"Global Research Platforms: Past, Present, Future"

1st Global Research Platform Workshop
Calit2's Qualcomm Institute
UC San Diego
September 17, 2019

Dr. Larry Smarr
Director, California Institute for Telecommunications and Information Technology
Harry E. Gruber Professor,
Dept. of Computer Science and Engineering
Jacobs School of Engineering, UCSD
http://lsmarr.calit2.net

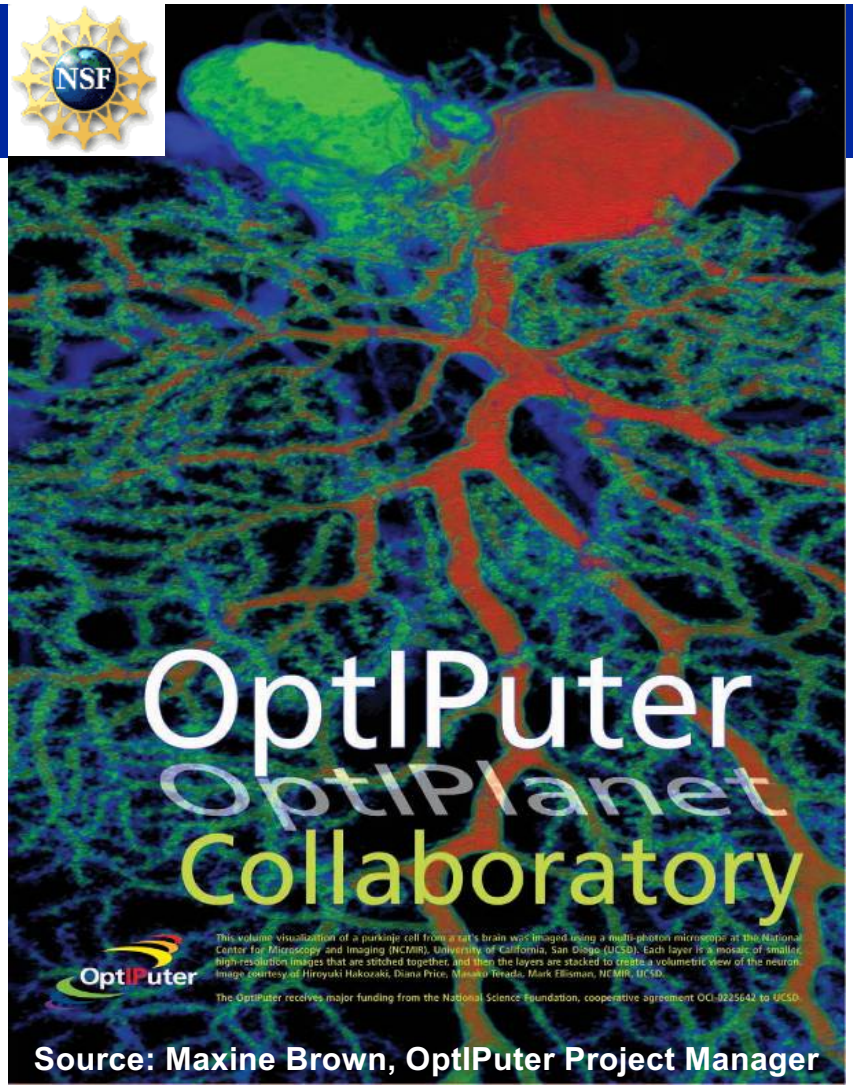

Source: Maxine Brown, OptIPuter Project Manager

The OptIPuter
Exploits a New World
in Which
the Central Architectural Element
is Optical Networking,
Not Computers.

Distributed Cyberinfrastructure
to Support
Data-Intensive Scientific Research
and Collaboration

PI Smarr,
2002-2009

PACIFIC RESEARCH PLATFORM

First Global Telepresence Meeting
Using Digital Cinema 4k Streams Over 1 Gbps Optical Fiber

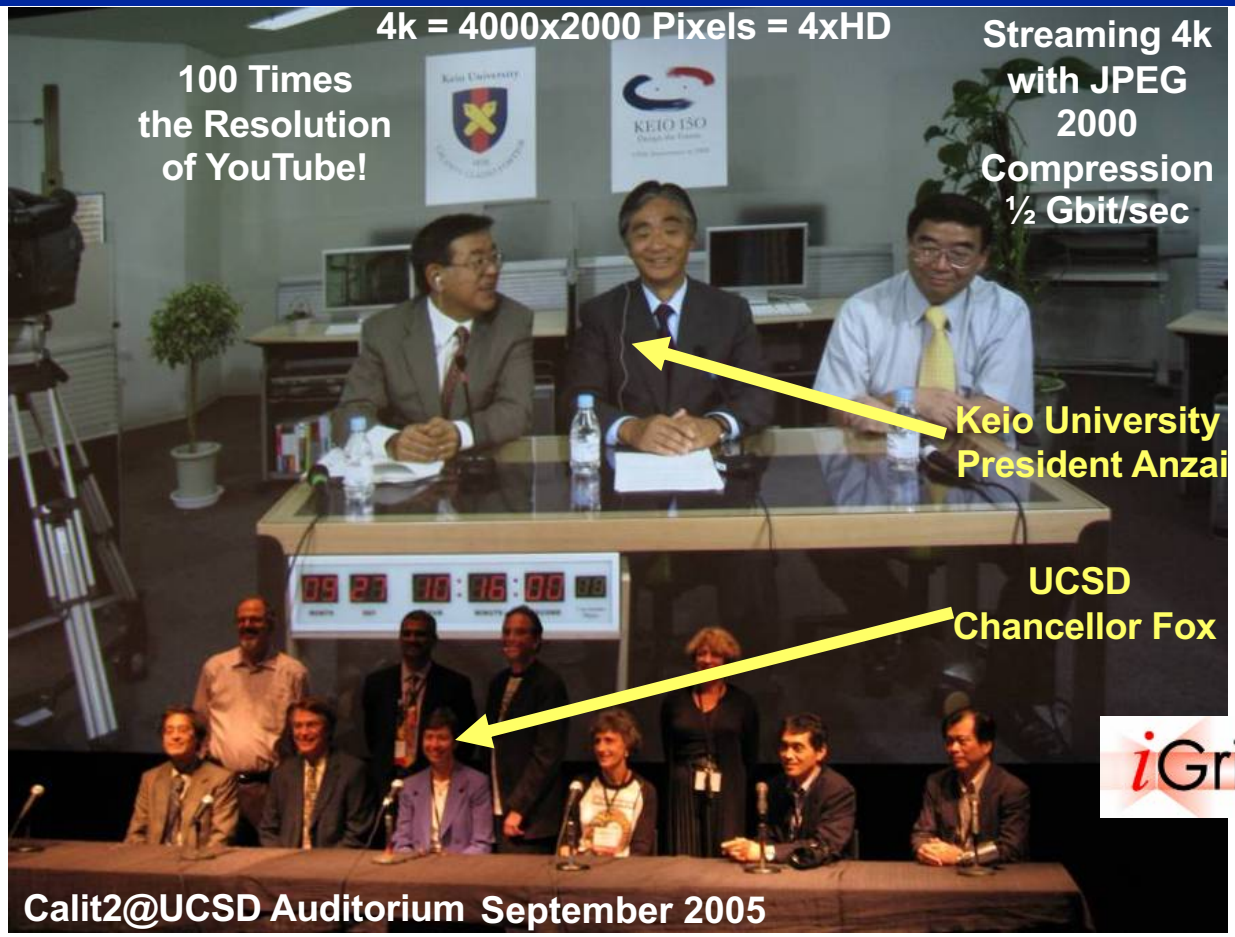

CineGrid

Lays Technical Basis for Global Digital Cinema

Sony
NTT
SGI

iGrid 2005

PRP PACIFIC RESEARCH PLATFORM

it²

Integrated "OptIPlatform" Cyberinfrastructure System: A 10Gbps Lightpath Cloud

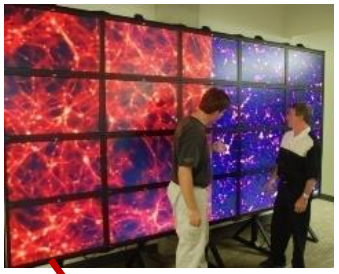

End User OptIPortal

Osaka University

Instruments

HD/4k Telepresence

HD/4k Video Cams

HPC

TeraGrid

LS 2009 Slide

Campus Optical Switch

10G Lightpath

Data Repositories & Clusters

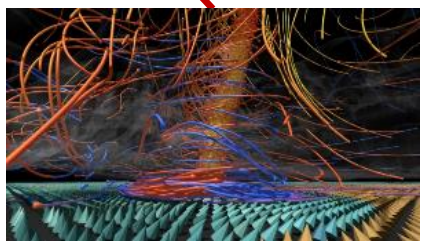

HD/4k Video Images

PRP PACIFIC RESEARCH PLATFORM

OptIPuter

SDSC

it²

Launch of the 100 Megapixel OzIPortal Kicked Off a Rapid Build Out of Australian OptIPortals

No Calit2 Person Physically Flew to Australia to Bring This Up!

