Nautilus cluster

Dmitry Mishin
Applications Developer
San Diego Supercomputer Center / Calit2, UCSD
dmishin@ucsd.edu

PACIFIC RESEARCH PLATFORM

Nautilus is a distributed hyperconverged cluster

- **Multi-Institution**
  - .edu
  - .gov
  - .org
  - .com
- **Hyper-Converged**
  - CPU
  - GPU
  - Storage
  - non von Neumann
- **Global Federation**
  - CILogon.org Identity
  - Namespace Isolation
  - OIDC refresh token service

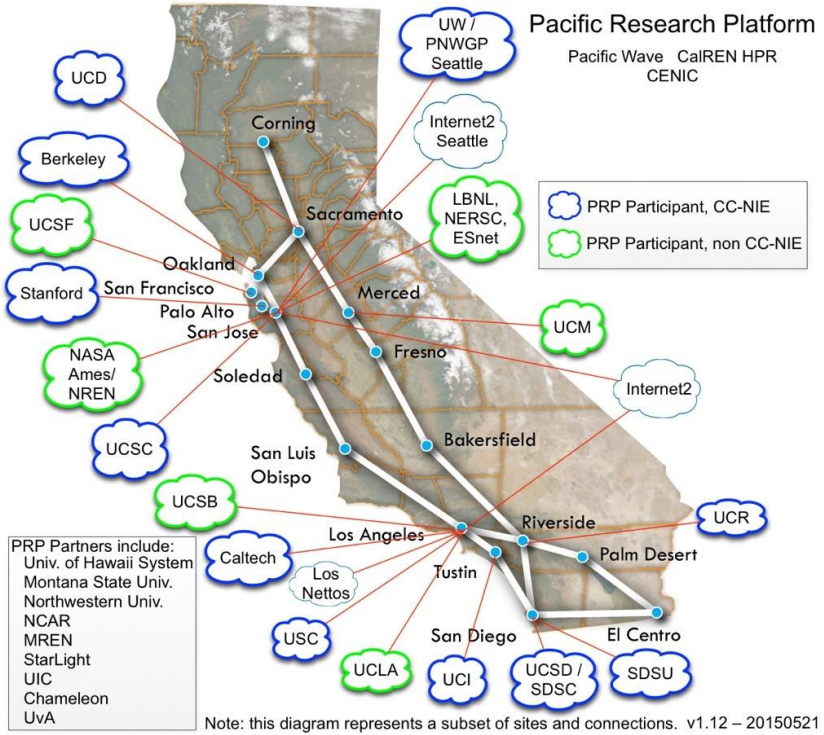

PACIFIC RESEARCH PLATFORM

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

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

MSU, UCM, UCB, UCSC, UCI, UCR, UCSD, SDSU, Stanford, Caltech

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

JACOBS SCHOOL OF ENGINEERING

SDSC UC San Diego it²

CHASE-CI

This project, called the **Cognitive Hardware And Software Ecosystem Community Infrastructure (CHASE-CI)**, will build a cloud of hundreds of affordable Graphics Processing Units (GPUs), networked together with a variety of neural network machines to facilitate development of next generation cognitive computing.

This cloud will be accessible by 30 researchers assembled from 10 universities via the NSF-funded Pacific Research Platform. These researchers will investigate a range of problems from image and video recognition, computer vision, contextual robotics to cognitive neurosciences using the cloud to be purpose-built in this project.

256 FP16/32 GPUs in 32 2U Intel Scalable Dual 12 core with Optane Memory

PRP PACIFIC RESEARCH PLATFORM

it²

FIONA8: a FIONA with 8 GPUs
Supports PRP Data Science Machine Learning--4M GPU Core Hours/Week

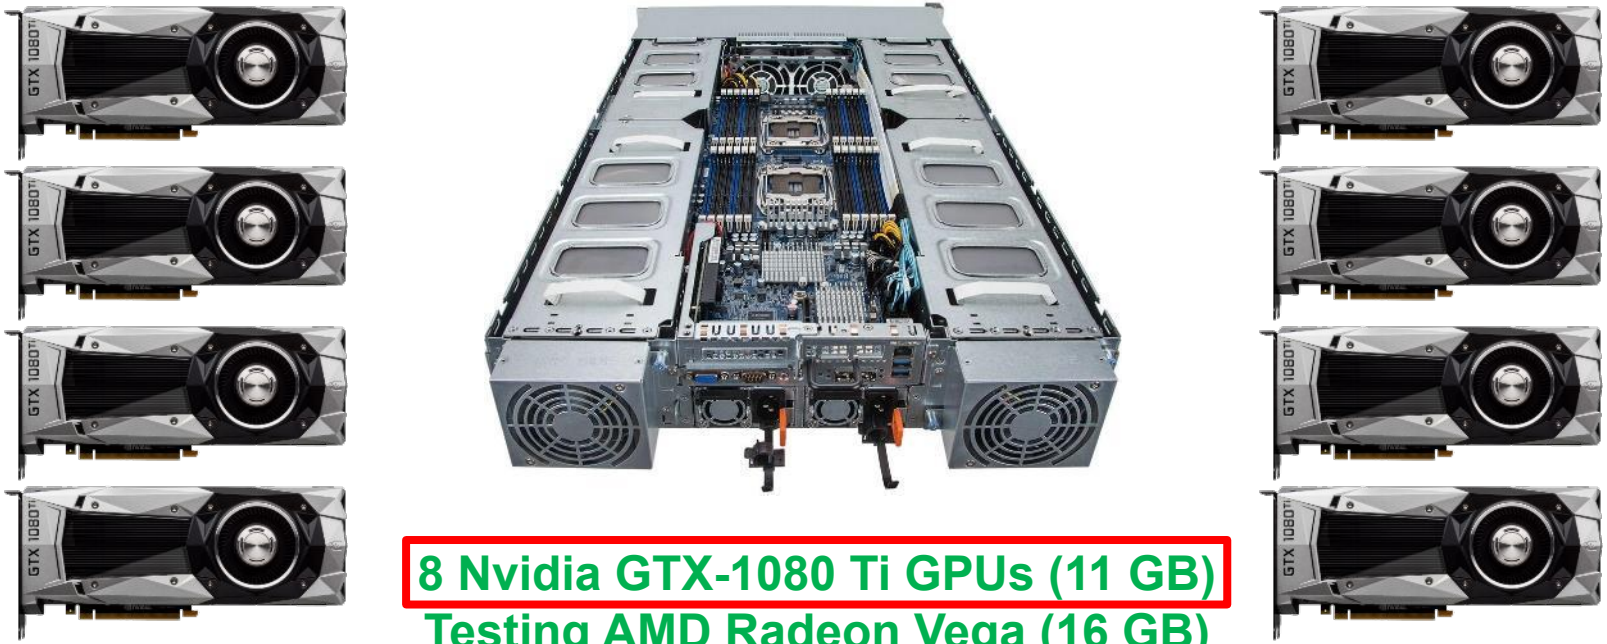

8 Nvidia GTX-1080 Ti GPUs (11 GB)

Testing AMD Radeon Vega (16 GB)

24 CPU Cores, 32,000 GPU cores, 96 GB RAM, 2TB SSD, Dual 10Gbps ports

3" High; ~$16,000

Nautilus is a distributed hyperconverged cluster

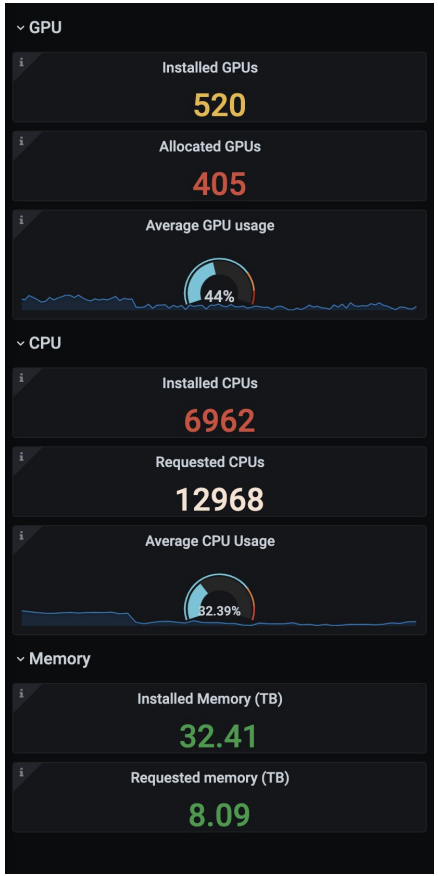

Ceph storage:

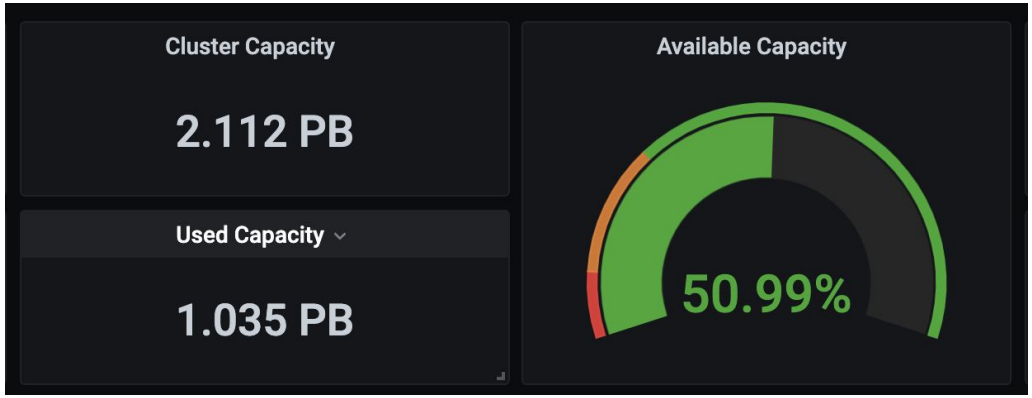

PACIFIC RESEARCH PLATFORM

it²

Typical use of Nautilus

Admins - Network monitoring tools (PerfSonar, sflow, tstat, artemis)

Users - Mostly Machine Learning, GPUs + CPUs (Tensorflow, PyTorch, etc)

Additional services - Jupyter, GitLab, Nextcloud, WebODM, RocketChat

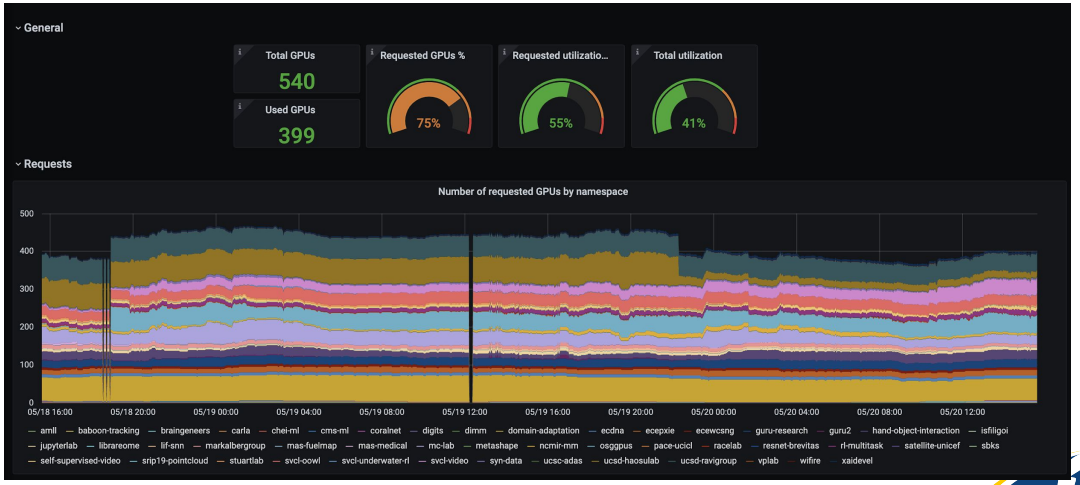

PACIFIC RESEARCH PLATFORM

Federated Identity

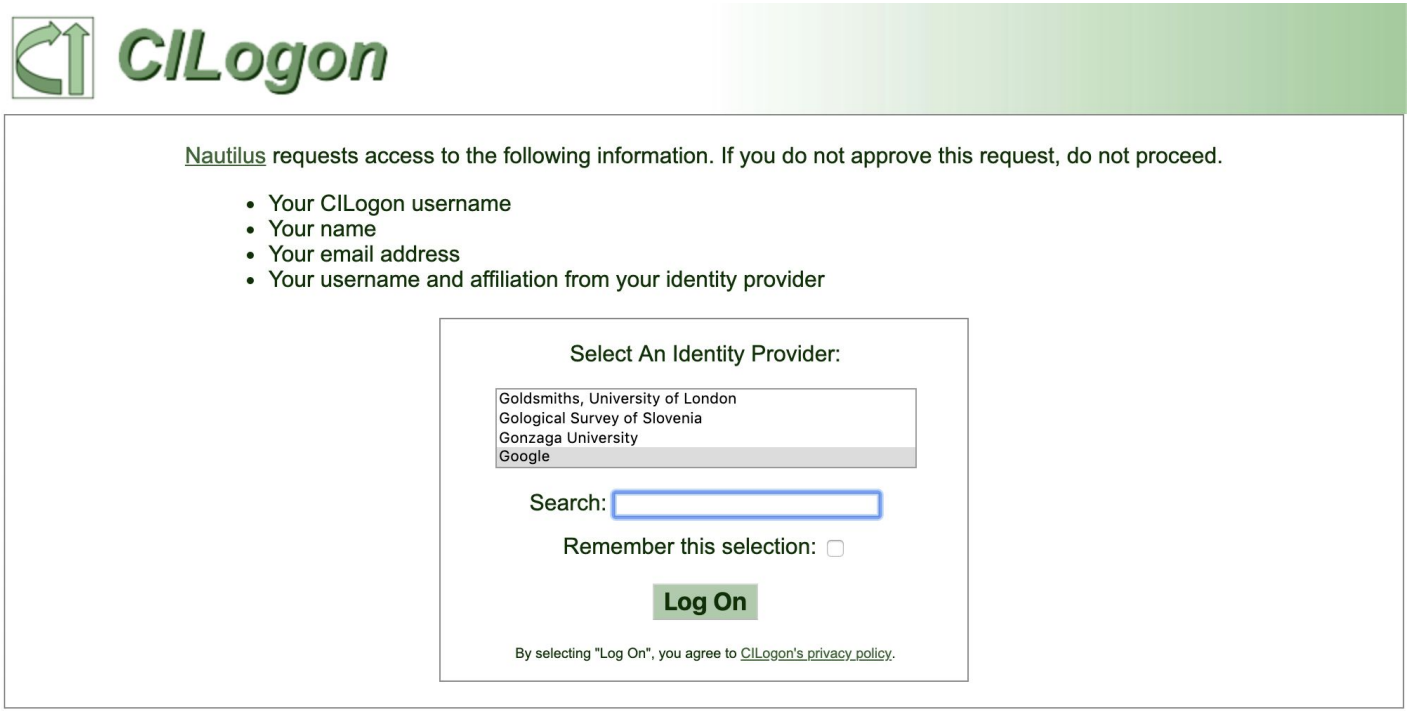

PACIFIC RESEARCH PLATFORM

it²

Rook is Cloud Native Ceph in our Hyper-converged cluster

Design

With Rook running in the Kubernetes cluster, Kubernetes applications can mount block devices and filesystems managed by Rook, or can use the S3/Swift API for object storage. The Rook operator automates configuration of the Ceph storage components and monitors the cluster to ensure the storage remains available and healthy. There is also a REST API service for configuring the Rook storage and a command line tool called `rook`.

