



**measurement  
operations  
engagement**

18 October 2016

CANS Houston

*John Hess <[jhess@pacificwave.net](mailto:jhess@pacificwave.net)>*



# Pacific Research Platform

- NSF CC-NIE and similar projects represent significant investments in campus infrastructure including SDN, Science DMZ's (~130 projects)
- But scientists are still struggling with the complexity of using the network and interoperability between different implementations of Science DMZ's
- PRP focuses on enabling the science communities across the Pacific region to make effective use of the high performance infrastructure
- Kick-off meeting at Stanford (December 2014):
  - take advantage of the regional infrastructure
  - perfSONAR\_ for measurement / analysis
  - MaDDash for visualization
- Include DTN's: use a common software suite for data movement; reflect disk-to-disk performance on MaDDash
- PRPv0 demonstrated as a proof-of-concept at the CENIC Conference (March 2015)

# 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/>



Science DMZ  
Coined 2010

# Initially Proposed PRP Science Driver Teams

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

# One Year In: The Pacific Research Platform is Now a Working End-to-End Science-Driven DMZ-Connector



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, UC San Diego SDSC
- Frank Wuerthwein, UC San Diego Physics SDSC



Source:  
John Hess, CENIC

# Pacific Research Platform

“For the hardest problems—not just in physics but in climate science and genomics—there are massive teams working around the world,” says ESnet director Greg Bell. “Our job is to make geography irrelevant.”