Covise, Phil Weber, Jurgen Schulze, Calit2
CGLX, Kai-Uwe Doerr , Calit2
http://www.calit2.net/newsroom/release.php?id=1421

OptIPortals
Were Adopted Globally

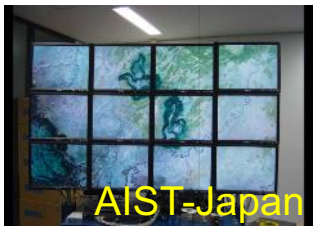

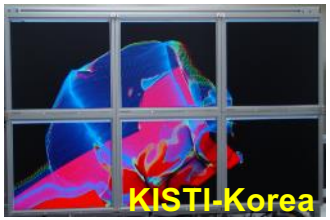

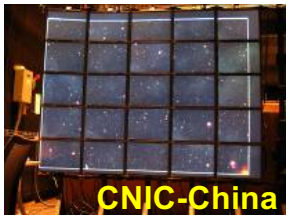

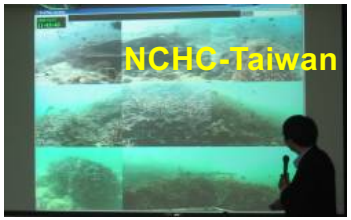

PRP
PACIFIC RESEARCH PLATFORM

Powered by Rocks ROCKS

evl electronic visualization laboratory

it²

We Have Been Working Towards Distributed Big Data for 15 Years: NSF OptIPuter, Quartzite, Prism Awards

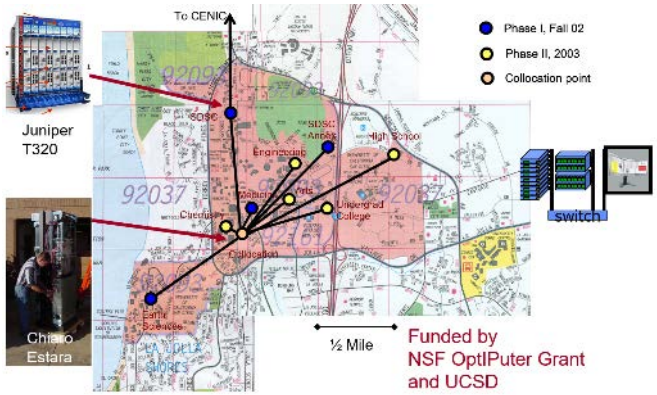

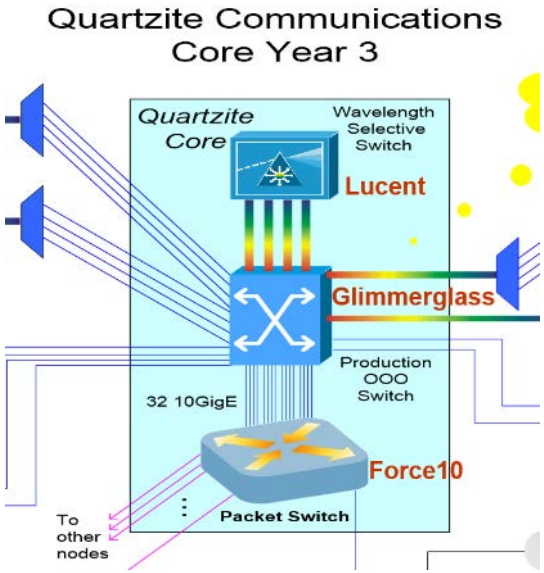

Goals by 2008:
>= 50 endpoints at 10 GigE
>= 32 Packet switched
>= 32 Switched wavelengths
>= 300 Connected endpoints

See Following Talk by UCSD's Valerie Polichar

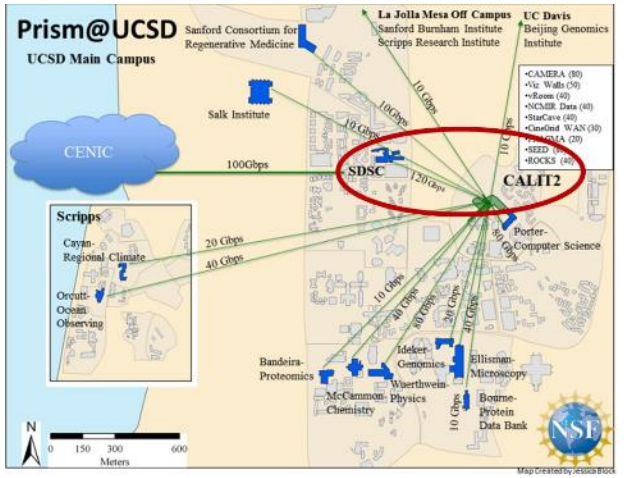

PRP PACIFIC RESEARCH PLATFORM

PI Smarr, 2002-2009

PI Papadopoulos, 2004-2007

PI Papadopoulos, 2013-2015

it²

Before the PRP: ESnet's ScienceDMZ Accelerates Science Research: DOE & NSF Partnering on Science Engagement and Technology Adoption

ScienceDMZ Coined in 2010 by ESnet
Basis of PRP Architecture and Design

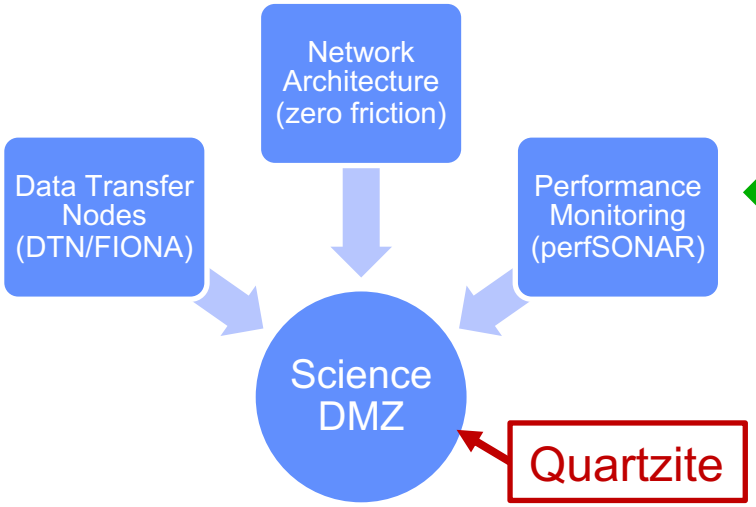

http://fasterdata.es.net/science-dmz/

DOE

NSF

NSF Campus Cyberinfrastructure Program
Has Made Over 250 Awards

Prism

2012 2013 2014 2015 2016 2017 2018

Slide Adapted From Inder Monga, ESnet

PRP PACIFIC RESEARCH PLATFORM

ESnet FASTERDATA KNOWLEDGEBASE

NSF

it²

2015 Vision: The Pacific Research Platform Will Connect Science DMZs Creating a Regional End-to-End Science-Driven Community Cyberinfrastructure

PACIFIC RESEARCH PLATFORM

Source: John Hess, CENIC

NSF CC*DNI Grant
$6.3M 10/2015-10/2020
Year 5 Starts in 3 Weeks!

PI: Larry Smarr, UC San Diego Calit2
Co-PIs:

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

Letters of Commitment from:

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

ESnet: Given Fast Networks, Need DMZs and Fast/Tuned DTNs

SDSC CITRIS BANATAO INSTITUTE it²

PRP Engineers Designed and Built Several Generations of Optical-Fiber Big-Data Flash I/O Network Appliances (FIONAs)

UCSD-Designed FIONAs Solved the Disk-to-Disk Data Transfer Problem *at Near Full Speed* on Best-Effort 10G, 40G and 100G Networks

Two FIONA DTNs at UC Santa Cruz: 40G & 100G
Up to 192 TB Rotating Storage

Add Up to 8 Nvidia GPUs Per 2U FIONA
To Add Machine Learning Capability

PRP PACIFIC RESEARCH PLATFORM

NSF

FIONAs Designed by UCSD's Phil Papadopoulos, John Graham, Joe Keefe, and Tom DeFanti

SDSC

it²

2018/2019: PRP Game Changer!
Using Kubernetes to Orchestrate Containers Across the PRP

CADE METZ BUSINESS 06.10.14 01:15 PM

GOOGLE OPEN SOURCES ITS SECRET WEAPON IN CLOUD COMPUTING

WIRED

"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

"Everything at Google runs in a container."
--Joe Beda,Google

How Kubernetes Conquered 2017 (and is Positioned for 2018)

Dan Meyer
December 25, 2017

sdx central

Kubernetes, an open-source software project that started at Google, has exploded in popularity and is now used by at least 54% of the Fortune 500. BUSINESS INSIDER **Rosalie Chan** Jan. 27, 2019,

PRP PACIFIC RESEARCH PLATFORM

CLOUD NATIVE COMPUTING FOUNDATION

it²

PRP Has Adopted Rook Cloud-Native Storage Orchestrator, Which Runs 'Inside' Kubernetes

Battle-tested, production storage

Rook is based on an embedded version of Ceph, which has 10+ years of production deployments and runs some of the worlds largest clusters.