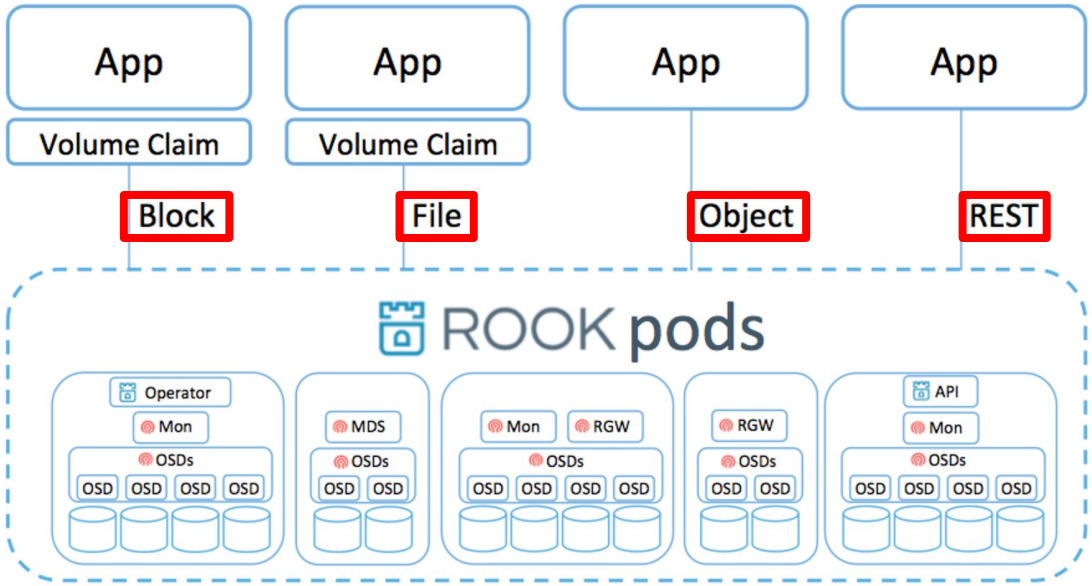

PACIFIC RESEARCH PLATFORM

Nautilus 2PB+ Storage upgrade

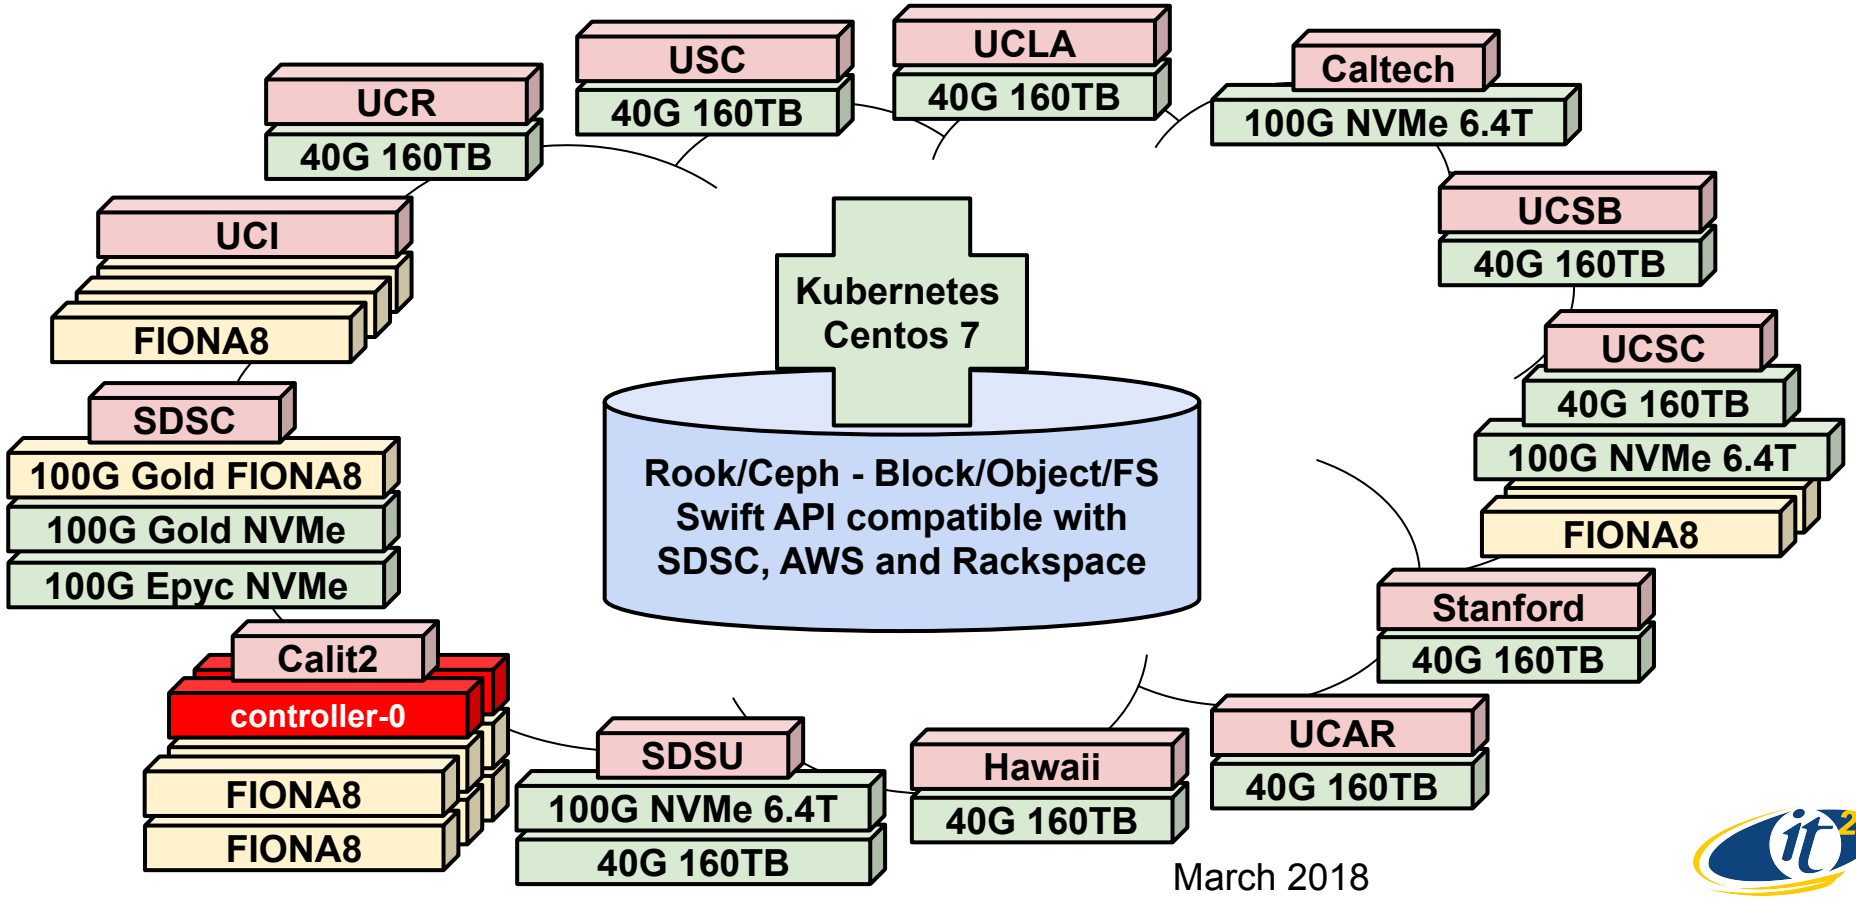

March 2018

Schedule FPGAs

The Xilinx FPGA device plugin for Kubernetes is a Daemonset deployed on the kubernetes(a.k.a k8s) cluster which allows you to:

Discover the FPGAs inserted in each node of the cluster and expose info of the FPGAs such as quantities, DSA(shell) type and timestamp, etc
Run FPGA accessible containers in the k8s cluster
More info about k8s device plugin, please refer to https://kubernetes.io/docs/concepts/extend-kubernetes/compute-storage-net/device-plugins/

https://github.com/Xilinx/FPGA_as_a_Service

PRP PACIFIC RESEARCH PLATFORM

it²

Walls, Caves and Waves

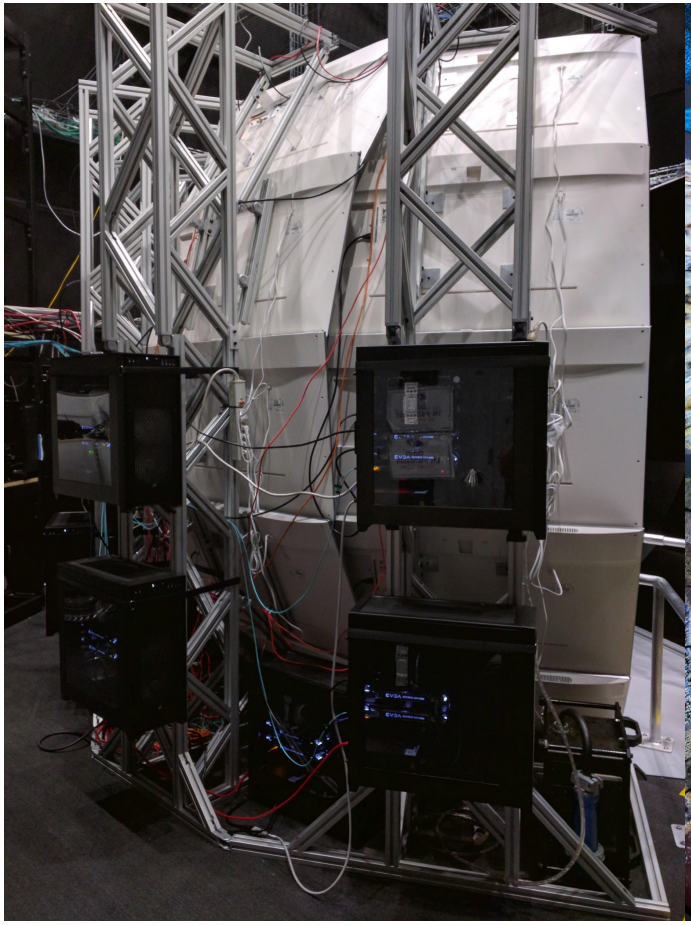

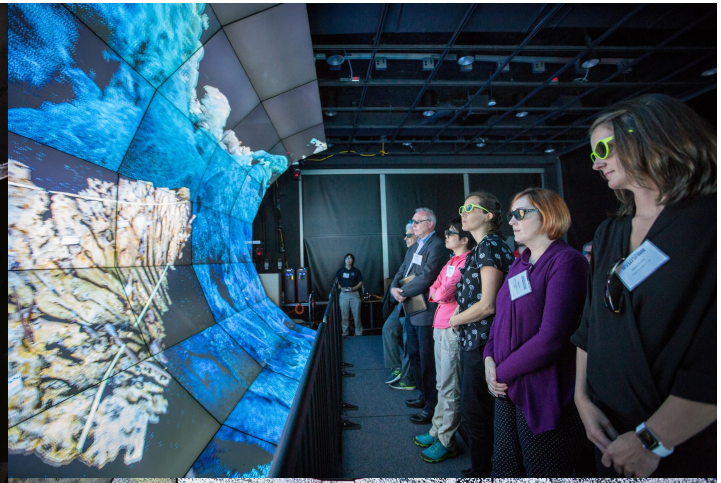

UCSD Adding >350 Game GPUs to Data Sciences Cyberinfrastructure - Devoted to Data Analytics and Machine Learning

UC San Diego

IT SERVICES

88 GPUs
for Students

Open Science Grid

48 GPUs for
OSG Applications

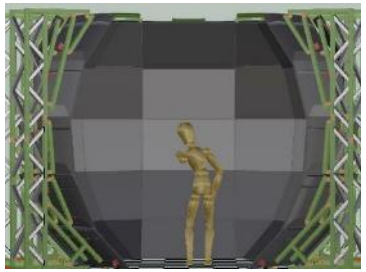

SunCAVE 70 GPUs
WAVE + Vroom 48 GPUs

UCSD Cognitive Science

JACOBS SCHOOL OF ENGINEERING

SDSC it²

QUALCOMM INSTITUTE

FIONA with
8-Game GPUs

CHASE-CI Grant Provides
96 GPUs at UCSD
for Training AI Algorithms on Big Data

it²

Project Calico

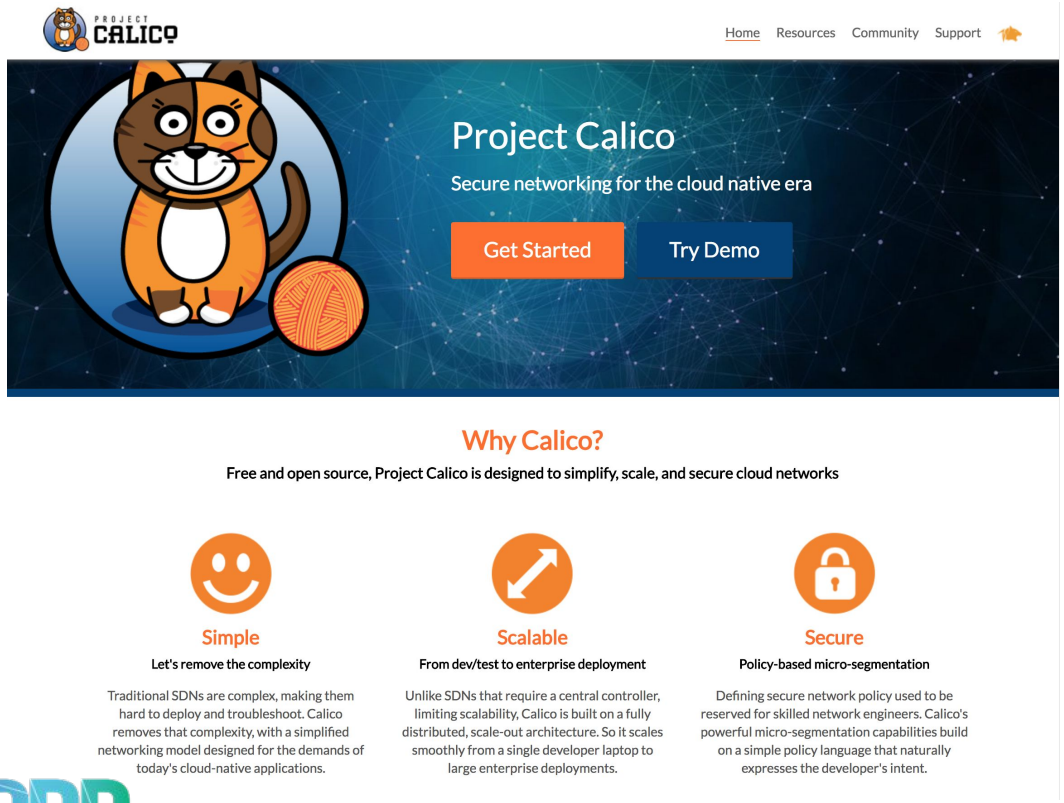

- CNI plugin
- Full cluster network security
  a. Nodes firewall
  b. Namespaces isolation

