

# **“Creating a Science-Driven Big Data Superhighway”**

**Remote Briefing to the Ad Hoc Big Data Task Force  
of the NASA Advisory Council Science Committee  
NASA Goddard Space Flight Center  
June 28, 2016**

**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>**



**Vision:  
Creating a Pacific Research Platform**

**Use Optical Fiber Networks to Connect  
All Data Generators and Consumers,  
Creating a “Big Data” Freeway System**

*“The Bisection Bandwidth of a Cluster Interconnect,  
but Deployed on a 20-Campus Scale.”*

**This Vision Has Been Building for 15 Years**



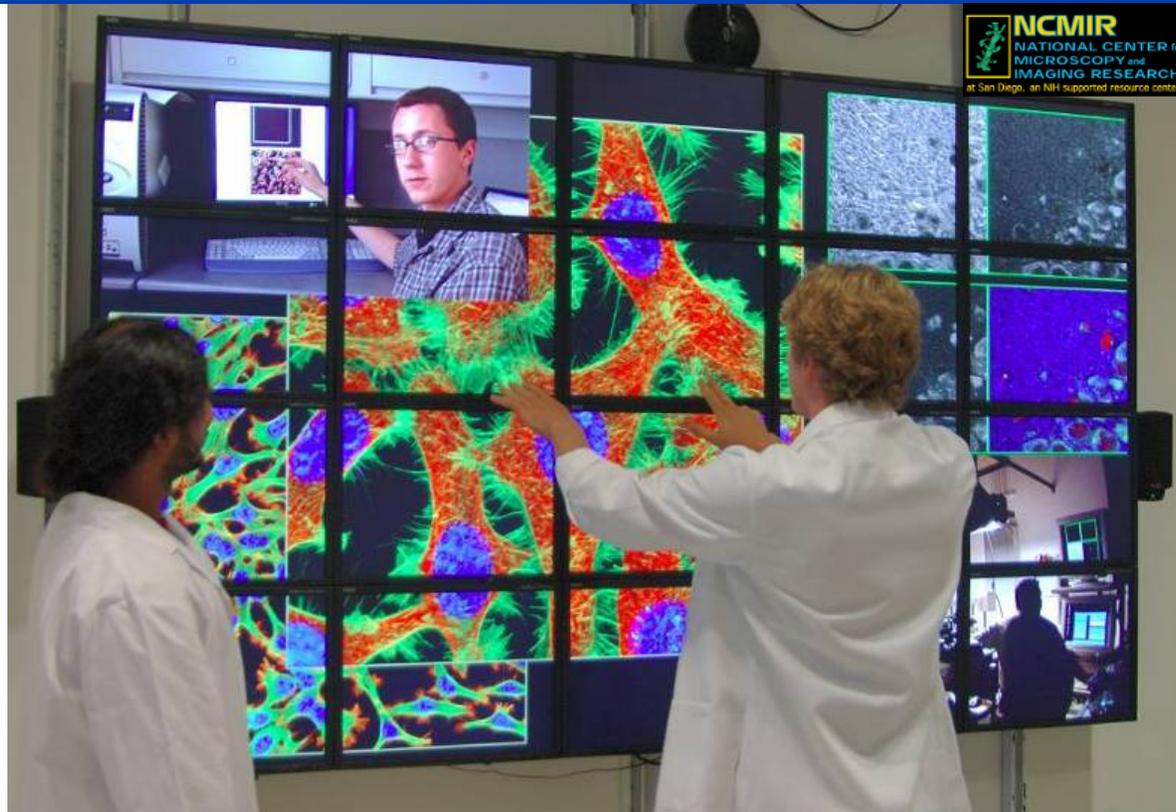
# NSF's OptIPuter Project: Demonstrating How SuperNetworks Can Meet the Needs of Data-Intensive Researchers

LS Slide 2005

OptIPortal-Termination Device for the OptIPuter Global Backplane



evl electronic visualization laboratory



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



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

SDSC  
SAN DIEGO SUPERCOMPUTER CENTER



## 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 as data transfer nodes (DTNs)
  - 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

The DOE ESnet Science DMZ and the NSF “Campus Bridging” Taskforce Report Formed the Basis for the NSF Campus Cyberinfrastructure Network Infrastructure and Engineering (CC-NIE) Program

Science DMZ  
Coined 2010



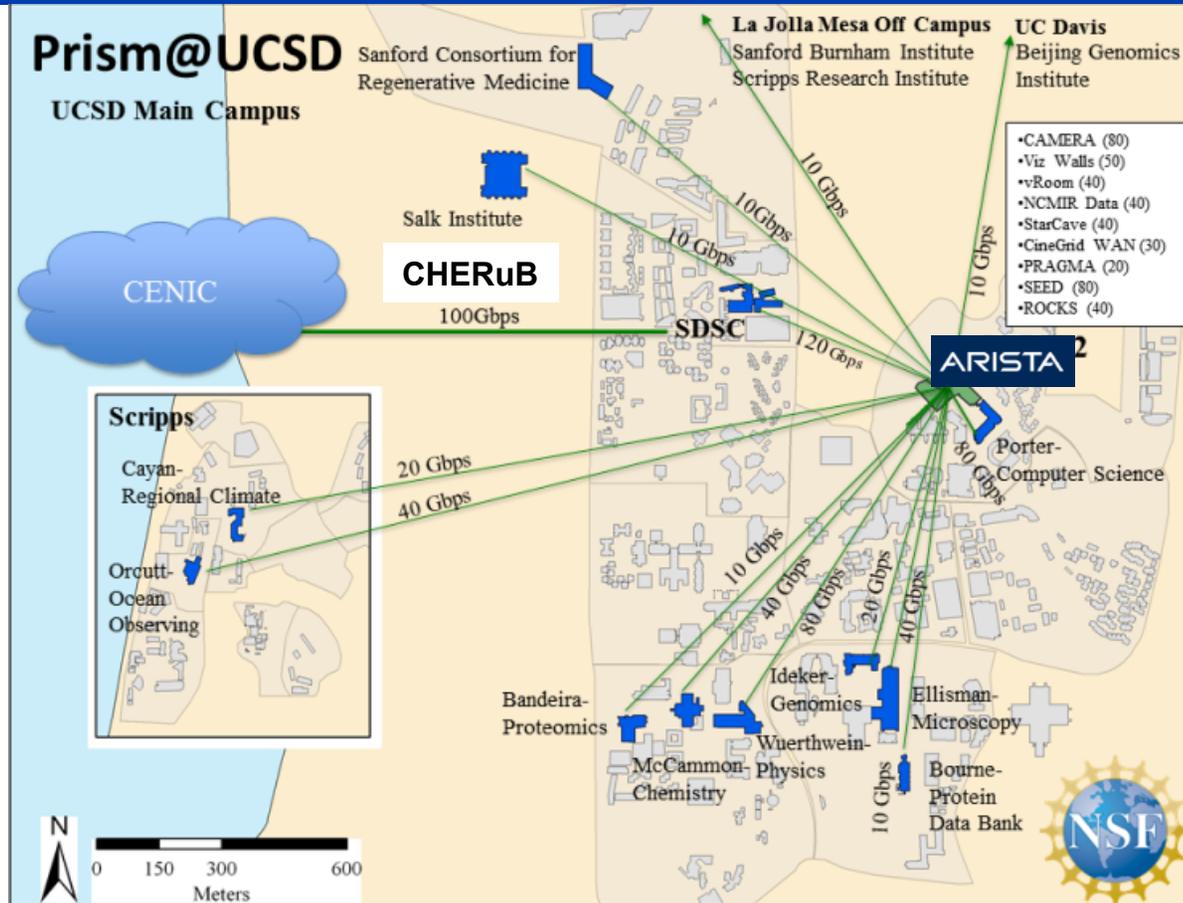
<http://fasterdata.es.net/science-dmz/>



# Creating a “Big Data” Freeway on Campus: NSF-Funded Prism@UCSD and CheruB Campus CC-NIE Grants

Prism@UCSD,  
PI Phil Papadopoulos,  
SDSC, Calit2,  
(2013-15)