Cloud-native environment integration

Rook runs as a cloud-native service for optimal integration with applications in need of block, object, or file storage.

PRP PACIFIC RESEARCH PLATFORM

Source: John Graham, Calit2/QI

it²

PRP's Nautilus Hypercluster Adopted Kubernetes to Orchestrate Software Containers and Manage Distributed Storage

Nautilus

kubernetes

Kubernetes (K8s) is an open-source system for automating deployment, scaling, and management of containerized applications.

Planet Scale

Designed on the same principles that allows Google to run billions of containers a week, Kubernetes can scale without increasing your ops team.

"Kubernetes with Rook/Ceph Allows Us to Manage Petabytes of Distributed Storage and GPUs for Data Science, While We Measure and Monitor Network Use."
--John Graham, Calit2/QI UC San Diego

PRP
PACIFIC RESEARCH PLATFORM

it²

PRP's California Nautilus Hypercluster Connected by Use of CENIC 100G Network

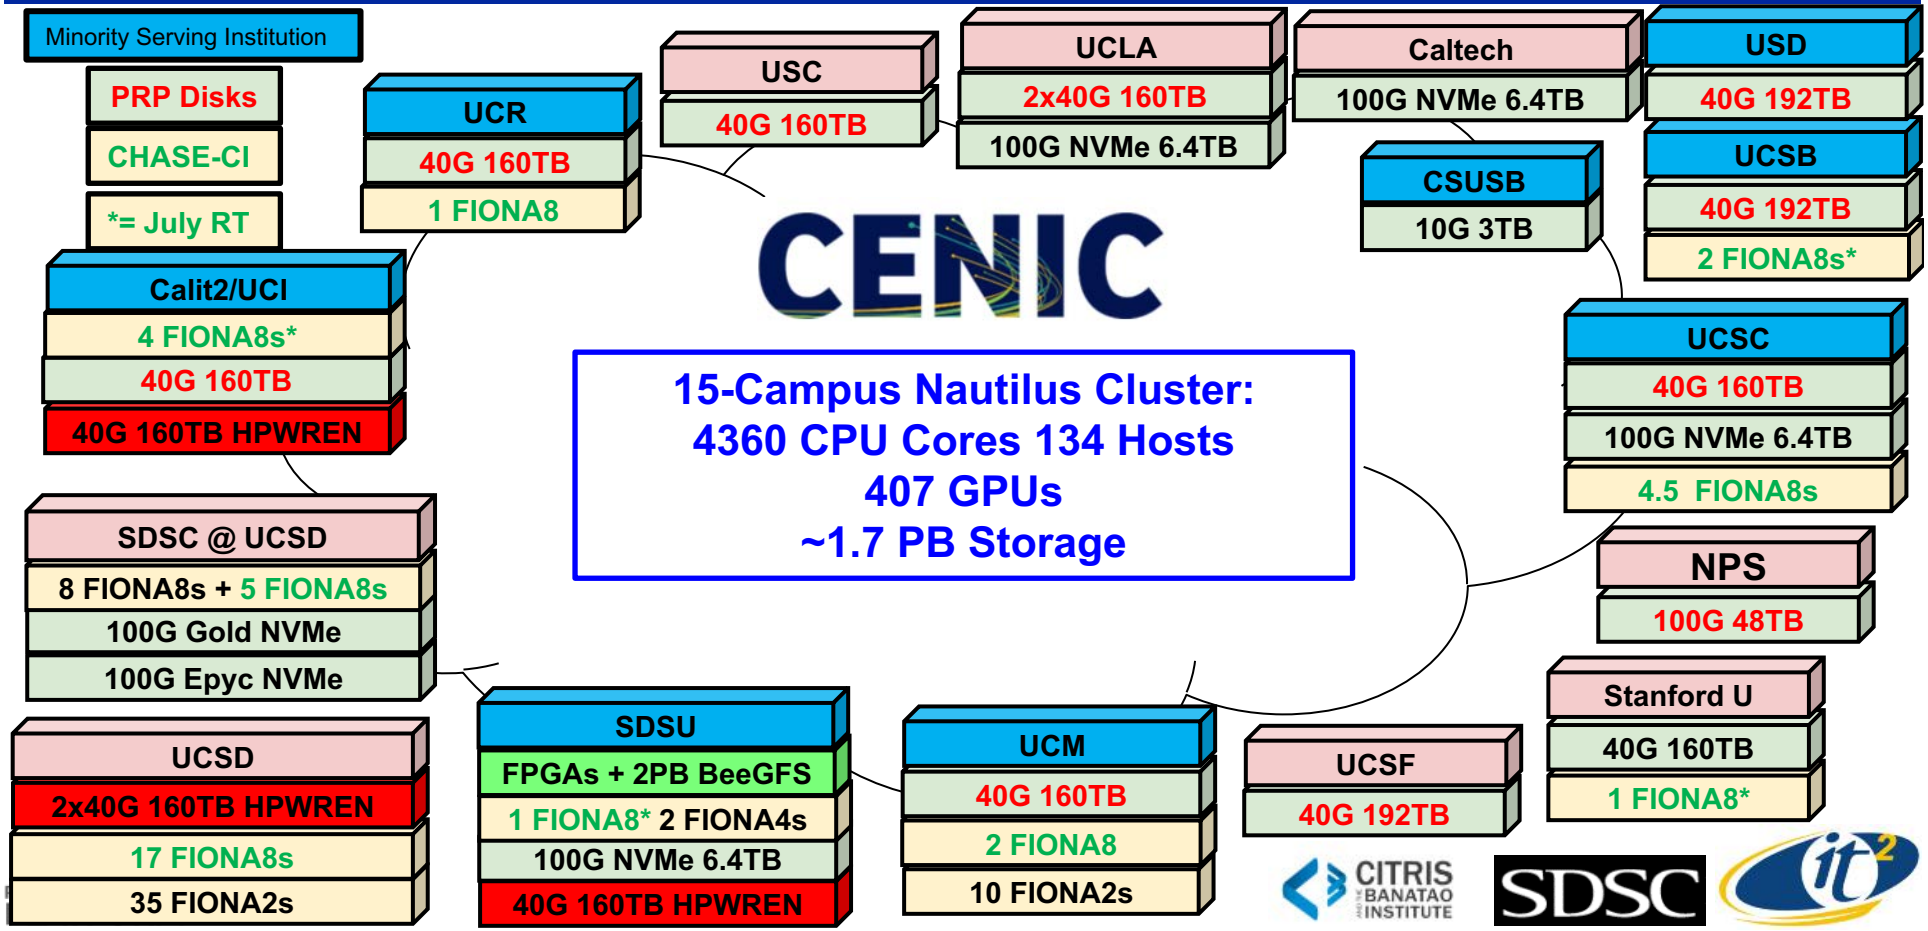

PRP's United States Nautilus Hypercluster FIONAs
We Now Connect 3 More Regionals and 3 Internet2 sites