https://www.projectcalico.org/

Pacific Research Platform

http://pacificresearchplatform.org

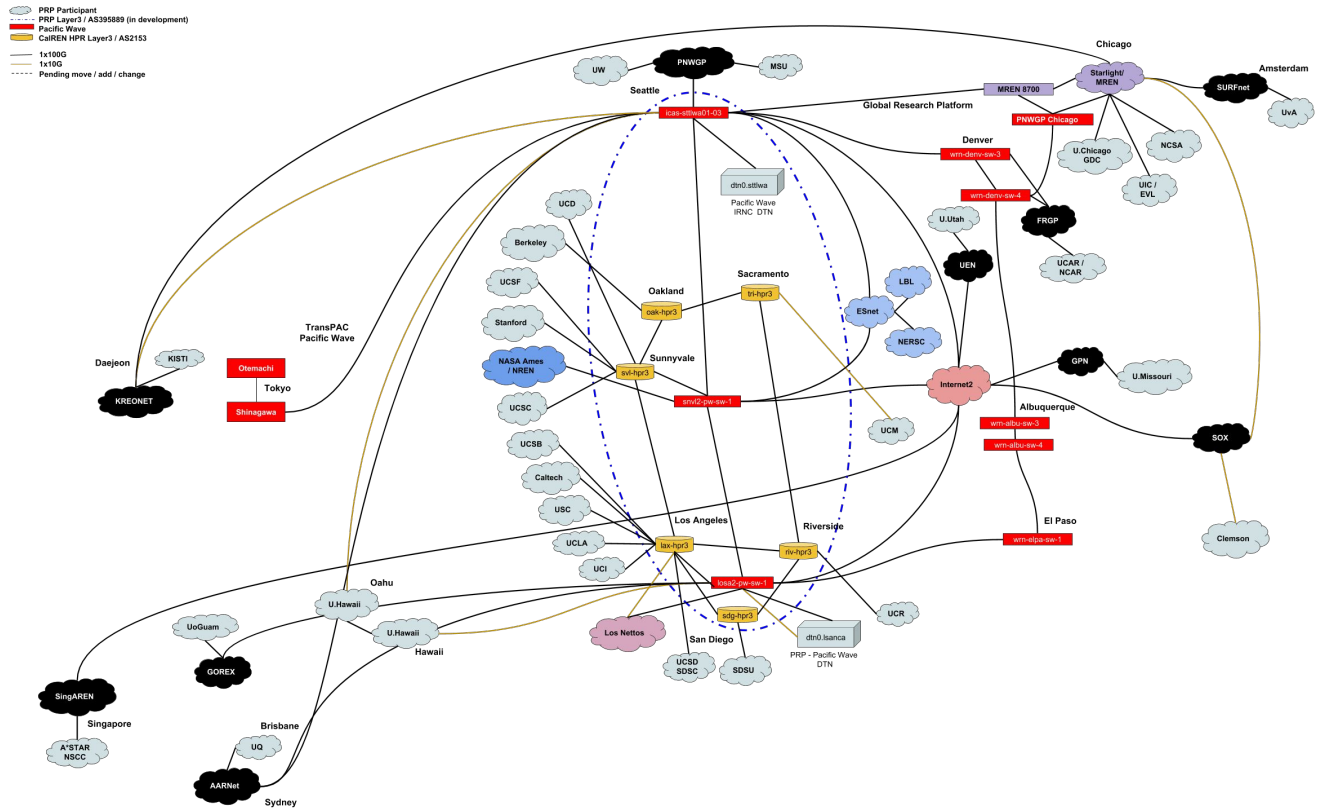

NOTE: this diagram represents a subset of sites, devices, and connections

v0.40.16
20180716

PerfSonar for network monitoring

AUTOMATION !!!

- Perfsonar deploys automatically as new nodes are joined to the cluster.
- No Human is needed to configure the MaDDash.
- The Meshconfig is handled by a golang webservice.

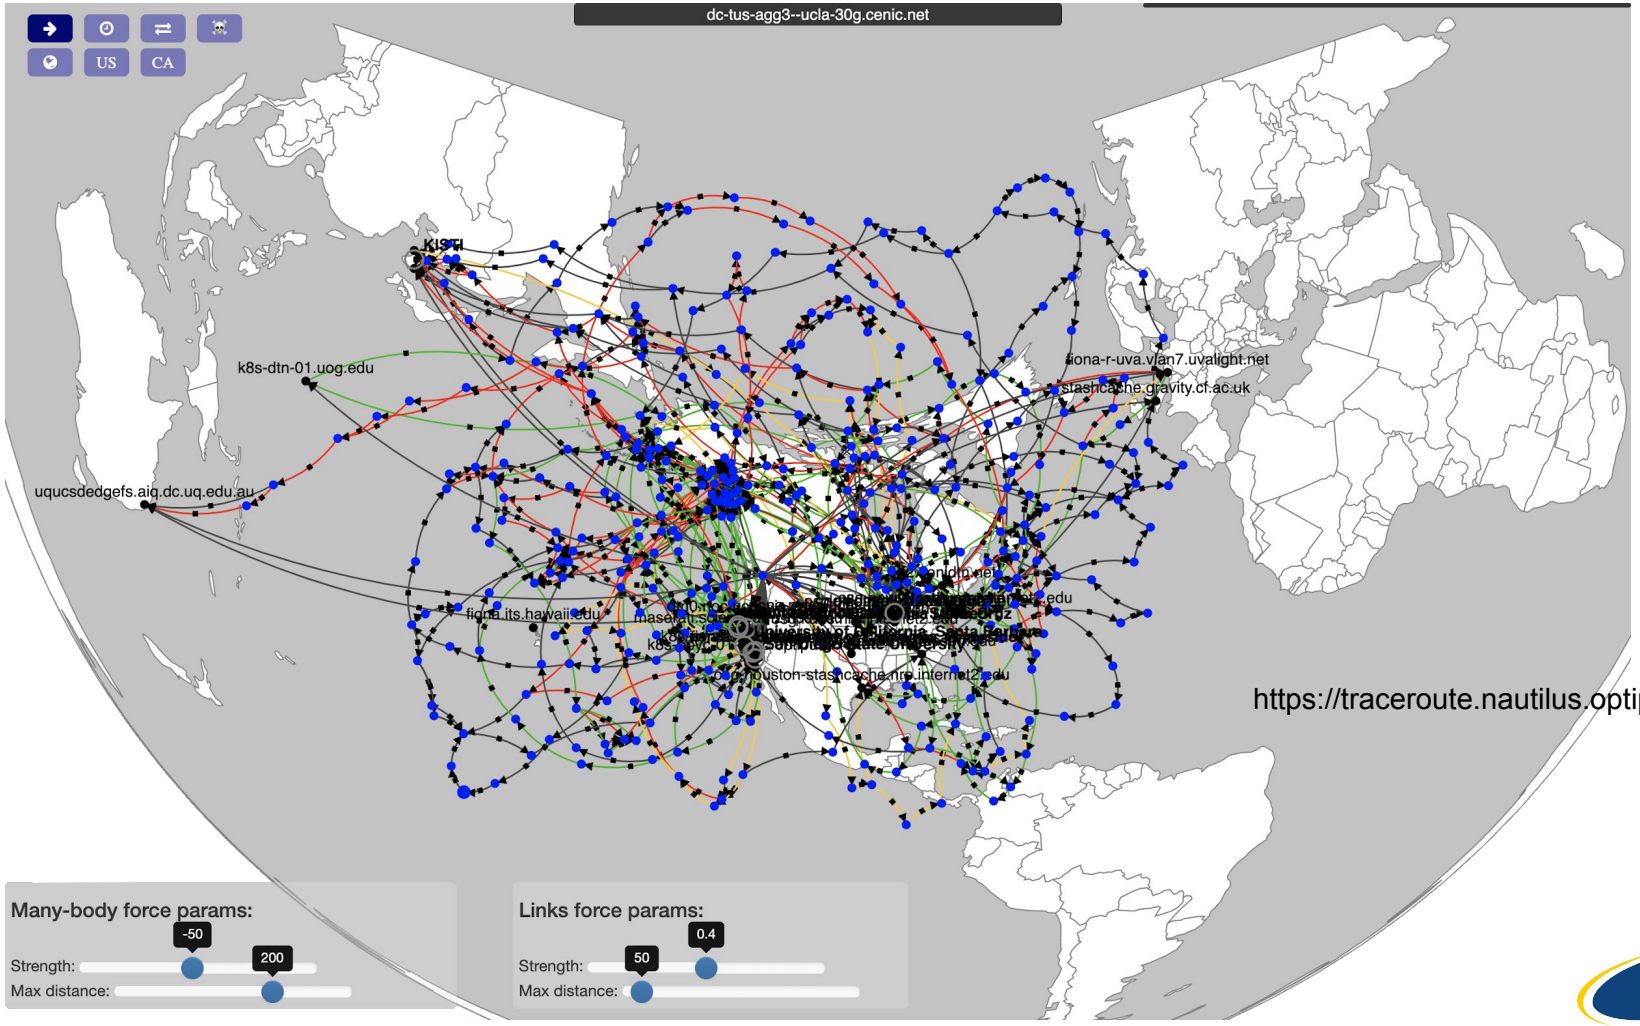
dc-tus-agg3--ucla-30g.cenic.net
US
CA
KISTI
k8s-dtn-01.uog.edu
uqucsdedgefs.aiq.dc.uq.edu.au
fiona-r-uva.vlan7.uvalight.net
stashcache.gravity.cf.ac.uk
fiona.its.hawaii.edu
https://traceroute.nautilus.optiputer.net
Many-body force params:
-50
Strength:
200
Max distance:
Links force params:
0.4
Strength:
50
Max distance:

ElastiFlow sflow visualization

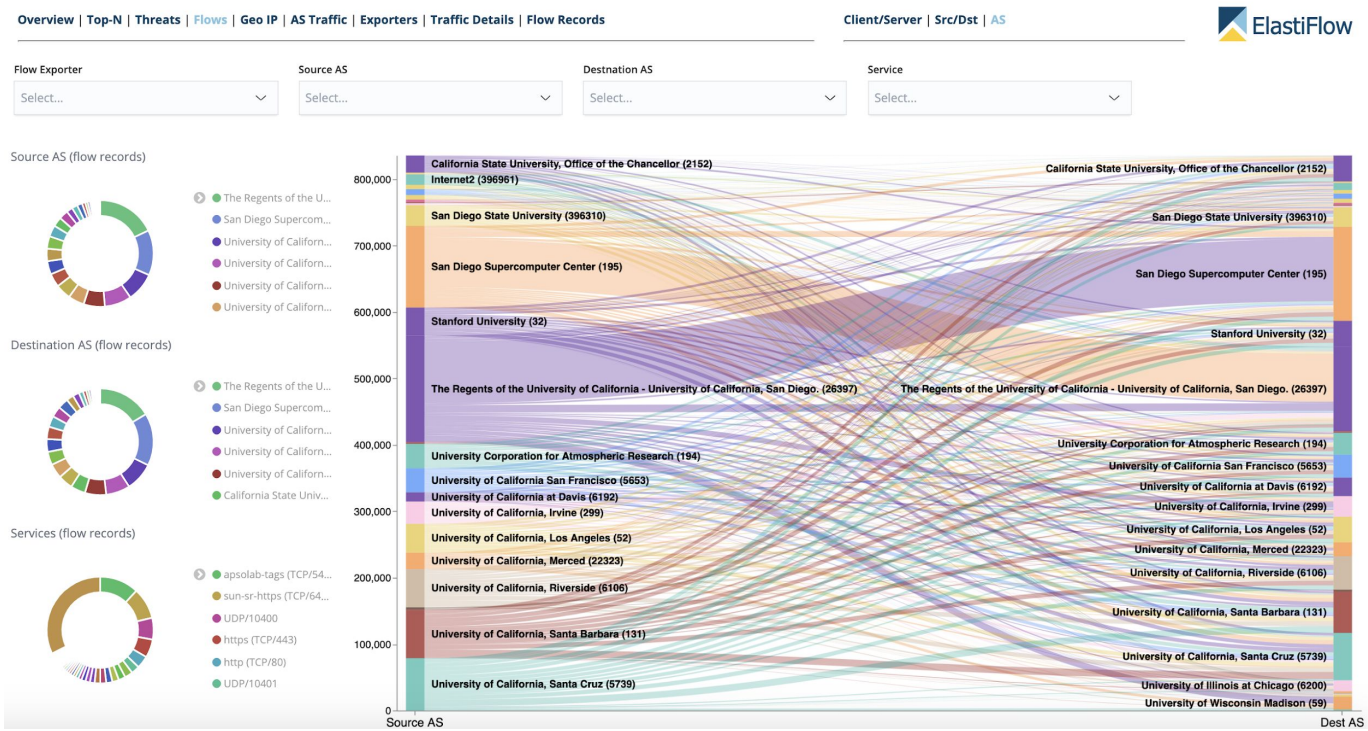

ElastiFlow™ provides network flow data collection and visualization using the Elastic Stack (Elasticsearch, Logstash and Kibana). It supports Netflow v5/v9, sFlow and IPFIX flow types (1.x versions support only Netflow v5/v9).

JupyterLab running as a service

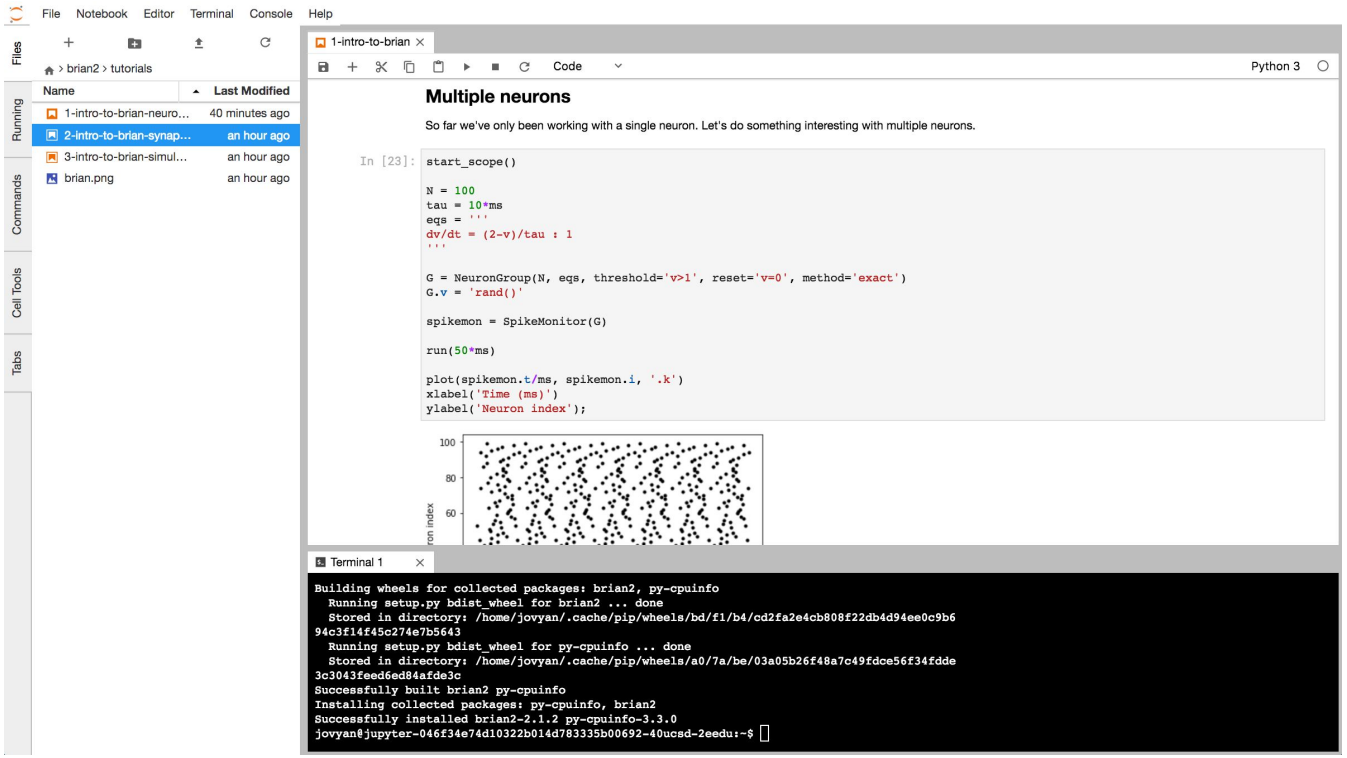

Kubernetes provides Jupyter Notebooks direct access to the GPU.