CheruB,  
PI Mike Norman,  
SDSC



# FIONA – Flash I/O Network Appliance: Linux PCs Optimized for Big Data on DMZs

**FIONAs Are  
Science DMZ Data Transfer Nodes (DTNs) &  
Optical Network Termination Devices**

UCSD CC-NIE Prism Award & UCOP  
Phil Papadopoulos & Tom DeFanti  
Joe Keefe & John Graham



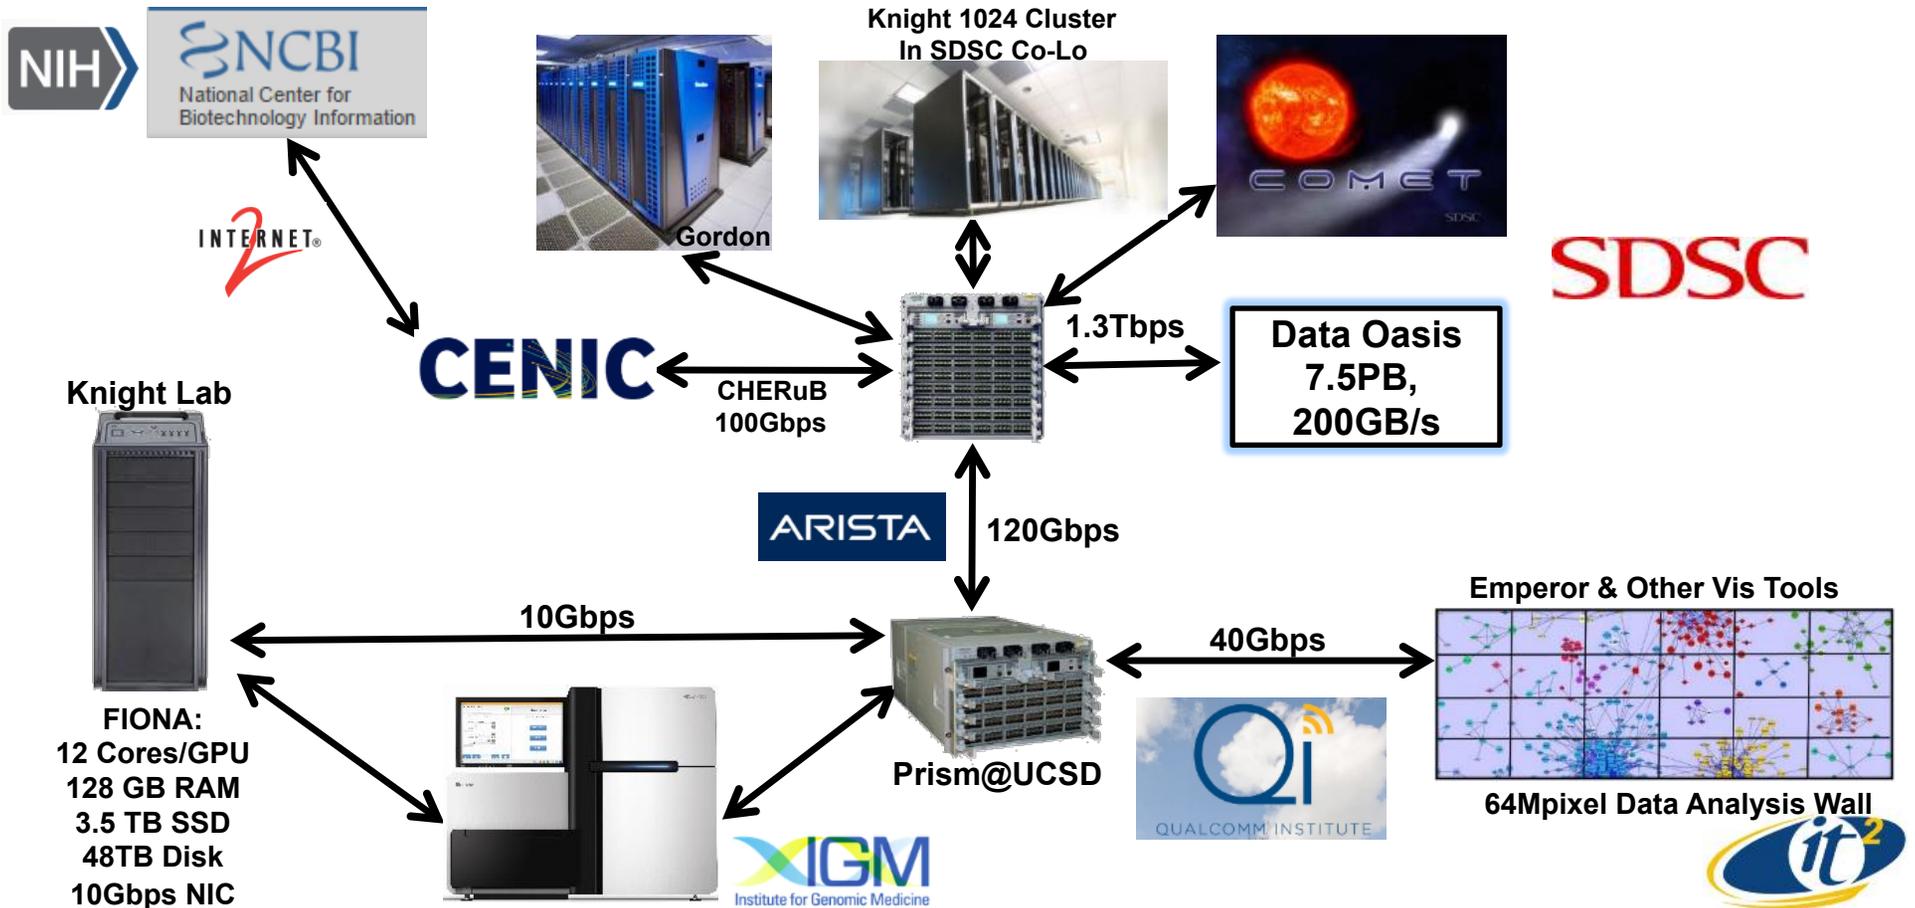
## Rack-Mount Build:

2

	Cost	\$8,000	\$20,000
Intel Xeon Haswell		E5-1650 v3 6-Core	2x E5-2697 v3 14-Core
RAM		128 GB	256 GB
SSD		SATA 3.8 TB	SATA 3.8 TB
Network Interface		10/40GbE Mellanox	2x40GbE Chelsi+Mellanox
GPU			NVIDIA Tesla K80
		RAID Drives 0 to 112TB (add ~\$100/TB)	

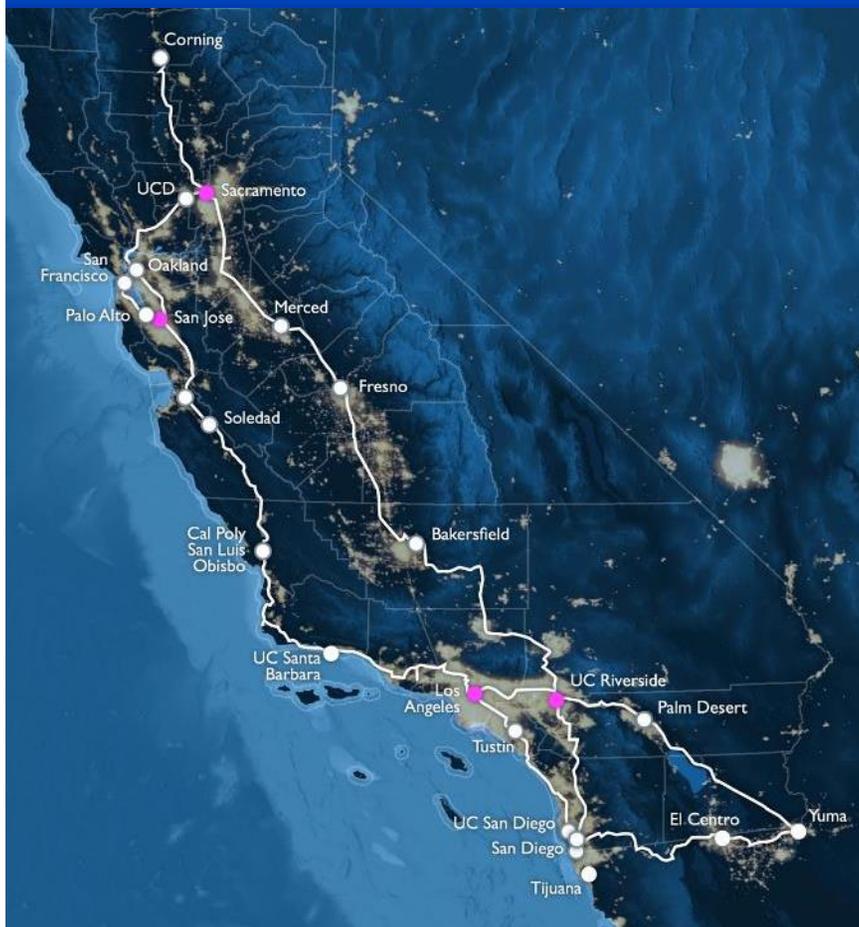


# How Prism@UCSD Transforms Big Data Microbiome Science: Preparing for Knight/Smarr 1 Million Core-Hour Analysis





## We Are Building on 15 Years of Member Investment in CENIC: California's Research & Education Network

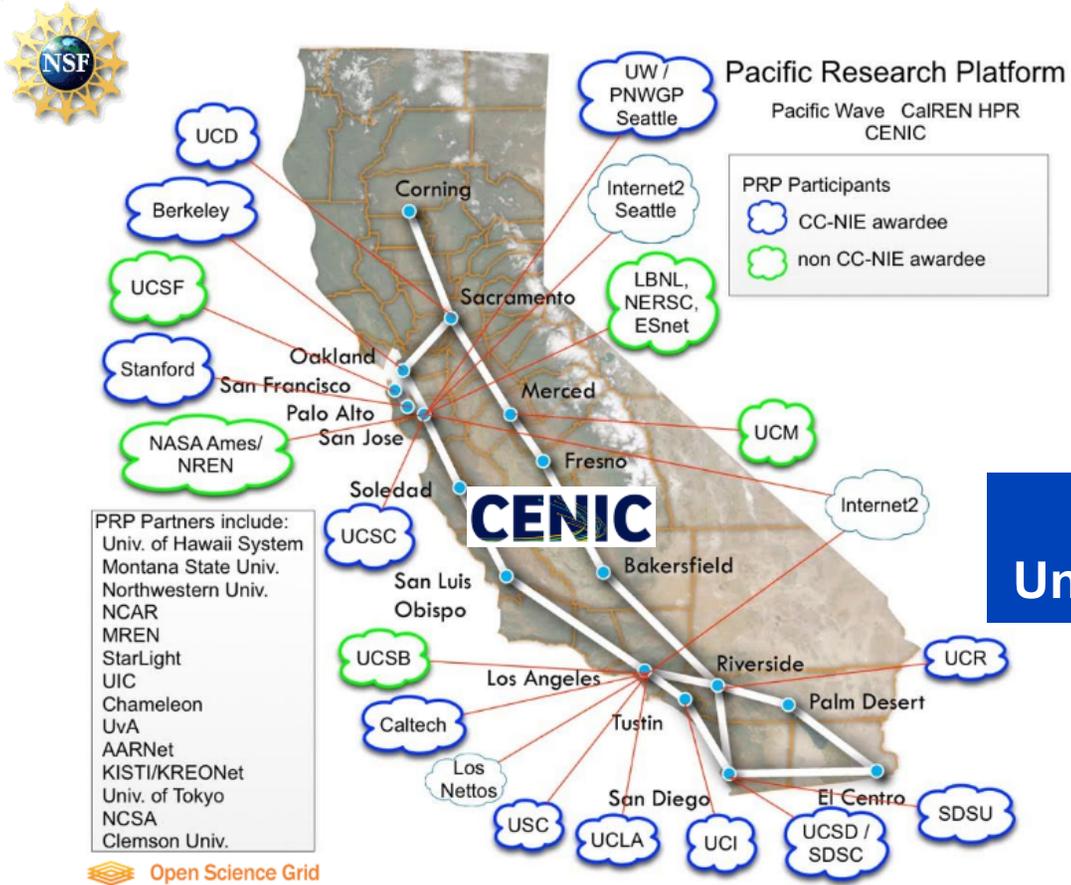


- **Members in All 58 Counties Connect via Fiber-Optics or Leased Circuits**
  - 3,800+ Miles of Optical Fiber
  - Over 10,000 Sites Connect to CENIC
  - 20,000,000 Californians Use CENIC
- **Funded & Governed by Segment Members**
  - UC, Cal State, Stanford, Caltech, USC
  - Community Colleges, K-12, Libraries
  - Collaborate With Over 500 Private Sector Partners
  - 88 Other Peering Partners
    - (Google, Microsoft, Amazon ...)

