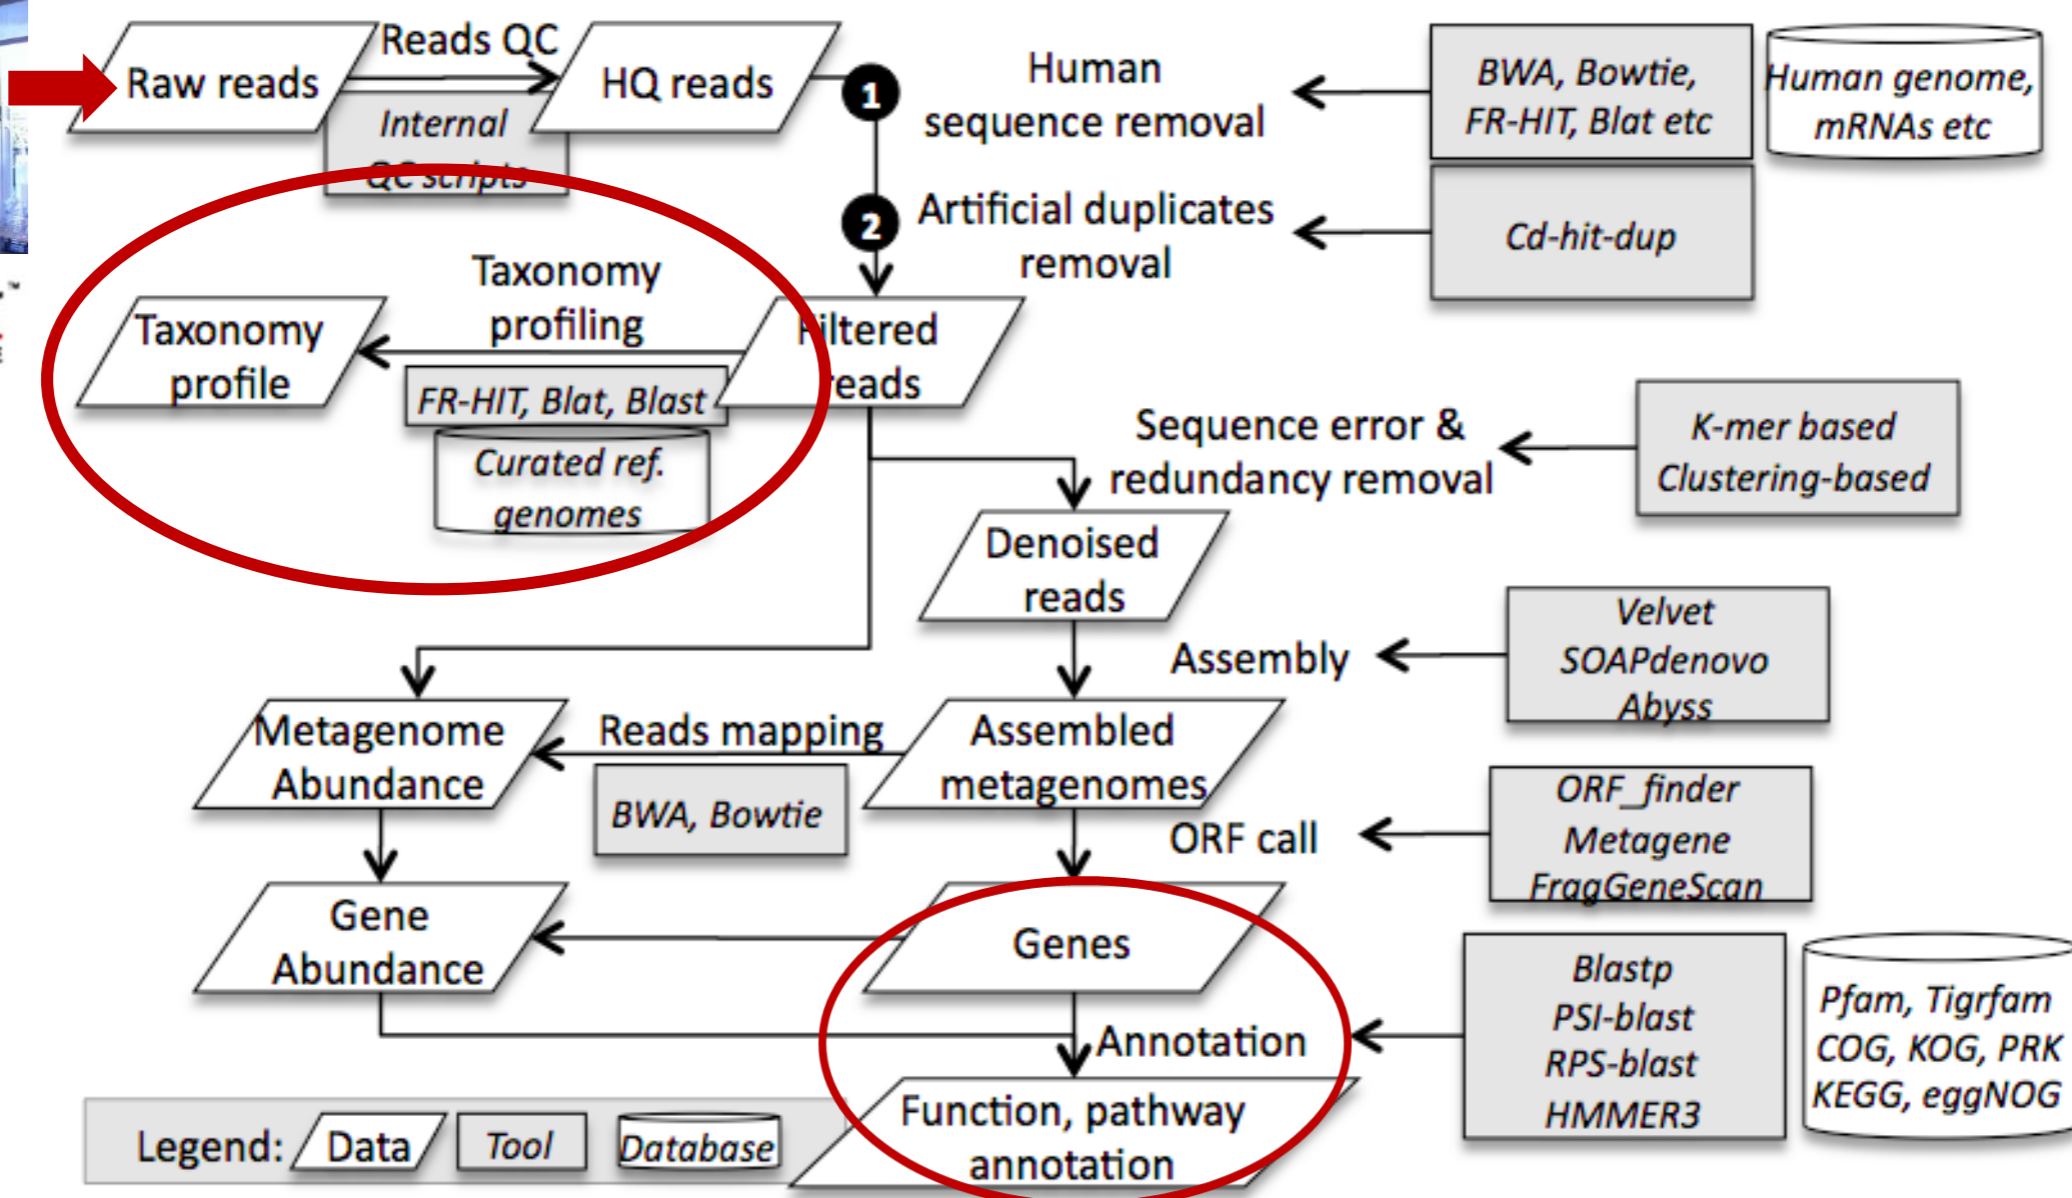
Dell Solutions Center
Industry Solutions Lab
SANGER DSU



