PRP at Work:
Measuring and Monitoring Our Networks' Performance & Reliability
And Prototyping Affordable Cloud Storage/Compute

Tom DeFanti

Research Scientist, Co-PI The Pacific Research Platform and CHASE-CI

California Institute for Telecommunications and Information

University of California San Diego

Distinguished Professor Emeritus, University of Illinois at Chicago

via

John Hess

CENIC – Pacific Wave

PROVIDING NEEDED PERFORMANCE

- **R&E networks must meet performance needs of a diverse mix of users & apps**
  - **That includes everything from:**
    - **Low latency for high volume interactive K-12 student testing**
    - **Through interactive video performance**
    - **To “big data” flows among labs**

PRP
PACIFIC RESEARCH PLATFORM

it²

THAT'S BEEN HARD FOR NETS TO DO

- **Internet2 has tried multiple times via End-2-End performance projects (e.g. the 'user expectations' efforts), as have others.**
- **And, there are many network performance tools, and they do lots of useful things**
- **But tools used by most of R&E don't do direct, active, deterministic, measurement of actual traffic between sites**

PRP
PACIFIC RESEARCH PLATFORM

WITH NSF, CENIC & PNWGP SUPPORT

- Active measurement is a good way to get, and help sustain true end-to-end user performance
- Based on ESnet's fasterdata.es.net
- Supported by the large NSF funded 'Science DMZ' project called 'Pacific Research Platform' (PRP)
- And a new one called 'Towards the NRP'

PRP
PACIFIC RESEARCH PLATFORM

The Pacific Research Platform Networks Connects Campuses to Create a Regional End-to-End DMZ Connector

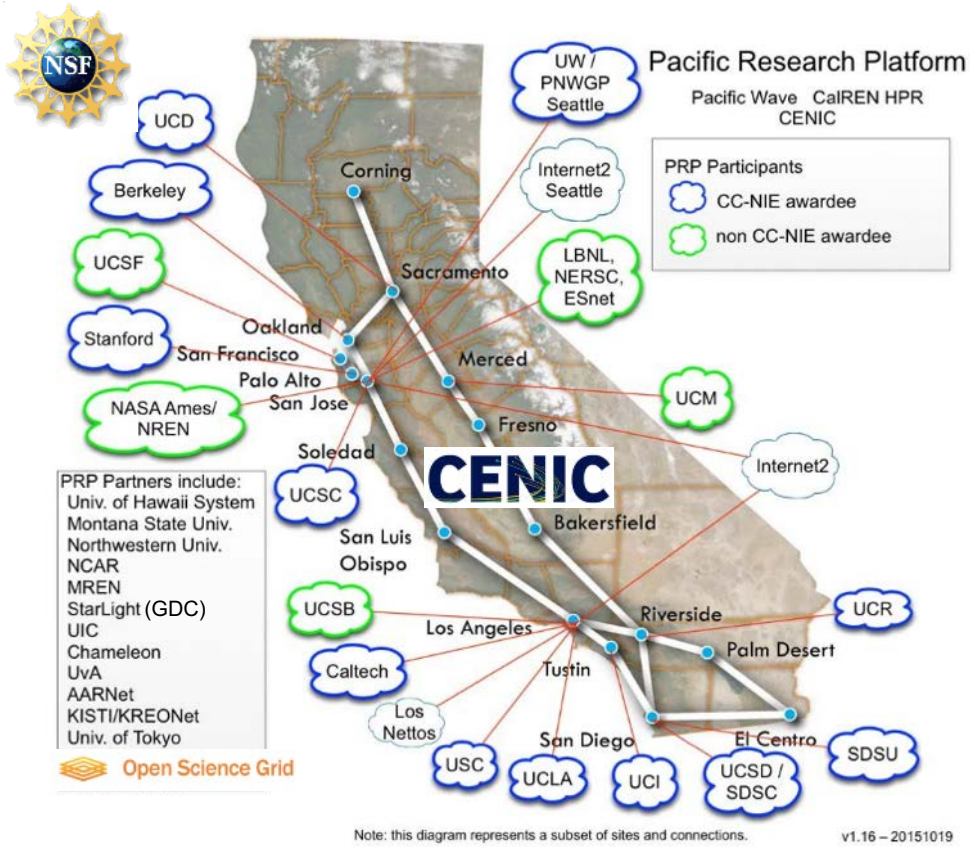

PACIFIC RESEARCH PLATFORM

Source: John Hess, CENIC

NSF Grant
10/2015-10/2020

PI: Larry Smarr, UC San Diego Calit2

Co-PIs:

- **Camille Crittenden, UC Berkeley CITRIS,**
- **Tom DeFanti, UC San Diego Calit2/QI,**
- **Philip Papadopoulos, Emeritus (pending)**
- **Frank Wuerthwein, UCSD Physics and SDSC**

Letters of Commitment from:

- 50 Researchers from 15 Campuses
- 32 IT/Network Organization Leaders

SDSC CITRIS AND THE BANATAO INSTITUTE it²

PERFORMANCE INSTRUMENTATION

- **We use purpose-built 'FIONA' PCs that are tuned to test end-to-end 1G, 10G, 40G and 100G connections, our version of ESnet's Data Transfer Nodes (DTNs). (FIONA=Flash I/O Network Appliance)**
- **perfSONAR and GridFTP logs are then turned into visualizations**
- **Disk-to-disk transfers of 10GB were performed 4x a day until the PRP networks and the campus endpoints were tuned and capable of full bandwidth utilization with essentially no TCP backoff**
- **New efforts are expanding the types of testing and visualizations, using Kubernetes as an orchestration engine to automate this distributed cluster of FIONAs**

PERFormance Service Oriented Network monitoring ARchitecture

PRP Measured Disk-to-Disk Throughput with 10GB File Transfer 4 Times Per Day in Both Directions for All PRP Sites until they Worked