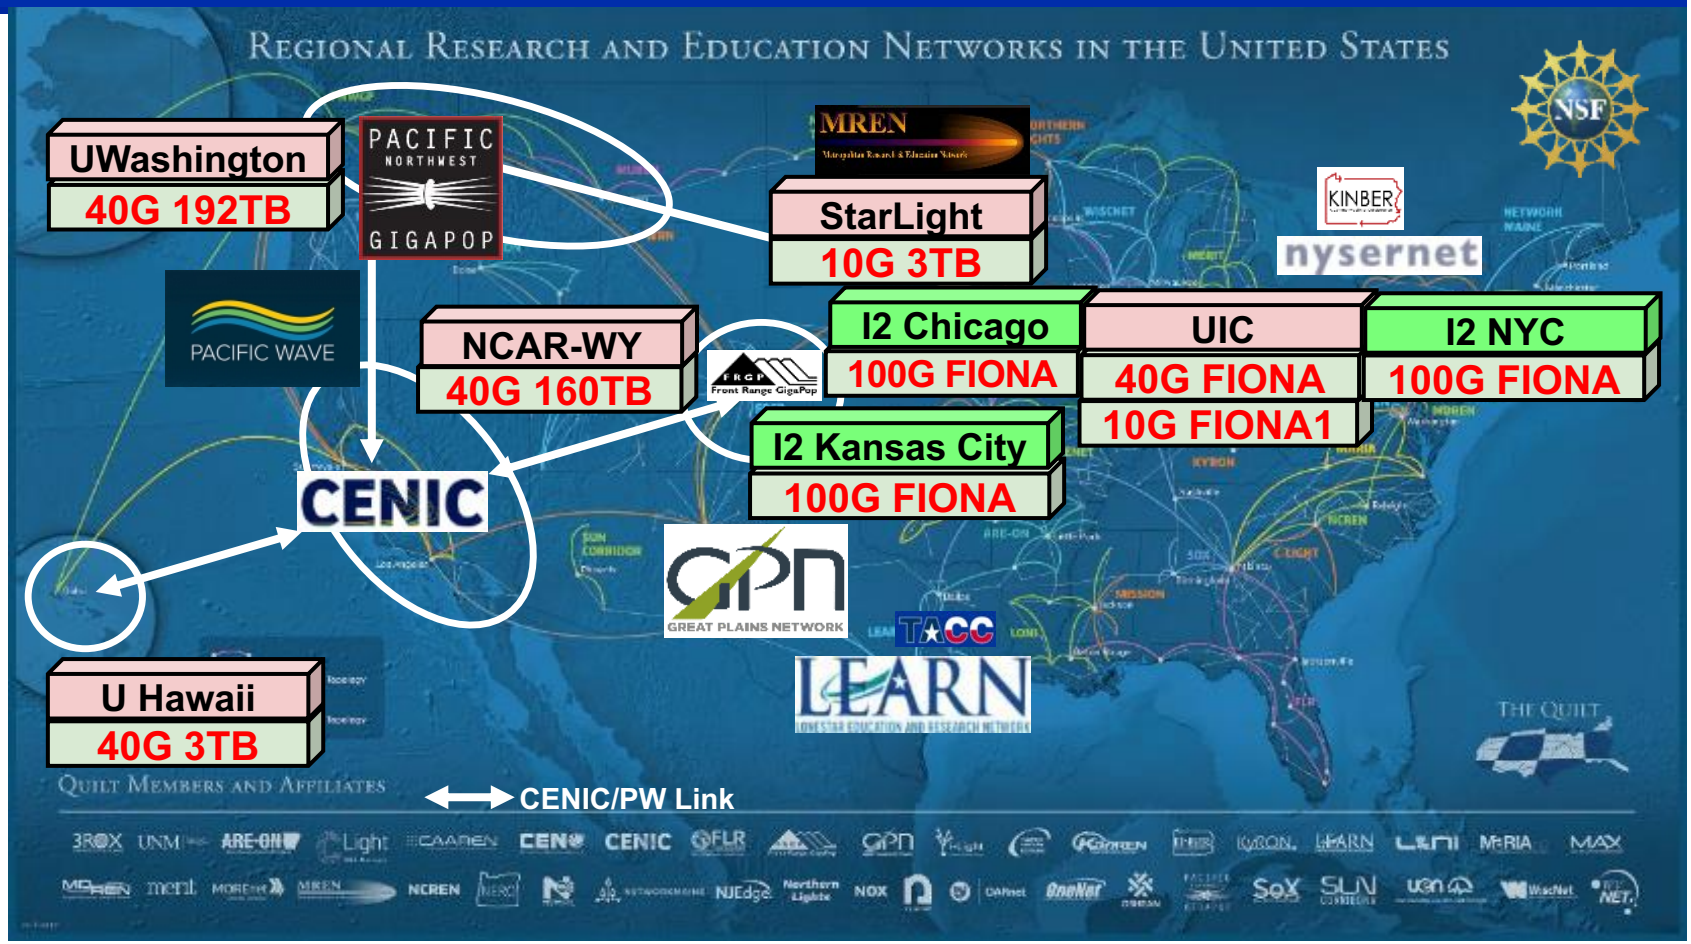

PACIFIC RESEARCH PLATFORM

Global PRP Nautilus Hypercluster Is Rapidly Adding International Partners Beyond Our Original PRP Partner in Amsterdam

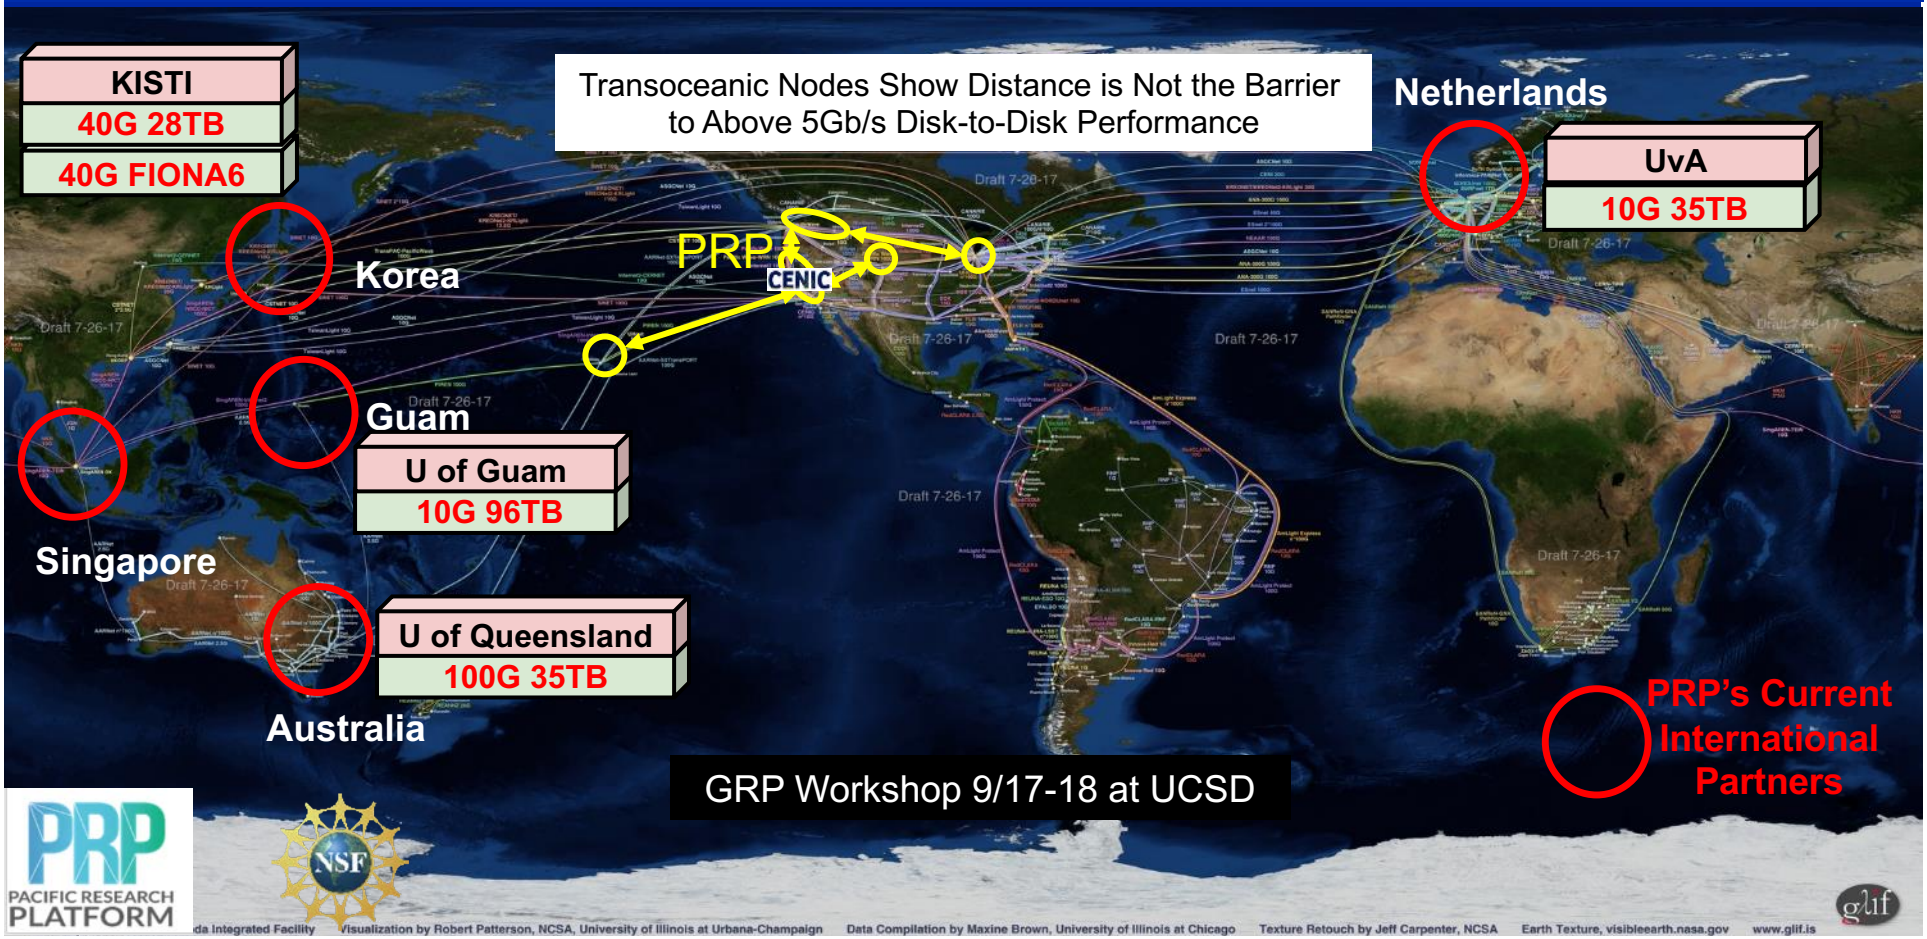

Operational Metrics: Containerized Trace Route Tool Allows Realtime Visualization of Status of PRP Network Links on a National and Global Scale