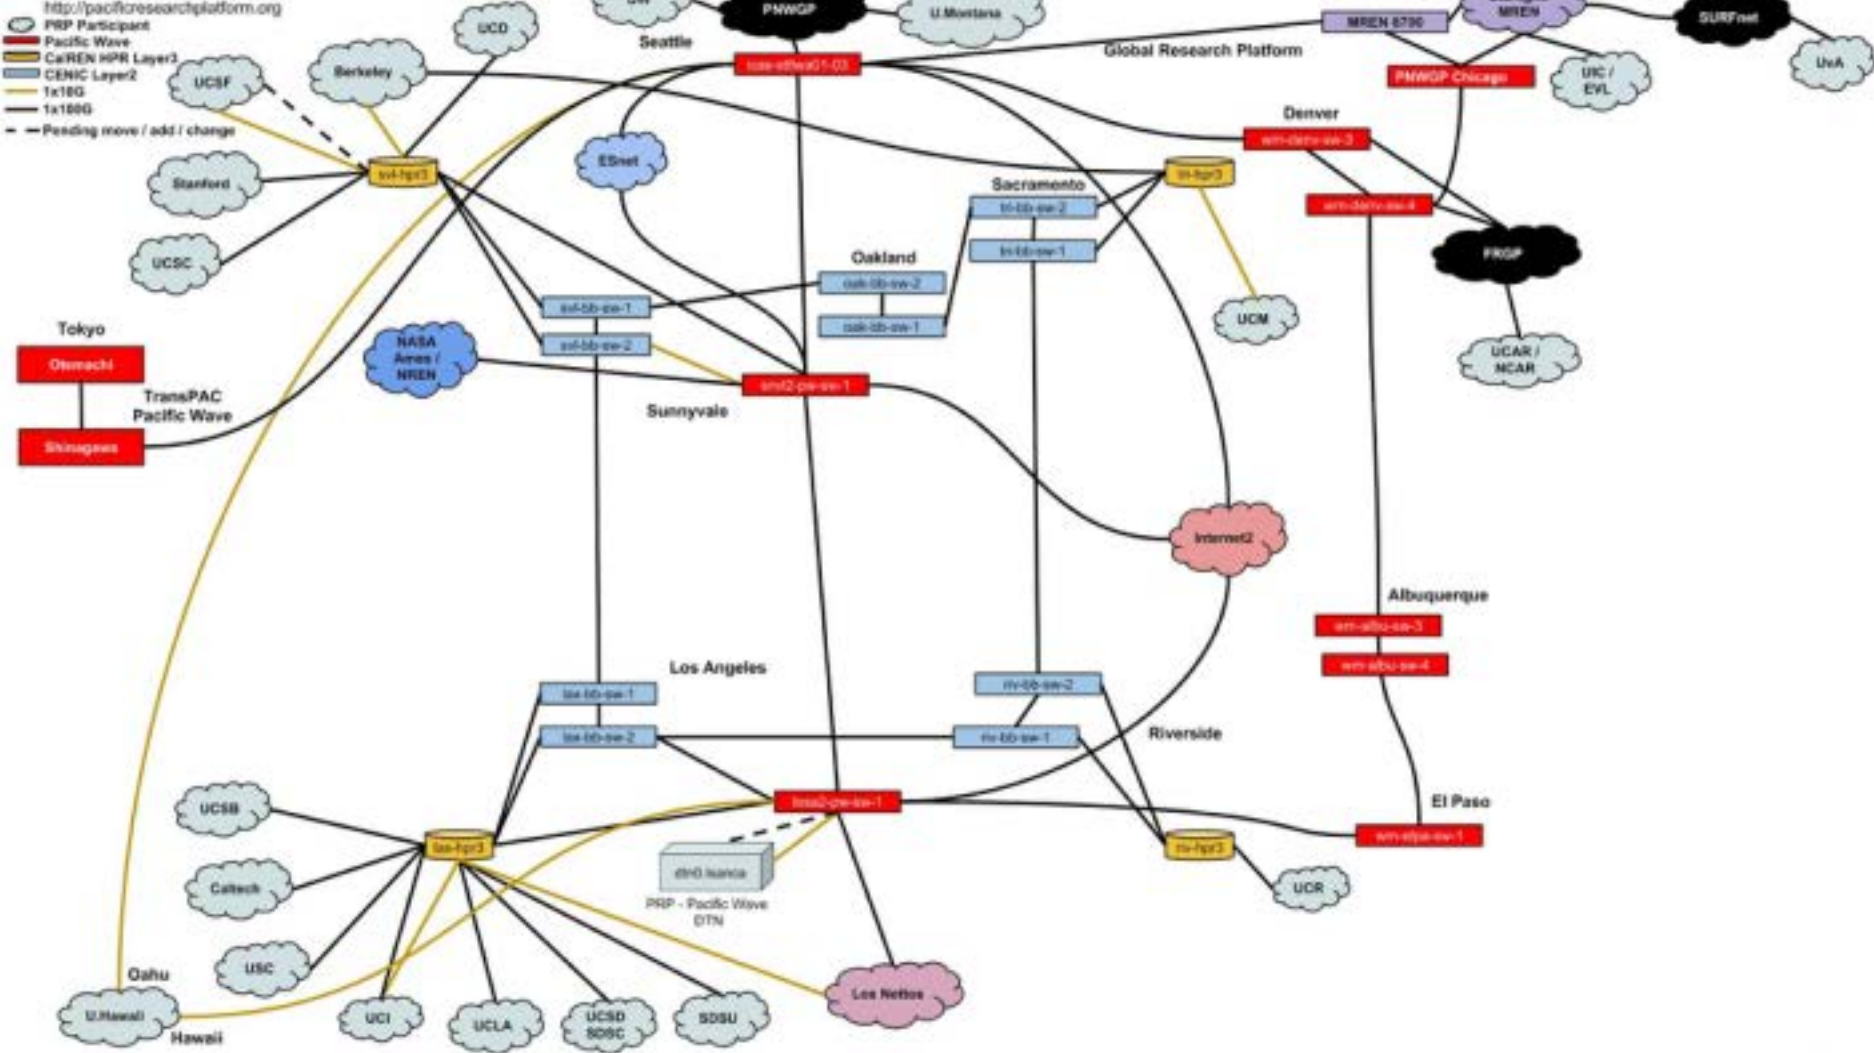
Inter-Institutional ScienceDMZ (PRP) kick-off meeting, Stanford, November 2014

2014 “WHAT TO EXPECT IN 2015: ULTRAFAST DATA TRANSFER SPEEDS UP SCIENCE” Popular Science, December 2014