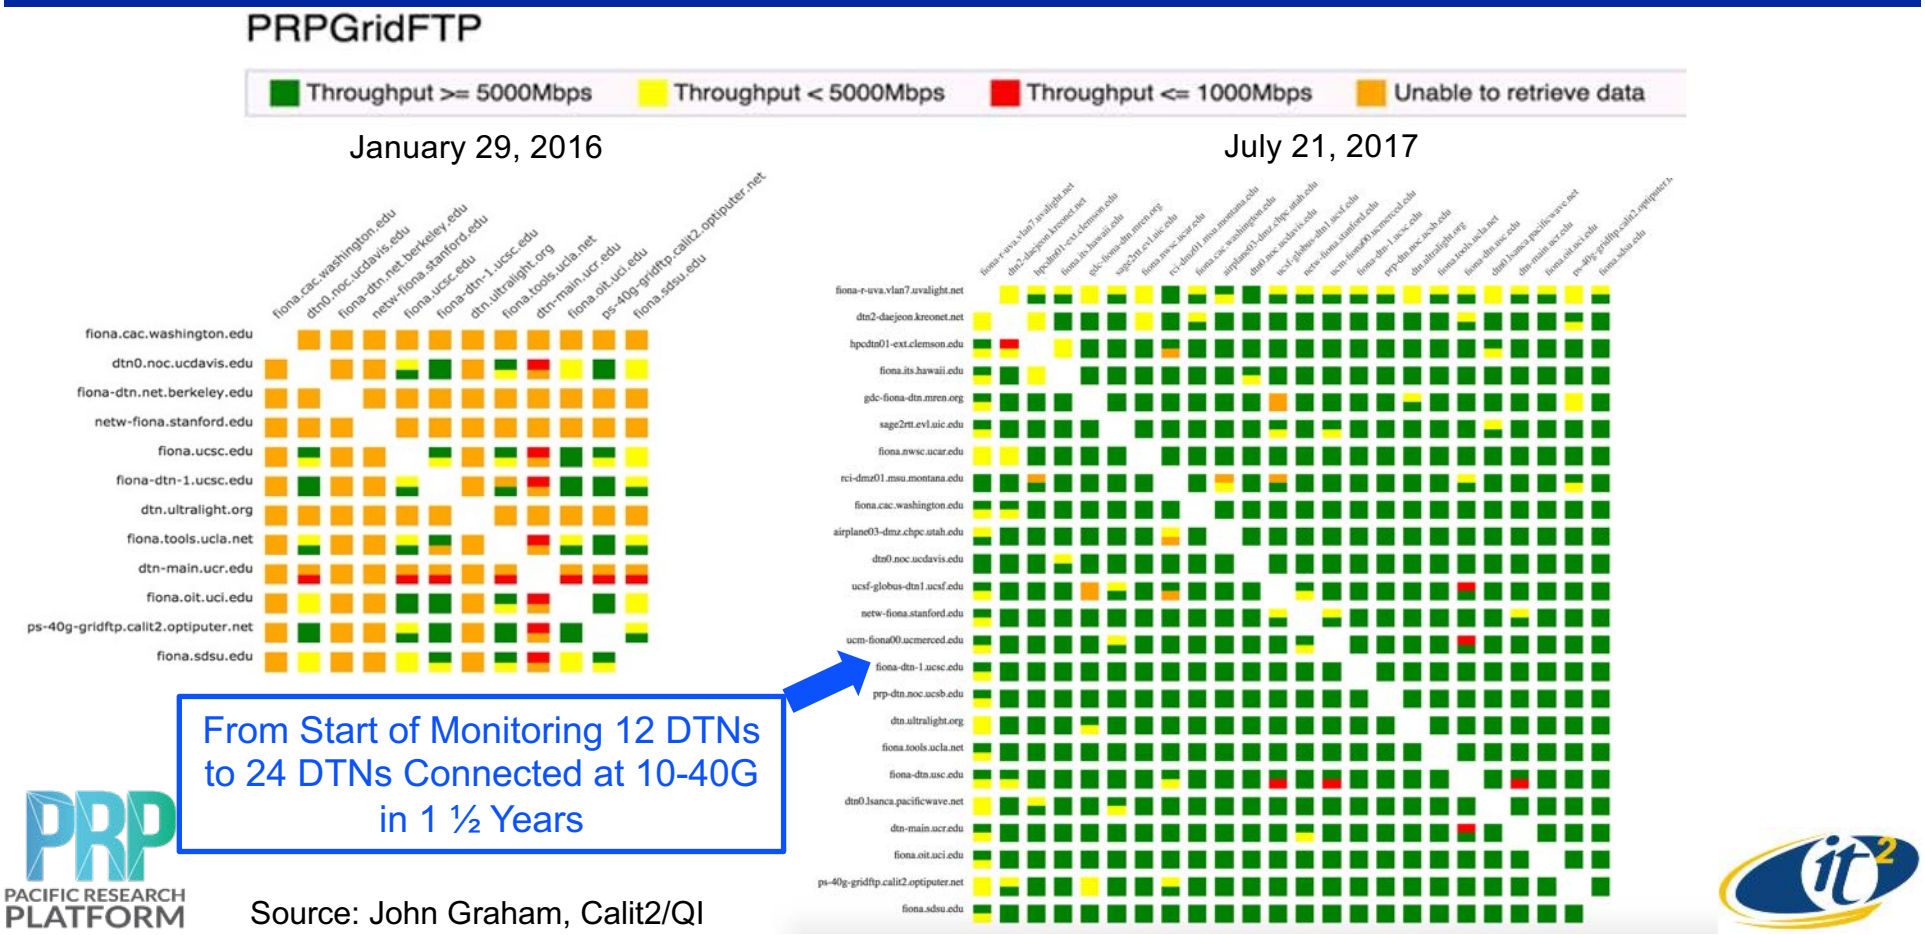

From Start of Monitoring 12 DTNs to 24 DTNs Connected at 10-40G in 1 ½ Years

Source: John Graham, Calit2/QI

PRP traceroute

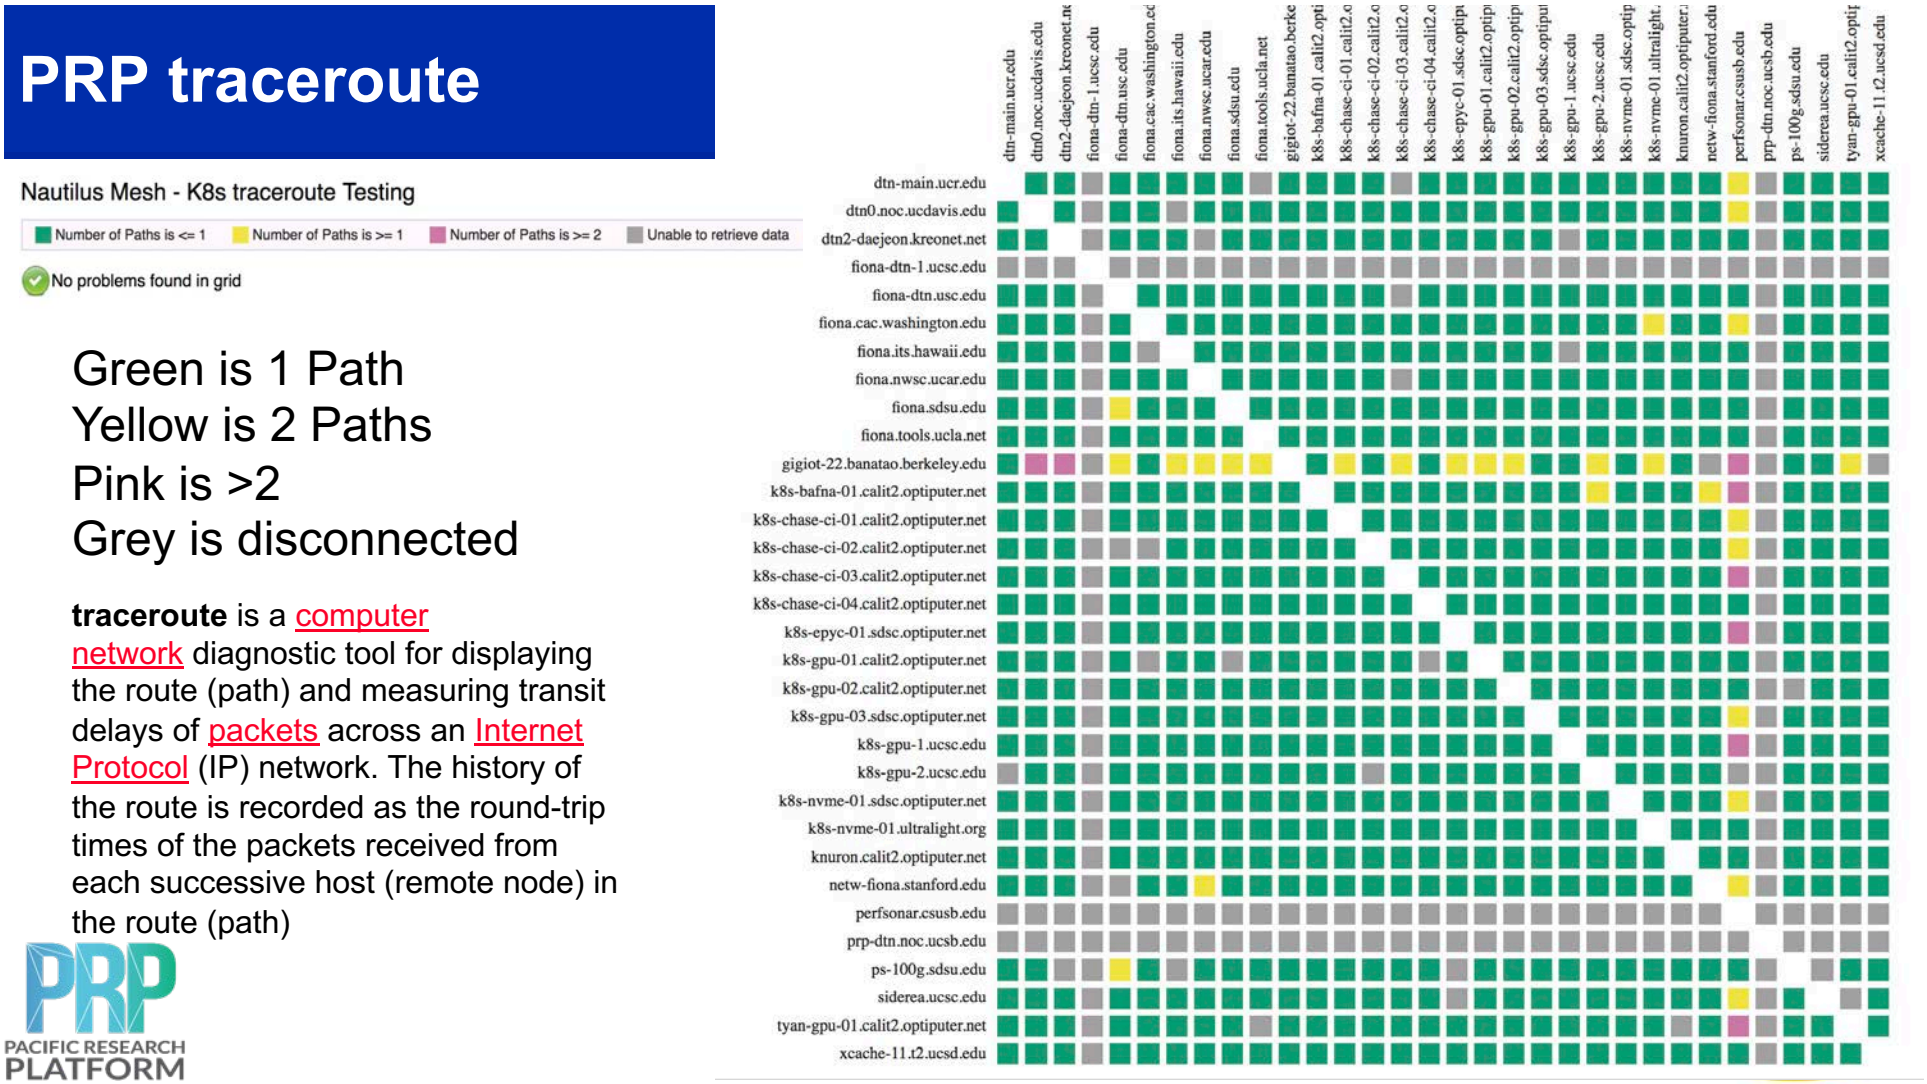

Green is 1 Path
Yellow is 2 Paths
Pink is >2
Grey is disconnected

traceroute is a computer network diagnostic tool for displaying the route (path) and measuring transit delays of packets across an Internet Protocol (IP) network. The history of the route is recorded as the round-trip times of the packets received from each successive host (remote node) in the route (path)

PACIFIC RESEARCH PLATFORM

Operational Metrics: Traceroute Tool Allows Real-time Visualization of Status of Network Links

Presets

https://traceroute.nautilus.optiputer.net/

Source: Dmitry Mishin(SDSC), John Graham (Calit2)

