

Open Cyberinfrastructure for an Open Society

Frank Würthwein
Director, San Diego Supercomputer Center

December 20th 2023



Three Challenges

- The gap between those who have and those who can't afford is becoming wider.
- How do we bridge the gap between classroom and research scale of any concept we teach?
- The end of Moore's Law is leading to a proliferation of "architectures" ... domain science adoption is at risk.

- I will switch back and forth between
 - Vision & Ideas
 - R&D
 - Existing “production” systems

Idea

R&D

Production

To minimize confusion, I'll tag my slides with these labels as much as possible

Towards an Open Infrastructure

Horizontally open => institutions can integrate their resources

Vertically open => projects can build on the infrastructure

Production


- Create an Open National Cyberinfrastructure that allows the federation of CI at all ~4,000 accredited, degree granting higher education institutions, non-profit research institutions, and national laboratories.
 - Open Science
 - Open Data
 - Open Source
 - Open Infrastructure
 - ← Open Compute
 - ← Open Storage & CDN
 - ← Open devices/instruments/IoT, ...?

Openness for an Open Society

A background image showing a diverse group of young people, likely students, smiling and looking towards the right. The image is slightly blurred, with the focus on a young man in the foreground wearing glasses and smiling.

The Minds We Need

- **Connect every community college, every minority serving institution, and every college and university, including all urban, rural, and tribal institutions** to a world-class and secure R&E infrastructure, with particular attention to institutions that have been chronically underserved;
- **Engage and empower every student and researcher** everywhere with the opportunity to join collaborative environments of the future, because we cannot know where the next Edison, Carver, Curie, McClintock, Einstein, or Katherine Johnson will come from; and



Lot's of funded projects that contribute to this **shared vision** in different ways.

Hardware funded by NSF, DOD, DOE, ...

Services building on top include:

Open Science Data Federation

Open Science Compute Federation

National Data Platform

Fusion Data Platform

R&E platforms for multiple campuses

R&D

Production

Open Infrastructure is “owned” and “built” by the community for the community



NATIONAL RESEARCH PLATFORM (NRP)

OUR ATTEMPT TO EXECUTE ON THIS VISION



Flexible Architecture to build on horizontally and vertically



- Depending on effort available and control desired, you can build on NRP both vertically and horizontally at different layers of the stack.



- NRP is a non-local extendable container deployment platform, thus allowing many uses unthinkable for a SLURM cluster in a data center.



