



SwitchON Workshop

16 October 2015

São Paulo, Brasil

---

# **US West-Coast Future Internet Infrastructure**

## **Pacific Wave Update**

## **Pacific Research Platform**

## **International Routing Research Collaboration**

---

*John Silvester, [silvester@usc.edu](mailto:silvester@usc.edu)*



# PACIFIC WAVE UPDATE

# A Brief History of Pacific Wave

- **Late 1990's:** Exchange points in Seattle (UW and PNWGP) and Los Angeles (USC and Los Nettos)
- **2004:** Birth of Pacific Wave as a joint project of CENIC and PNWGP – introduction of the “distributed exchange” concept
- **2005-2010** US West Coast dark fiber acquisition; addition of Sunnyvale (SF Bay Area), mostly to allow connection to some US Federal R&E networks; backbone 10G; connectors 100M - 10G; NSF funding (NSF funding OCI-IRNC TransLight Pacific Wave)
- **2011-2015** Multiple 10G to 100G; NSF funding OCI-IRNC TransLight Pacific Wave II)



- Operates US Western Region R&E Exchange Points
  - PW has exchange points in the Western United States at Seattle, SF Bay Area (Sunnyvale), Los Angeles, Denver, Albuquerque and El Paso
  - Includes a diverse ring fabric to allow participating networks to peer across that fabric with any other participant
  
- 17 major R&E networks directly connect to the Pacific Wave fabric     [Current PW Participants](#)

# Current Developments

## Redundancy and Span

- To improve resiliency and to allow peering with European Networks in Chicago and other networks and entities led to development of a ring topology
- Additional PW access points
  - Denver, CO – The Front Range GigaPoP (FRGP) serves universities in Colorado as well as NOAA, UCAR and other major national facilities
  - Albuquerque, NM – The University of New Mexico operates the Albuquerque GigaPoP serving universities in New Mexico and connects to both ESnet and Los Alamos National Lab
  - El Paso, TX – Operated by University of New Mexico, El Paso is expected to be a major junction for US-Mexico connectivity
  - Hawaii – The University of Hawaii is a long time collaborator with Pacific Wave and has recently added additional capacity between Hawaii and mainland US