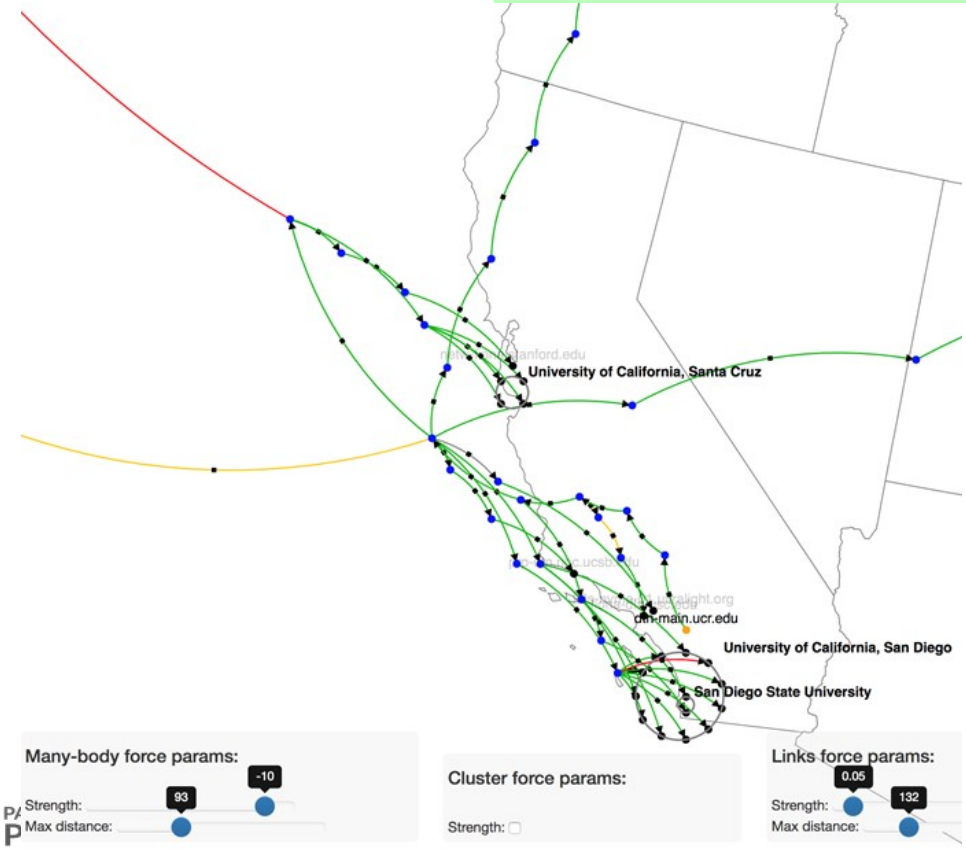

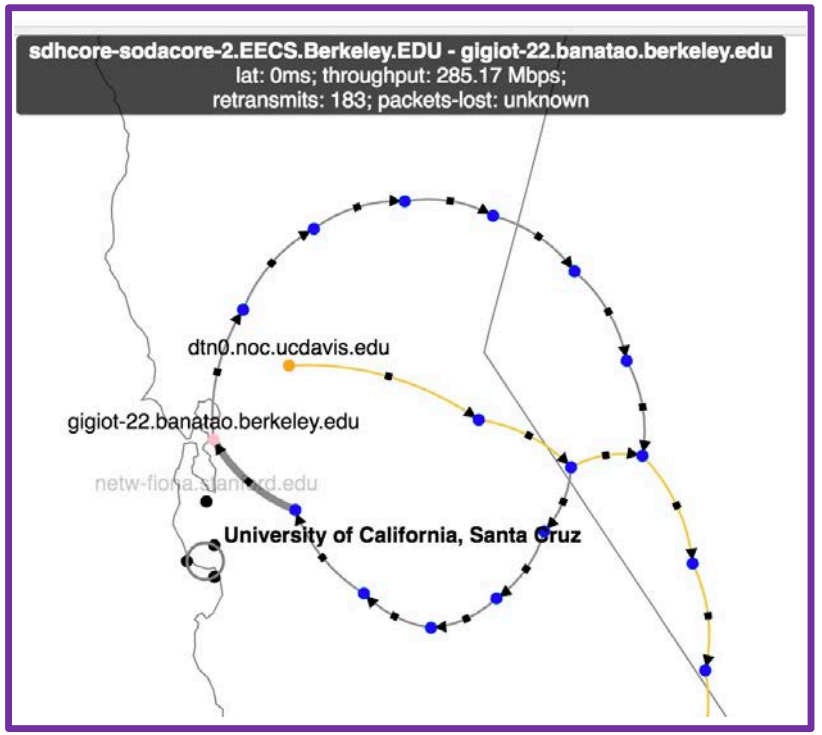

PRP Traceroute

FIONette—1G, $275
6TB/day

Green is 1 Path
Yellow is 2 Paths
Pink is >2

Shows flows that took different paths to the same destination
(Grey is not connected/firewalled)

PACIFIC RESEARCH PLATFORM

The **One-Way Active Measurement Protocol** (**OWAMP**) Measures Unidirectional Characteristics such as One-Way Delay and One-Way Loss

Nautilus Mesh - K8s OWAMP Testing

Loss rate is <= 0 | Loss rate is >= 0 | Loss rate is >= 0.01 | Unable to retrieve data

Found a total of 4 problems involving 4 hosts in the grid

Nautilus Mesh - K8s OWAMP Testing

perfsonar.csusb.edu

Site is down

Category: CONFIGURATION

Potential Solutions:

- Verify the host is up
- If recently added to the mesh, verify the mesh config file
- If recently removed from the mesh, verify that the perfSO
- Verify the local and remote sites allow access to TCP po

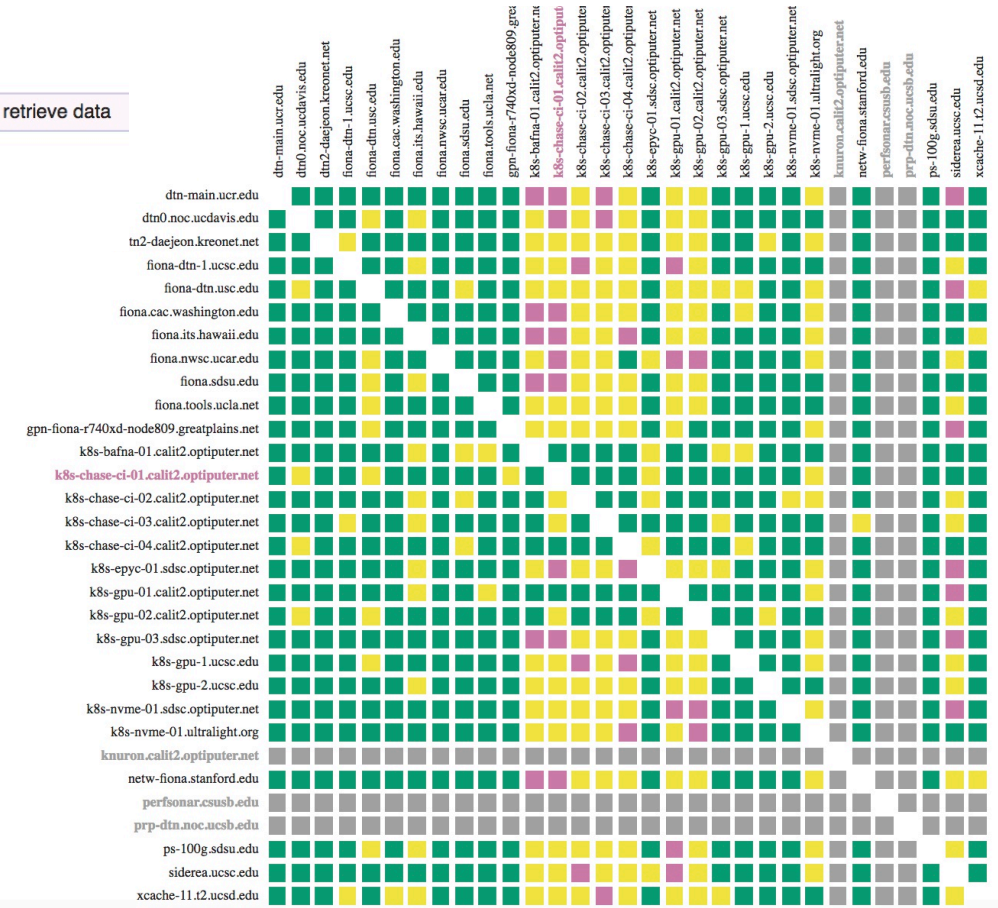

PACIFIC RESEARCH PLATFORM

ACTIVE, FACT-BASED, NET MANAGEMENT

- **Our approach gives:**
  - **Proactive measurements of actual performance**
  - **Early warning of issues to NOC and engineers**
- **And, the FIONA devices can run other software to monitor for security incursions and other issues**
- **Further, the FIONA platform allows really inexpensive network node & end-site based R&E Cloud capabilities**

