

# Pacific Research Platform

## the US West Coast Big Data Freeway System

presented by: Cees de Laat  
many, but not all, slides from:  
**Tom DeFanti & Larry Smarr et al.**



Pacific Research Platform leadership team: (L-R) Larry Smarr, Camille Crittenden, Philip Papadopoulos, Thomas DeFanti, and Frank Würthwein



# CdL Sabbaticals in spring 2015



## 1 - Berkeley - San Diego

### All Photos sabb-1



## 2 - Washington - Chicago - Ottawa

### All Photos sabb-2



# Why me?



### Activity

work

leisure

### Purpose

This personal website contains logs of my sabbatical activities in the spring of 2015. I went on these sabbaticals to be able to engage the community based on content in the more and more managerial stuff I was sucked into at UvA. Also my participation in the SURF scientific advisory board made it necessary to study future directions of cyber infrastructure. Apart from those I always keep on my toes if the research questions we as a group at UvA are studying, are still novel and valid. Directions change very rapidly particular field of study in System and Network Engineering.

### Goals

- Discuss with peers the 5 and 10 year outlook for Cyber Infrastructure. Budgets are shrinking, networks become mature, Supercomputing and High Throughput data compressing is now also done with huge commercial cloud centers, so what is the position of specific mission resources in this wild ocean of public capabilities. What do science community need to do ourselves and what can we just buy in the future from the (cloud) market. What do and what do we not need to do on leadership supercomputing. How do we relate to public cloud. We also do not have scientific water. What do we need to do on data at the central level? Do we need a national data service including data stewardship, etc.
- Next year we are chairing the eIRG workshop. The themes will be CI directions and funding models on national and european level. During the sabbatical seeking out potential speakers.
- Reflection on our own research questions

### Personal

- Take some holiday
- Have Emelie joining me for a couple of days
- Drive the HWY1
- Visit the desert
- Go to Chicago theatre, music bars, poetry



# Was there on sabbatical



# CC-DNI – DIBBs program

- Data Infrastructure Building Blocks (DIBBs) program.
- The DIBBs program encourages development of robust and **shared data-centric cyberinfrastructure** capabilities to accelerate **interdisciplinary and collaborative research** in areas of inquiry stimulated by data.
- Effective solutions will **bring together cyberinfrastructure expertise and domain researchers**, to ensure that the resulting cyberinfrastructure components address the researchers' data needs.
- **→→→→ Science Drivers!**
- **→→→→ CI plans!**

# Science Drivers

---

## Particle Physics Data Analysis

- The Large Hadron Collider (LHC). Run 2 will have ~2x the energy, generating ~10x the data volume of Run 1.

## Astronomy and Astrophysics Data Analysis

- Includes two data-intensive telescope surveys that are precursors to the Large Synoptic Survey Telescope (LSST)

*Intermediate Palomar Transient Factory (iPTF)*

*Dark Energy Spectroscopic Instrument (DESI)*

- Galaxy Evolution

*Southern California Center for Galaxy Evolution (CGE)*

*Assembling Galaxies of Resolved Anatomy (AGORA)*

- Gravitational Wave Astronomy

*The Laser Interferometer Gravitational-Wave Observatory (LIGO)*

## Biomedical Data Analysis

*Cancer Genomics Hub (CG Hub) and Cancer Genomics Browser*

*Microbiome and Integrative 'Omics*

*Integrative Structural Biology*

# Science Drivers (2)

---

## Earth Sciences Data Analysis

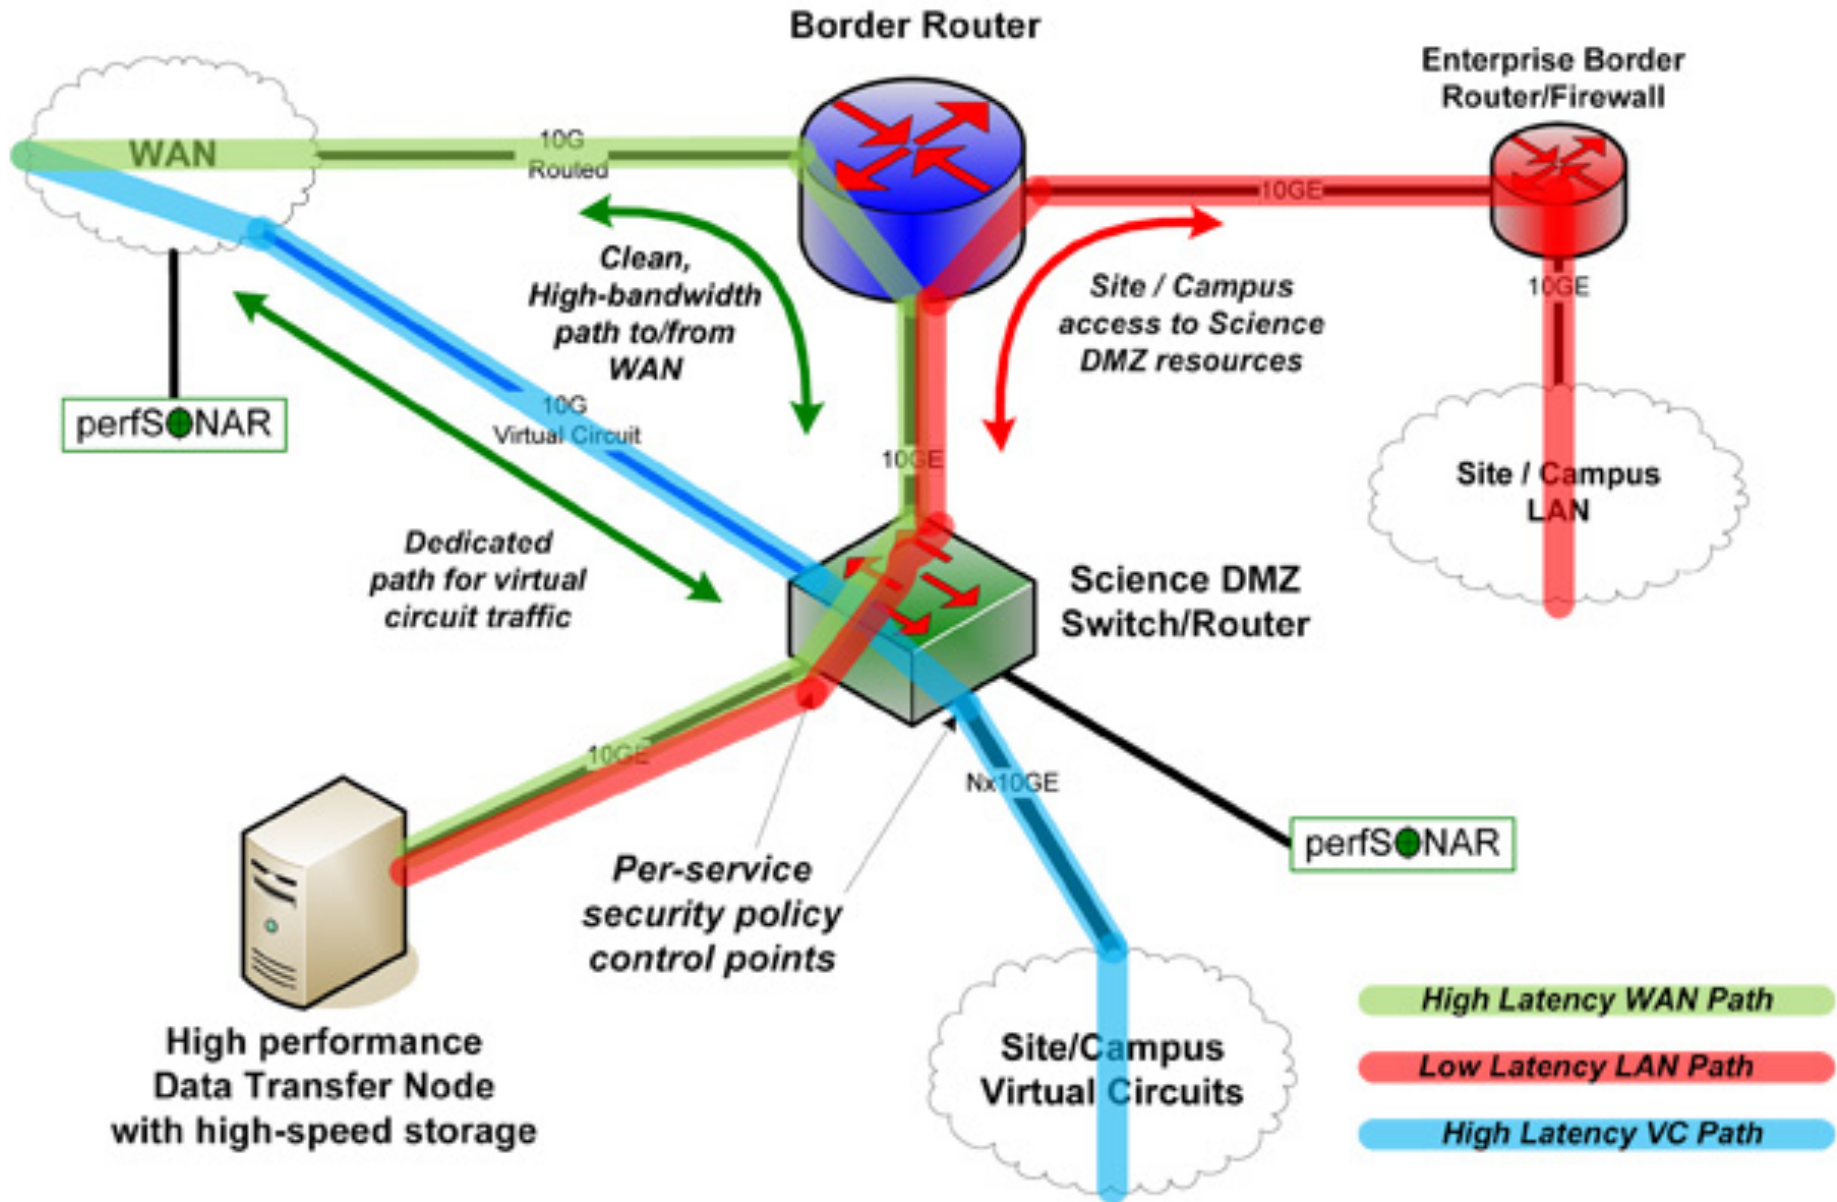
- Data Analysis and Simulation for Earthquakes and Natural Disasters  
*Pacific Earthquake Engineering Research Center (PEER)*
- Climate Modeling  
*National Center for Atmospheric Research (NCAR)*  
*University Corporation for Atmospheric Research (UCAR)*
- California/Nevada Regional Climate Data Analysis  
*California Nevada Climate Applications Program (CNAP)*
- CO2 Subsurface Modeling