9/16/2019

Source: Dima Mishin, SDSC

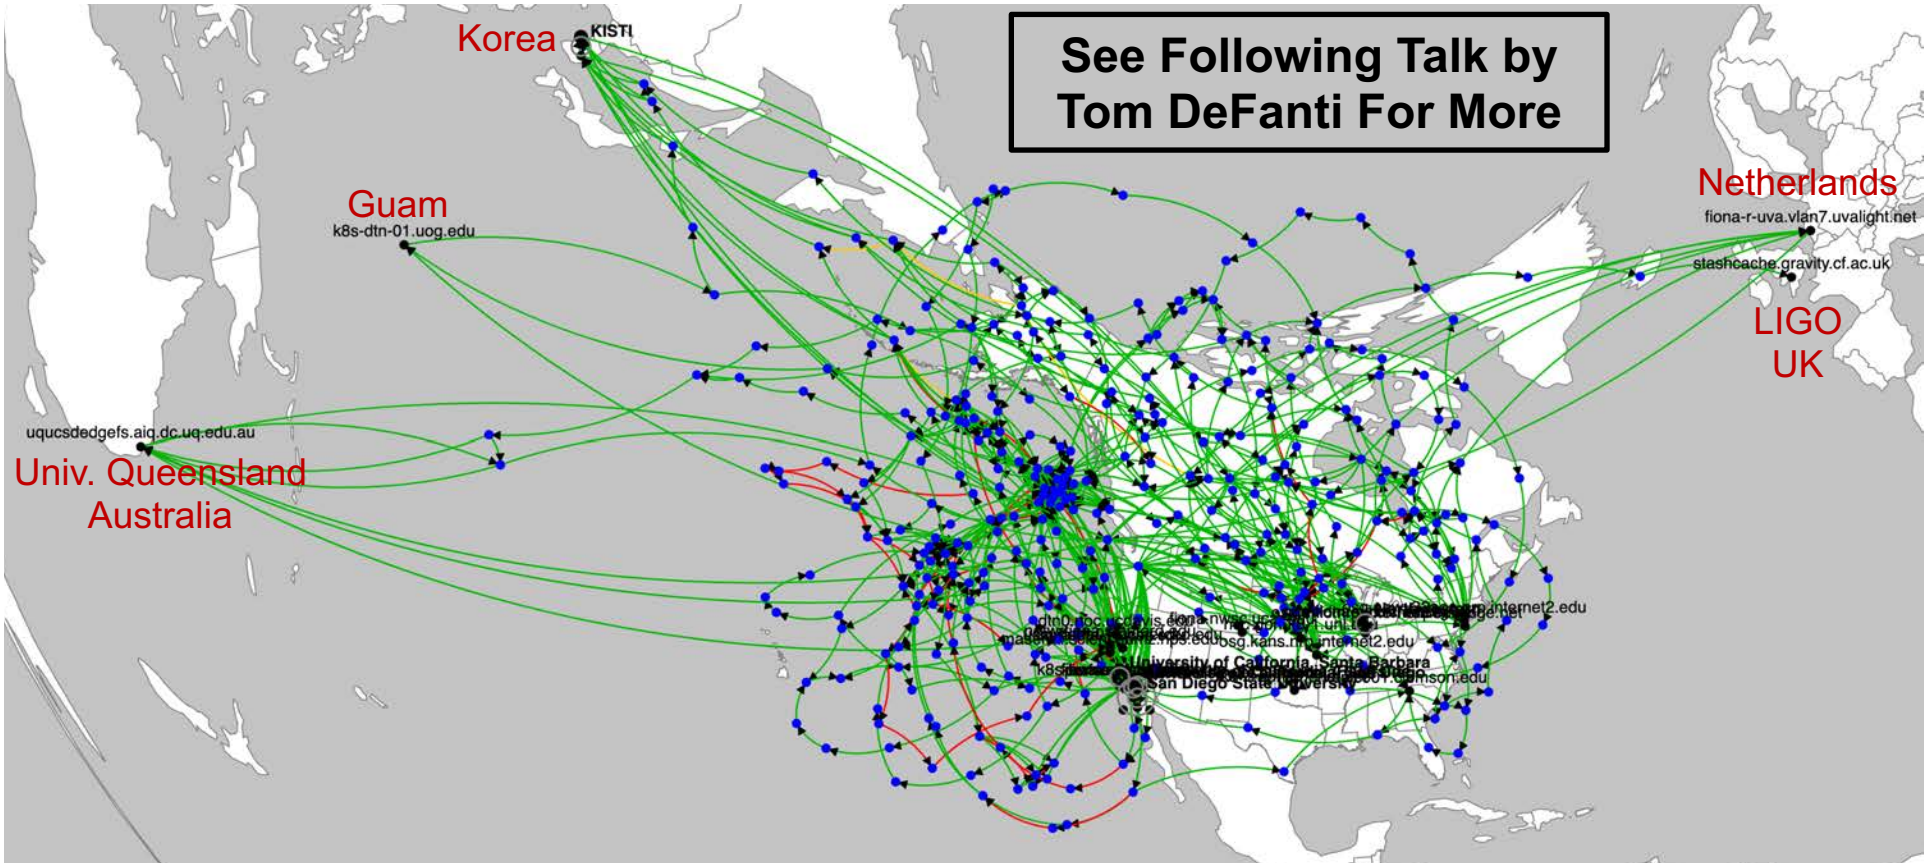

PRP is Science-Driven:
Connecting Multi-Campus Application Teams and Devices

Particle Physics

CMS

Biomedical 'omics

NSF

LIGO

IceCube
South Pole Neutrino Observatory

Telescope Surveys

Earth Sciences

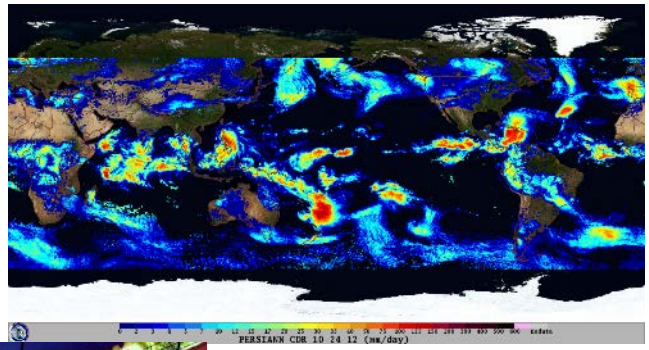

Visualization, Virtual Reality, Collaboration

CITRIS
SDSC

Scott Sellars Rapid 4D Object Segmentation of NASA Water Vapor Data - "Stitching" in Time and Space

NASA *MERRA v2 – Water Vapor Data Across the Globe

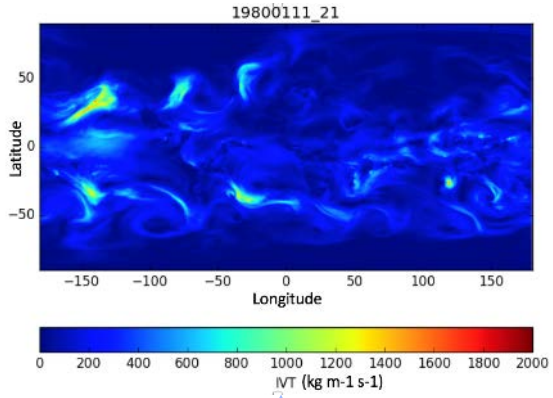

Object Detection, Segmentation and Tracking

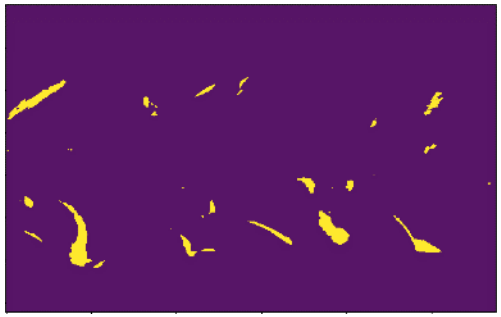

4D Object Constructed (Lat, Lon, Value, Time)

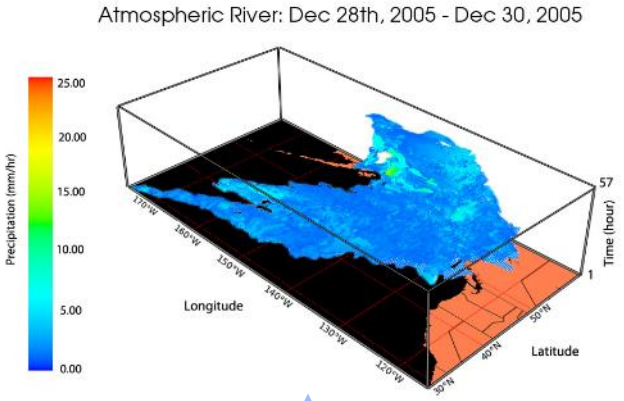

Scott L. Sellars[1], John Graham[1], Dima Mishin[1], Kyle Marcus[2] , Ilkay Altintas[2], Tom DeFanti[1], Larry Smarr[1], Joulien Tatar[3], Phu Nguyen[4], Eric Shearer[4], and Soroosh Sorooshian[4]

[1]Calit2@UCSD; [2]SDSC; *[3]Office of Information Technology, UCI; [4]Center for Hydrometeorology and Remote Sensing, UCI*

PRP Enabled Scott's Workflow to Run 532 Times Faster!