PRP
PACIFIC RESEARCH PLATFORM

it²

Installing 16 10&12 TB Drives in June at UC Merced, UC Riverside, and Stanford

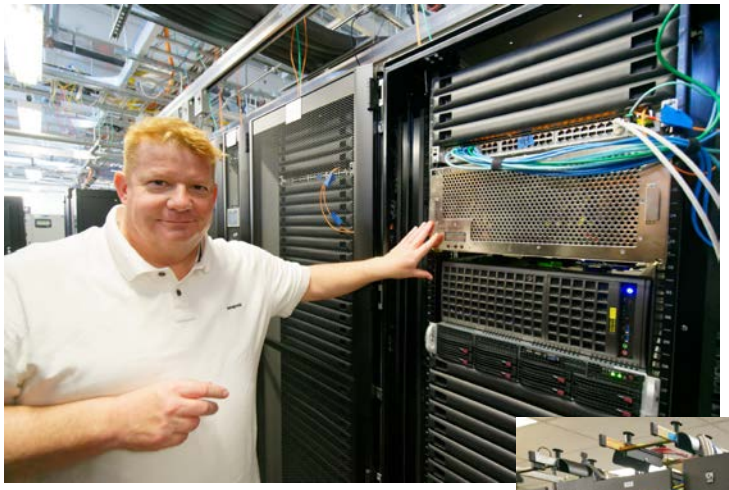

PRP
PACIFIC RESEARCH PLATFORM

it²

Now We Use Kubernetes to Manage FIONAs Across PRPv2

"Kubernetes is a way of stitching together
a collection of machines into, basically, a big computer,"
--Craig Mcluckie, Google
and now CEO and Founder of Heptio

Allows the PRP to Deploy Petabytes at $10/TB/yr of Distributed Storage and GPUs for Data Science as well as Measure and Monitor Use

PRP PACIFIC RESEARCH PLATFORM

CLOUD NATIVE COMPUTING FOUNDATION

it²

PRP is Deploying Distributed Petabytes of DMZ FIONA Storage for Posting/Staging Data at $10/TB per Year by Leveraging our Base of Already Installed FIONAs