# Pacific Wave: Supporting collaboration

## Pacific Research Platform v1



NOTE: This diagram represents a subset of sites, devices, and connections

v1.00  
20160629

# Pacific Wave: Measurement and Visualization

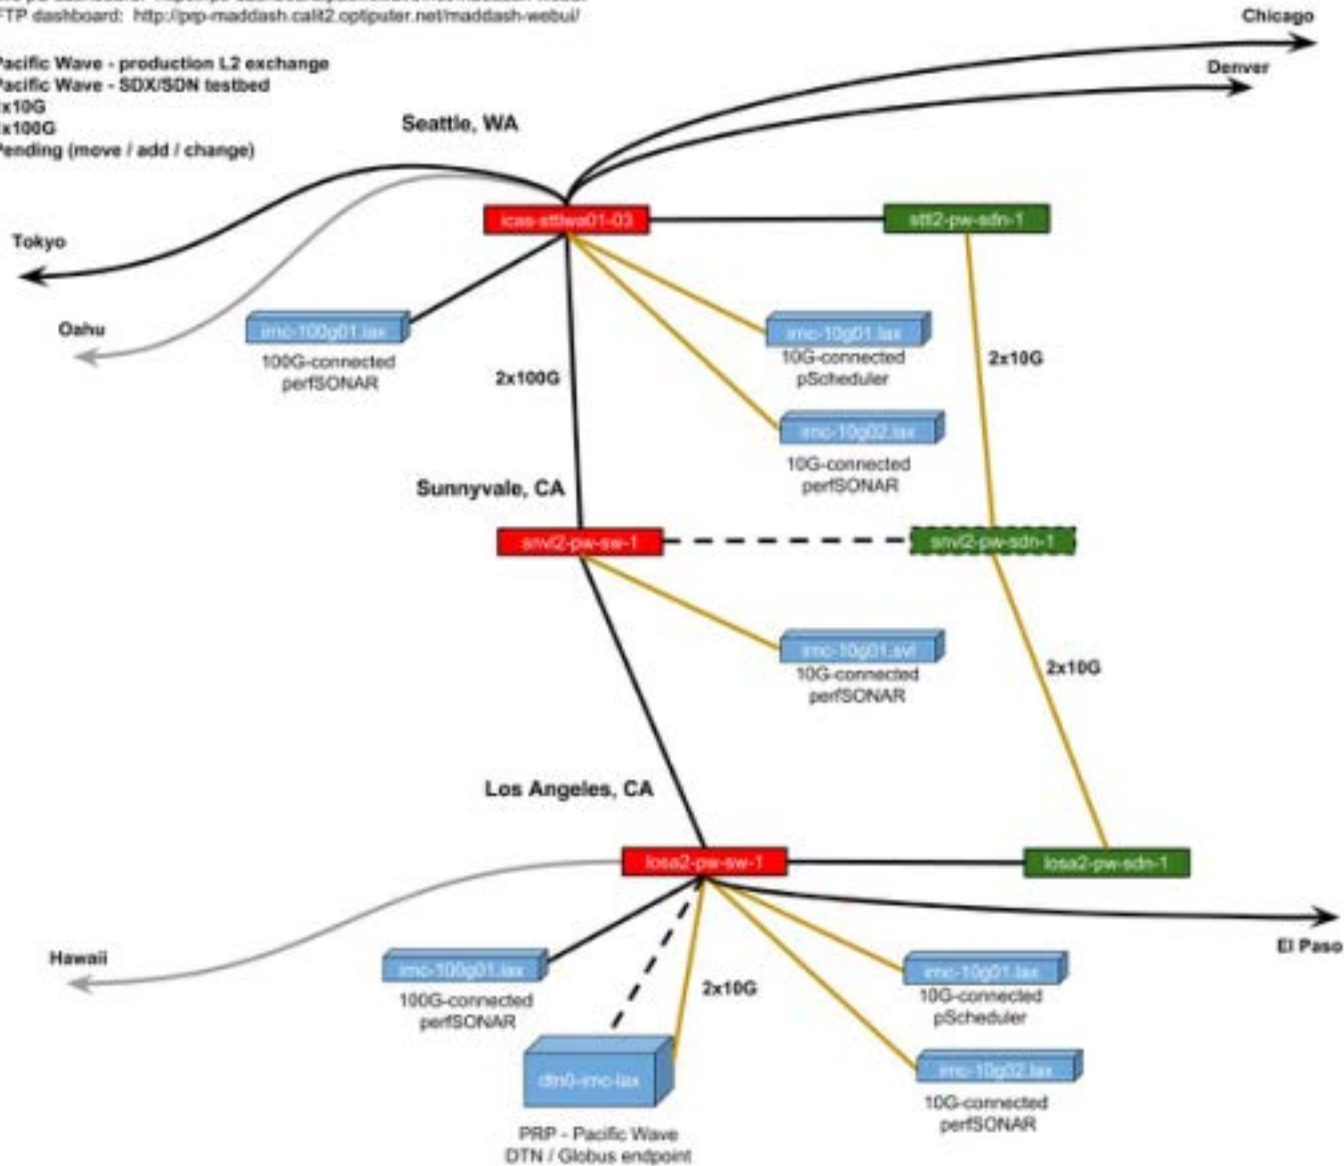
- 10G-connected pS nodes deployed in Seattle, Sunnyvale, and Los Angeles
- 10G-connected pScheduler node in testing at Los Angeles: moving to pS 4.0RC1
- 2x10G-connected PRP-contributed DTN in testing at Los Angeles
  - GridFTP testpoint on PRP MaDDash
  - Globus-managed endpoint for science data movement
  - platform for exploring ICN, NDN, ...
- 100G-connected perfSONAR nodes in testing at Seattle and Los Angeles
  - currently Centos 6.8 w/pS 3.5.x
  - moving to CentOS 7.2 w/pS 4.0RC1
- Pacific Wave perfSONAR dashboard beta testing MaDDash 2.0RC1
  - central esmond MA
  - a collaboration of Pacific Wave participants and partners
- IRNC NOC Performance Engagement Team