Pacific Research Platform (10-100 Gb/s)

Complete workflow time: 19.2 days→52 Minutes!

UC, Irvine

GPUs

Calit2's FIONA

SDSC's COMET

Source: Scott Sellers, CW3E

UC, San Diego

GPUs

Calit2's FIONA

Moving a Supercomputer Application to a Distributed Computation on PRP:

Simulating the Injection of CO_2 in Brine-Saturated Reservoirs:

Poroelastic and Pressure-Velocity Fields Solved In Parallel Over the PRP, Using MPI for Domain Decomposition, Orchestrated by Kubernetes Containers on FIONA Nodes

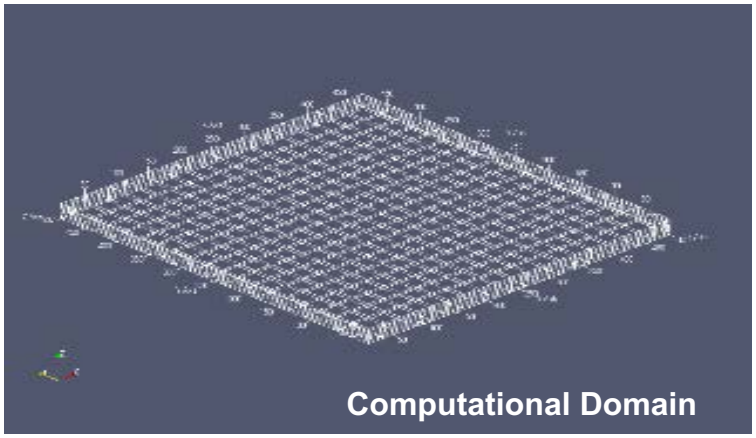

PACIFIC RESEARCH PLATFORM

Numerical methods:
Finite Elements for Solid Phase,
Finite Volume for Aqueous Phase

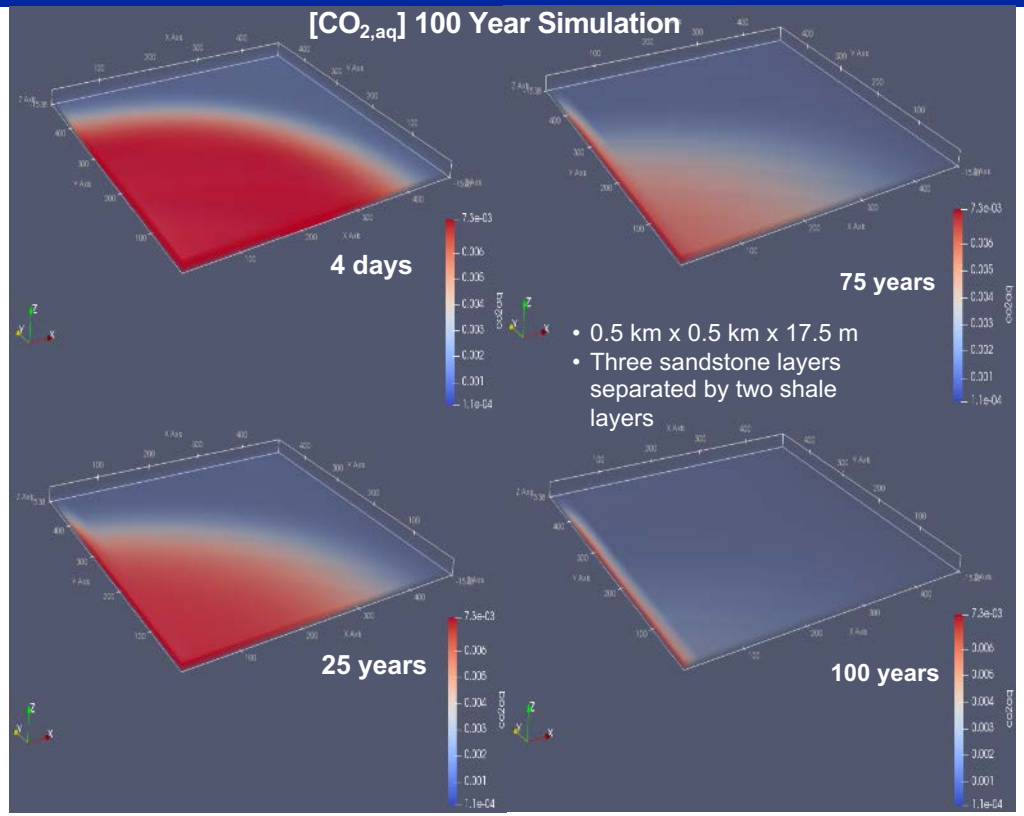

Evolution of Formation Water Carbonic Acid Concentration in Time, After 4 Days of Injection

PRP Distributed-Parallel Computation: Using OpenMPI on Kubernetes to Create Distributed Pods

```
[paolini@fiona k8s]$ kubectl get pod -n sds
NODE [Creating Kubernetes Distributed Pods]
k8s-nvme-01.sdsc.optiputer.net
fiona.cac.washington.edu
dtn-main.ucr.edu
siderea.ucsc.edu
dtn2-daejeon.kreonet.net
fiona-dtn-1.ucsc.edu
fiona.nwsc.ucar.edu
k8s-epyc-01.sdsc.optiputer.net
```

Source: Chris Paolini, SDSU

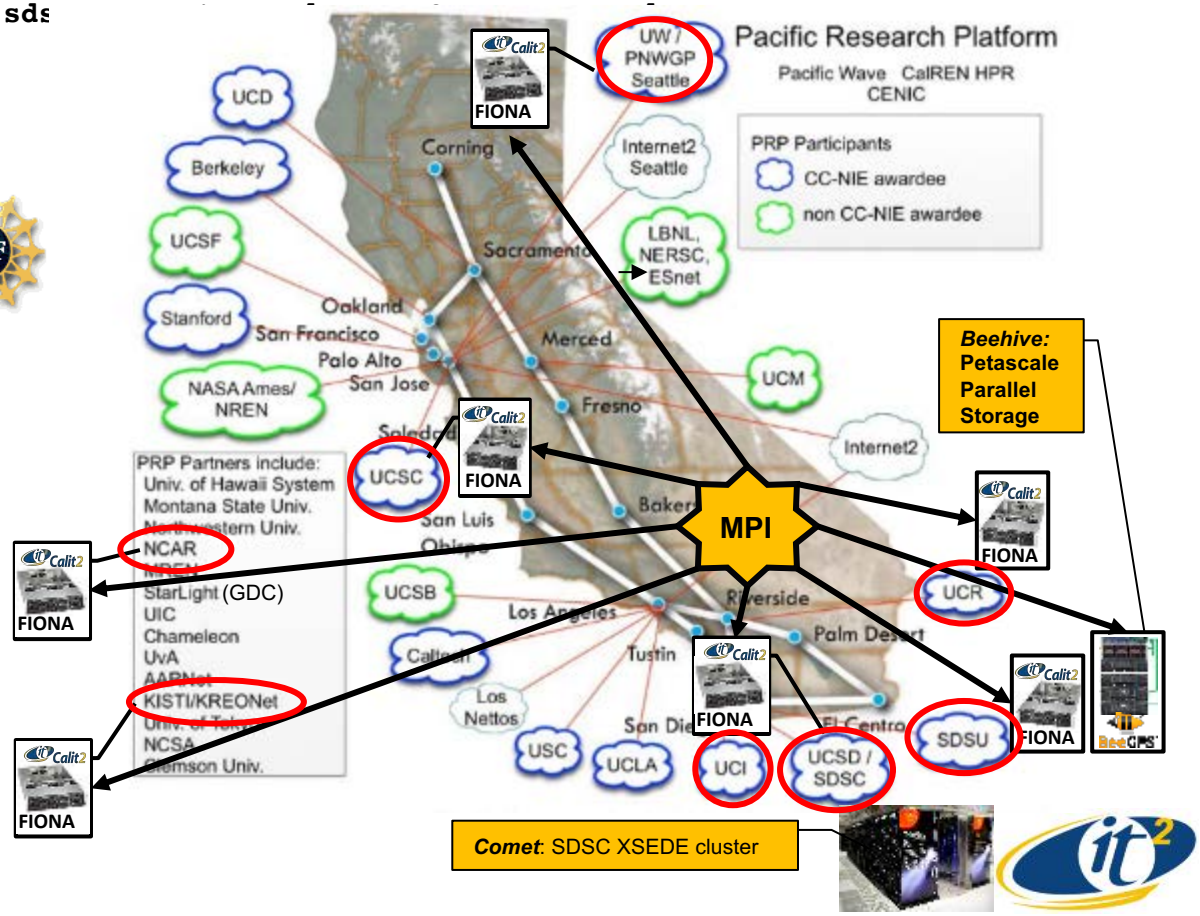

The IceCube Science Program Needed GPUs to Improve Its Pointing Accuracy

GPU Simulations Needed *to Improve Ice Model*.
=> Results in Significant Improvement in Pointing Resolution for Multi-Messenger Astrophysics