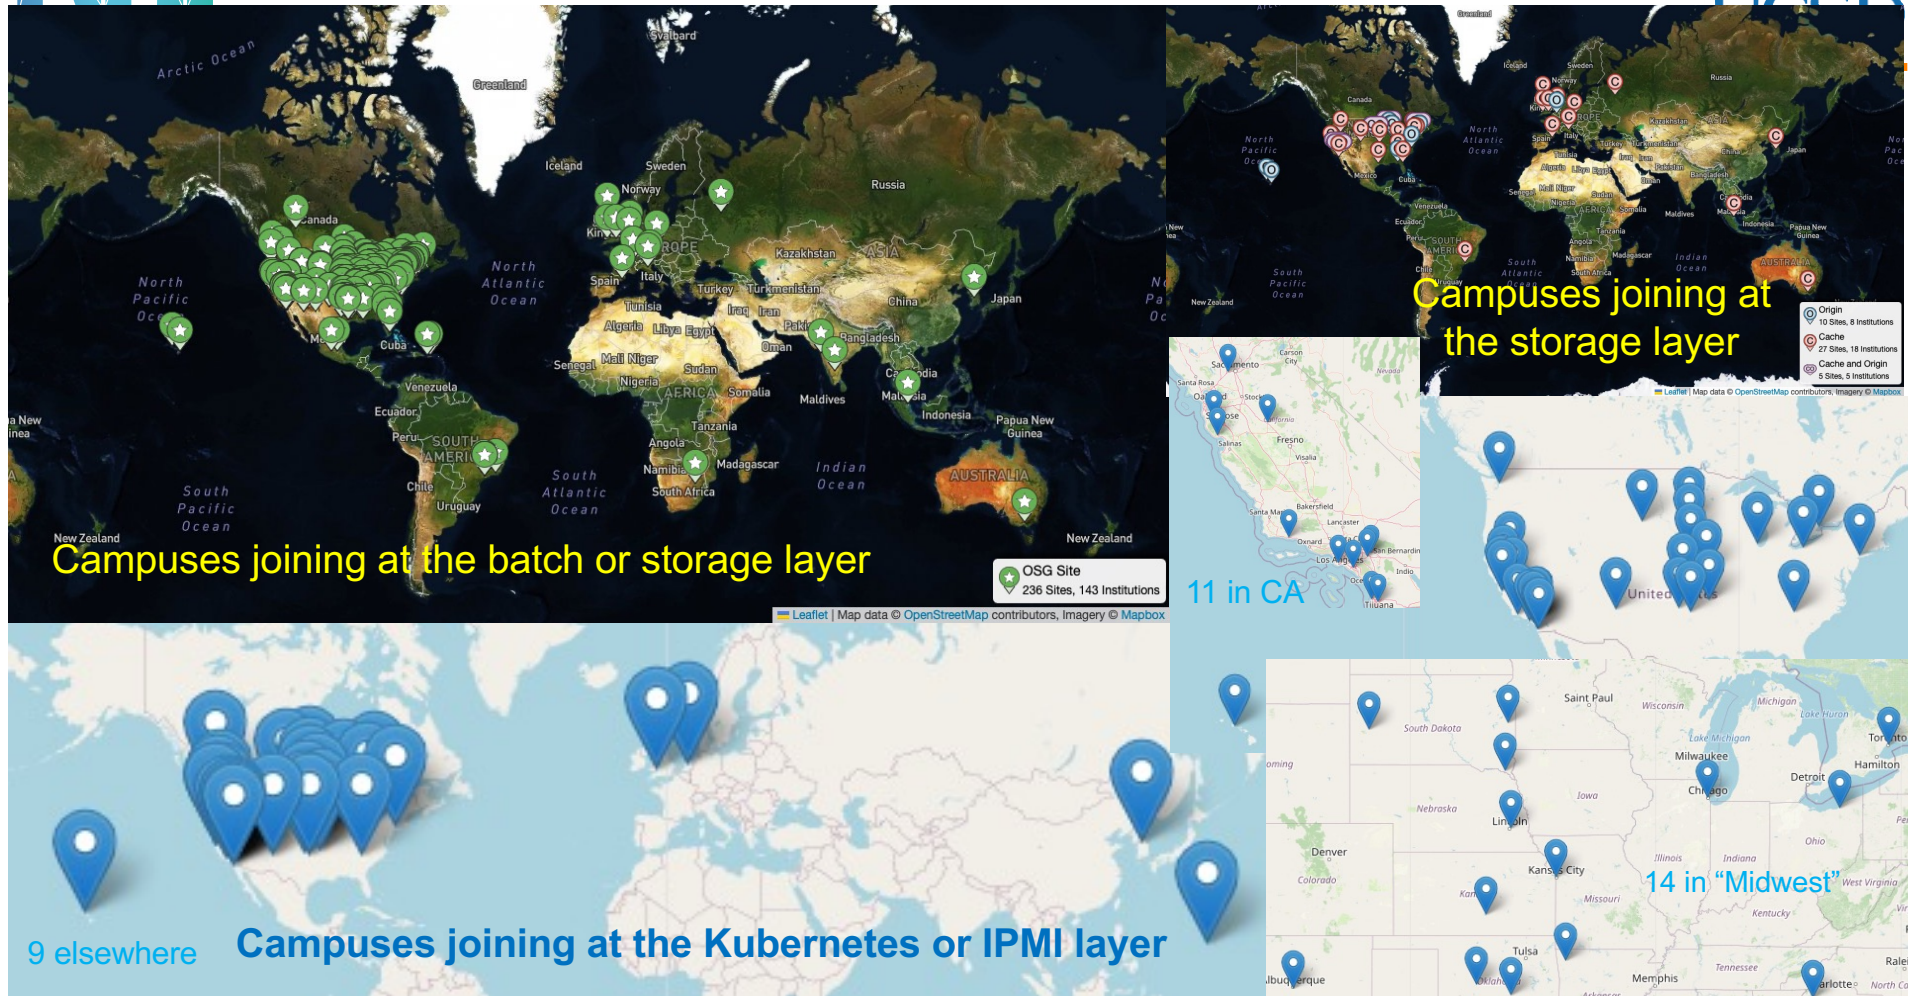
Open Science Data Federation

deployed on NRP



Data is stored at origins
and accessed via caches





**More than 150 Institutions across 5 continents
are presently integrating resources this way**

Bridging Education & Research by having them co-exist on one platform

A lot of the smaller colleges care more about educational use than research use.

All colleges struggle with applying classroom concepts to research scale problems.

There are presently half a dozen colleges that use NRP for education



UCSD's Information Technology Services has "cloned" NRP for use in education



SDSC Racked FIONAs:

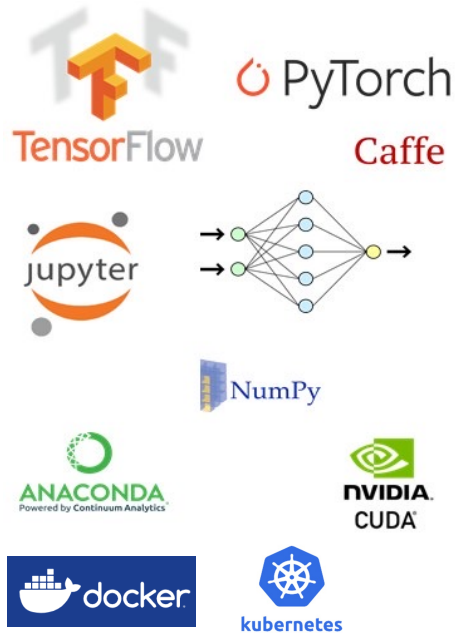
139 32-bit GPUs (~10% NRP)

2048 CPU-cores

10/25/40G Networking

Not Federated With NRP

Software Used By Students



- Student-Focused Platform For:
 - Undergraduate & Graduate Coursework
 - For-Credit Independent Study
 - Capstones & Projects
- Research-Driven Architecture
- Managed by UCSD IT Services

Source: Adam Tilghman, UCSD ITS
Data Science/ML Platform - 5/2023 Update






UCSD's DS/ML Platform Supported Nearly 60 Courses Spring Quarter 2023



#Courses by School/Division, Quarter




Source: Adam Tilghman, UCSD ITS
Data Science/ML Platform - 5/2023 Update



UCSD's DS/ML Platform Supported Over 6,000 Students Spring Quarter 2023



#Students by School/Division, Quarter



Source: Adam Tilghman, UCSD ITS
Data Science/ML Platform - 5/2023 Update



Selected Courses, Spring 2023



-
- Advanced Computer Vision
 - Bioinformatics for Immunologists
 - Computational Physics: Probabilistic Models/Sim.
 - Data Analysis/Design for Biologists
 - Data Science/Spatial Analysis
 - Deep Learning and Applications
 - Intro to Causal Inference
 - Neural Networks/Pattern Recognition
 - Numerical Analysis for Multiscale Biology
 - Robot Manipulation and Control



SDSU's VERNE Expands the UCSD Instructional Model by Federating with NRP: Visionary, Education, Research, Network, Ecosystem