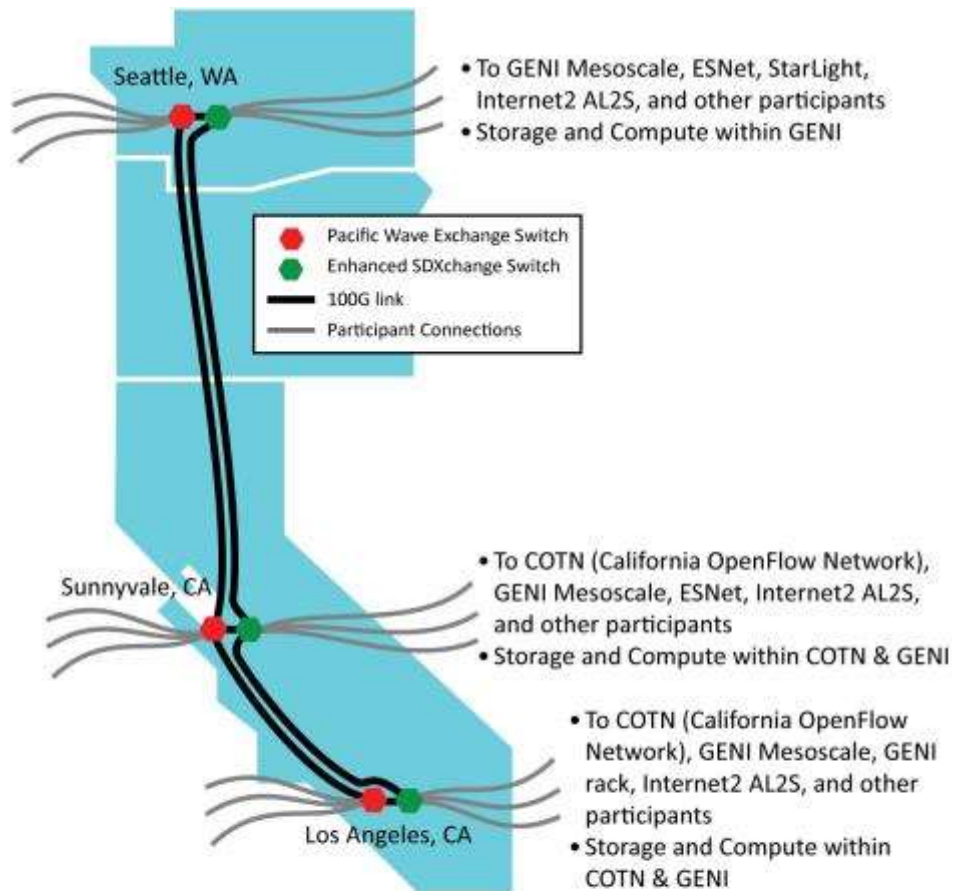
# CENIC/Pacific Wave NSF Award

- Pacific Wave Expansion Supporting SDX & Experimentation (ACI-1451050, September 2015)
  - Continued enhancement, upgrade and evolution of Pacific Wave to support more 100G connections
  - Additional 100G capacity between exchanges points along West Coast
  - SDN/SDX deployment on parallel infrastructure to enable experimentation while maintaining production use of the Pacific Wave exchange
  - Collaboration with other IRNC awardees on SDX development, measurement and monitoring.
  - perfSONAR deployment on Pacific Wave



# Distributed Exchange or Network?

- The result looks more like a (layer 2) network rather than an extruded linear distributed exchange but we prefer the term “Distributed Exchange” since it focuses on how the fabric is used to enable peering between networks.



**Traditional Peerings** between all Pacific Wave participants can remain on the Pacific Wave Exchange Switches.

**SDX Services** can be accessed via direct connection to Enhanced SDXchange Switches *OR* via connections to the Pacific Wave Exchange Switches.

**Mixed connections** can also take place between participants where one is connected to the Pacific Wave Exchange Switch and the other to an Enhanced SDXchange Switch.



# PACIFIC RESEARCH PLATFORM

# The Pacific Research Platform (PRP)

- NSF CC-NIE and similar projects represent significant investments in campus infrastructure including SDN, DMZ's (~130 projects)
- But the scientists are still struggling with the complexity of using the network and interoperability between different implementations of DMZ's
- PRP focuses on enabling the science communities to make effective use of the high performance infrastructure that is available.
- The idea was hatched in December 2014 – take advantages of the infrastructure, including a [PERFSONAR grid](#) for measurement.
- And DTN's and common software suite to demonstrate a proof of concept for the PRP
- Demonstrated at the CENIC Spring meeting (March 2015)

# Ten Week Sprint to Demonstrate the Interoperability of Science DMZs with CENIC/PW Backplane

Presented at CENIC 2015  
March 9, 2015

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.”**



2015 Conference  
SHAKING THINGS UP

March 9-11 2015



# Pacific Research Platform Project

- Led to a successful multi-year NSF proposal to move forward with the further development of PRP.
- PI: Prof Larry Smarr – CALIT2, UCSD  
Co-PI's: Camille Crittenden, CITRIS UCB  
Tom DeFanti – CALIT2, UCSD  
Philip Papadopoulos, SDSC, UCSD  
Frank Wuerthwein, UCSD and SDSC
- Press releases  
[CALIT2 PR](#); [CENIC PR](#)

# Pacific Research Platform - Proposal

## Abstract

The Pacific Research Platform is a project to forward the work of researchers and their access to technical infrastructure, with a vision of connecting all the National Science Foundation Campus Cyberinfrastructure grants (NSF CC-NIE & CC-IIE) to research universities within the region, as well as the Department of Energy (DOE) national labs and the San Diego Supercomputer Center (SDSC).