## Scalable Visualization, Virtual Reality, and Ultra-Resolution Video

*Cultural Heritage Data*  
*Networked Scalable Visualization*  
*Virtual Reality Systems*  
*Ultra-Resolution Video Systems*

**Note: different classes of users needing different CI properties, ideally under their own control!**

# Science-DMZ



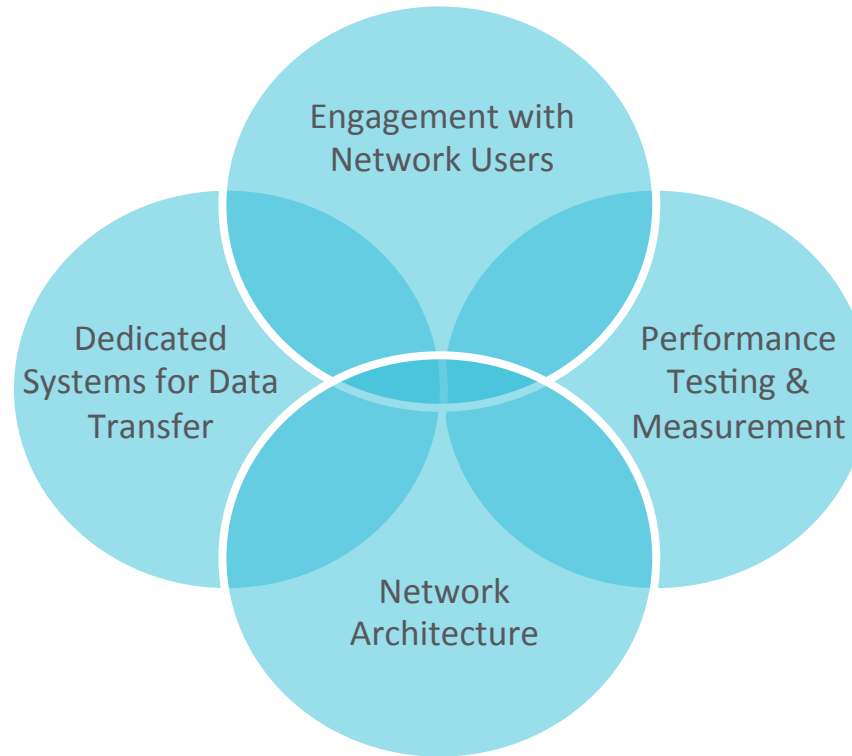
# Science DMZ Superfecta: Engagement

## 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



# PRPv0 - An experiment including:

Caltech

CENIC / Pacific Wave

ESnet / LBNL

NASA Ames / NREN

San Diego State University

SDSC

Stanford University

University of Washington

USC

UC Berkeley

UC Davis

UC Irvine

UC Los Angeles

UC Riverside

UC San Diego

UC Santa Cruz



# Science DMZ Data Transfer Nodes - Optical Network Termination Devices: Linux PCs Optimized for Big Data

- **FIONA – Flash I/O Network Appliance**

- UCOP Rack-mount build: US\$7.7K (Blue), \$21K (Red)
- Intel Xeon E5-1650 v3 Haswell, 2xIntel 2.60GHz Xeon E5-2697 v3 14-Core
- 1TB RAM (expandable to 16TB)
- Flash Drives: 4TB (up to 16TB)
- RAID Drives 0 to 112TB
- NVIDIA Tesla K80 24GB GPU
- 10GbE/40GbE Adapter (100GbE)
- Tested speed 37Gbs disk-to-disk
- UCSD CC-NIE Prism Award & UCOP
  - Phil Papadopoulos & Tom DeFanti
  - Joe Keefe & John Graham