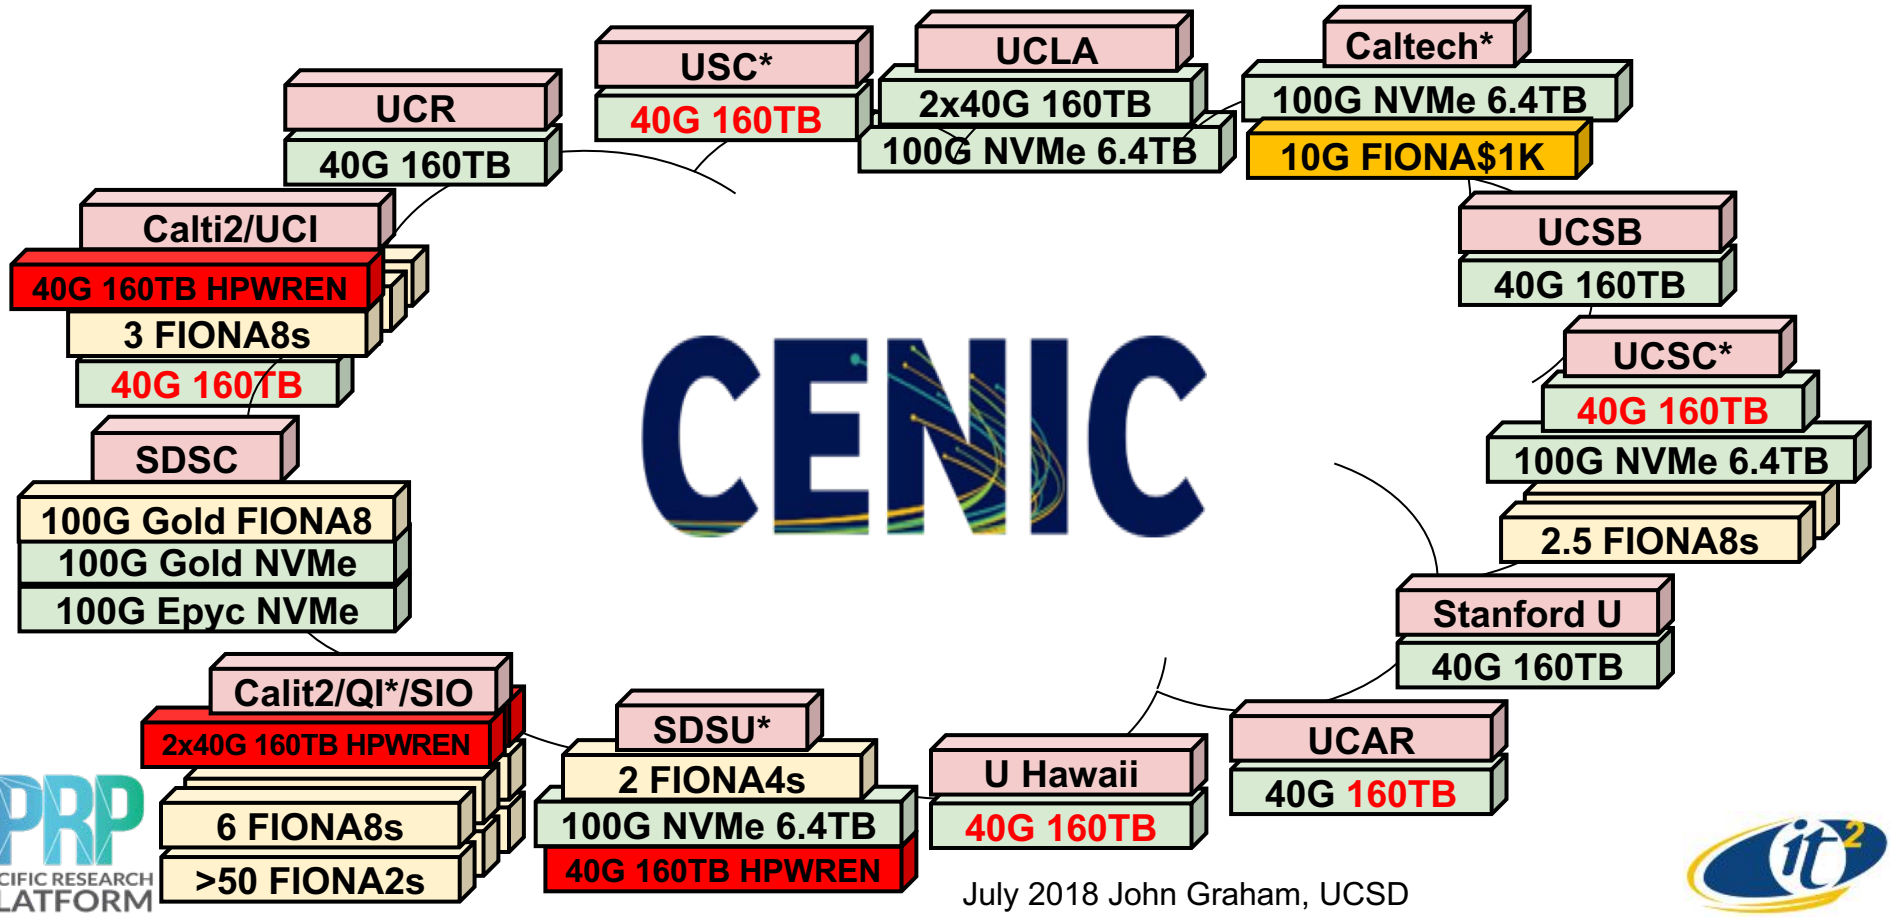

July 2018 John Graham, UCSD

Grafana Plot of First 730 TB at UCSD, Stanford, UCD, UCM

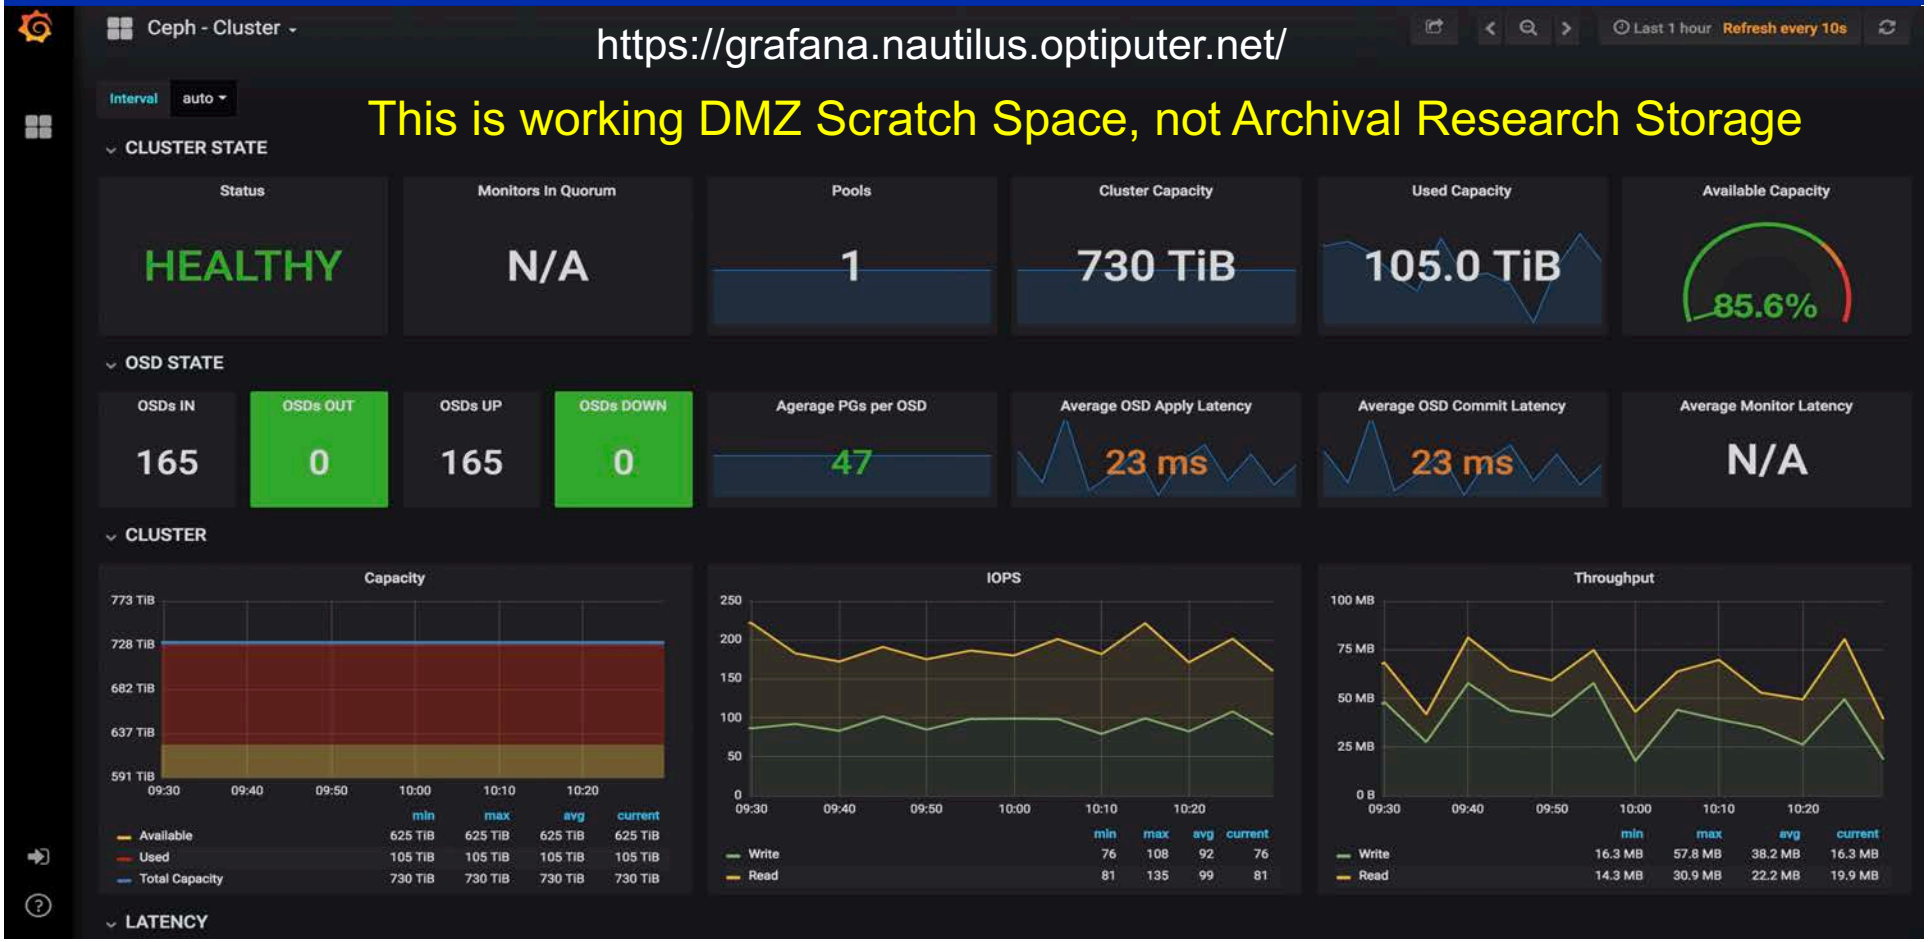

And the Community Computing Needed to Process the Data at a 10-1 Price Advantage vs. AWS (For Researchers, its 30-1)