## Science Drivers

Particle Physics

Astronomy and Astrophysics

Biomedical

Earth Sciences

Scalable Visualization, Virtual Reality, and Ultra-Resolution Video

# Pacific Research Platform Driven by Data-Intensive Research



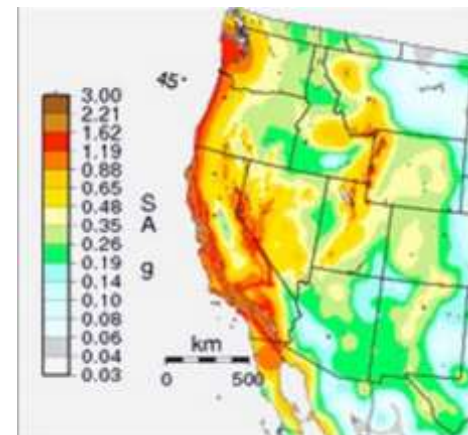
Particle  
Physics

Biomedical  
'omics



Telescope  
Surveys

Earthquake  
Engineering



Visualization,  
Virtual Reality,  
Collaboration



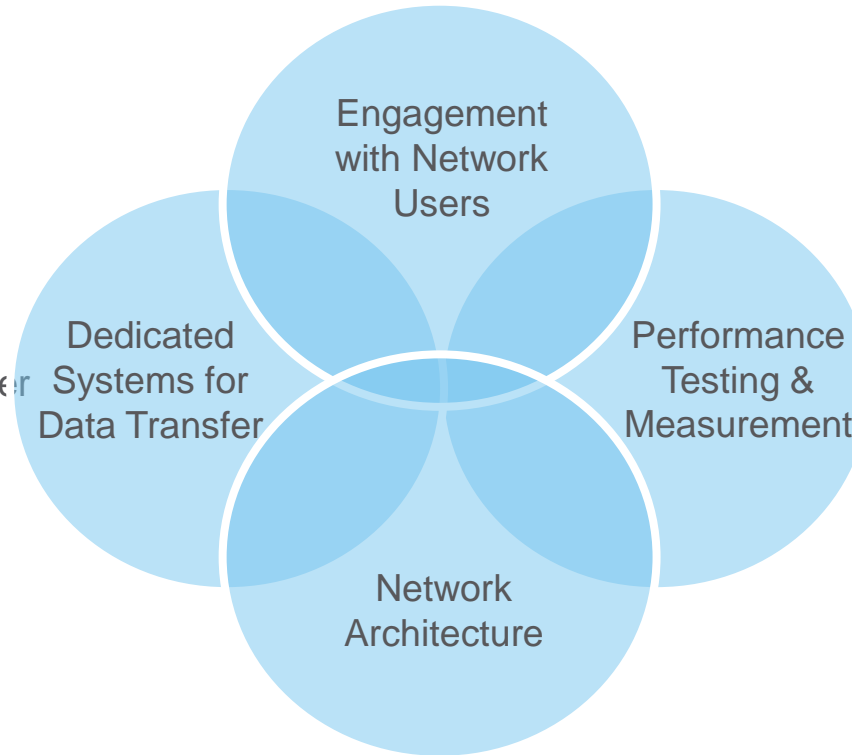
# Science DMZ Superfecta: Engagement

## Engagement

- Partnerships
- Education & Consulting
- Resources & Knowledgebase

## Data Transfer Node

- High performance
- Configured for data transfer
- Proper tools



## perfSONAR

- Enables fault isolation
- Verify correct operation
- Widely deployed in ESnet and other networks, as well as sites and facilities