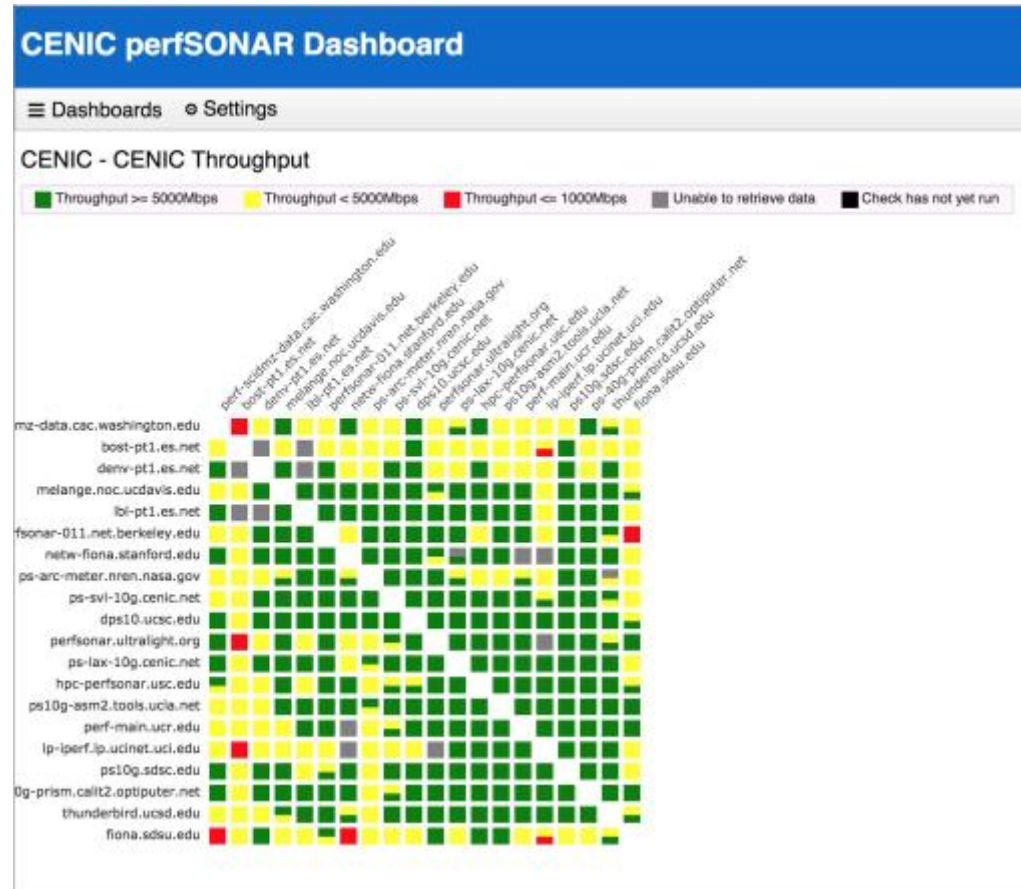
# What have we done

PRPv0 concentrated on the regional aspects of the problem. There are lots of parts to the research data movement challenge. This experiment mostly looked at the inter-campus piece . Over a 10-week period, lots of network and HPC staff at lots of sites collaborated to

- Build a mesh of perfSONAR instances to instrument the network
- Implement MaDDash -- Measurement and Debugging Dashboard
- Deploy Data Transfer Nodes (DTNs)
- Perform GridFTP file transfers to quantify throughput
- Activate an ad-hoc, partial BGP peering mesh across a fabric of 100G links to demonstrate the potential of networks with burst capacity greater than that of a single DTN
- Identify some specific optimizations needed
- Fix a few problems in pursuit of gathering illustrative data
- Identify anomalies for further investigation

# MaDDash of perfSONAR throughput and loss

- Performance for nodes that are close is better than for nodes that are far away
- Network problems that manifest over a distance may not manifest locally