# Pacific Wave: Performance Measurement

Pacific Wave pS dashboard: <https://ps-dashboard.pacificwave.net/maddash-webui/>  
 PRP GridFTP dashboard: <http://prp-maddash.ca12.cptiputer.net/maddash-webui/>

- Pacific Wave - production L2 exchange
- Pacific Wave - SDX/SDN testbed
- 1x10G
- 1x100G
- - - Pending (move / add / change)



NOTE: this diagram represents a subset of sites, devices, and connections

v0.11  
20160929



# Pacific Wave: Measurement Visualization

## Pacific Wave perfSONAR Dashboard: a collaborative project of Pacific Wave partner and participant networks

[Dashboards](#)
[Reports](#)
[Settings](#)
[Participants & Partners, and Affiliations](#)


Last page refresh time: September 30, 2016 04:43:01 AM PDT

### Pacific Wave Dashboard

#### Pacific Wave - IPv4 Throughput

■ Throughput >= 900Mbps
 ■ Throughput < 900Mbps
 ■ Throughput <= 500Mbps
 ■ Unable to retrieve data
 ■ Check has not yet run


⚠ Found a total of 10 problems involving 8 hosts in the grid



#### Pacific Wave - IPv6 Throughput

■ Throughput >= 900Mbps
 ■ Throughput < 900Mbps
 ■ Throughput <= 500Mbps
 ■ Unable to retrieve data
 ■ Check has not yet run

⚠ Found a total of 11 problems involving 8 hosts in the grid



#### Pacific Wave - IPv6 Packet Loss

■ Loss rate is <= 0
 ■ Loss rate is >= 0
 ■ Loss rate is >= 0.01
 ■ Unable to retrieve data
 ■ Check has not yet run

✅ No problems found in grid



# Weekly PRP Network Engineer Call and Notes Subscription List (60)

kmayeshiro@ucdavis.edu (Kevin Mayeshiro)  
drapp@ucsd.edu (Dave Rapp)  
dart@es.net (Eli Dart)  
ppapadopoulos@eng.ucsd.edu (Phillip Papadopoulos)  
timothy.carlson@pnnl.gov (Timothy Carlson)  
rpwagner@sdsc.edu (Rick Wagner)  
robert.tannenbaum@ucsf.edu (Tannenbaum, Robert)  
jgraham@ucsd.edu (John Graham)  
doan@usc.edu (Thanh Doan)  
mvn@ucla.edu (Michael Van Norman)  
dave@cenic.org (Dave Reese)  
geoff@ucsd.edu (Geoff Davis)  
pmurray@stanford.edu (Paul Murray)  
jah@ucsc.edu (John Haskins)  
darrell@cenic.org (Darrell Newcomb)  
jvanreij@stanford.edu (Johan van Reijendam)  
michael.duff@stanford.edu (Michael Duff)  
mark.foster@nasa.gov (Mark Foster)  
azher@hep.caltech.edu (Azher Mughal)  
russ-harvey@ucr.edu (Russ Harvey)  
ghidley@ucsd.edu (Hidley, Gregory)  
jnielsen@ucsc.edu (Jason Nielsen)  
dan.spanner@pnnl.gov (Dan Spanner)  
warner@ucsc.edu (Jim Warner)  
addlema@iu.edu (Hans Addleman)  
ccrittenden@berkeley.edu (Camille Crittenden)  
tdefanti@ucsd.edu (DeFanti, Tom)  
jshnell@berkeley.edu (Jack Schnell)  
bvincen@stanford.edu (Bruce Vincent)  
schylerb@uw.edu (Schyler Batey)