Nvidia Card	~Cost	32-bit GF	GB	per GF	per GB	cores	8-GPU PC	160 GPU rack
GTX 1080 Ti 11GB	$726	10609	11	$0.07	$66	3584	$13,804	$276,090
P100 16GB	$8,304	8071	16	$1.03	$519	3584	$74,432	$1,488,640
AWS p2.xlarge EC2 (8) K-80 GPUs+disk for 3 years +55% ICR							$370,512	$7,410,240
AWS p2.xlarge EC2 (8) K-80 GPUs+disk for 3 years							$239,040	$4,780,800

40 GPUs in 5 FIONA8s for ML in 15” of rack space 2.4 Million GPU hours per day

PRP PACIFIC RESEARCH PLATFORM

it²

Grafana Plot of First 83 GPUs Adding Weekly/Weekley

https://grafana.nautilus.optiputer.net/

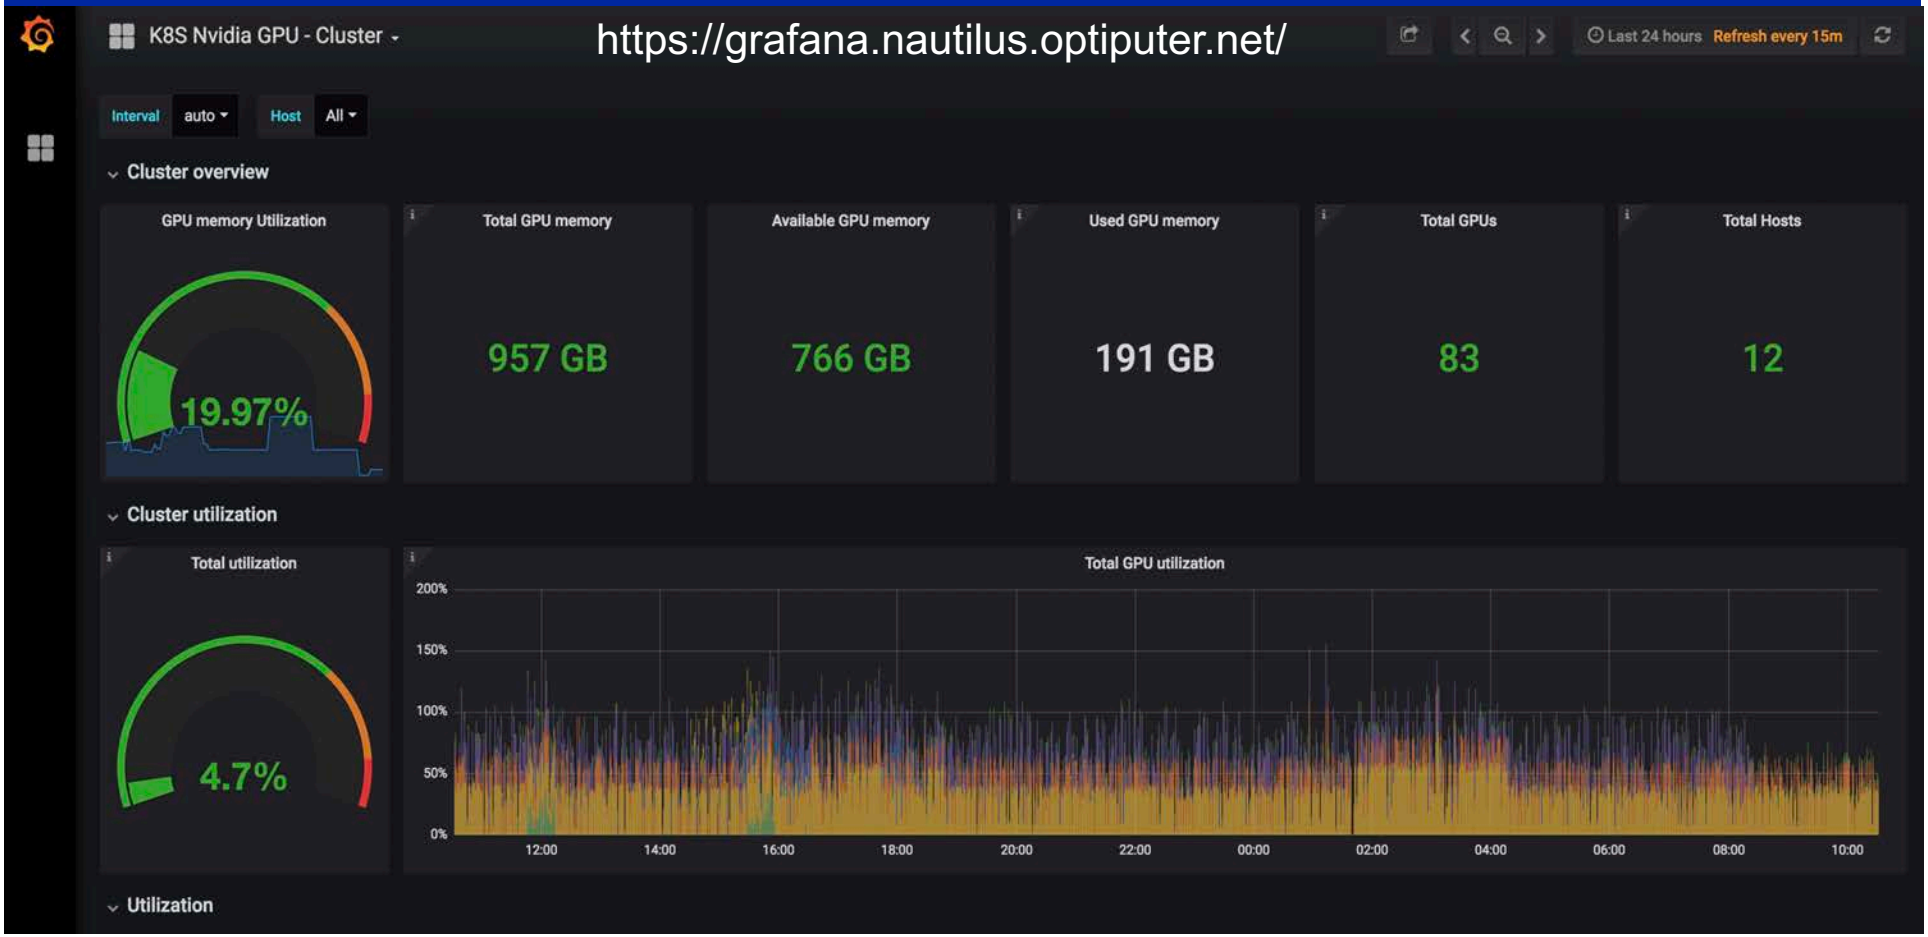

9 Xilinx Virtex 7 FPGA PCIe Development Platforms

- **Provided as a $250K gift from Calit2 partner KnuEdge**
- **Plugs into any FIONA in the extra 8-lane slot; can mix w/GPUs**
- **Almost 2M logic cells each card**