# John Graham's Network Results

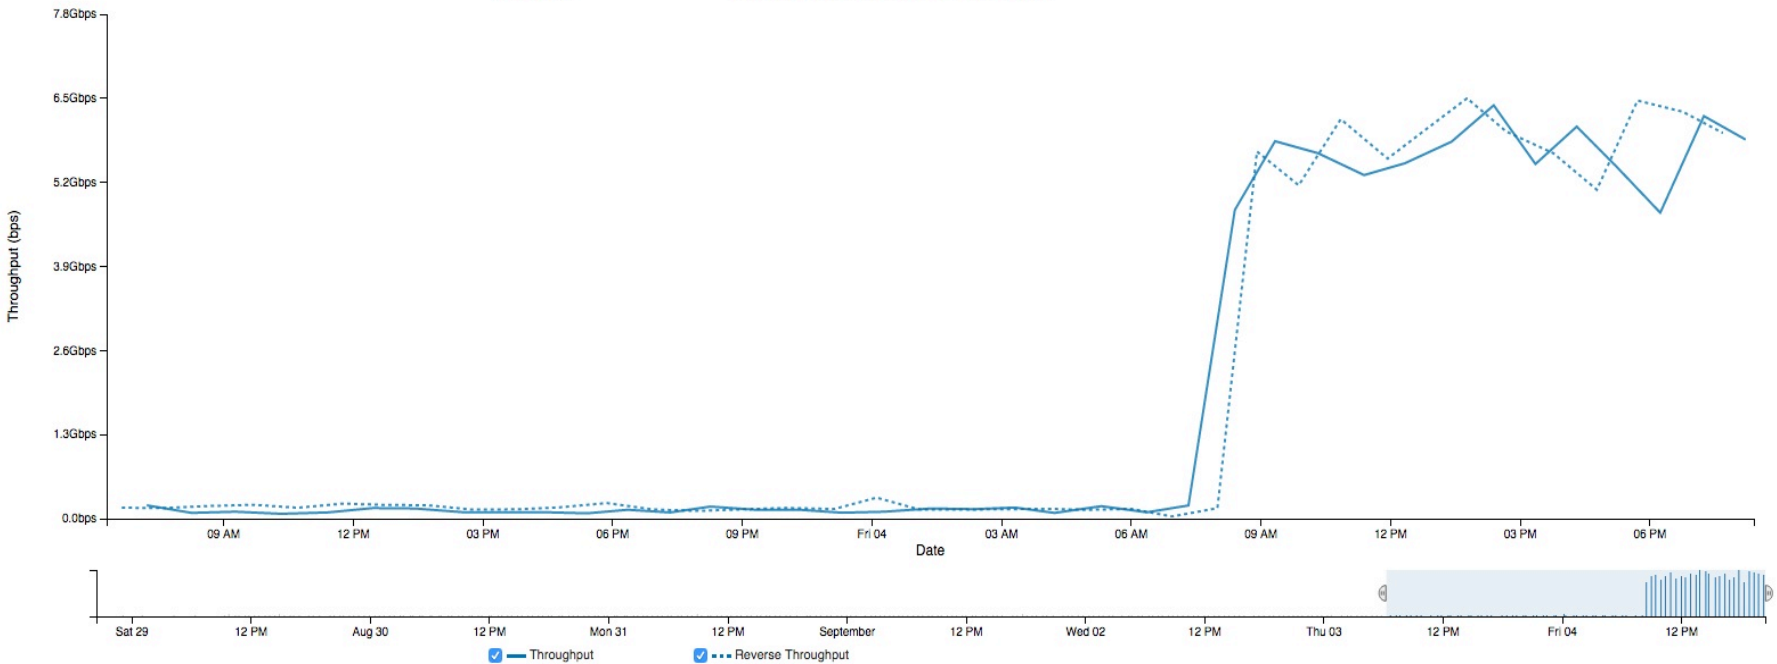
## Moving the CineGrid Exchange 30TB

Source: 10.19.21.50 - 10.19.21.50  
Capacity: Unknown MTU: Unknown

Destination: 10.19.21.51 - 10.19.21.51  
Capacity: Unknown MTU: Unknown

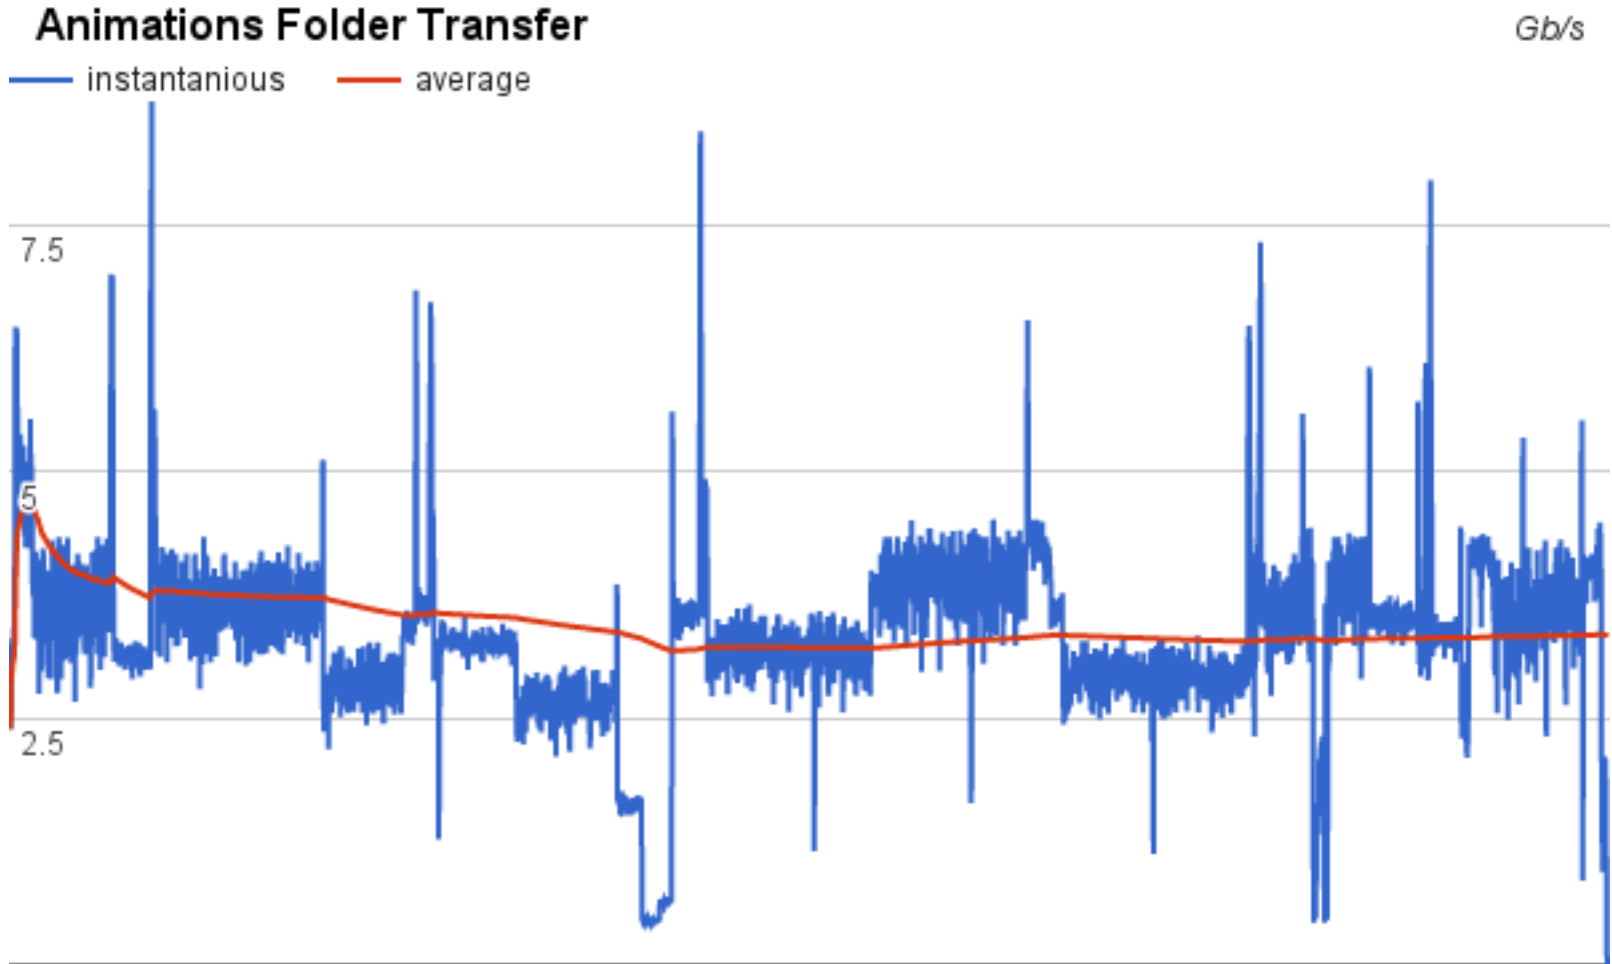
[Link to this chart](#)

Zoom: 1d 3d 1w 1m 1y  
Previous 1w  
Fri Aug 28 20:25:26 2015 -- Fri Sep 4 20:25:26 2015



# UCSD < -- > UvA

Iperf3 mem to mem : 32 Gbps



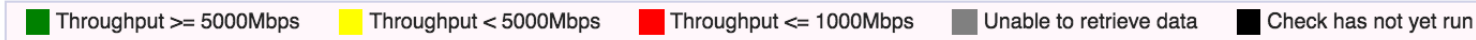
Limited by many 25 Mbyte 4k frame files, file system, ZFS, sata interfaces, etc.

# Data Transfer Testing SSD to SSD (Using Caltech's Fast Data Transfer FDT)

streams	from	to	GB	seconds	Mb/s	method
8	16 x 120GB SSD in a 2 vdev raidz1 pool at UCSD 16 x 120GB SSD in a 2 vdev raidz1 pool at UCLA	8x 800GB PCIe NVME	1600	240	53,333	FDT
	16 x 120GB SSD in a 2 vdev raidz1 pool at Stanford	/dev/null	100	54	14,834	GFTP
1	16 x 120GB SSD in a 2 vdev raidz1 pool at Stanford	16 x 4T SATA in a 2 vdev raidz1 pool with 16GB ZIL 960GB ARC on ssd	200	120	13,333	FDT
2	16 x 120GB SSD in a 2 vdev raidz1 pool at Stanford	16 x 4T SATA in a 2 vdev raidz1 pool with 16GB ZIL 960GB ARC on ssd	400	240	13,333	FDT
1	16 x 120GB SSD in a 2 vdev raidz1 pool at Stanford	16 x 4T SATA in a 2 vdev raidz1 pool with 16GB ZIL 960GB ARC on ssd	100	66	12,061	GFTP
4	4 x 800GB PCIe NVME at Caltech	16 x 1T SATA in a 2 vdev raidz1 pool with 16GB ZIL 960GB ARC on ssd	800	1080	5,926	FDT
1	1 x 800GB PCIe NVME at Caltech	16 x 1T SATA in a 2 vdev raidz1 pool with 16GB ZIL 960GB ARC on ssd	200	300	5,333	FDT
	ESNET bnl-diskpt1.es.net	16 x 4T SATA in a 2 vdev raidz1 pool with 16GB ZIL 960GB ARC on ssd	1	8	1,039	WGET
	<a href="https://cloud.sdsc.edu/auth/v1.0">https://cloud.sdsc.edu/auth/v1.0</a>	UCSD server with 1 SSD	10.74	113	763	swift.py
	USGS edcftp.cr.usgs.gov	16 x 4T SATA in a 2 vdev raidz1 pool with 16GB ZIL 960GB ARC on ssd	11.56	244	379	LFTP