# CENIC



# Next Step: The Pacific Research Platform Creates a Regional End-to-End Science-Driven “Big Data Superhighway” System



**NSF CC\*DNI Grant**  
**\$5M 10/2015-10/2020**

- PI: Larry Smarr, UC San Diego Calit2**  
**Co-Pis:**
- Camille Crittenden, UC Berkeley CITRIS,
  - Tom DeFanti, UC San Diego Calit2,
  - Philip Papadopoulos, UCSD SDSC,
  - Frank Wuerthwein, UCSD Physics and SDSC

**FIONAs as Uniform DTN End Points**



Note: this diagram represents a subset of sites and connections. v1.16 – 20151019

# 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.”**

Presented at CENIC 2015  
March 9, 2015





# Pacific Research Platform Driven by Multi-Site Data-Intensive Research



**Particle  
Physics**

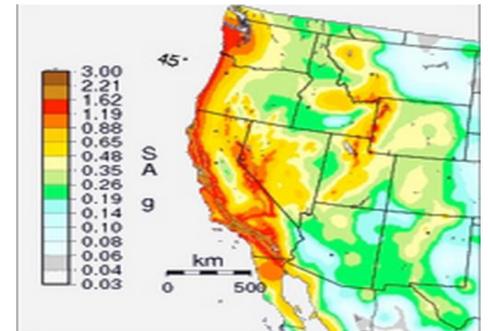


**Biomedical  
'omics**



**Telescope  
Surveys**

**Earthquake  
Engineering**



**Visualization,  
Virtual Reality,  
Collaboration**



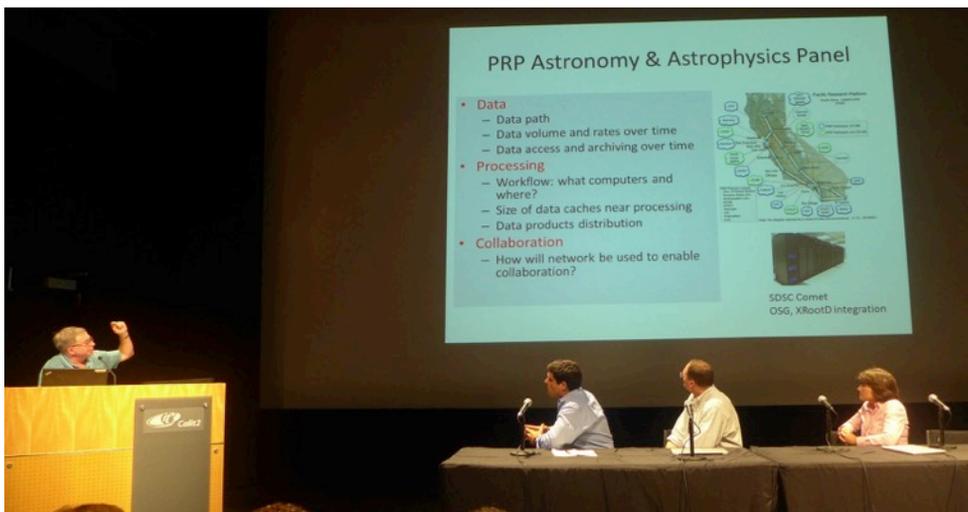
## PRP Timeline

- **PRPv1**
  - A Routed Layer 3 Architecture
  - 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**
  - Incorporating SDN/SDX, AutoGOLE / NSI
  - 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
  - Cooperating Research Groups



## Invitation-Only PRP Workshop Held in Calit2's Qualcomm Institute October 14-16, 2015

- **130 Attendees From 40 organizations**
  - Ten UC Campuses, as well as UCOP Plus 11 Additional US Universities
  - Four International Organizations (from Amsterdam, Canada, Korea, and Japan)
  - Five Members of Industry Plus NSF



## PRP First Application: Distributed IPython/Jupyter Notebooks: Cross-Platform, Browser-Based Application Interleaves Code, Text, & Images

IJulia  
IHaskell  
IFSharp  
IRuby  
IGo  
IScala  
IMathics  
Ialdor  
LuaJIT/Torch  
Lua Kernel  
IRKernel (for the R language)  
IErlang  
IOCaml  
IForth  
IPerl  
IPerl6  
loctave  
Calico Project

- kernels implemented in Mono, including Java, IronPython, Boo, Logo, BASIC, and many others



Evolved from the IPython Project

Source: John Graham, QI

IScilab  
IMatlab  
ICSharp  
Bash  
Clojure Kernel  
Hy Kernel  
Redis Kernel  
jove, a kernel for io.js  
IJavascript  
Calysto Scheme  
Calysto Processing  
idl\_kernel  
Mochi Kernel  
Lua (used in Splash)  
Spark Kernel  
Skulpt Python Kernel  
MetaKernel Bash  
MetaKernel Python  
Brython Kernel  
IVisual VPython Kernel



