GLIF Americas Working Group

University of Sydney
25 September, 2017

john hess, CENIC – Pacific Wave

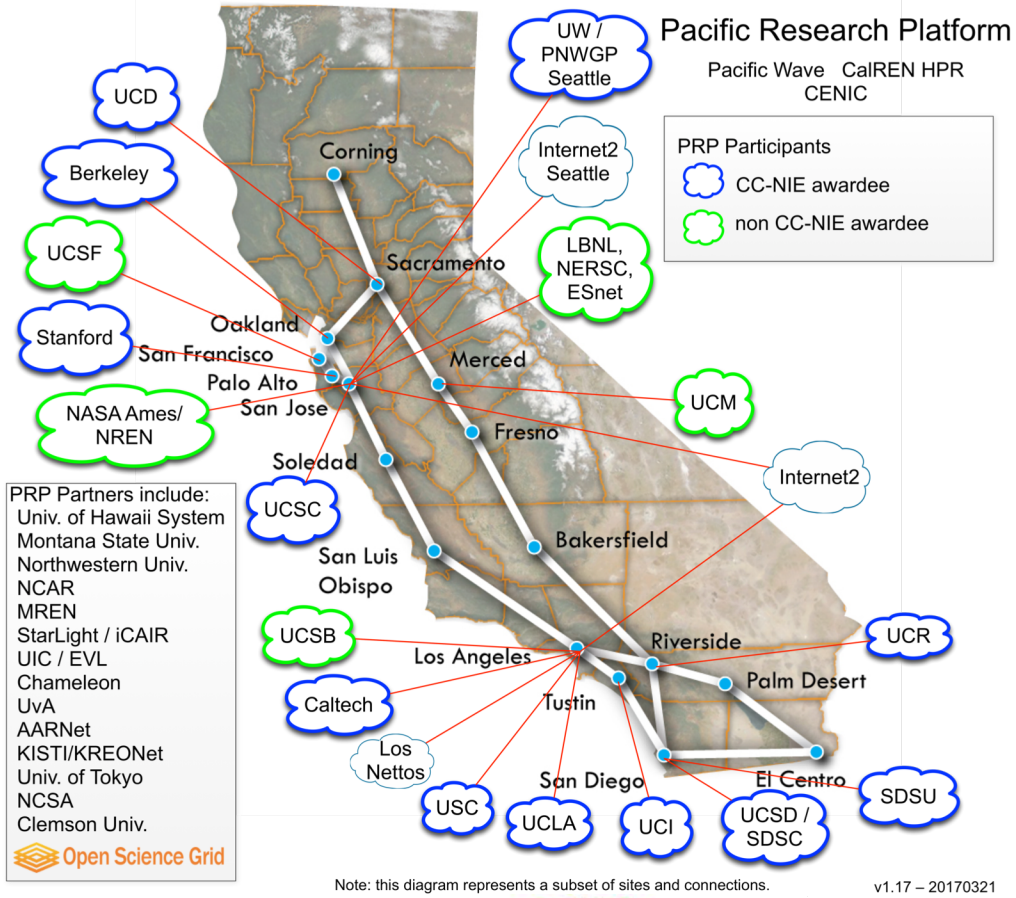

NSF CC*DNI Grant
$5M 10/2015-10/2020

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**

Science Teams:

- **Visualization and Virtual Reality**
- **Biomedical**
- **Earth Sciences**
- **Particle Physics**
- **Astronomy and Astrophysics**
- **Cryo-EM**
- **Deep Learning & Robotics**
- **High-Performance Wireless**

PACIFIC RESEARCH PLATFORM

Our Prototype System – Built for Scientists Out of a Bunch of Independently Managed Networks

- **Challenge:**
  - **Campus DMZs, Regional (e.g., CENIC), National (Internet2), International Networks (e.g., GLIF) are Individually-Architected Systems**
- **How Do <u>They Work Together</u> with Predictable Performance?**
- **➔ PRP is Focused on Disk-to-Disk Data Movement**
  - **From the Eyes of Domain Scientists**
    - **End-to-End for *Their Data* is Their Only Real Metric of Concern (As it Should Be)**

Source: Phil Papadopoulos

PRP PACIFIC RESEARCH PLATFORM

PRP Science DMZ Data Transfer Nodes (DTNs) - Flash I/O Network Appliances (FIONAs)