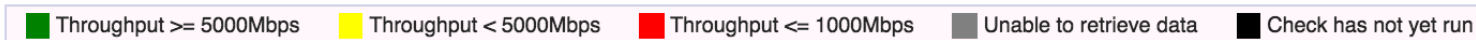
# CENIC Dashboard

## CENIC - CENIC Throughput



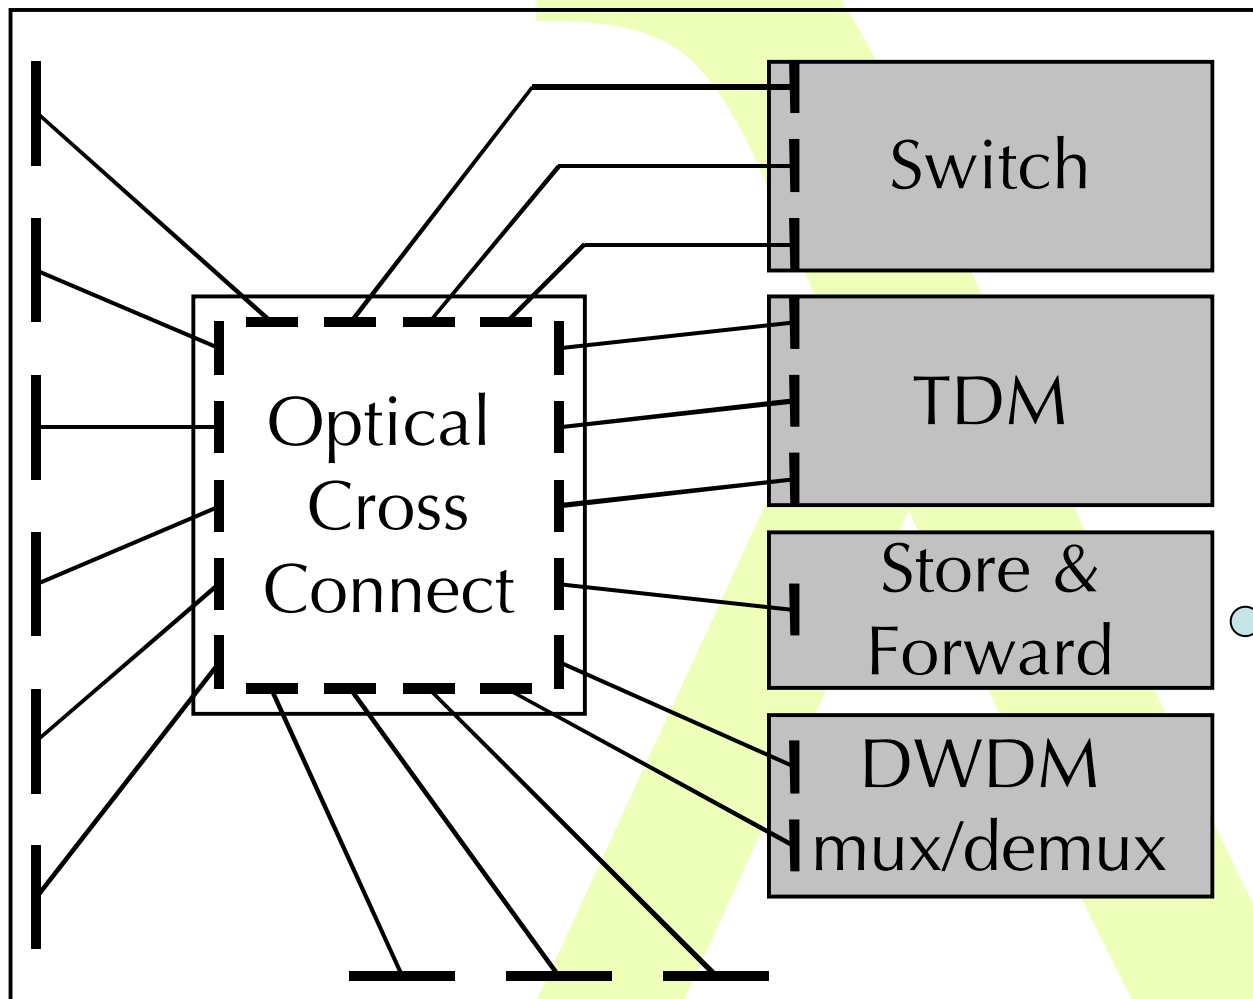
Next Gen VR Booth

## CENIC - CENIC IPv6 Throughput



# Optical Exchange as Black Box


## Optical Exchange



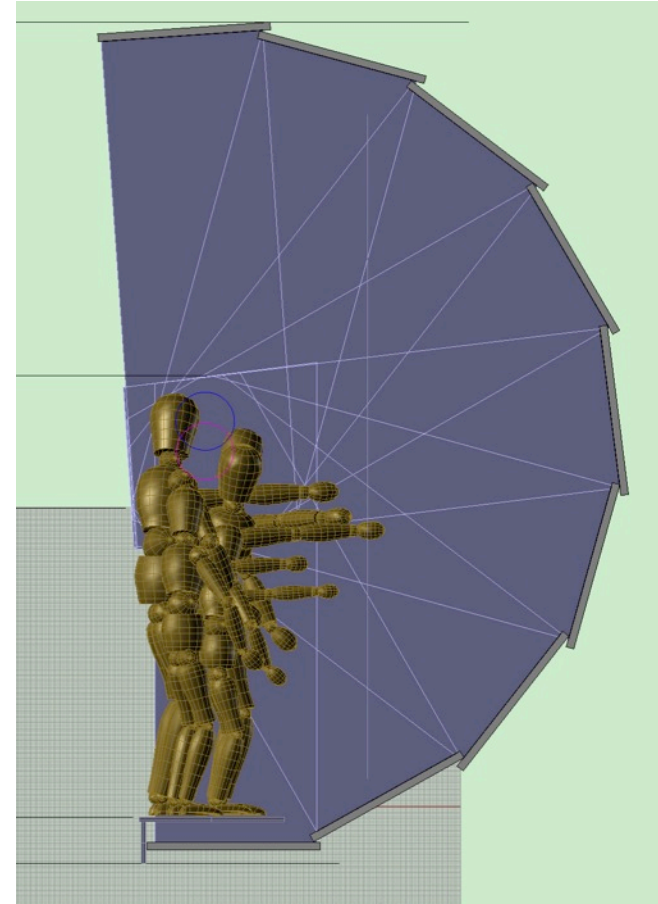
TeraByte  
Email  
Service



# WAVE Cluster—20x40G PRP-connected



# The WAVE—2013 ~\$400,000 35 Megapixels/eye \$0.01/pixel



WAVE VR display 7 high by 5 wide HD panels built for the SME Building, UCSD  
125,000 Cores (50 nvidia 780s), 200TF (single precision), 1 Terabit Network

# Pacific Research Platform Funded UCSD/UCB \$1M/yr 5 years + CENIC \$\$

## Abstract

The Pacific Research Platform is a project to forward the work of advanced 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

US NSF ACI-1540112 and ACI-1541349

# Pacific Research Platform Strategic Arc

---

Build upon Pacific Wave as a backplane for data-intensive science

- High performance data movement provides capabilities that are otherwise unavailable to scientists
- Integrating Science DMZs across the West Coast
- This capability is extensible, both regionally and nationally
- SDSC's Andrea Zonca is adapting CILogon for secure access

Goal: scientists can get the data they need, where they need it, when they need it

- PRPv0: a proof of concept experiment to develop and inform requirements for future work.
- Engage with scientists to map their research on to the Pacific Research Platform