CENIC



TensorFlow



PyTorch



kubernetes



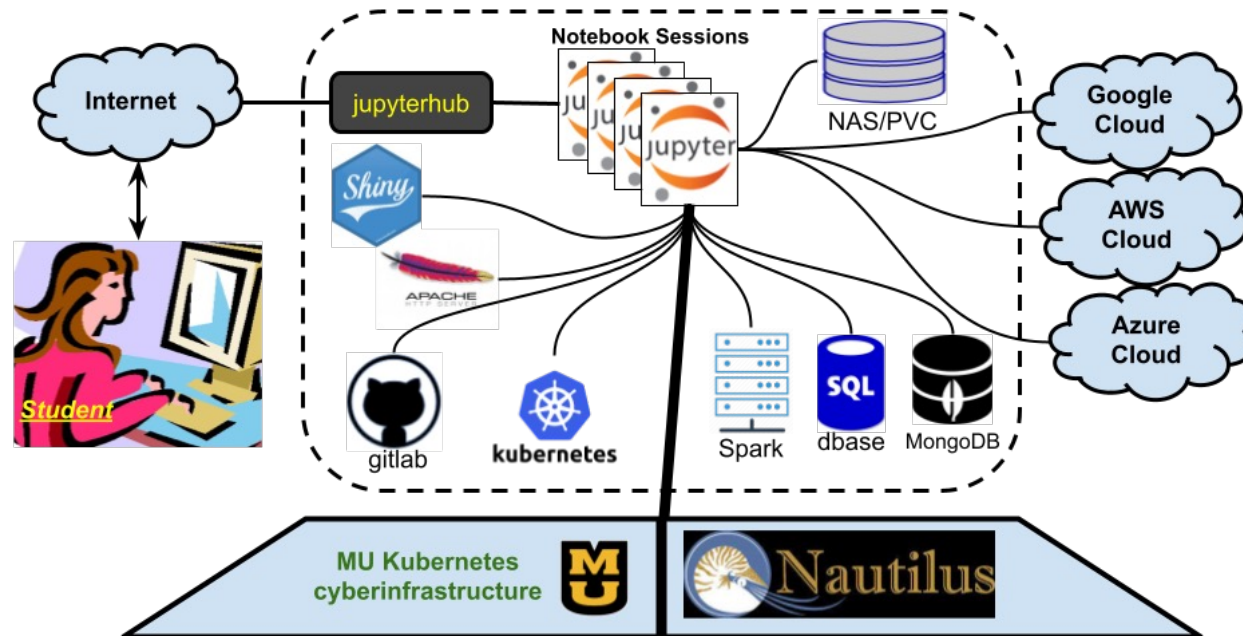
docker

- Dell PowerEdge Cluster for Instructional Use, with 15 Nodes with 32 A100 GPUs (managed as 232 independent GPUs), 768 CPU cores, and 960 TB Storage.
- A Learning Resource that *Surges* for Research and Instruction
- System Administration as a Service from the “Cat II: Prototype NRP project”
- Friendly User mode Spring 2023 with
 - ~90 Users in Computational Engineering 361 & Astronomical Techniques ASTR 350
 - Research in Developing Synthetic Medical Data for Machine Learning Training

Slide courtesy of Jerry Sheehan, SDSU

Starting to expand this to the entire CSU system:
23 colleges with 440,000 students
21 of them are Hispanic Serving Institutions

University of Missouri is Leveraging Jupyter Hubs on Nautilus for STEM & AI/ML Experiential Learning



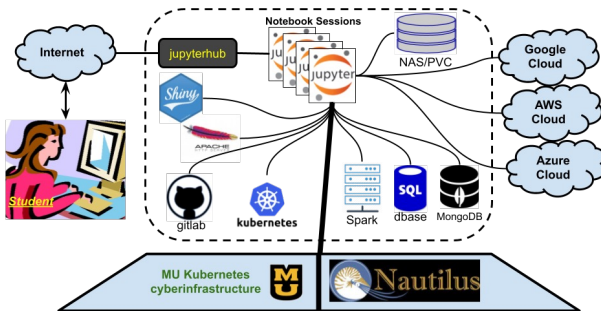
Has 16 GPUs On-Site Now -
Adding 44 with a New CC*

Grant Scott

<https://scottgs.mufaculty.umsystem.edu/>




Customized Jupyter Hubs on Nautilus Support Multiple Courses and Certificates



- Computer Science Undergraduate Classes
- Computer Science HPC Classes
- Undergraduate Data Science
- HPC Emphasis Graduate Data Science
- State & Federal Government Training Programs