UCSD Designed FIONAs
To Solve the Disk-to-Disk
Data Transfer Problem
at Full Speed
on 10G, 40G and 100G Networks

FIONette -- 1Gb/s,
2TB, $1,000USD

Sources: Phil Papadopoulos, SDSC &
Tom DeFanti, Joe Keefe & John Graham, Calit2

Low-cost perfSONAR nodes

ODROID-C2 - 1Gb/s, 32GB, $100:
* Amlogic ARM® Cortex®-A53(ARMv8) 1.5Ghz quad core CPUs
* Mali™-450 GPU (3 Pixel-processors + 2 Vertex shader processors)
* 2Gbyte DDR3 SDRAM
* Gigabit Ethernet
* HDMI 2.0 4K/60Hz display
* H.265 4K/60FPS and H.264 4K/30FPS capable VPU
* 40pin GPIOs + 7pin I2S
* eMMC5.0 HS400 Flash Storage slot / UHS-1 SDR50 MicroSD Card slot
* USB 2.0 Host x 4, USB OTG x 1 (power + data capable)

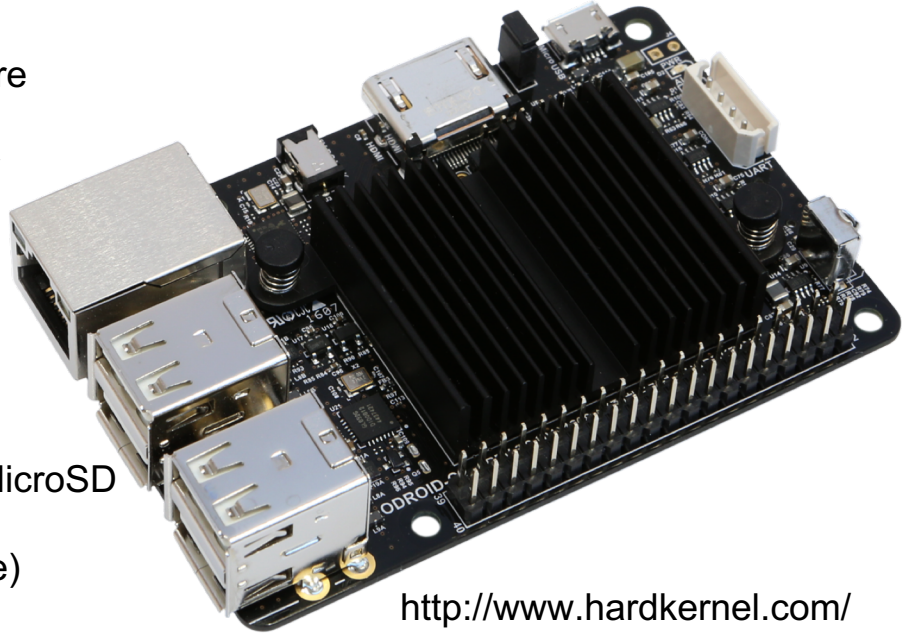

http://www.hardkernel.com/

Project tracking candidate, low-cost hardware:
https://github.com/perfsonar/project/wiki/Low-Cost-perfSONAR-Nodes

More Than 30 PRP Installed FIONAs: Customized to the Needs of Application Teams

- **Data Transfer Nodes**
  - 1, 10, 40, and 100Gb/s NICs
- **perfSONAR Nodes**
  - 10, 40, and 100Gb/s NICs
- **Storage Transfer Nodes**
  - Up to 160TB of Rotating Disks
  - Nonvolatile Memory Disks (NVMe - 10x Faster than Flash)
- **Compute Transfer Nodes**
  - 12-48 Intel CPU Cores
  - 1-8 GPUs (Delivers Up to 500,000 GPU Core Hours/Day)
    - **NSF CISE::CHASE-CI (**Award #1730158)
- **Visualization Transfer Nodes**
  - 3-45 Tiled displays (up to 180 Megapixels, 2D & 3D)
  - 360-Megapixel SunCAVE Coming Soon