# Next Steps and Near-term Goals

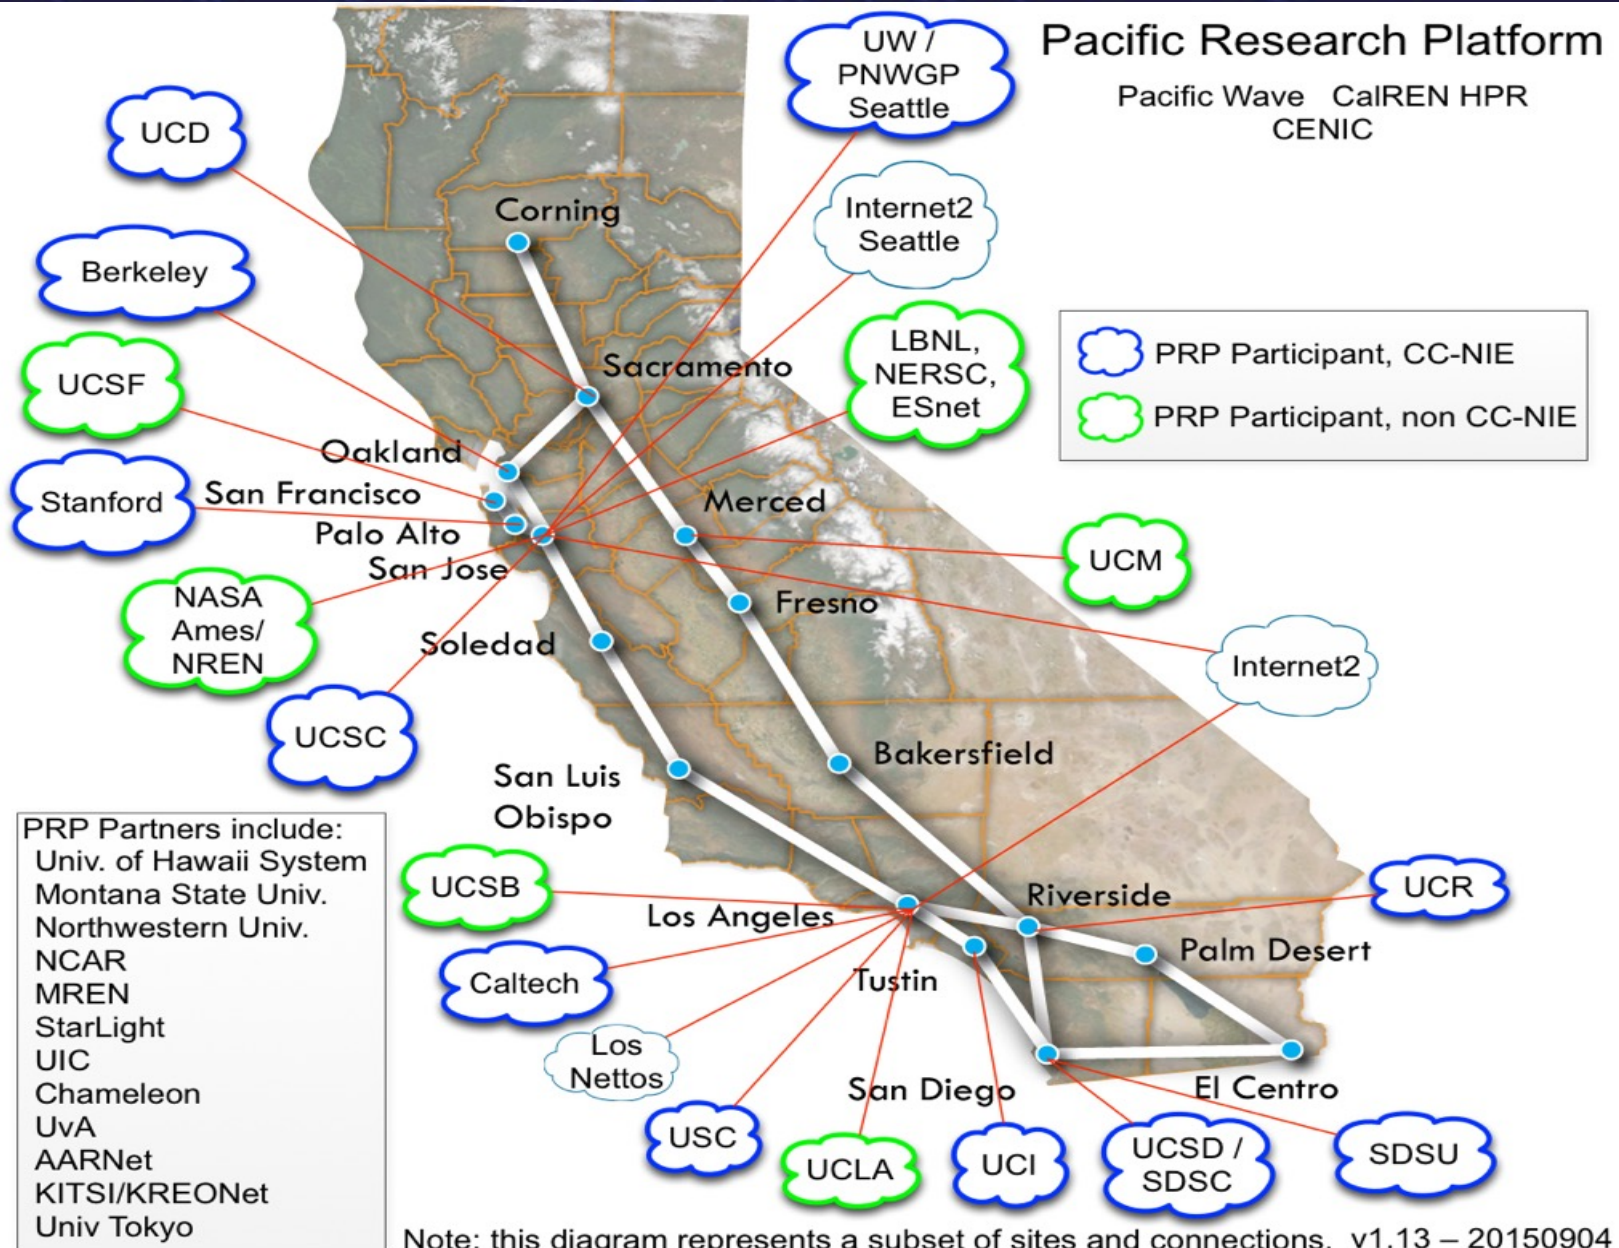
- Migrate from experiment to persistent infrastructure as part CalREN HPR 100G Layer 3 upgrade; becomes PRPv1
- Sending out the \$7,700 FIONA/DTNs & \$2,300 PerfSONARs to all 10 UC campuses in September with UC funds; more to non-UC partners in October after NSF contract starts
  - Working with CENIC to identify the people and the connections
  - Upgrading the early FIONAs from March to the new PRPv1 standard
  - Collecting base-line tests State-wide, to Seattle, Chicago, and Amsterdam
  - Identifying and installing security mechanisms
- The PRP will be rolled out in two phases.
  - First, the PRPv1 platform will focus on deploying/maximizing data-sharing to include all member campuses and their science apps
  - Second, the PRPv2 will be rolled out as an advanced, IPv6-based version with robust security and software-defined networking (SDN) features
- Work with campus IT organizations to make “last mile” connections between researchers and the Science DMZ
- PRP Workshop, October 14-16; 60 people signed up/session
- Demos at SC’15, tbd.