Grant Scott

<https://scottgs.mufaculty.umsystem.edu/>




NRP brings CS R&D and Domain R&D onto the same platform

NRP blurs the lines between "testbed" and "production" CI

**Create social cohesion to accelerate domain science
adoption of new programming paradigms & architectures**

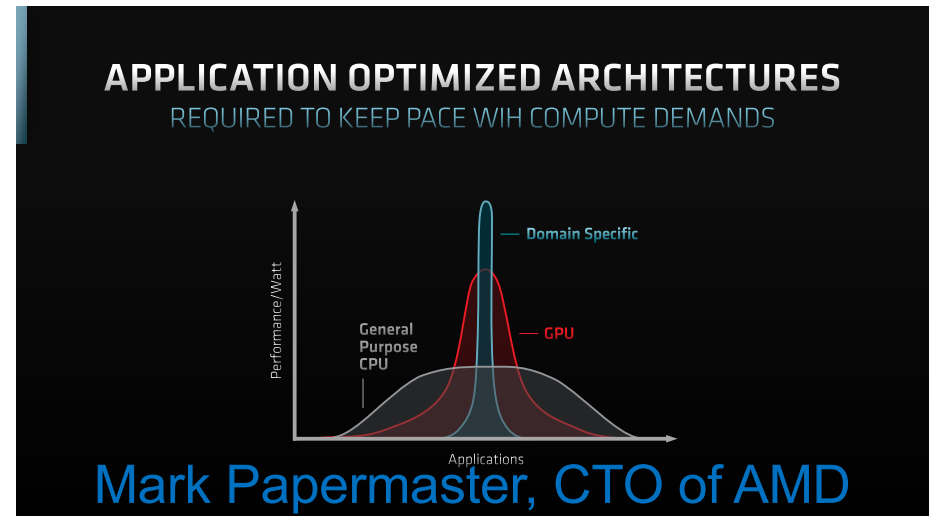
“end of Moore’s law” motivates new architectures



<https://doi.org/10.1098/rsta.2019.0061>

R&D

PI, Tajana Rosing



PRISM, a Jump 2.0 project
funded by SRC
is early user of FPGAs@NRP


21

NRP supports FPGAs (Xilinx & Intel), P4 switches, NVIDIA DPUs & HGXs


Committed to be a “Playground” of technologies, easily deployed and operated.

This is HPCs future if we continue business as usual!
... and “scale” might not be the answer...

AVERAGE PERFORMANCE IMPROVEMENT PER 11 YEARS FOR SUM OF TOP500 LIST SYSTEMS



Performance improvements vs time
slowed down by $O(100)$




- **Programmable computational capabilities emerged in devices of all kinds**
 - Storage devices with embedded FPGAs => "Computational Storage"
 - GPUs on Network Interface Cards => "Data Flow Programming"
 - Programmable switches, down to individual ports => "Programmable Networks"
- We innovate nextGen systems in ATL to solve grand challenges of science
- Innovations made available to all of open science via our Open Infrastructure

Idea

Strategic Objective is to bring CS Research closer to Domain Research in the hope of decreasing time to adoption of new technologies & ideas