IceCube is Now a Major PRP CPU/GPU Application Through the OSG Portal

PRP PACIFIC RESEARCH PLATFORM

See Following GRP Talk by Igor Sfiligoi, SDSC/UCSD

it²

In 2019 a Growing Number of Applications and Campuses Have Driven a 4-Fold Increase in Requested Nautilus GPUs

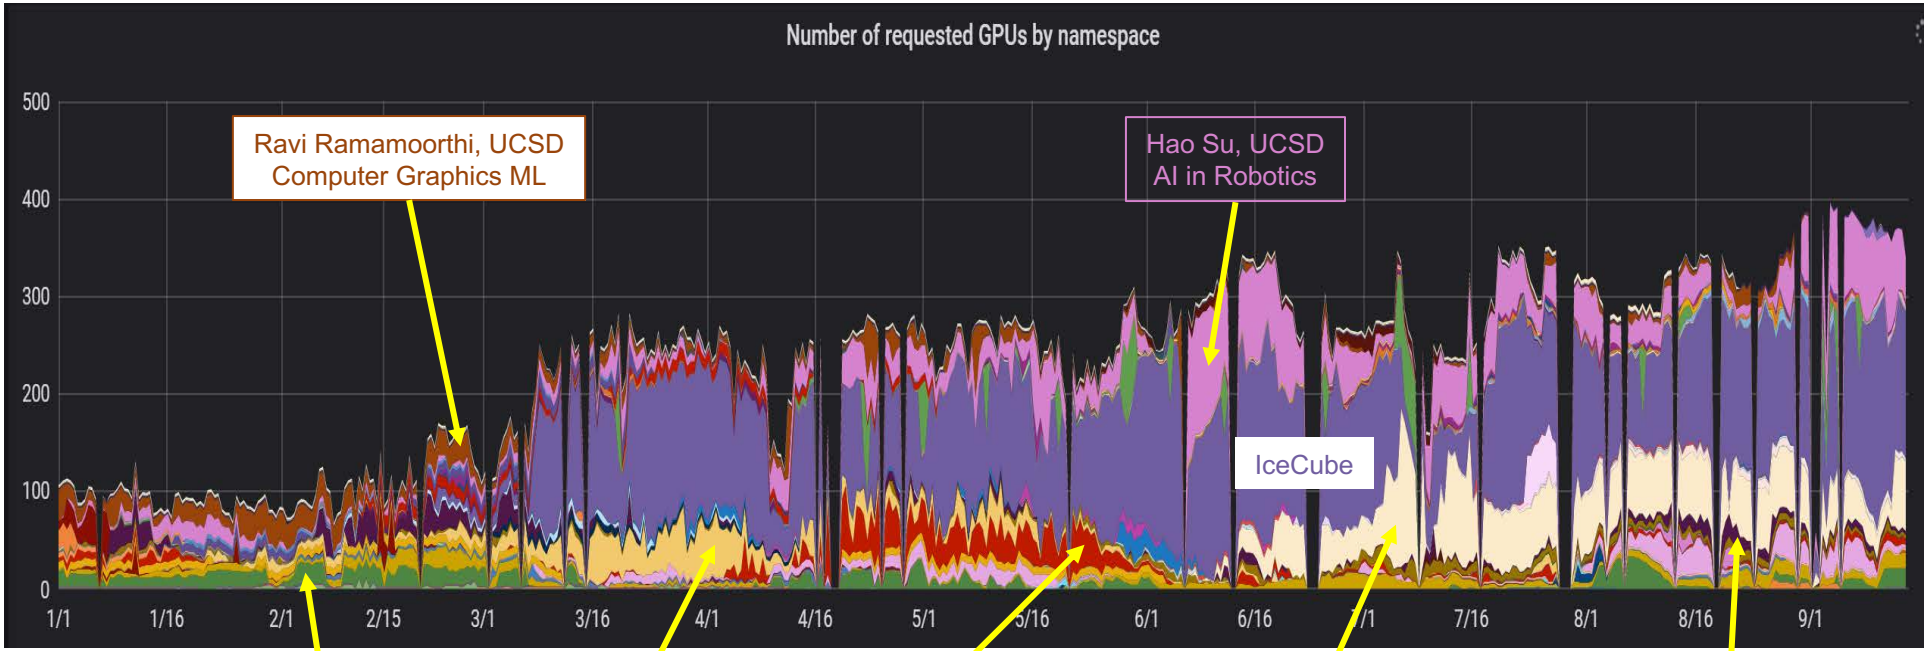

PACIFIC RESEARCH PLATFORM

NSF CHASE-CI Grant Creates a Community Cyberinfrastructure Adding a Machine Learning Layer Built on Top of the Pacific Research Platform

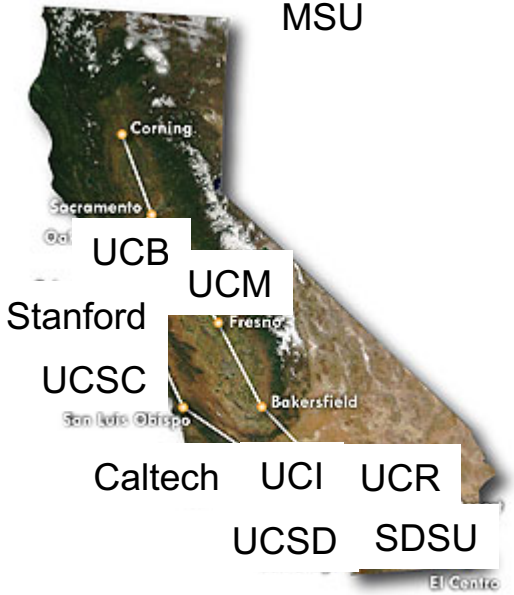

NSF Grant for High Speed "Cloud" of 256 GPUs
For 30 ML Faculty & Their Students at 10 Campuses
for Training AI Algorithms on Big Data

CI-New: Cognitive Hardware and Software Ecosystem Community Infrastructure (CHASE-CI)

For the Period September 1, 2017 – August 31, 2020

SUBMITTED – January 18, 2017

PI: Larry Smarr, Professor of Computer Science and Engineering, Director Calit2, UCSD
Co-PI: Tajana Rosing, Professor of Computer Science and Engineering, UCSD
Co-PI: Ken Kreutz-Delgado, Professor of Electrical and Computer Engineering, UCSD
Co-PI: Ilkay Altintas, Chief Data Science Officer, San Diego Supercomputer Center, UCSD
Co-PI: Tom DeFanti, Research Scientist, Calit2, UCSD

PACIFIC RESEARCH PLATFORM

JACOBS SCHOOL OF ENGINEERING
SDSC UC San Diego

Three UC Irvine Information and Computer Science Professors Using CHASE-CI Over PRP This Year