Source: John Graham, Calit2

Network Measurement and Analysis: perfSONAR

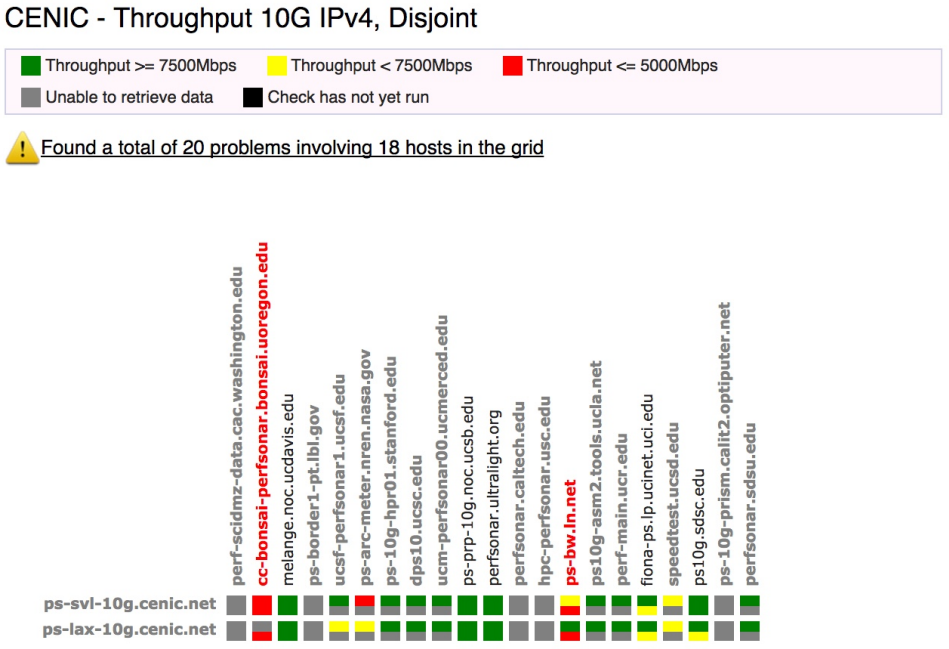

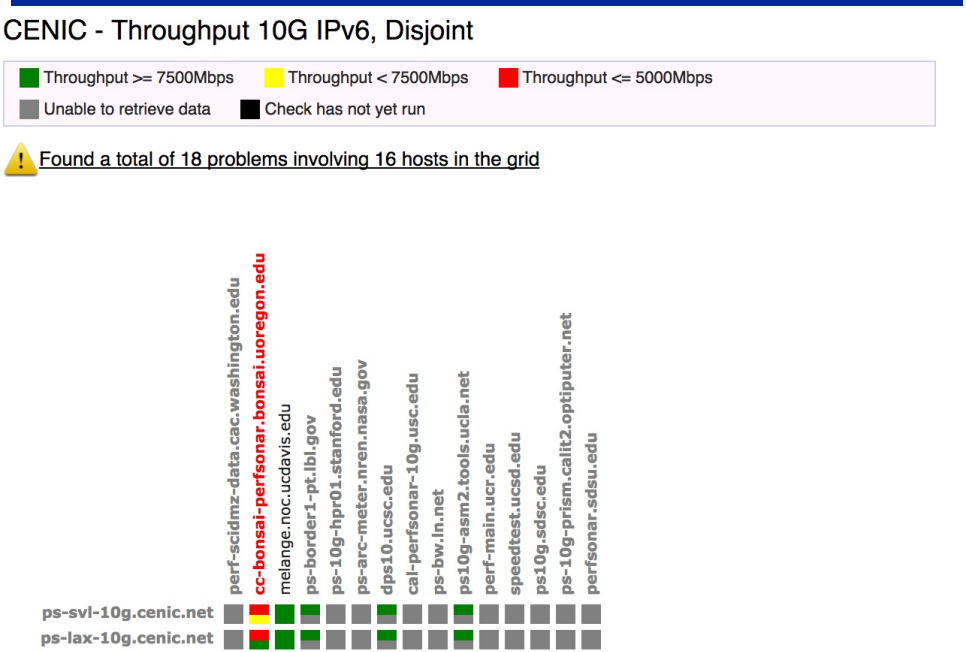

- Disjoint: a_member hosts test to all b_member hosts; no tests between b_member hosts*
- Grids for reflecting results IPv4-only and IPv6-only tests

Visualizing disk-to-disk throughput

source: prp-maddash.calit2.optiputer.net/maddash-webui/

Prototype Traceroute Visualization tool

• Allow filtering of paths by choosing one or more sources and targets or showing all available data when there are no sources and targets selected

• Color-code the graph edges with information from an esmond database (Measurement Archive) and switch between them: latency, bandwidth, retries (retransmits)