R&D


NVIDIA BlueField DPU



P4 programmable switches




Xilinx SmartSSD



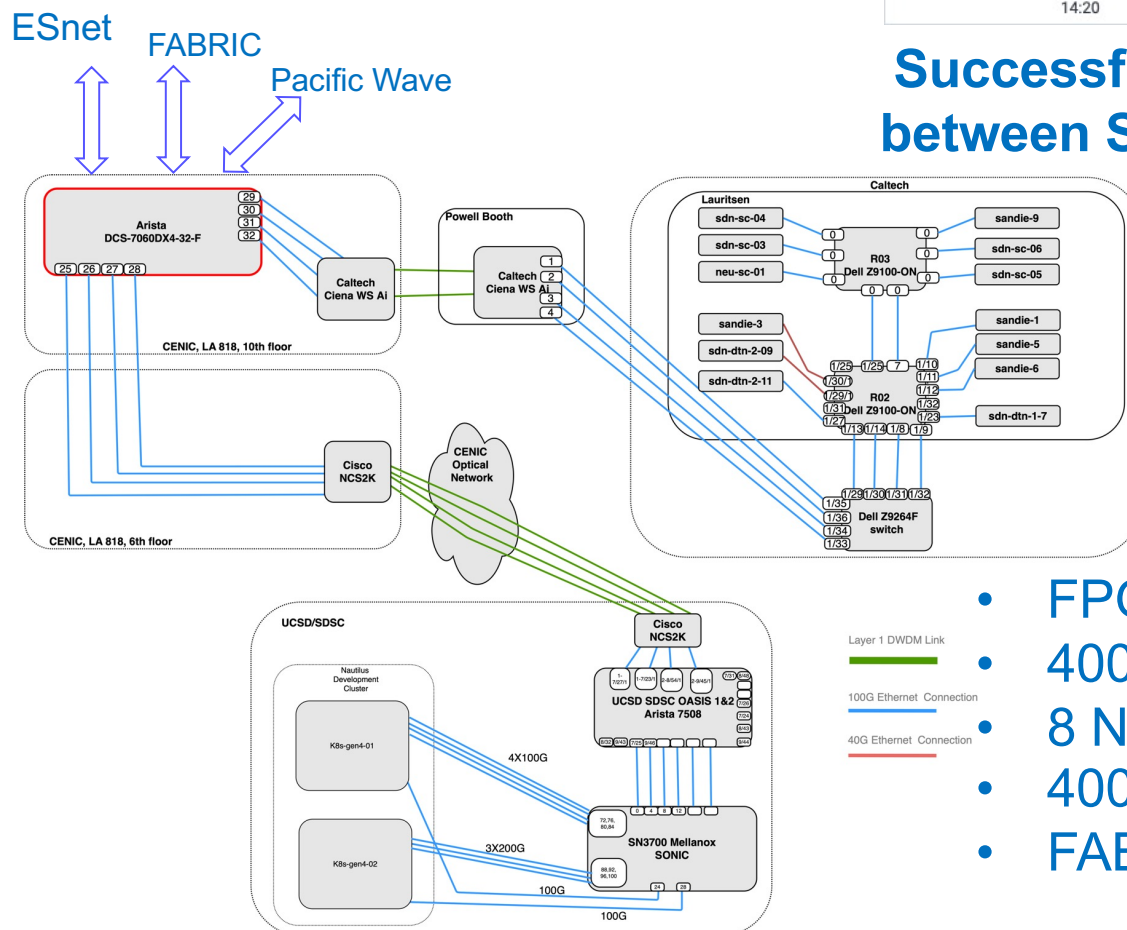
400G WAN Infrastructure

Thank you to CENIC

We can peer at 400G in LA
with multiple networks
via our 400G Arista switch



Successfully sustained 400G transfers
between SDSC & Caltech using XRootD



Infrastructure at SDSC:

- FPGAs: 32 U55C, 24 Bitware 520
- 400G P4 programmable switches
- 8 NVIDIA HGX w 8 80G A100s each
- 400TB of NVMe
- FABRIC node

- NRP has a very ambitious vision
 - Horizontally open
 - Today about 3x # of GPUs total than what was part of Cat-II PNRP award
 - PNRP award started testbed operations phase on 3/27/23
 - Vertically Open
 - We have built the “Open Science Data Federation” on top of NRP,
 - ... and are starting to build “Fusion Data Platform for AI” on top of NRP
 - ... and are starting to build elements of the National Discovery Cloud for Climate on top of NRP (Pelican, National Data Platform, NCAR integration, ...)
 - “Playground” for CS R&D on the same platform as “Production” system for Domain Scientists
- We are off to an excellent start ... but there is lot's more to come over the course of the next 5 years.

Acknowledgements

- This work was partially supported by the NSF grants OAC-1541349, OAC-1826967, OAC-2030508, OAC-1841530, OAC-2005369, OAC-21121167, CISE-1713149, CISE-2100237, CISE-2120019, OAC-2112167