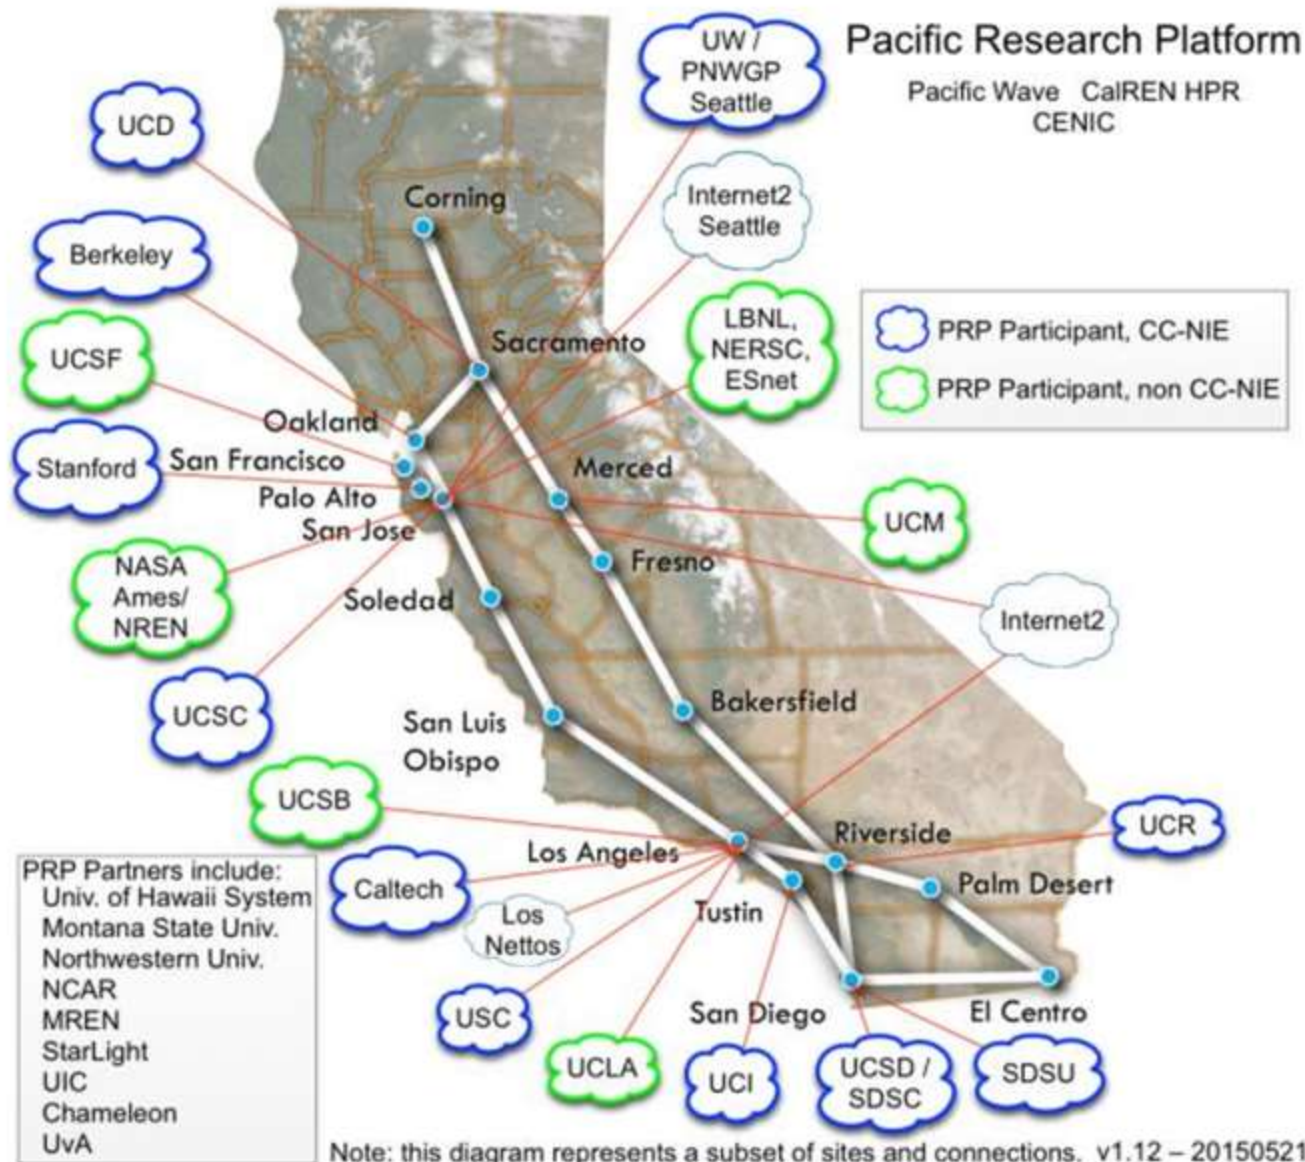
## Science DMZ

- Dedicated location for DTN
- Proper security
- Easy to deploy - no need to redesign the whole network