# Longer-term Goals

---

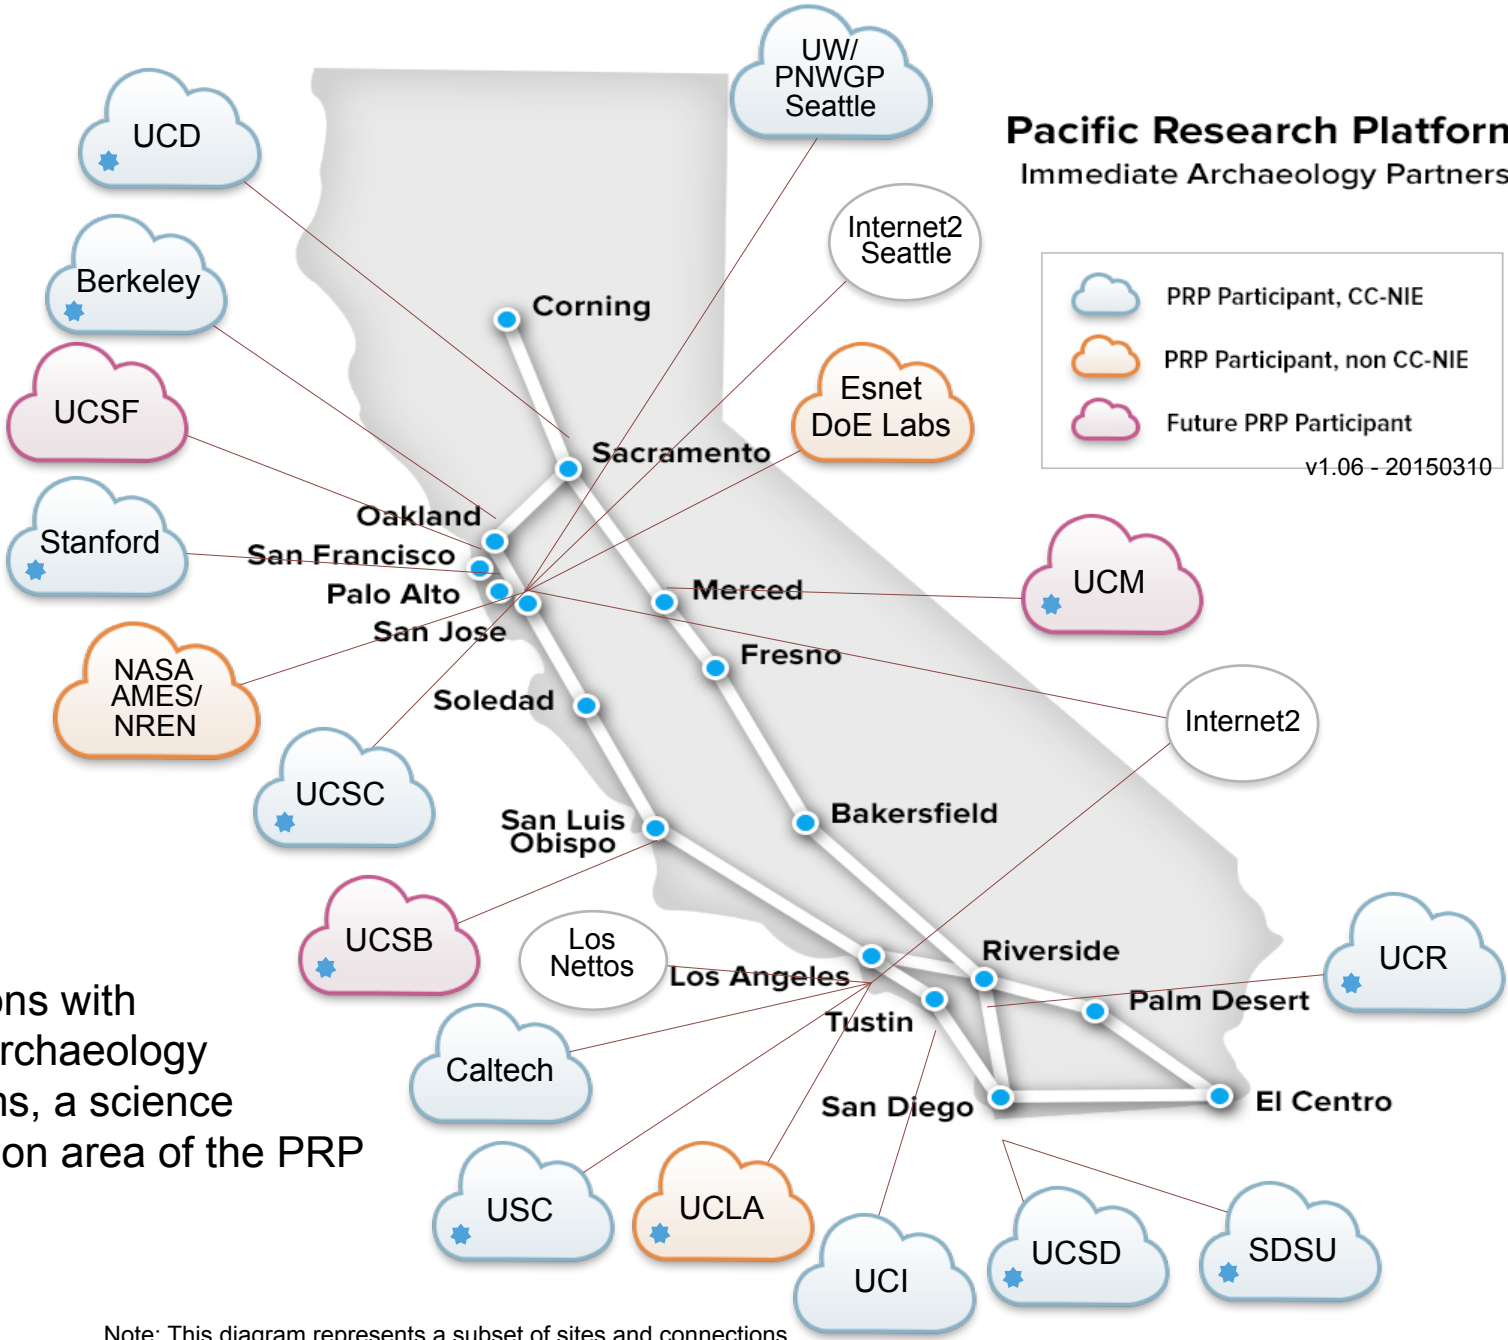
- An Integrated West Coast Science DMZ for Data-Intensive Research
- Science DMZ interoperability / integration across regions, nationally, and internationally
- SDN/SDX, ...
- Commercial services – Amazon AWS, Microsoft Azure, Google, ...

# Pacific Research Platform: A Regional Science DMZ



# Pacific Research Platform

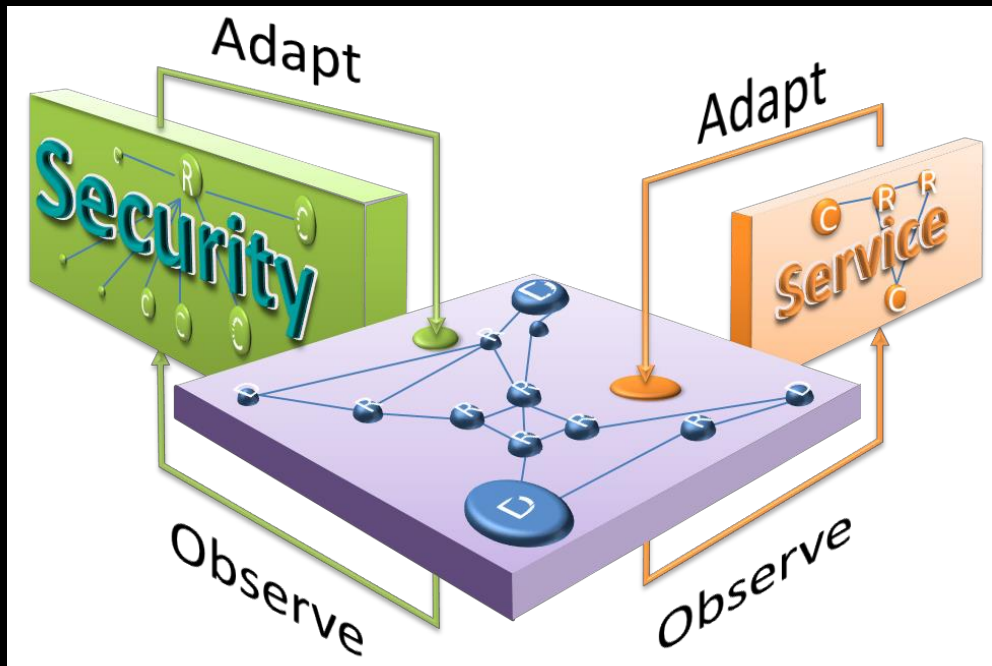
## Immediate Archaeology Partners



v1.06 - 20150310

★ Institutions with Active Archaeology Programs, a science application area of the PRP





# SARNET

Security Autonomous Response  
with programmable NETWORKS

Cyber Security program

PI: CdL

Co-Pi's: RM, LG, RW

- 400 + 285 + 300 kEuro:
- 3 PhD's
- Prog & Eng manpower

- Network virtualizations and SDN
- Reasoning
- Risk evaluation
- Trust groups
- Execute response & adaptation




[delaat.net/sarnet](http://delaat.net/sarnet)

# Service Provider Group framework

*A Service Provider Group (SPG) is an organisation structure providing a defined service only available if its members collaborate.*