WebODM

Drone Data Processing Summary

05 March, 2018

Treasure Island, Flight 1, High Quality Preset, No resize

Treasure Island (urban area), nadir grid mapping mission with 80/70 overlap, 250 ft, 107 images, 17 acres.

Flight Info
Location: Treasure Island
Date: 2018-02-17
Start time: 10:13
End time: 10:19
UAV: Phantom III Standard
Sensor: Phantom III Standard

Capture Settings
Flight Control App: Pix4Dcapture
Mission Type: grid
Num images: 107
Front overlap: 80%
Side overlap: 70%
Altitude (ft): 250
Image GSD (in): 1.31
Drone Speed: fast
Trigger mode: fast
White balance: sunny
Exposure: auto

Processing Settings
Platform: WebODM
StitchID: ti_2018-02-17_flt01_stch01
Parameter Preset: High Quality
Processing Time (min): 192
Description: High Quality Preset, no resize

Images

Locations | Footprints

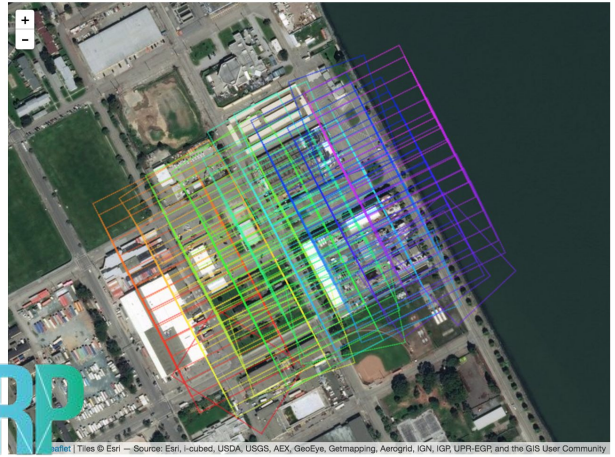

WebODM Results

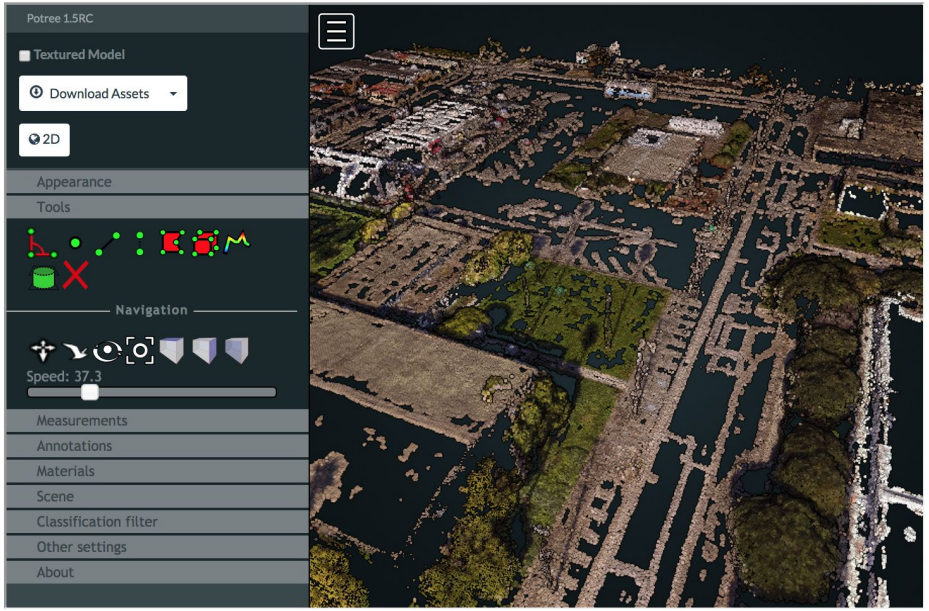

Full screen

Assessment: Significant warping of cars and building tops. Mild warping of parking lot lines.

PACIFIC RESEARCH PLATFORM

GitLab code and container repo

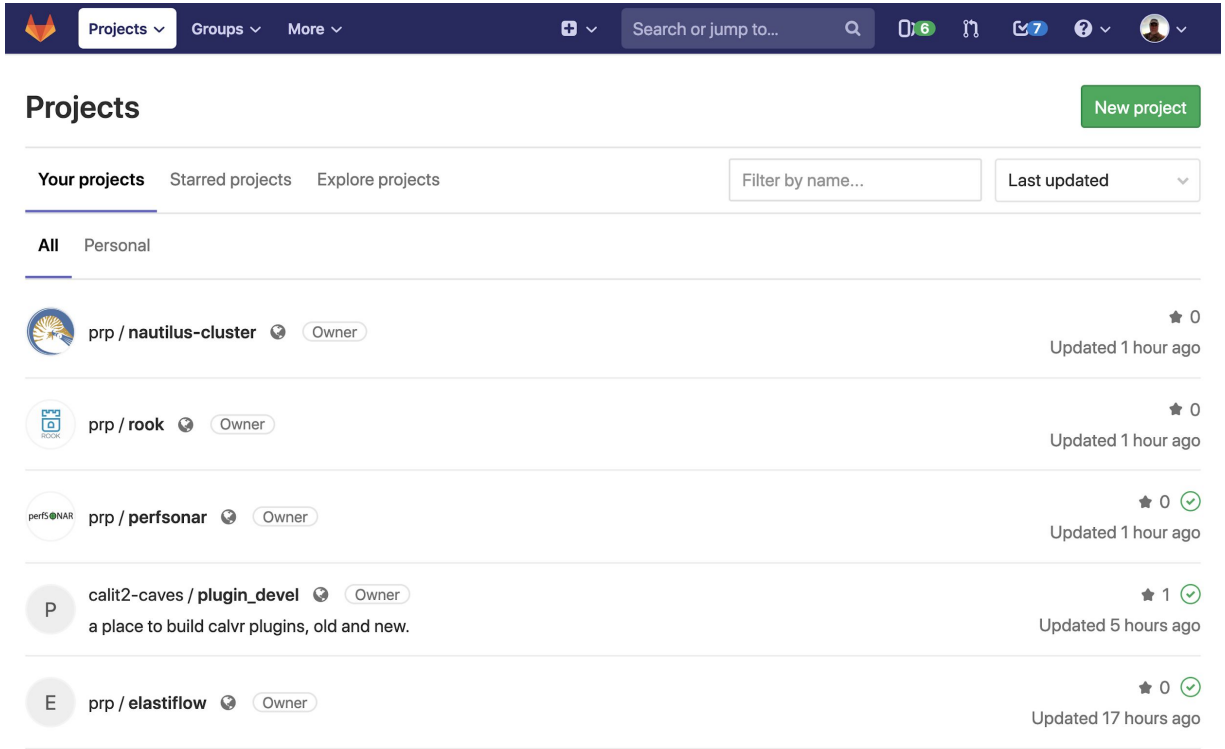

https://gitlab.nautilus.optiputer.net

Projects Groups More Search or jump to... 6 7

perfSONAR

.gitlab-ci.yml 4.04 KB

Edit Web IDE Replace Delete

```yaml
image: gitlab-registry.nautilus.optiputer.net/prp/golang-docker-gitlabci

before_script:
  - docker login -u gitlab-ci-token -p $CI_BUILD_TOKEN gitlab-registry.nautilus.optiputer.net

stages:
  - build-and-push

build-and-push-perfsonar:
  stage: build-and-push
  tags:
  - build-as-docker
  only:
    changes:
      - perfsonar-toolkit/**/*
      - perfsonar-toolkit/*
  script:
  - cd perfsonar-toolkit
  - docker build -t gitlab-registry.nautilus.optiputer.net
  - docker tag gitlab-registry.nautilus.optiputer.net/${C
  - docker push gitlab-registry.nautilus.optiputer.net/${

build-and-push-testpoint:
  stage: build-and-push
  only:
    changes:
```