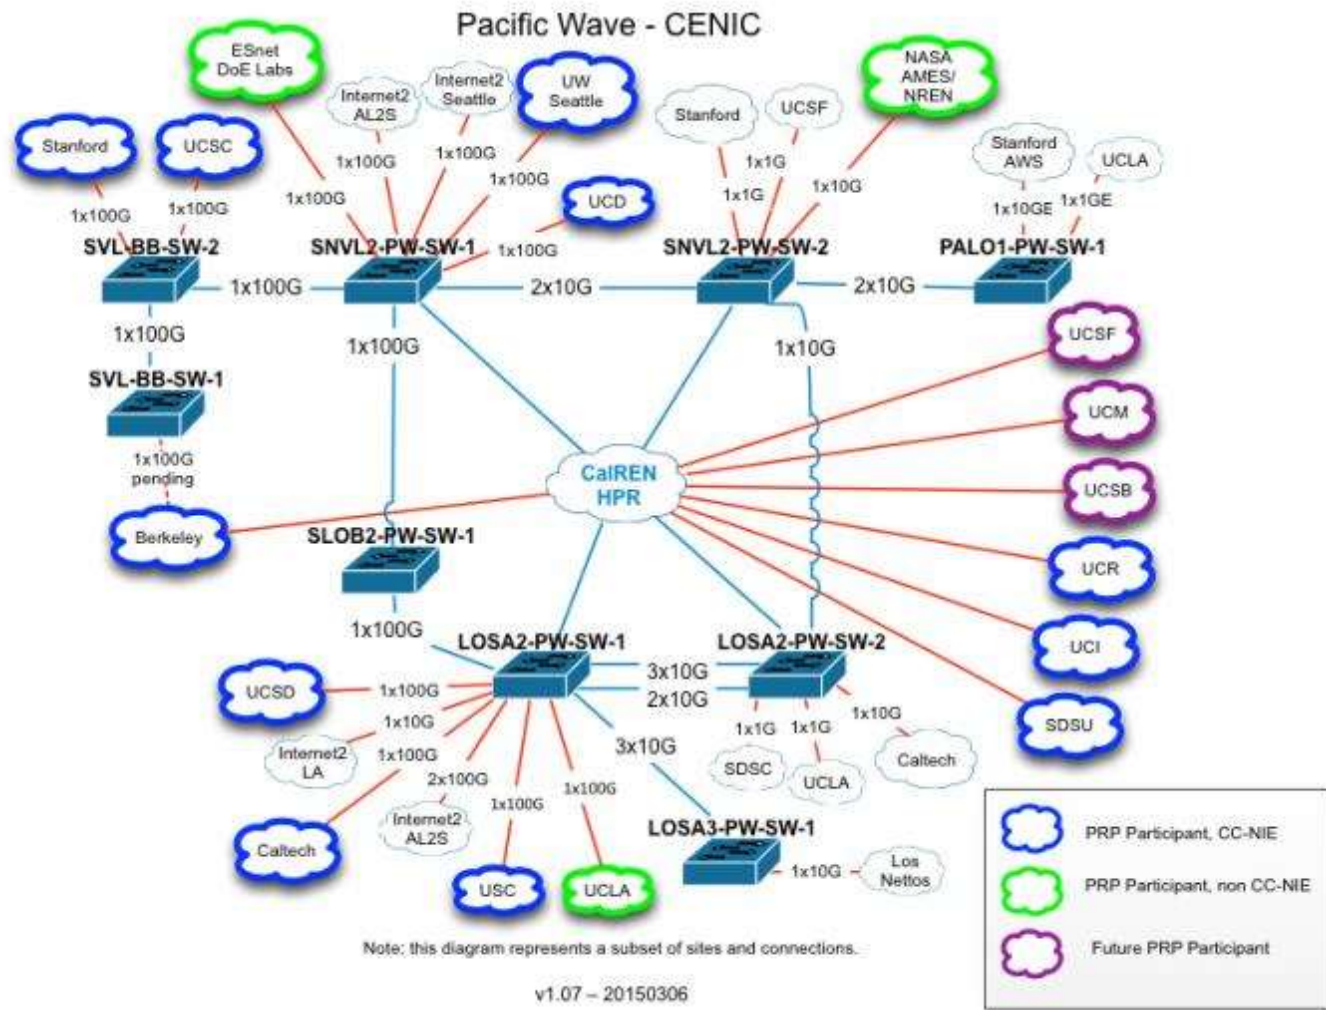


# PRP: A Regional Science DMZ





# CENIC/Pacific Wave Network supporting PRP prototype demo



# Next Steps

- Migrate from experiment to persistent infrastructure
- Expand DTN and perfSONAR deployment to include all 10 UC campuses, Caltech, Stanford, USC, and University of Washington
- Incorporate federated access to resources
- Engage with scientists to begin to map their research collaborations across the Pacific Research Platform
- Engage with PRAGMA and perfSonar grid
- Work with campus IT organizations to make “last mile” connections between researchers and the Science DMZ

# Longer-term Goals

---

- PRP Evolution An Integrated West Coast Science DMZ for Data-Intensive Research
- Science DMZ interoperability / integration across regions, nationally, and internationally
- Managing Domain Specific Science Communities - we anticipate that as it matures, SDN is a natural technology to manage the multiple science domain structures (including security) envisioned by PRP and that we envision that in the next few years the PW fabric will be run using SDN for that purpose
- IPv6
- Expand to Commercial services – Amazon AWS, Microsoft Azure, Google?

**USC-UFGM Collaboration**

**INTERNET ROUTING RESEARCH**

**PEERING PLATFORM**

# USC & UFMG collaborate on routing research

Active collaboration since 2011

- USC: Ethan Katz-Bassett and his students
  - co-leads Networked Systems lab with Ramesh Govindan, Wyatt Lloyd, and Minlan Yu
- UFMG: Italo Cunha and his students
- Research focus: Systems for measuring Internet routing
  - Motivation: Routing contributes to performance, availability, and security problems.
    - Lack of visibility hinders operations and research.
  - Goal: Develop distributed measurement systems to enable better operation & novel research.

# PEERING: An AS for us

- Traditionally, researchers are outside Internet's routing ecosystem, so rely on measurement of existing routes and on modeling