bill.strossman@ucr.edu (Bill Strossman)  
hutton@ucsd.edu (Hutton, Thomas)  
jeff.haferman@gmail.com (Jeff Haferman)  
jhess@cenic.org (John Hess)  
kollross@illinois.edu (Matthew Kollross)  
jmadden@ucsd.edu (Jim Madden)  
jmschopf@iu.edu (Jennifer Schopf)  
celestea@usc.edu (Celeste Anderson)  
delaat@uva.nl (Cees de Laat)  
jyy@uci.edu (Jessica (Jie Yun) Yu)  
brad@ucsc.edu (Brad Smith)  
r-butler@illinois.edu (Randy Butler)  
jmeyer@ucsd.edu (Meyer, Jon)  
michael.kilpatrick@ucsf.edu (Michael Kilpatrick)  
manish.sompura@ucsf.edu (Sompura, Manish)  
pepin@clmson.edu (James Pepin)  
pradulov@indiana.edu (Predrag Radulovic)  
gpeek@ucsc.edu (George Peek)  
erik.mccroskey@berkeley.edu (Erik McCroskey)  
paolini@engineering.sdsu.edu (Chris Paolini)  
newman@hep.caltech.edu (Harvey Newman)  
sbellamine@cenic.org (Sana Bellamine)  
jsilvest@usc.edu (John A. Silvester)  
jsonstro@ucsc.edu (Joshua Sonstroem)  
jweekley@ucmerced.edu (Jeffrey Weekley)  
tboerner@ncsa.illinois.edu (Tim Boerner)  
vpolichar@ucsd.edu (Polichar, Valerie)  
ismarr@ucsd.edu (Larry Smarr)  
mleeoh@gmail.com (Jysop Lee, KISTI)  
pschmitz@berkeley.edu (Patrick Schmitz)





# The Campus FIONAs – Flash I/O Network Appliance: Linux PCs Optimized for DMZs over Distance

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

**Phil Papadopoulos & Tom DeFanti  
Joe Keefe & John Graham**



## Typical FIONA Rack-Mount Build:

<b>Cost</b>	<b>\$8,000</b>	<b>\$20,000</b>
<b>Intel Xeon Haswell</b>	<b>E5-1650 v3 6-Core</b>	<b>2x E5-2697 v3 14-Core</b>
<b>RAM</b>	<b>128 GB</b>	<b>256 GB</b>
<b>SSD Flash</b>	<b>SATA 3.8 TB</b>	<b>SATA 3.8 TB</b>
<b>Network Interface</b>	<b>10/40G</b>	<b>2x40G or 100G</b>
<b>GPU</b>	<b>none</b>	<b>NVIDIA Tesla K80</b>
<b>RAID Drives 0 to 200TB (add ~\$100/TB); \$\$ NVMe for 100G</b>		




# PRP use cases

## NCMIR

National Center for Microscopy and Imaging Research  
<https://ncmir.ucsd.edu>

- Pacific Wave
- CalREN HPR Layer3
- CENIC Layer2
- 1x10G
- 1x100G
- Pending (move / add / change)



NOTE: This diagram represents a subset of sites, devices, and connections

v4.07  
20160829




# PRP MaDDash: Point-to-Point Bandwidth Map GridFTP File Transfers-Note Huge Improvement in Five Months

■ Throughput  $\geq$  5000Mbps   
 ■ Throughput  $<$  5000Mbps   
 ■ Throughput  $\leq$  1000Mbps   
 ■ Unable to retrieve data   
 ■ Check has not yet run

January 29, 2016 PRPV1 (L3)


Green is Disk-to-Disk  
In Excess of 5Gbps

June 6, 2016 PRPV1 (L3)






# 9/23/16: More DMZs



PRPGridFTPdataMovers



# PRP Roadmap

## ■ PRPv1

- A routed Layer 3 architecture
- Tested, measured, optimized, with multi-domain science data
- Bring many of our science teams up
- Each community with 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
- Trusted Platform Module (TPM) hardware
- Support bit-rates up to 100Gbps in bursts and streams
- Develop means to operate a shared federation of caches

# Resources:



**Pacific Wave**  
<http://www.pacificwave.net/>  
<https://ps-dashboard.pacificwave.net>



**CENIC**  
<http://www.cenic.org/>  
<https://ps-dashboard.cenic.net>



**PNWGP**  
<http://www.pnwgp.net/>



**Pacific Research Platform**  
<http://pacificresearchplatform.org/>  
[http://cenic.org/files/publications/PRP\\_Overview\\_%C6%92.pdf](http://cenic.org/files/publications/PRP_Overview_%C6%92.pdf)  
<http://prp-maddash.calit2.optiputer.net/maddash-webui/>



**Calit2**  
<http://www.calit2.net/>



**CITRIS**  
<http://citris-uc.org/>



**ESnet**  
<http://www.es.net/>  
<http://fasterdata.es.net/>  
<http://ps-dashboard.es.net/>



**NSF**  
<http://www.nsf.gov/>



# Thanks!

Pacific Wave:

John Hess <jhess@pacificwave.net>

PacificWave NOC

24x7x365

noc@pacificwave.net

+1-888-PAC-WAVE

+1-206-722-9283