PRP
PACIFIC RESEARCH PLATFORM

it²

Reaching beyond CENIC/Pacific Wave: Using Internet2 to Connect Regional Networks—NRP Pilot

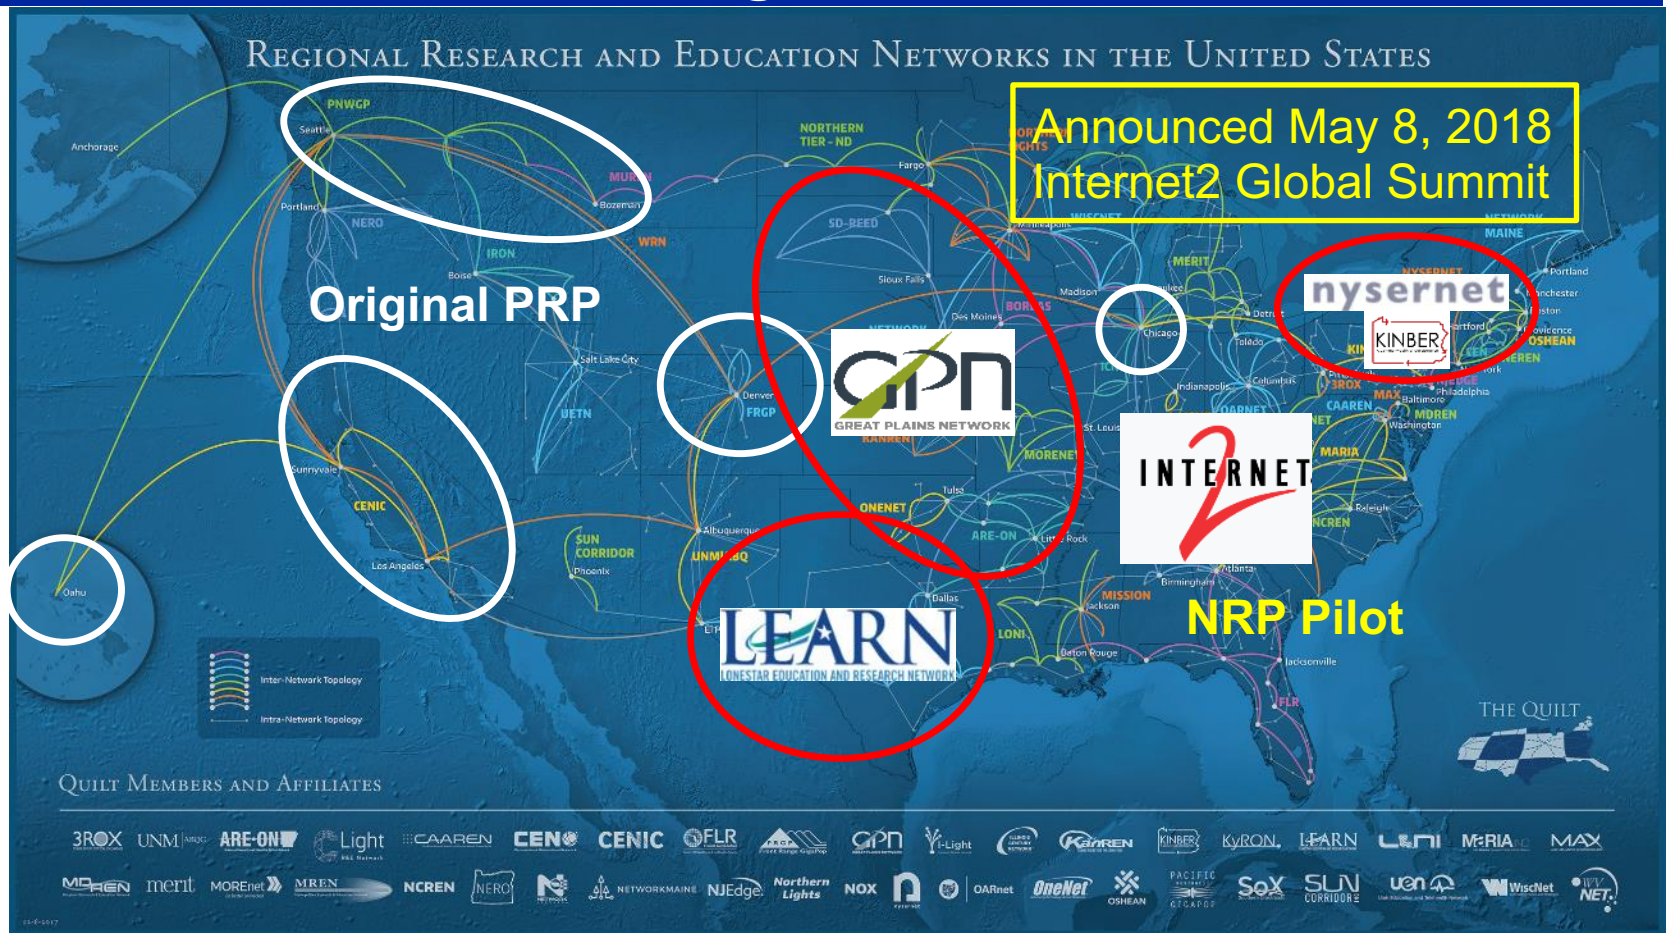

ESnet
CENIC
SDSC
CITRIS and the BANATAO INSTITUTE
PRP PACIFIC RESEARCH PLATFORM

PRP Has Great Support:

- US National Science Foundation (NSF) awards:

  CNS-0821155, CNS-1338192, CNS-1456638, CNS-1730158,
  ACI-1540112, & ACI-1541349
- CENIC, Pacific North West Gigapop (PNWGP), Pacific Wave & StarLight
- University of California Office of the President CIO
- UCSD Next Generation Networking initiative
- Calit2 and Calit2's Qualcomm Institute
- DOE ESnet

PRP
PACIFIC RESEARCH PLATFORM