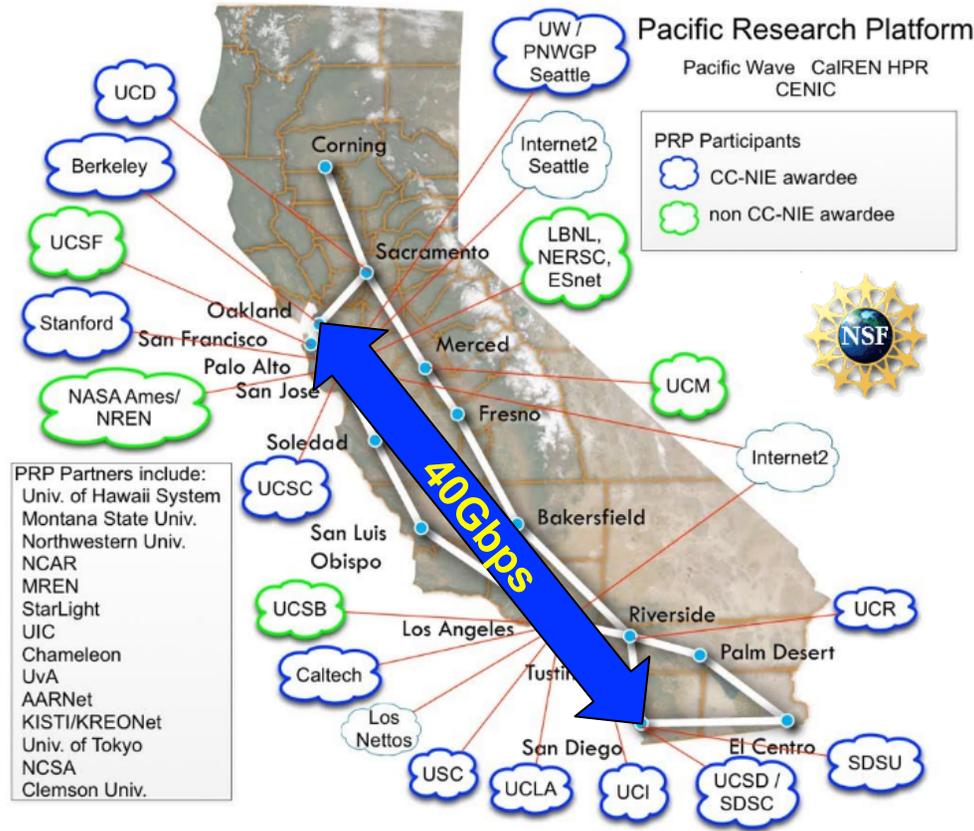
# PRP UC-JupyterHub Backbone

## Next Step: Deploy Across PRP

UCB

UCSD

Source: John Graham, Calit2



### GPU JupyterHub:

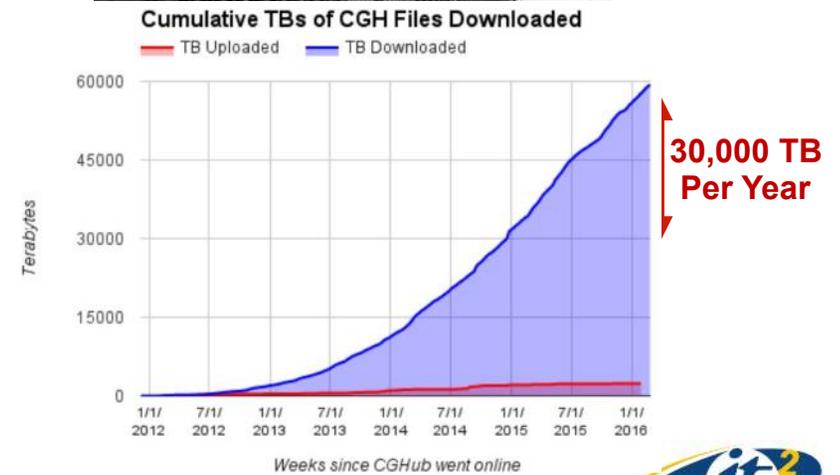
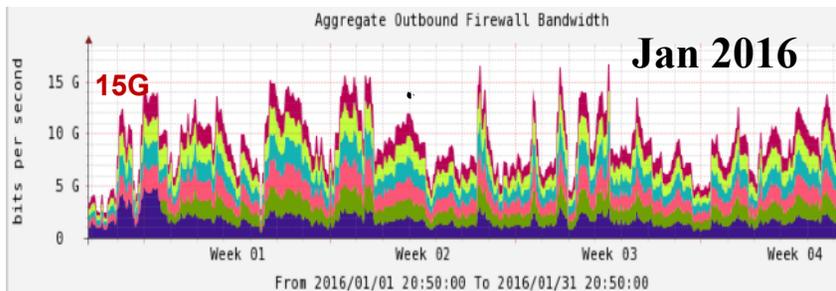
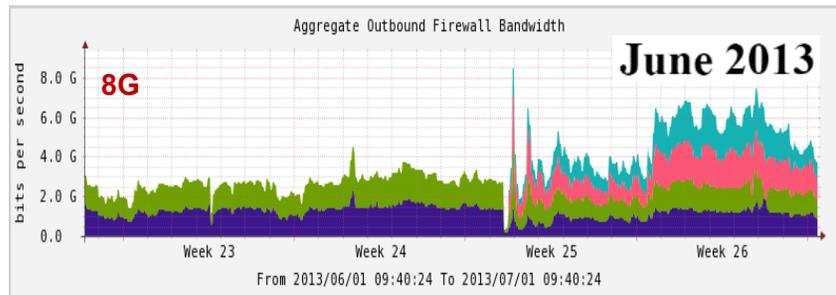
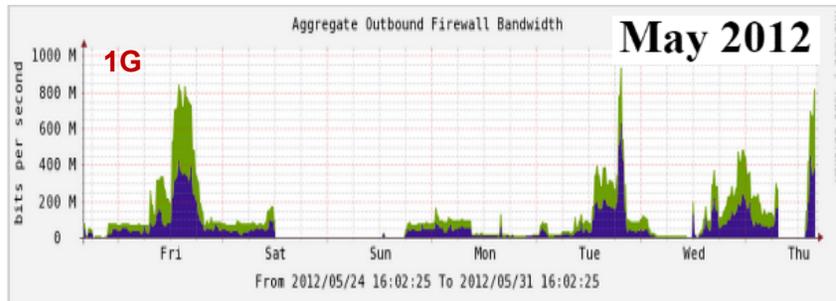
- 2 x 14-core CPUs
- 256GB RAM
- 1.2TB FLASH
- 3.8TB SSD
- Nvidia K80 GPU
- Dual 40GbE NICs
- And a Trusted Platform Module

### GPU JupyterHub:

- 1 x 18-core CPUs
- 128GB RAM
- 3.8TB SSD
- Nvidia K80 GPU
- Dual 40GbE NICs
- And a Trusted Platform Module



# Cancer Genomics Hub (UCSC) is Housed in SDSC: Large Data Flows to End Users at UCSC, UCB, UCSF, ...



Data Source: David Haussler, Brad Smith, UCSC



## Two Automated Telescope Surveys Creating Huge Datasets Will Drive PRP



**INTERMEDIATE PALOMAR TRANSIENT FACTORY**

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

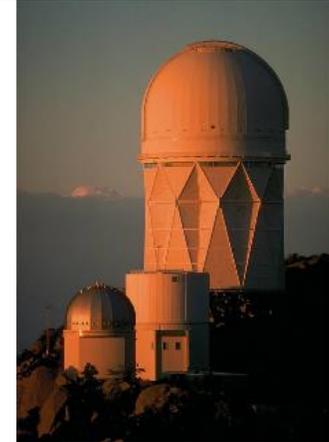
**30GB per night**

**120GB per night**

**Precursors to  
LSST and NCSA**

**PRP Allows Researchers  
to Bring Datasets from NERSC  
to Their Local Clusters  
for In-Depth Science Analysis**

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



**Dark Energy Spectroscopic Instrument**

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

**150 GB per night**

**800GB per night**

Source: Peter Nugent, Division Deputy for Scientific Engagement, LBL  
Professor of Astronomy, UC Berkeley



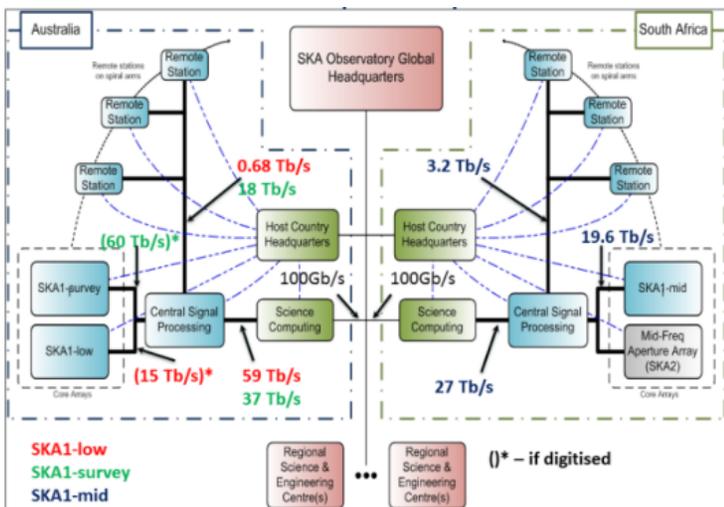
# Global Scientific Instruments Will Produce Ultralarge Datasets Continuously Requiring Dedicated Optic Fiber and Supercomputers

## Square Kilometer Array



IBM to build exascale supercomputer for the world's largest, million-antennae telescope