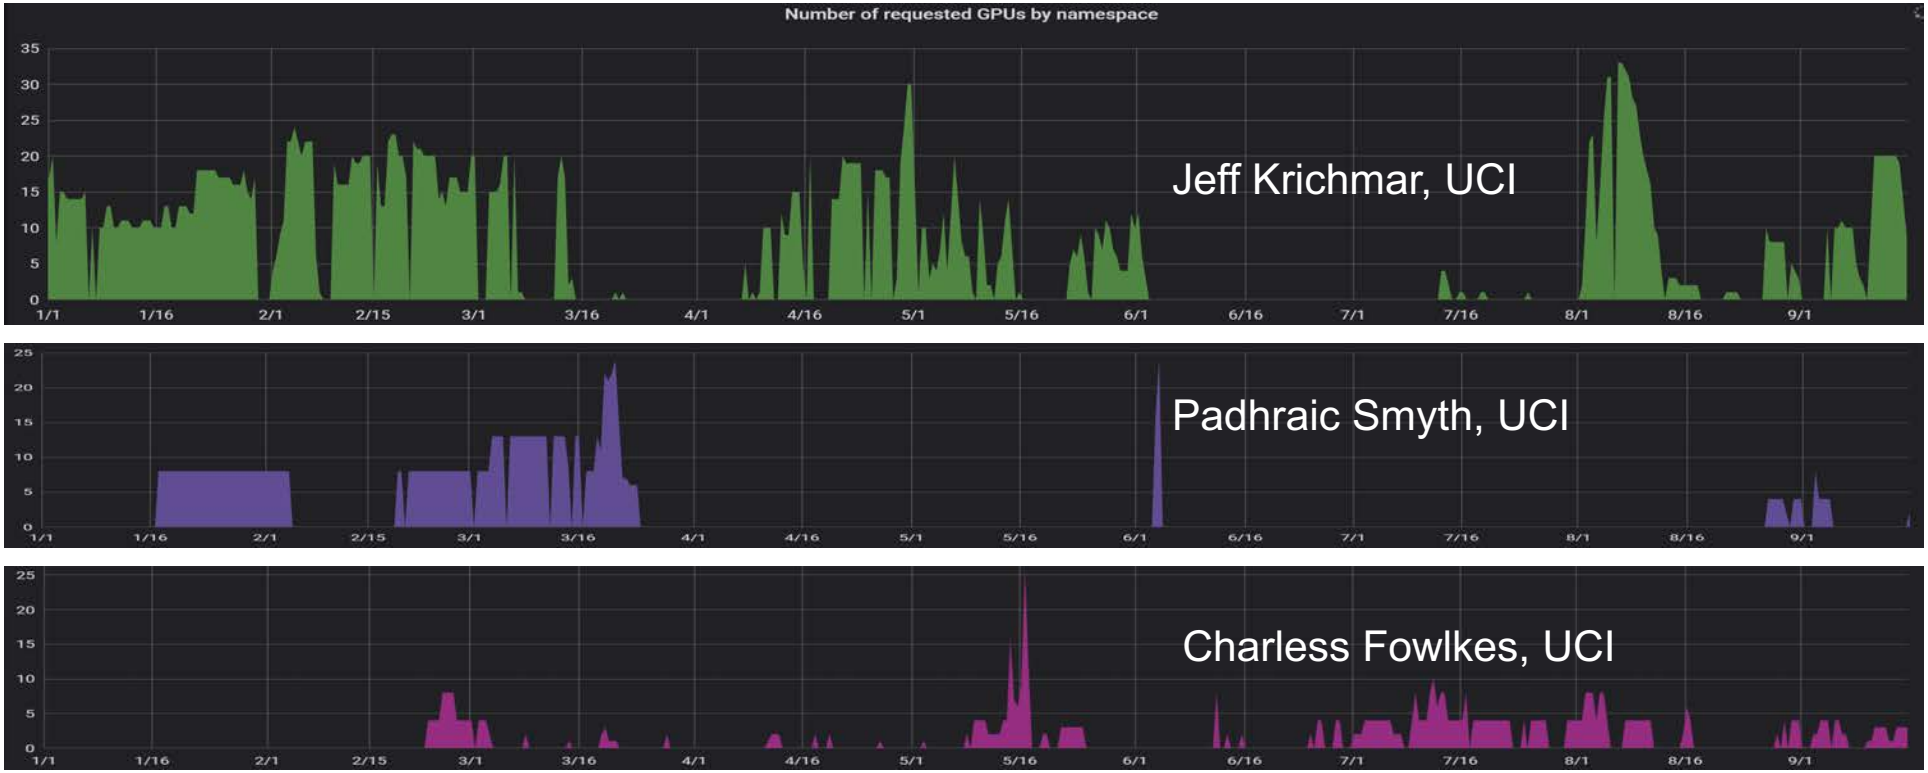

PACIFIC RESEARCH PLATFORM

Using Kubernetes to Surround the PRP/CHASE-CI Machine Learning Platform With Clouds of CPUs, GPUs and Non-Von Neumann Processors

288
64-bit GPUs

CHASE-CI

Field-Programmable Gate Arrays (FPGA)

XILINX

Micron

Amazon Boosts Cloud-Computing Performance With New, GPU-Accelerated AWS Instances

amazon web services

Google Cloud TPU

Microsoft Azure GPU-ACCELERATED MICROSOFT AZURE

Japan plans to build the fastest deep learning supercomputer

AIST ADVANCED INDUSTRIAL SCIENCE AND TECHNOLOGY (AIST)

4352x NVIDIA Tesla V100 GPUs

it²

AIST's Machine Learning Supercomputer Has Been Connected to PRP/CHASE-CI at 6.4 Gbps!

SDSC CTO Phil Papadopoulos and Satoshi Sekiguchi sign 5-Year MOU on behalf of UC San Diego and AIST.

"The two sides will collaborate on cyberinfrastructure projects, notably PRP/CHASE-CI and AIST's AI Bridging Cloud Infrastructure (ABCI)."

- **ABCI Has 4352 NVIDIA Tesla V100 GPUs**
  - **Peak Performance: 550 PFlops (FP16)**
  - **Linpack: 20 PFlops (#8 in TOP500)**
- **ABCI-PRP: Grand Challenge Project**
  - **Gerald Pao, Salk Institute**
  - **George Sugihara, SIO**
  - **Creation of Neuromorphic Deep-Learning Architectures by Large-Scale Dynamic Modeling of Transparent Fish Brains on ABCI**
  - **Disk-to-Disk data transfer from UCSD to ABCI**
    - **FDT version 0.26.1**
    - **574 GB Transferred in 12 Minutes**
    - **6.4 Gbs**

PRP PACIFIC RESEARCH PLATFORM

Source: Ryousei Takano, AIST

it²

Kubernetes Simplifies Moving Software Containers From Global PRP into Commercial Clouds

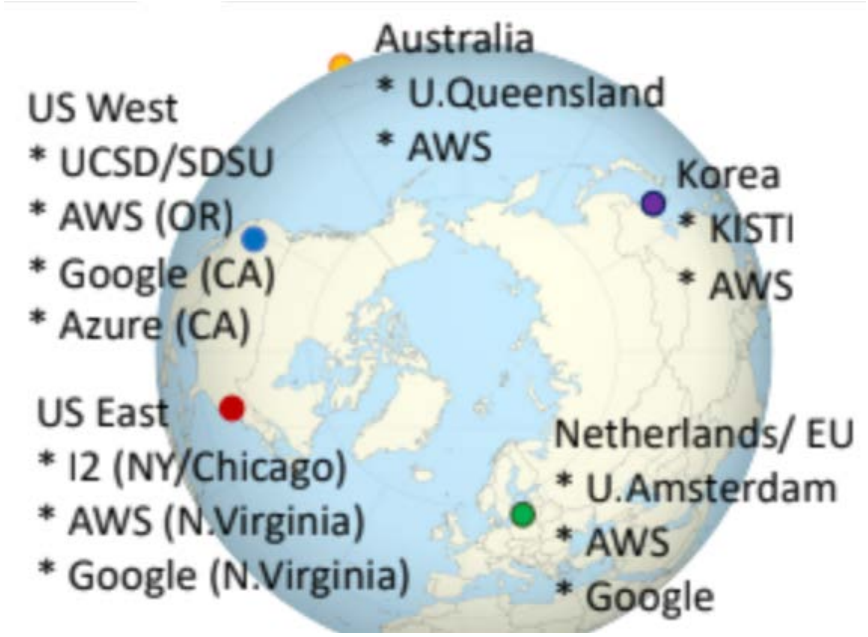

Objective:

Quantify network capabilities of major Cloud providers and compare to what's available through Nautilus

PACIFIC RESEARCH PLATFORM

See GRP Poster by Igor Sfiligoi, UCSD

Global Scientific Instruments Will Produce Ultralarge Datasets Continuously Requiring Dedicated Optic Fiber, Supercomputers, and Machine Learning

Square Kilometer Array – See Following Talk by Shaun Amy

Large Synoptic Survey Telescope (LSST) - See Following Talk by Jeff Kantor

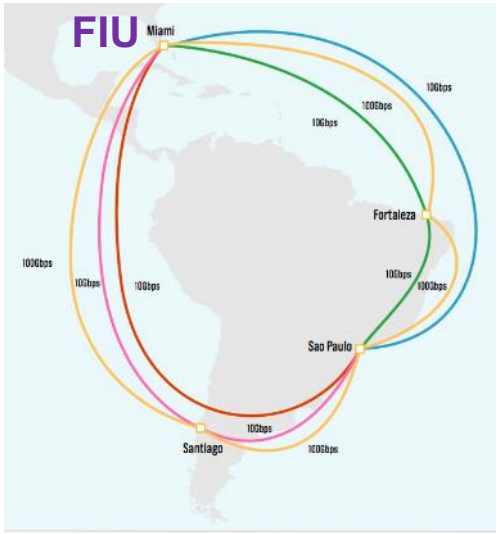

NCSA will be the global central hub for the LSST processing, archiving, and serving the terabytes of data that will be collected every night of the decade-long survey

PACIFIC RESEARCH PLATFORM

PRP/TNRP/CHASE-CI Support and Community:

- US National Science Foundation (NSF) awards to UCSD, NU, and SDSC
  - CNS-1456638, CNS-1730158, ACI-1540112, ACI-1541349, & OAC-1826967
  - OAC 1450871 (NU) and OAC-1659169 (SDSU)
- UC Office of the President, Calit2 and Calit2's UCSD Qualcomm Institute
- San Diego Supercomputer Center and UCSD's Research IT and Instructional IT
- Partner Campuses: UCB, UCSC, UCI, UCR, UCLA, USC, UCD, UCSB, SDSU, Caltech, NU, UWash UChicago, UIC, UHM, CSUSB, HPWREN, UMo, MSU, NYU, UNeb, UNC,UIUC, UTA/Texas Advanced Computing Center, FIU, KISTI, UVA, AIST
- CENIC, Pacific Wave/PNWGP, StarLight/MREN, The Quilt, Kinber, Great Plains Network, NYSERNet, LEARN, Open Science Grid, Internet2, DOE ESnet, NCAR/UCAR & Wyoming Supercomputing Center, AWS, Google, Microsoft, Cisco