CI/CD

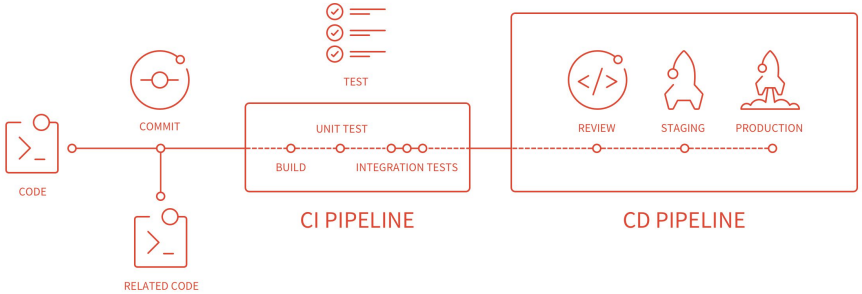

PACIFIC RESEARCH PLATFORM

it²

RocketChat for community support

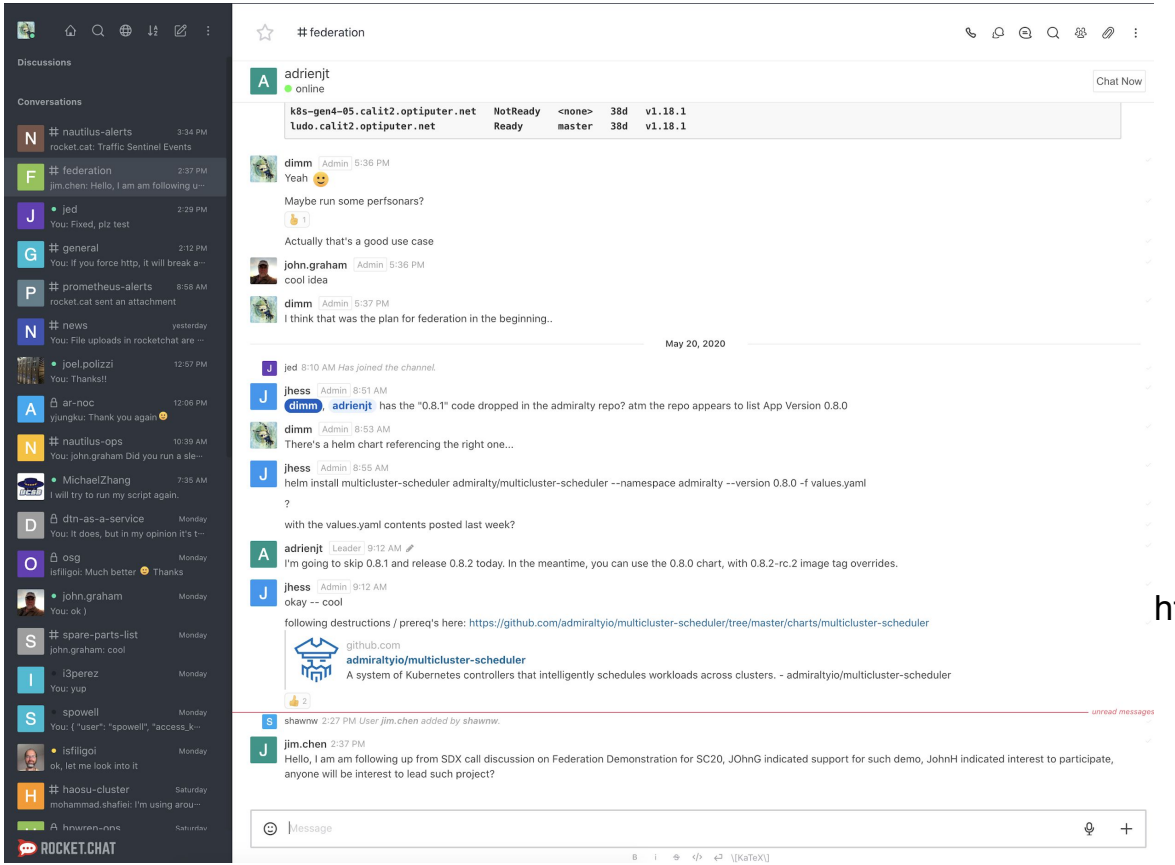

Currently:
>1K users
>300K messages

https://rocket.nautilus.optiputer.net

Namespaces as a collaboration environment

adalab

University of California, San Diego: Advanced Data Analytics

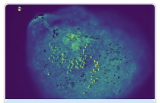

braingeneers

UC Santa Cruz & UC San Francisco: Machine Learning with the recordings derived from cortical organoids

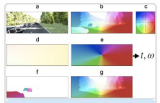

carl-uci

University of California, Irvine: Reinforcement learning and motion decomposition

chei-ml

University of California, San Diego and SIO: Deep learning for coral species segmentation

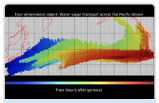

connect

University of California, San Diego: Machine Learning in Earth Sciences

coralnet

University of California, San Diego: Computer Vision and Machine Learning for Coral Ecology

deepgtex-prp

Clemson University: Deep Learning in Oncogenomics

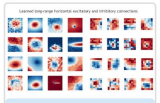

desalab

University of California, San Diego: Biologically plausible deep learning for vision

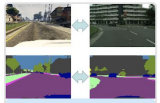

domain-adaptation

University of California, San Diego: Transfer and multitask learning

dreiman

University of California, Santa Cruz: Deep learning research with applications in cosmology, extragalactic astronomy and astrophysics

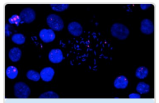

ecdna

University of California, San Diego: Deep Learning for medical imaging

ecewcsng

University of California, San Diego: Deep Learning for sensor data processing

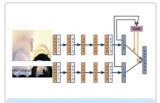

guru-research

University of California, San Diego: Applications of

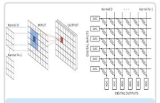

kube-environment

University of California, San Diego: ReRAM-based

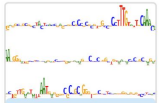

kundajelab

Stanford University: Deep Learning in Genomics

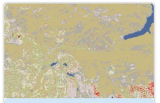

mas-fuelmap

University of California, San Diego: Generating

mesl

University of California, San Diego: Convolutional

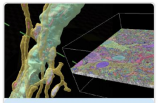

ncmir-mm

University of California, San Diego: Image

Currently:
>300 namespaces

https://ucsd-prp.gitlab.io/nautilus/namespaces/

Federation

Scheduling	• Schedule pods in a remote cluster, control execution, watch the state
Network	• Be able to communicate to pods in remote cluster
Storage	• Be able to store data in remote cluster

PRP
PACIFIC RESEARCH PLATFORM

it²

AARCH64 IoT node

Also used as AARCH64 Gitlab runner for building IoT container images

PACIFIC RESEARCH PLATFORM

Thanks to Our Support:

- US National Science Foundation (NSF) awards
  - CNS 0821155, CNS-1338192, CNS-1456638, CNS-1730158, ACI-1540112, & ACI-1541349
- University of California Office of the President CIO
- UCSD Chancellor's Integrated Digital Infrastructure Program
- UCSD Next Generation Networking initiative
- Calit2 and Calit2 Qualcomm Institute
- CENIC, PacificWave and StarLight
- DOE ESnet