By Sebastian Anthony on April 2, 2012 at 11:48 am | 8 Comments

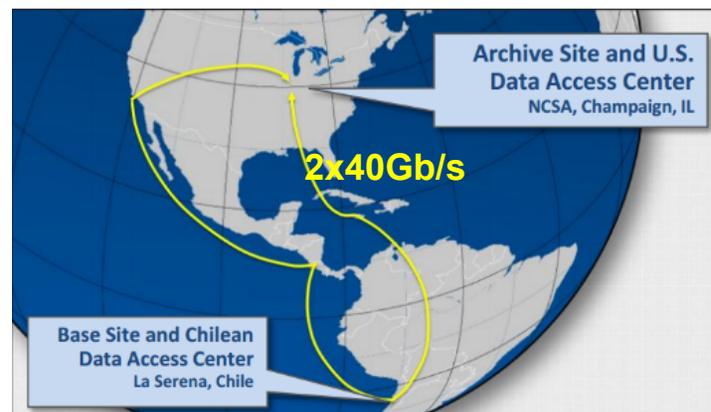


<https://tnc15.terena.org/getfile/1939>

## Large Synoptic Survey Telescope



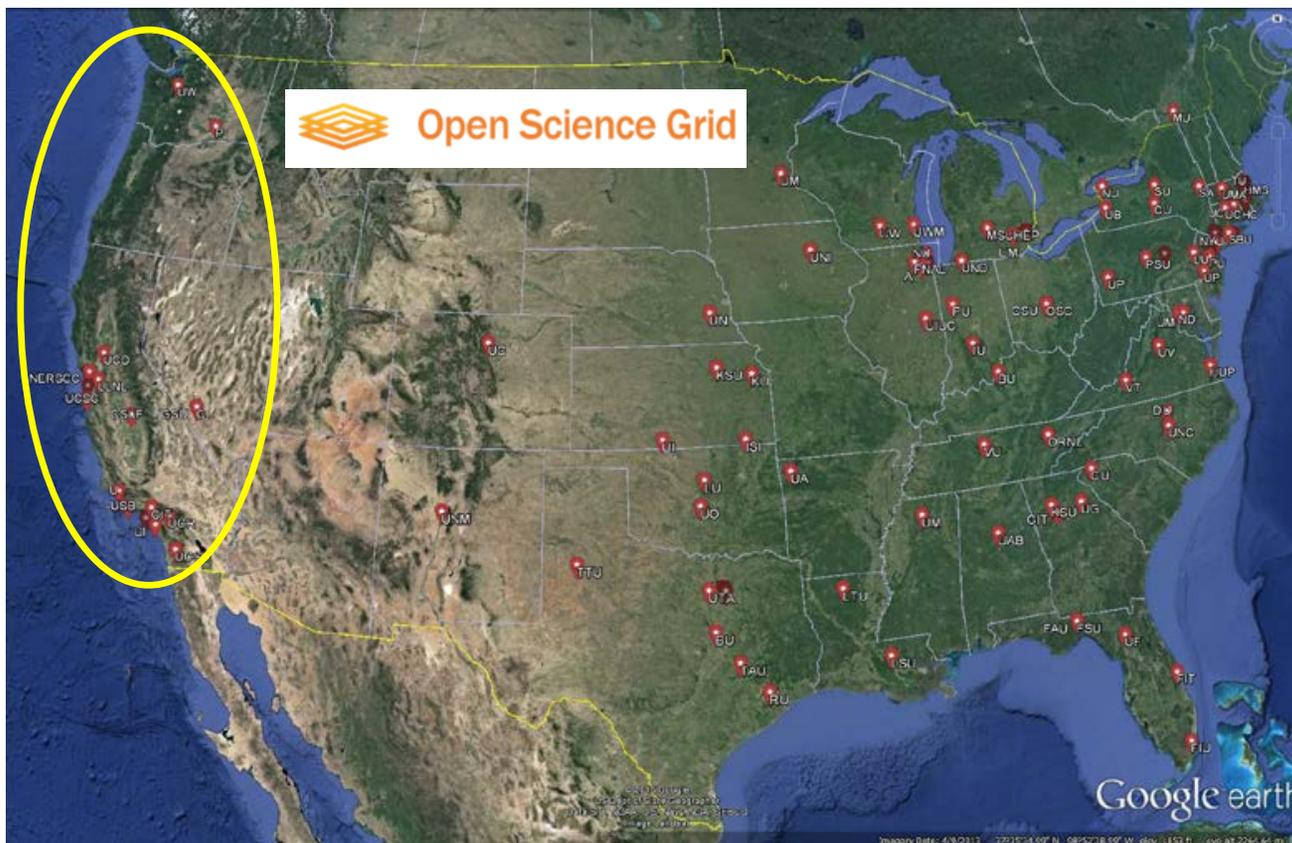
Tracks ~40B Objects,  
Creates 10M Alerts/Night  
Within 1 Minute of Observing



[www.lsst.org/sites/default/files/documents/DM%20Introduction%20-%20Kantor.pdf](http://www.lsst.org/sites/default/files/documents/DM%20Introduction%20-%20Kantor.pdf)