*Examples:*

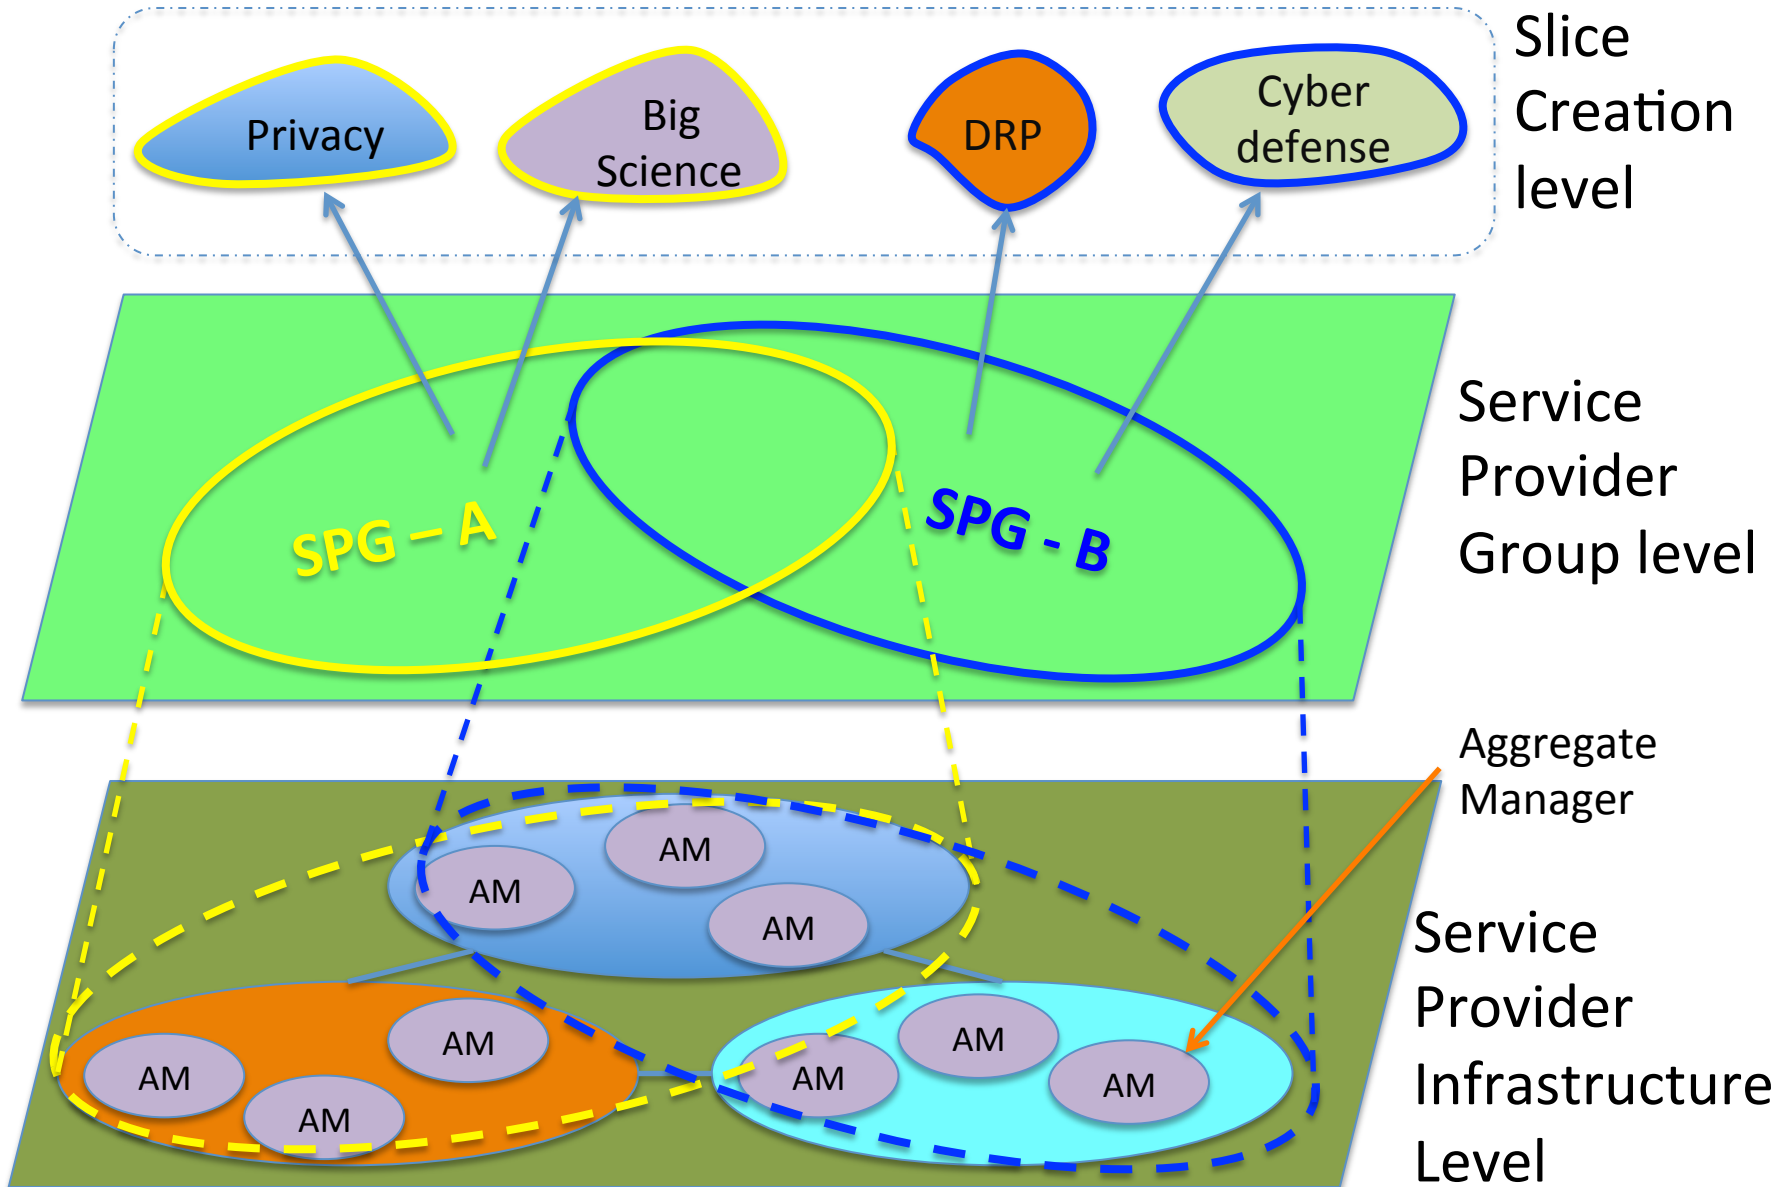
Internet2NET+



# Service Provider Group Characteristics

- **Autonomous members** acting together on a decision to provide a service none could provide on its own
- Appears as **a single provider** to a customer
- Appears as **a collaborative group** to members with standards, rules and policies that are defined, administered, enforced and judged by the group.
- Autonomy in the group: every member signs an agreement **declaring compliance** with common rules, unless local law determines otherwise.
- Membership rules **organizes trust** amongst members and manage group reputation and viability.

# Envisioned role of the SPG: define slice archetypes?



# Research Questions

- SARNET:
  - Is a cyber security alliance, allowing networks to join/leave freely, feasible?
  - What is needed to organize an alliance, considering the SPG concept?
- Considering future networking concepts:
  - Is a SPG a concept that should identify and arrange slice archetypes e.g. defining cyber-security assurance levels
  - What concerns should the SPG address (e.g. economical-, legal-, administrative-, etc. slice ownerships)?

Collaboration welcomed: [delaat@uva.nl](mailto:delaat@uva.nl)

# Links

- ESnet fasterdata knowledge base
  - <http://fasterdata.es.net/>
- Science DMZ paper
  - [http://www.es.net/assets/pubs\\_presos/sc13sciDMZ-final.pdf](http://www.es.net/assets/pubs_presos/sc13sciDMZ-final.pdf)
- Science DMZ email list
  - <https://gab.es.net/mailman/listinfo/sciencedmz>
- perfSONAR
  - <http://fasterdata.es.net/performance-testing/perfsonar/>
  - <http://www.perfsonar.net>
- Sample Campus & Regional Cyberinfrastructure Plans
  - <http://fasterdata.es.net/campusCiplanning/>

# Pacific Research Platform

---

Questions?