Computing on Data: Complex Software Pipelines - From Sequence to Taxonomy and Function



J. Craig Venter™
INSTITUTE



PI: (Weizhong Li, CRBS, UCSD):
NIH R01HG005978 (2010-2013, \$1.1M)



Planning for climate change in California

substantial shifts on top of already high climate variability

**SIO Campus Climate Researchers Need to Download
Results from Remote Supercomputer Simulations
to Make Regional Climate Change Forecasts**

Dan Cayan

USGS Water Resources Discipline

Scripps Institution of Oceanography, UC San Diego

much support from Mary Tyree, Mike Dettinger, Guido Franco and
other colleagues

Sponsors:

California Energy Commission

NOAA RISA program

California DWR, DOE, NSF



Earth Sciences: Pacific Earthquake Engineering Research Center



University of California, Berkeley -
Lead Institution



California Institute of Technology



Oregon State University



Stanford University

UC DAVIS

University of California, Davis



University of California, Irvine



University of California, Los Angeles



University of California, San Diego

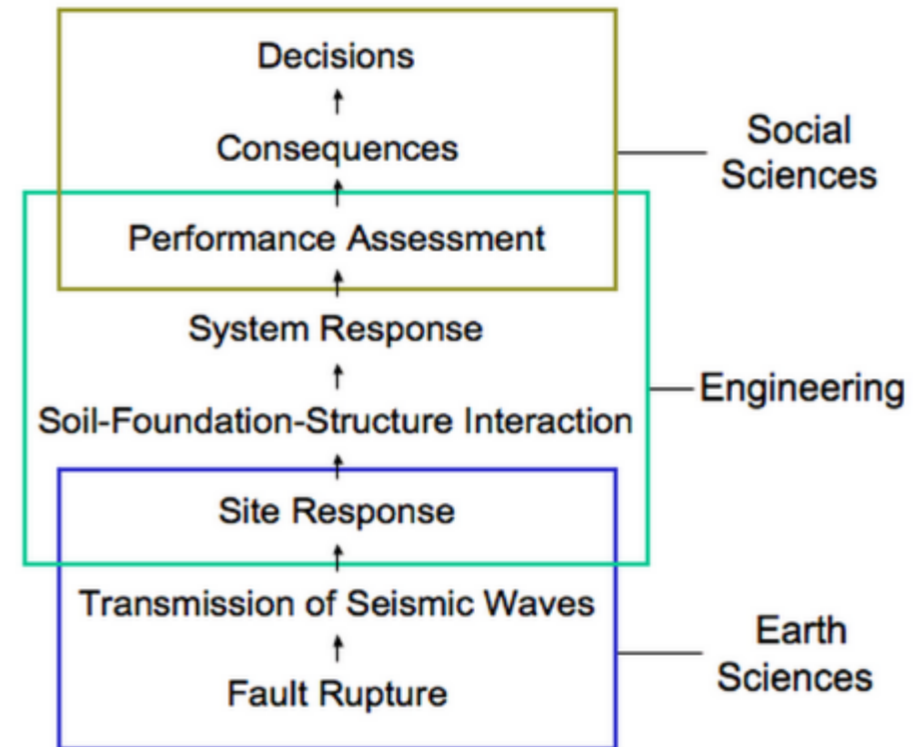


University of Southern California




University of Washington

**Enabling
Real-Time Coupling
Between
Shake Tables
and
Supercomputer
Simulations**



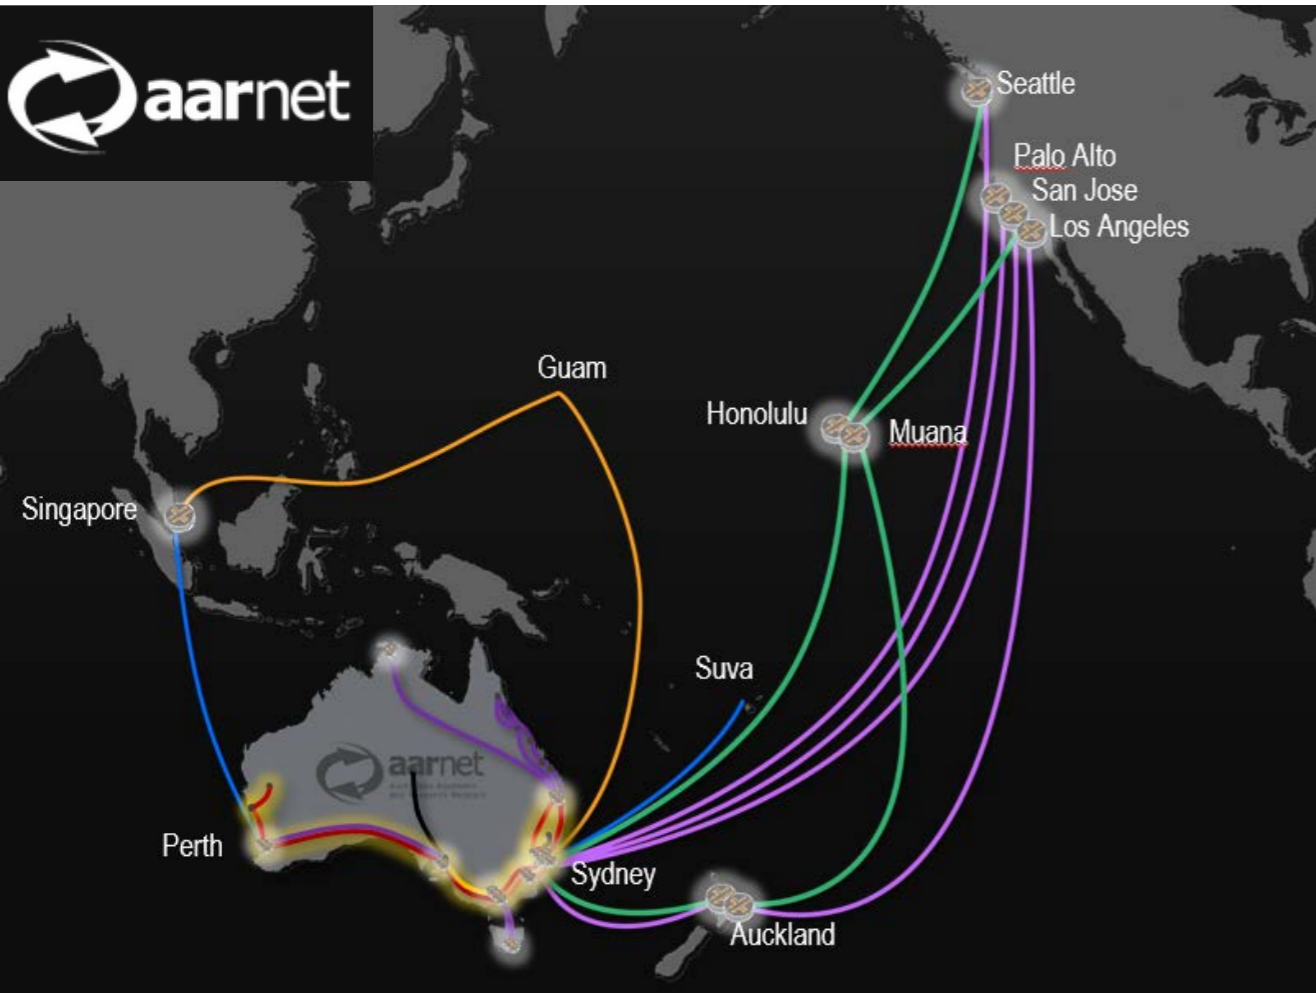
Collaboration Between EVL's CAVE2 and Calit2's VROOM Over 10Gb Wavelength