• Properly recognize flapping routes and display those as dotted edges

• Select the closest node or edge using Voronoi diagrams, eliminating the need to precisely aim at a thin edge, which helps in busy locations

• Display an item's information, such that it can be copied to clipboard, by clicking on it

• Retrieve AS numbers for all nodes, clustering the nodes based on their AS

source: Dmitry Mishin, SDSC

Visualizing paths

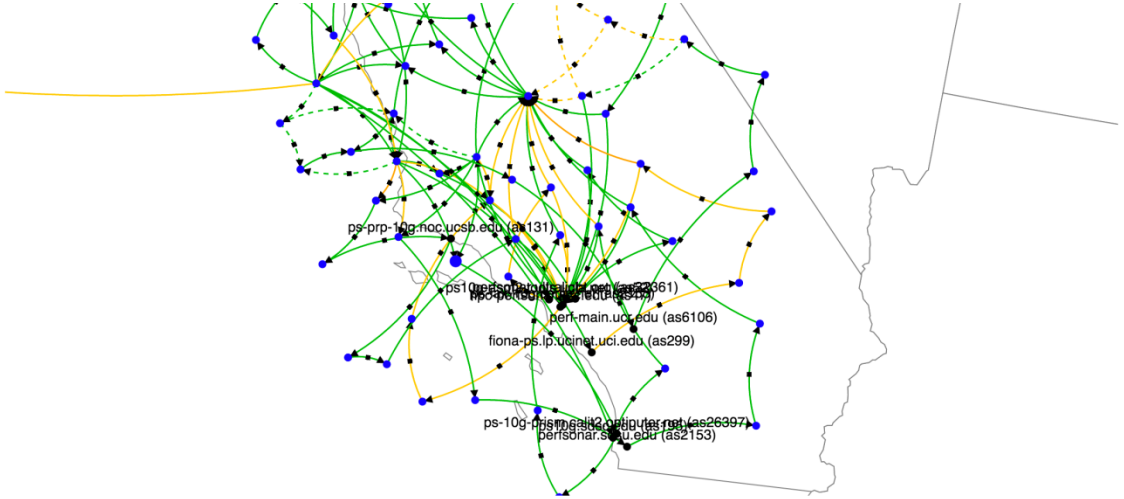

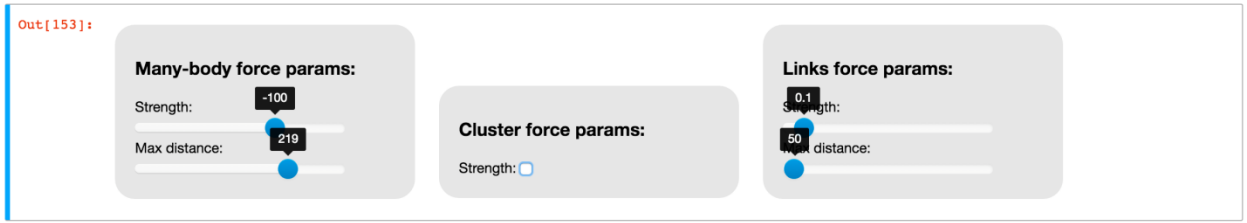

http://jupyter.calit2.optiputer.net:8080/url/jupyter.calit2.optiputer.net/Public/tracert/Traceroute-cenic14.ipynb?flush_cache=true