- PEERING is a testbed that allows researchers to emulate a network, connect it to real ISPs, then announce/select routes and send/receive traffic
  - Currently 11 PoPs around the world, including US and Brazil
  - Hundreds of peers, including Hurricane Electric, Google, etc
  - Used by us and others in papers at SIGCOMM, IMC, SIGMETRICS, SOSR, HotNets
  - Topics include: Measuring routing policies, testing software-defined Internet exchange (SDX), measure RPKI-based BGP security adoption
- Looking for people who can deploy a new site (especially at an IXP) or who want to run experiments
- See <http://peering.usc.edu> for more details

## Providing visibility into routing: *LIFEGUARD*

- **Locating Internet Failures Effectively and Generating Usable Alternate Routes Dynamically (SIGCOMM 2012)**
  - Long-lasting partial outages contribute much of the Internet unavailability
  - **LIFEGUARD**: techniques a provider can use to locate even unidirectional failures outside the provider's network and to steer traffic to the provider around the failures

## Providing visibility into routing - *PoiRoot*

- Locating the root cause of Internet path changes (SIGCOMM 2013)
  - When a path changes, how can we tell what triggered the change, given that changes cascade through the Internet according to policy and topology that may not be visible to us?
  - PoiRoot: a system a provider can use to isolate the network responsible for changes



## Providing visibility into routing: *Sibyl*

- Sibyl: A practical Internet route oracle (in progress)
  - Researchers and operators need route information to answer questions, but available tools like traceroute only answer "*What is the path from here to there?*"
  - Sibyl: A platform that combines available vantage points to answer rich queries like "*What paths traverse a Level3-Sprint peering other than LA or NYC?*"