Source: NTT Sponsored ON*VECTOR Workshop at Calit2 March 6, 2013

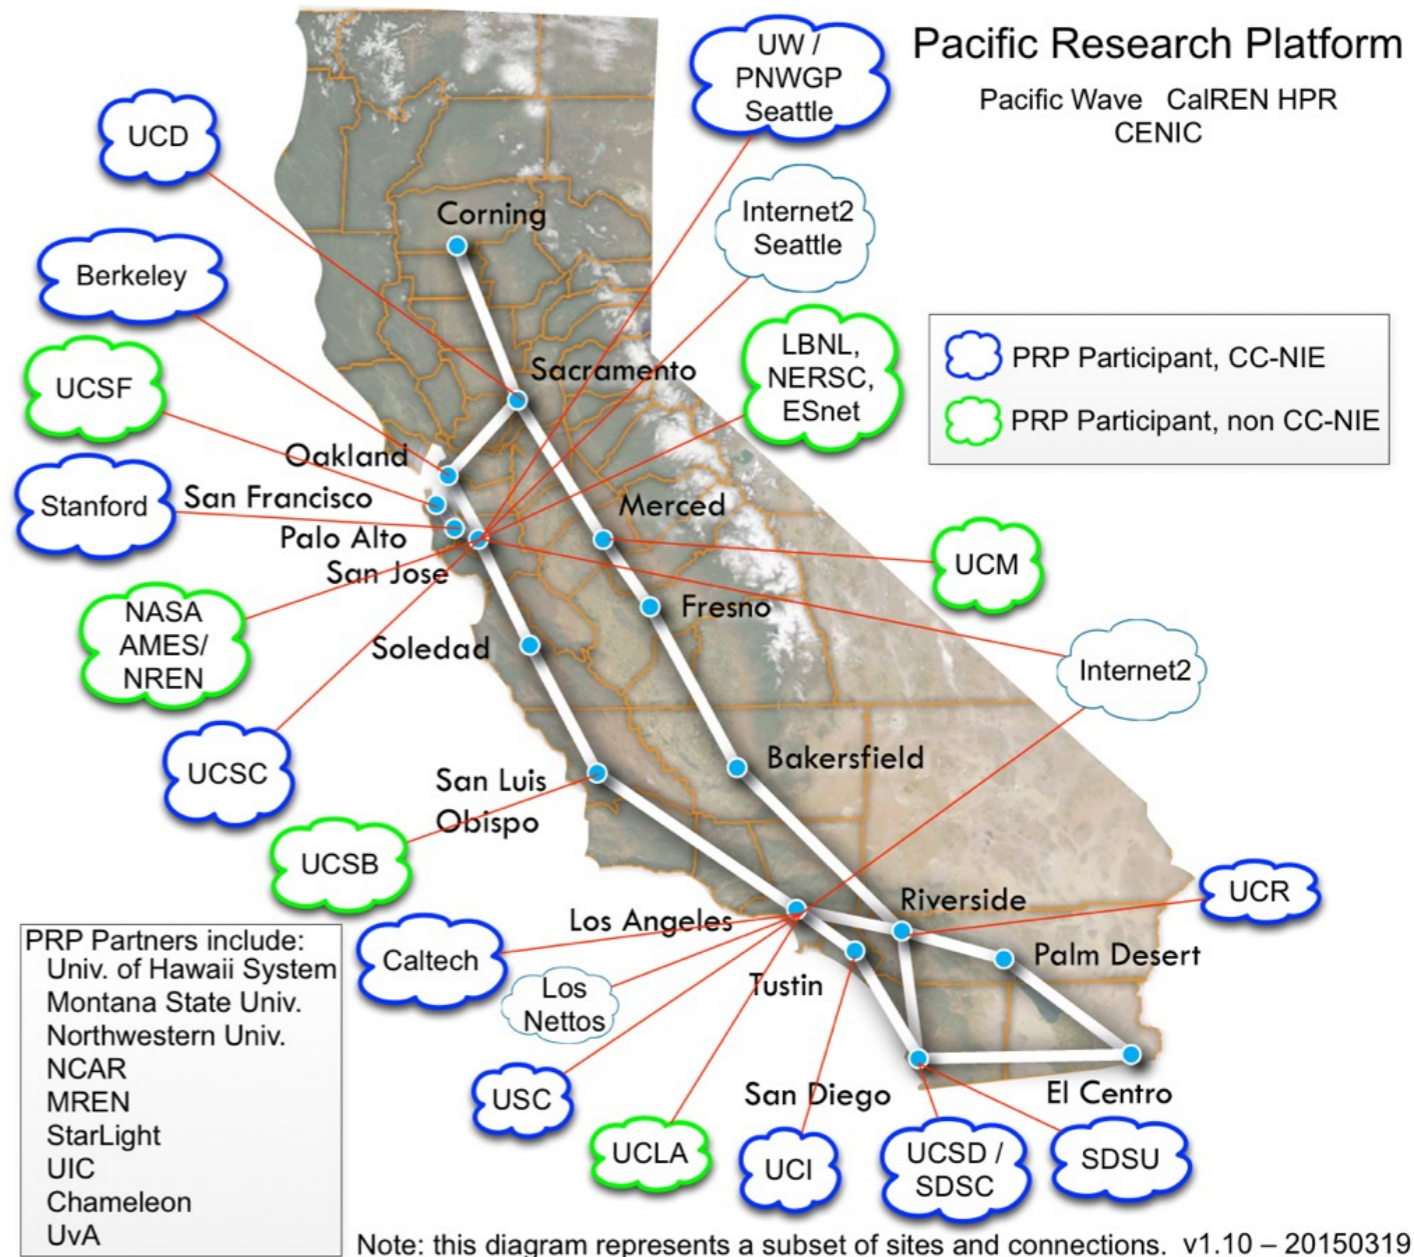


Optical Fibers Link Australian and US Big Data Researchers



Australia forges US partnership to accelerate scientific discovery and innovation