## OSG Federates Clusters in 40/50 States: Creating a Scientific Compute and Storage “Cloud”

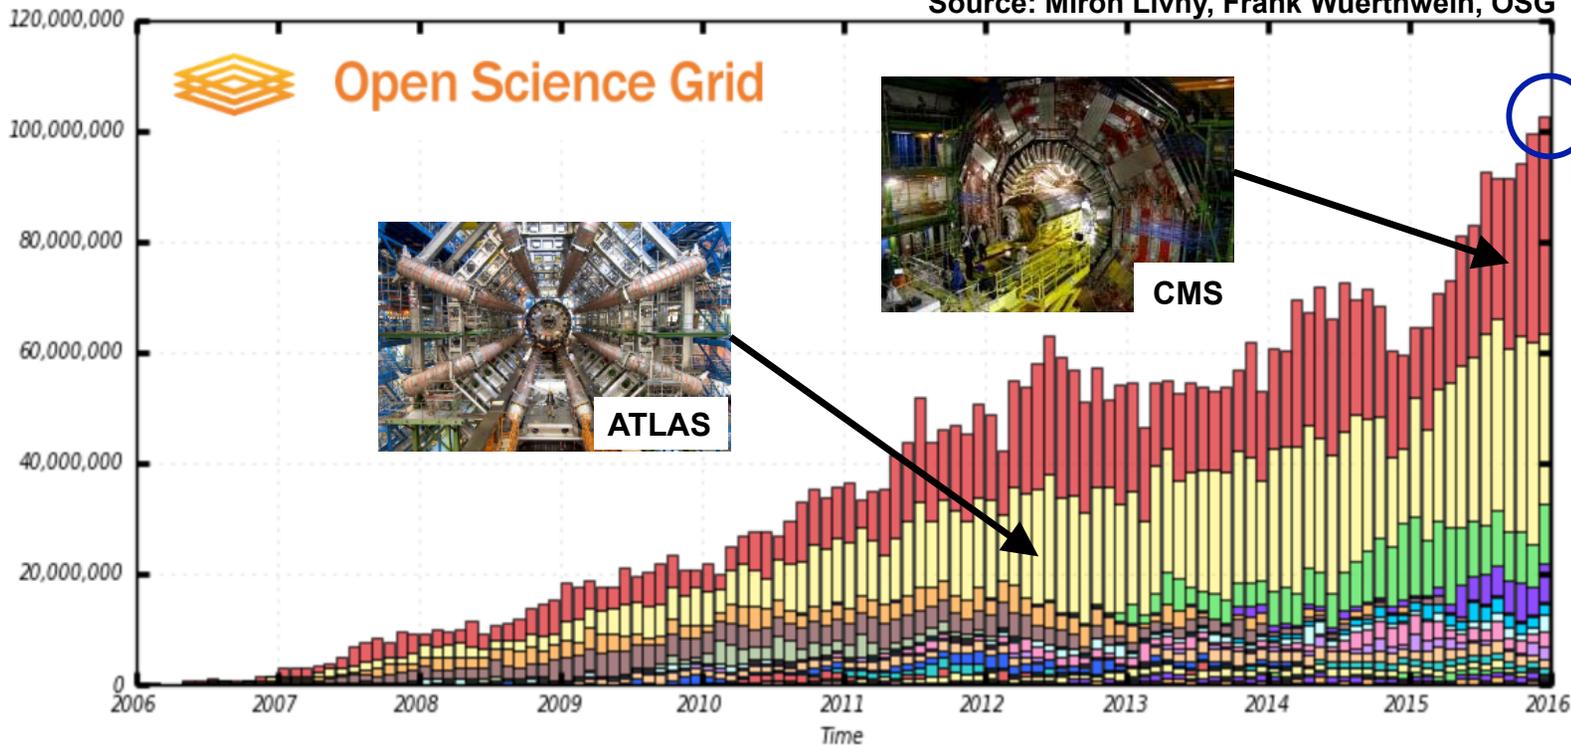


Source: Miron Livny, Frank Wuerthwein, OSG



# We are Experimenting with the PRP for Large Hadron Collider Data Analysis Using The West Coast Open Science Grid on 10-100Gbps Optical Networks

Source: Miron Livny, Frank Wuerthwein, OSG



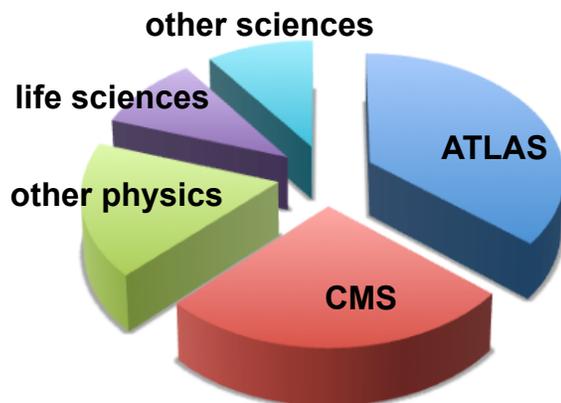
Crossed  
100 Million  
Core-Hours/Month  
In Dec 2015

Supported Over  
200 Million Jobs  
In 2015

Over 1 Billion  
Data Transfers  
Moved  
200 Petabytes  
In 2015



# PRP Prototype of Aggregation of OSG Software & Services Across California Universities in a Regional DMZ



OSG Hours 2015 by Science Domain



- Aggregate Petabytes of Disk Space & PetaFLOPs of Compute, Connected at 10-100 Gbps
- Transparently Compute on Data at Their Home Institutions & Systems at SLAC, NERSC, Caltech, UCSD, & SDSC

**PRP Builds on SDSC's LHC-UC Project**



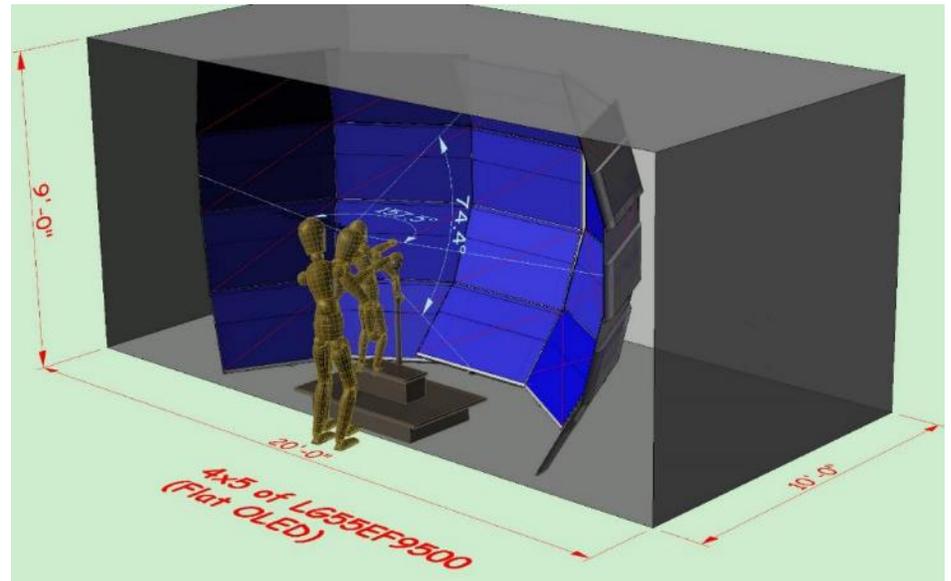
Source: Frank Wuerthwein, UCSD Physics; SDSC; co-PI PRP



# PRP Links Creates Distributed Virtual Reality



WAVE@UC San Diego



CAVE@UC Merced



# Planning for climate change in California

substantial shifts on top of already high climate variability

**UCSD Campus Climate Researchers Need to Download  
Results from NCAR 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