PRP Science Engagement Workshops

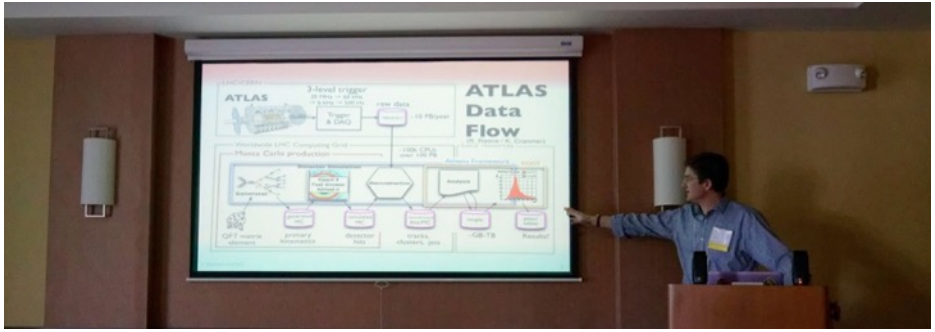

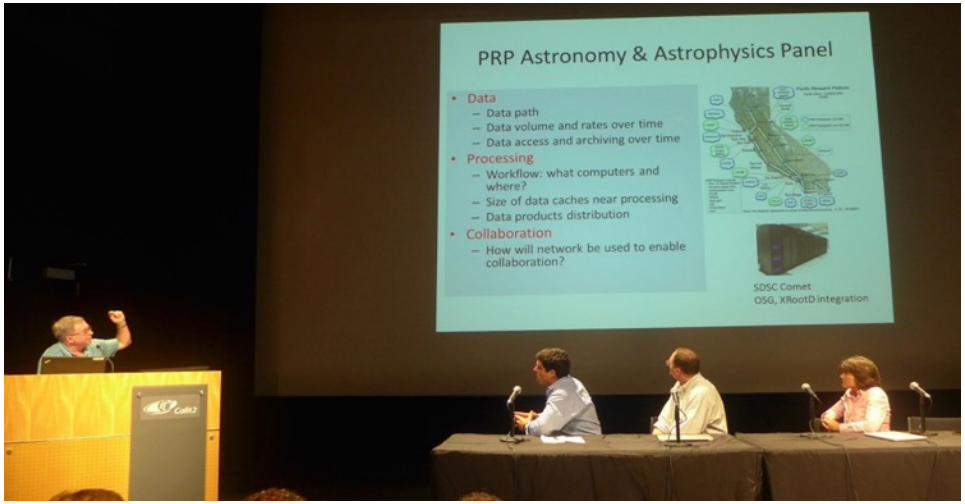

Source: Larry Smarr, Calit2

PRP's First Two Years:
Connecting Campus Application Teams and Devices

Particle Physics

CMS

Biomedical 'omics

INTERMEDIATE PALOMAR TRANSIENT FACTORY

Telescope Surveys

Earthquake Engineering

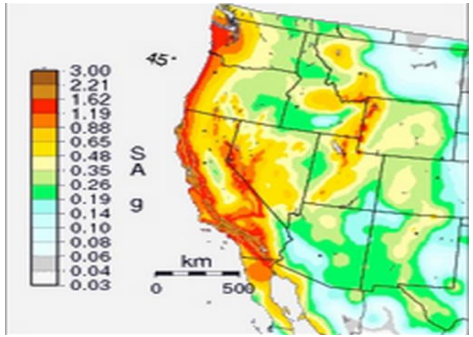

Visualization, Virtual Reality, Collaboration

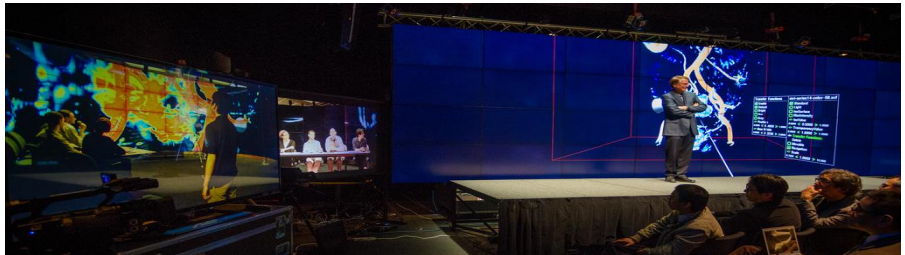

Source: Larry Smarr, Calit2

PRP
PACIFIC RESEARCH PLATFORM

100 Gbps FIONA at UCSC Connects the UCSC Hyades Cluster to the NERSC Supercomputer at LBNL

Shawfeng Dong, UCSC Cyber Infrastructure engineer

UCSC Feb 7, 2017

Source: Larry Smarr, Calit2

250 images per night

800GB per night

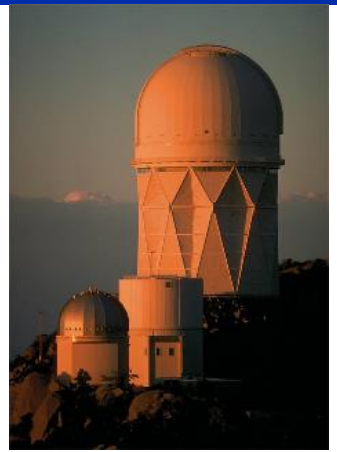

Dark Energy Spectroscopic Instrument

Supporting UCSC Remote Access
to Large Data Subsets
of the Dark Energy Spectroscopic Instrument (DESI) and
AGORA Galaxy Simulation Data
Produced at NERSC.