August 5, 2015



Next Step: Use AARnet/PRP to Set Up Planetary-Scale Shared Virtual Worlds



Digital Arena, UTS Sydney



CAVE2, Monash U, Melbourne



CAVE2, EVL, Chicago

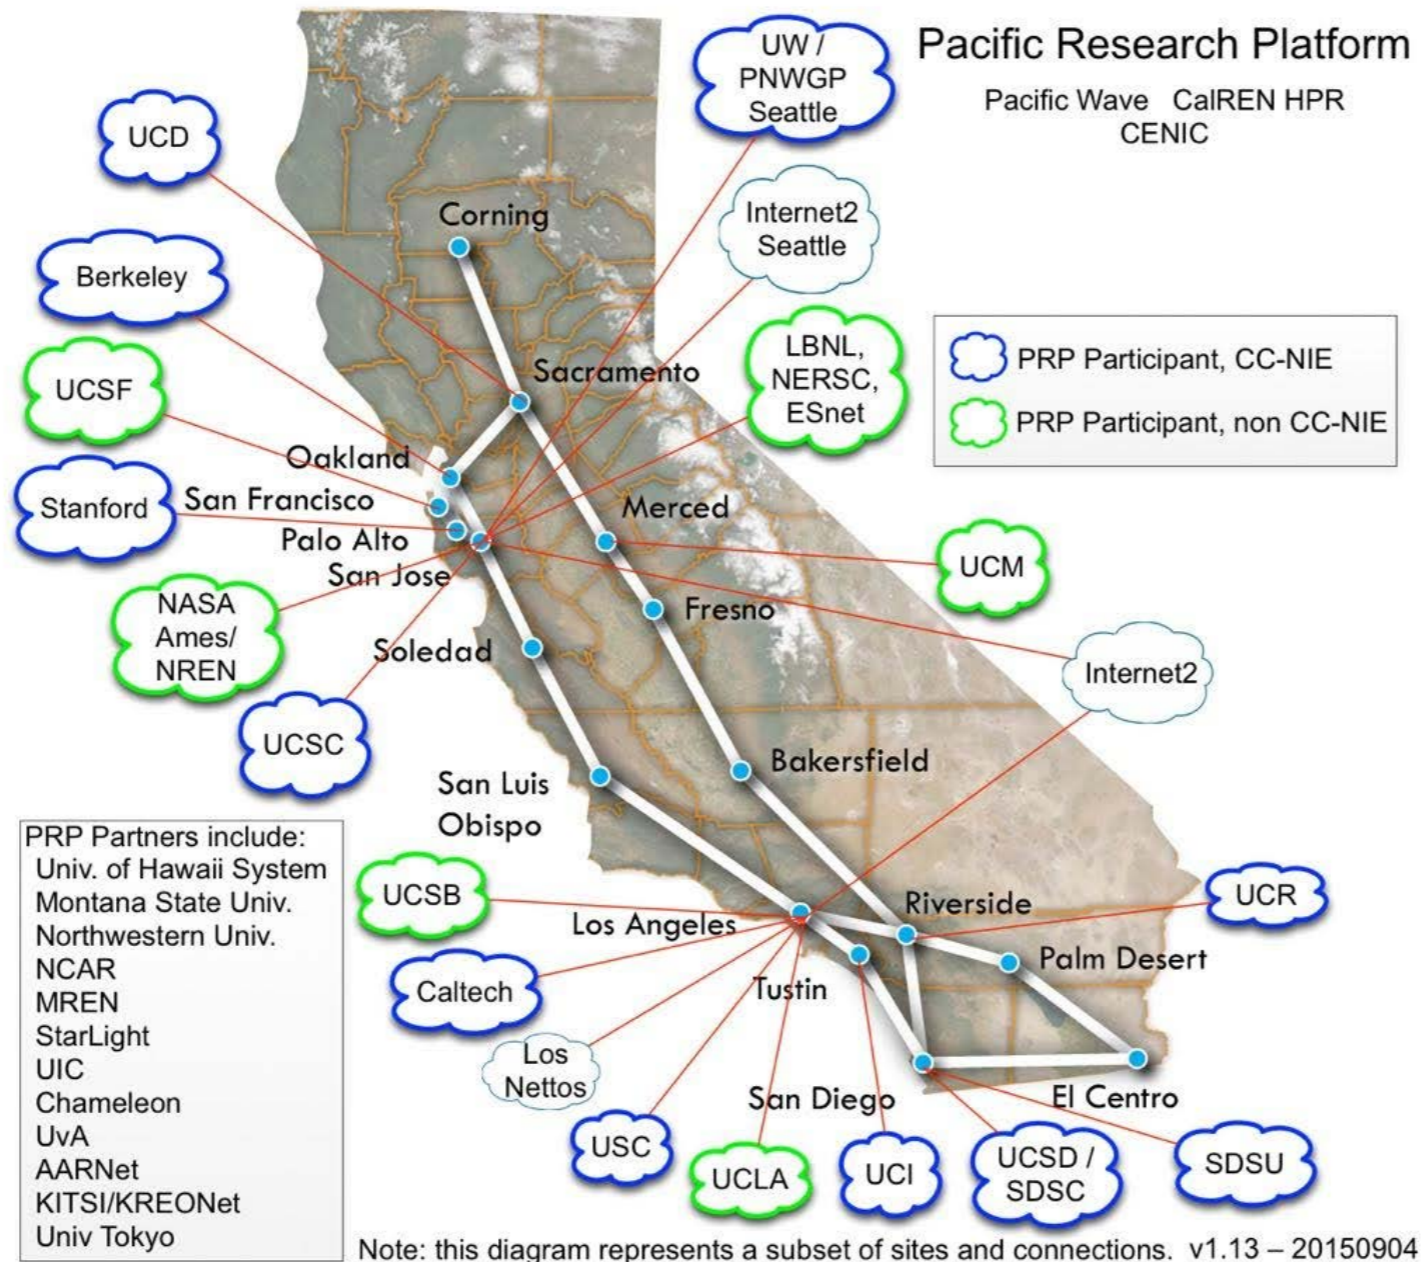
PRP Timeline

- **PRPv1**
 - A Layer 2 and Layer 3 System
 - Completed In 2 Years
 - Tested, Measured, Optimized, With Multi-domain Science Data
 - Bring Many Of Our Science Teams Up
 - Each Community Thus Will Have Its Own Certificate-Based Access To its Specific Federated Data Infrastructure.
- **PRPv2**
 - Advanced Ipv6-Only Version with Robust Security Features
 - E.G., Trusted Platform Module Hardware and SDN/SDX Software
 - Support Rates up to 100Gb/s in Bursts And Streams
 - Develop Means to Operate a Shared Federation of Caches



The Pacific Wave Platform

Creates an End-to-End Regional Science Big Data Freeway



**Opportunity:
 Connect NCSA
 to End Users
 on PRP Campuses
 @10Gbps**

Source:
John Hess, CENIC