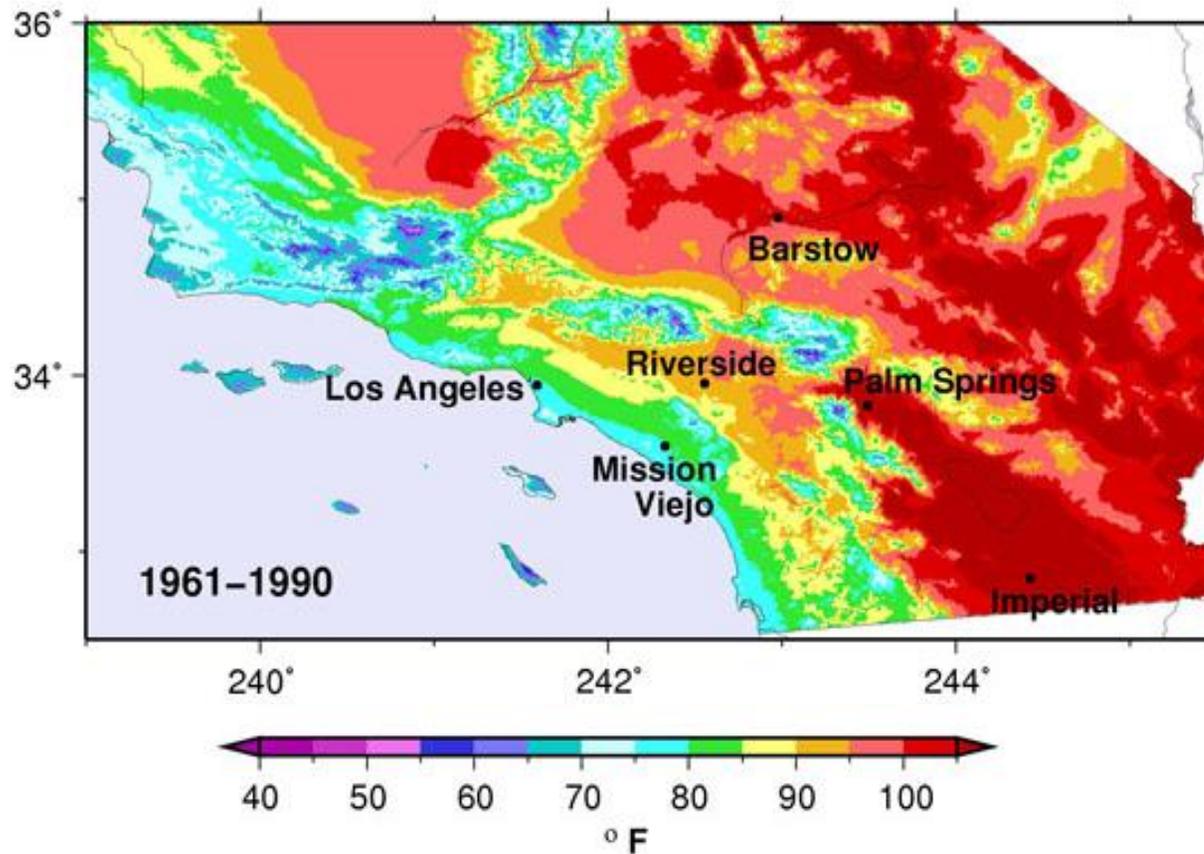
**NCAR Upgrading to 10Gbps Link Over Westnet  
from Wyoming and Boulder to CENIC/PRP**

*Sponsors:*

California Energy Commission  
NOAA RISA program  
California DWR, DOE, NSF

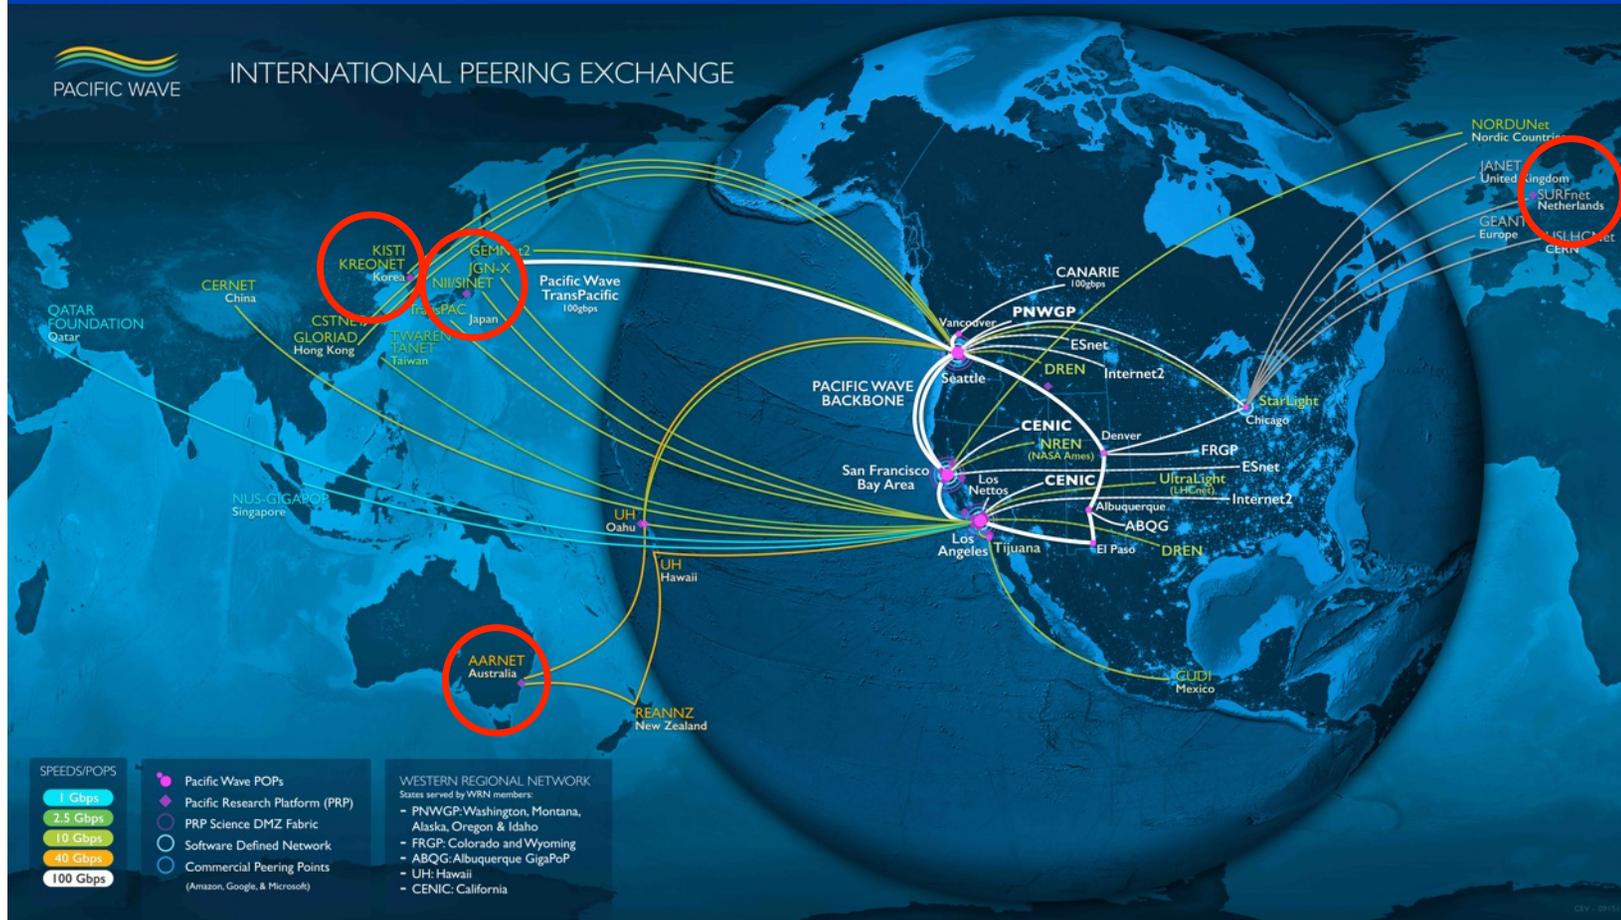


## Downscaling Supercomputer Climate Simulations To Provide High Res Predictions for California Over Next 50 Years



Source: Hugo Hidalgo, Tapash Das, Mike Dettinger

# Next Step: Global Research Platform Building on CENIC/Pacific Wave and GLIF



**Current  
International  
GRP Partners**