PRP Will Link the Laboratories of the Pacific Earthquake Engineering Research Center

PEER

PACIFIC EARTHQUAKE ENGINEERING RESEARCH CENTER

http://peer.berkeley.edu/

HOME | ABOUT PEER | NEWS | EVENTS | RESEARCH | PRODUCTS | LABORATORIES | PUBLICATIONS | NISEE | BIP MEMBERS | EDUCATION | FAQs | LINKS

PEER Mega Research Programs


Tall Buildings Initiative


Transportation Systems


Lifelines & NGL


Tsunami

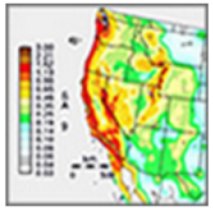
NGA-West 2

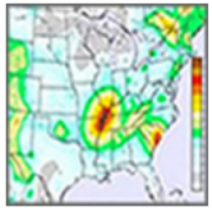
NGA-East


Global GMPE

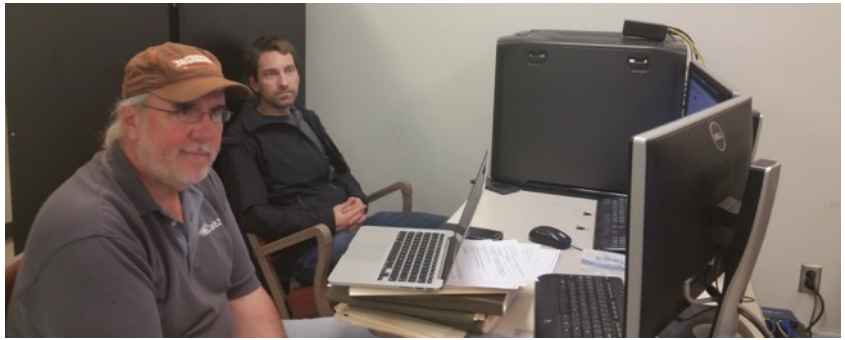

John Graham Installing FIONette at PEER Feb 10, 2017

PEER Labs: UC Berkeley, Caltech, Stanford, UC Davis, UC San Diego, and UC Los Angeles

Source: Larry Smarr, Calit2

PRP Now Enables Distributed Virtual Reality

PRP

WAVE@UC San Diego

WAVE @UC Merced

Transferring 5 CAVEcam Images from UCSD to UC Merced:
2 Gigabytes now takes 2 Seconds (8 Gb/sec)

Source: Larry Smarr, Calit2

The Prototype PRP Has Attracted New Application Drivers

Evolved from the IPython Project

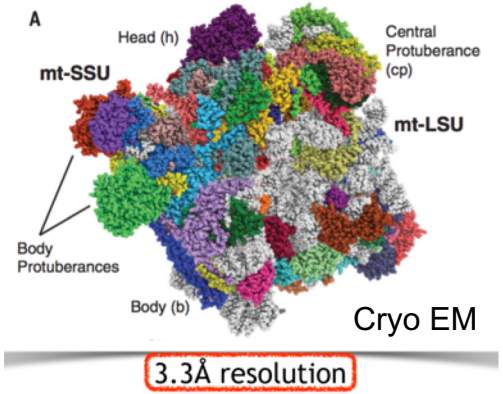

Cryo EM

Frank Vernon - Expansion of HPWREN

Tom Levy, Cultural Heritage

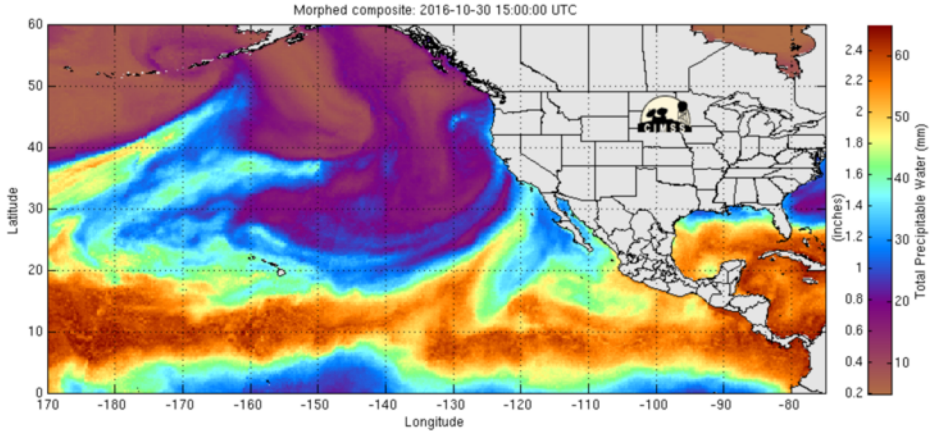

Scott Sellars, Marty Ralph
Center for Western Weather and Water Extremes

Source: Larry Smarr, Calit2

What do we envision for PRPv2

Among the technical challenges we have encountered with PRPv1: selectively announcing reachability of Science DMZ resources; choosing a traffic-engineered path; developing and implementing tools to ensure traffic fits within participants security model. We are now deploying a PRPv2 BGP pilot to explore solutions:

- An ARIN-assigned ASN: Pacific Research Platform / AS395889
- Route Servers at exchange points form the control-plane to determine reachability of Science DMZ resources, with traffic traversing high-speed data-plane
- BGP Communities for tagging classes of Science DMZ networks
- Peering will be native IPv6 only (may support IPv4 as transport)
- Initial phase will include UCSD, SDSC, UCSC, Stanford, NCSA, UIC/EVL
- Stretch goals: BGP + SDN/SDX for dynamically provisioned 'super-channels' supporting data movement among cooperating research groups

PRP PACIFIC RESEARCH PLATFORM

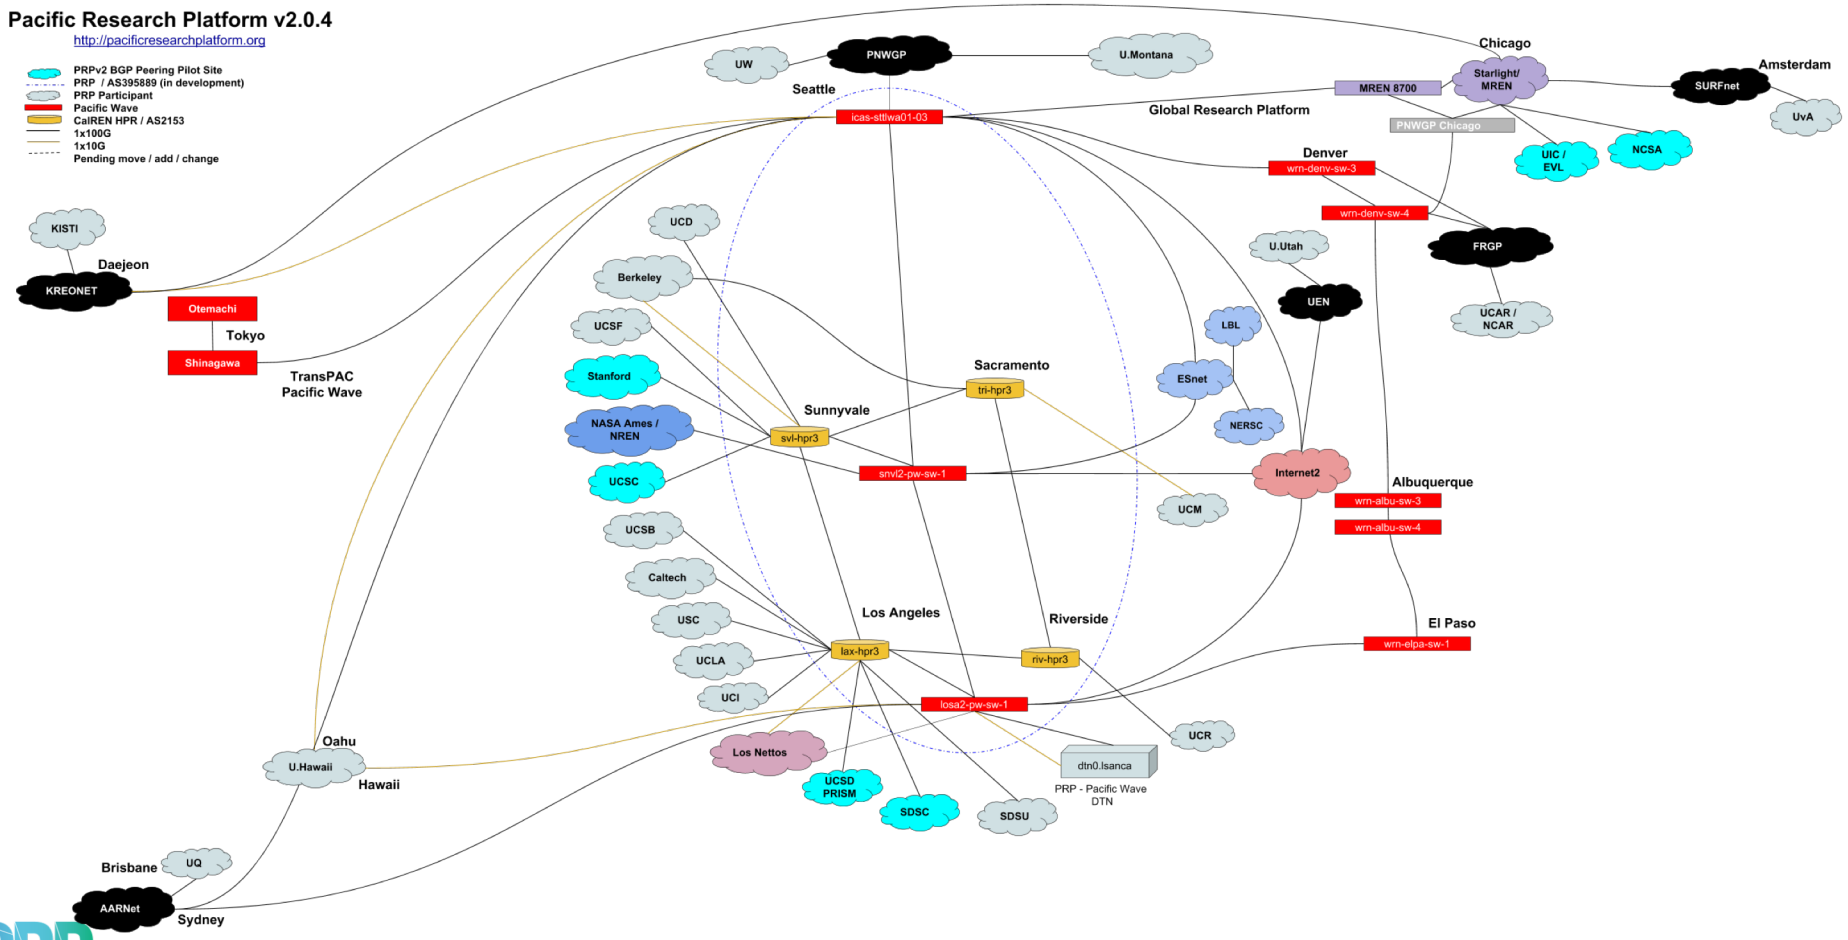
Pacific Research Platform v2.0.4
http://pacificresearchplatform.org
PRPv2 BGP Peering Pilot Site
PRP / AS395889 (in development)
PRP Participant
Pacific Wave
CalREN HPR / AS2153
1x100G
1x10G
Pending move / add / change
UW
PNWGP
U.Montana
Seattle
icas-sttlwa01-03
Chicago
Starlight/ MREN
MREN 8700
Global Research Platform
PNWGP Chicago
Amsterdam
SURFnet
UvA
UIC / EVL
NCSA
Denver
wrn-denv-sw-3
wrn-denv-sw-4
FRGP
U.Utah
UEN
UCAR / NCAR
KISTI
Daejeon
KREONET
Otemachi
Tokyo
Shinagawa
TransPAC
Pacific Wave
UCD
Berkeley
UCSF
Stanford
NASA Ames / NREN
UCSC
Sacramento
tri-hpr3
Sunnyvale
svl-hpr3
snvl2-pw-sw-1
LBL
ESnet
NERSC
Internet2
UCM
Albuquerque
wrn-albu-sw-3
wrn-albu-sw-4
UCSB
Caltech
USC
UCLA
UCI
Los Angeles
lax-hpr3
Riverside
riv-hpr3
losa2-pw-sw-1
El Paso
wrn-elpa-sw-1
UCR
Los Nettos
UCSD PRISM
SDSC
SDSU
dtn0.lsanca
PRP - Pacific Wave DTN
Oahu
U.Hawaii
Hawaii
Brisbane
UQ
AARNet
Sydney

PRPv2 BGP Pilot: route servers for control plane

http://pacificresearchplatform.org