

Open Infrastructure for an Open Society

Frank Würthwein
Director, San Diego Supercomputer Center

Lustre User Group Meeting

May 3rd 2023



Vision




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

- 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

National Research Platform (NRP)



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

We want you to ...
... grow NRP.
... build on NRP.


NRP is “owned” and “built” by the community for the community

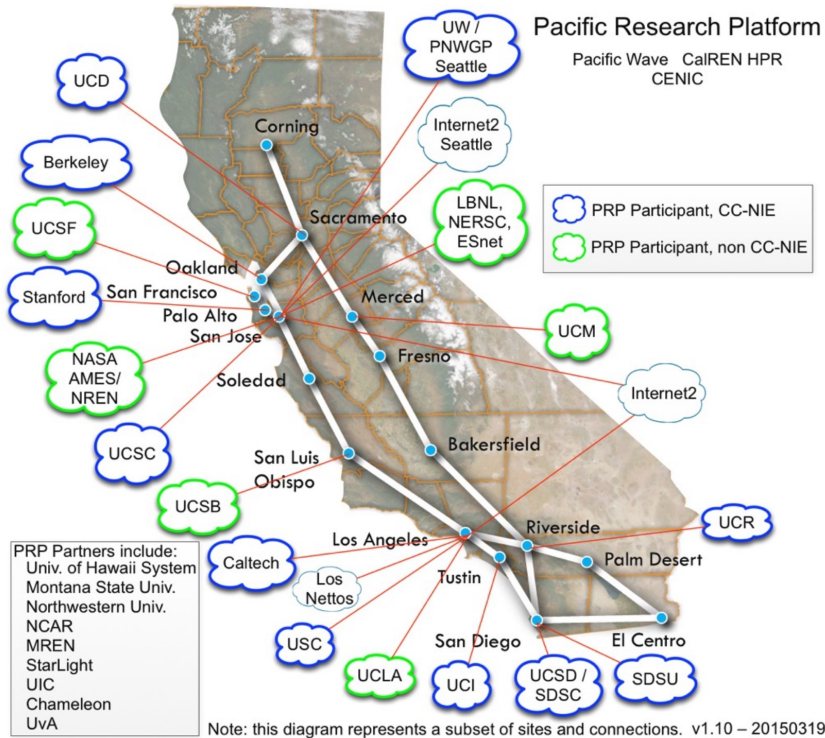
NRP operates at all layers of the stack, from IPMI up

- IPMI reduces TCO and lower threshold to entry
- Kubernetes allows service deployments
 - Also the natural layer for application container deployment
- Admiralty allows K8S federation with folks who want control
 - Including cloud integration to access TPUs & other cloud only architectures
- HTCondor allows NRP to show up as a “site” in OSG

The layer you integrate at depends on

- Control you want
- Effort you can afford





Pacific Research Platform

A distributed Kubernetes cluster
(2015 – 2023)

- **When the PRP Grant Was Funded in 2015, It Started With:**
 - 6 States ➡ Now 43 States
 - 19 Campuses ➡ Now 110 Campuses
 - 9 Minority Serving Institutions ➡ Now 23 MSIs
 - 2 NSF EPCoR States ➡ Now 19 EPSCoR States, 2 Territories, and Wash DC

Expanded Diversity & Inclusion Nautilus Campus Users Across the Country: West

California

- California Institute of Technology
- California State Fullerton (MSI)
- California State Monterey Bay (MSI)
- California State Polytechnic U., Pomona (MSI)
- California State U. San Bernardino (MSI)
- California State U., Northridge (MSI)
- Harvey Mudd College
- Naval Postgraduate School
- San Diego State U. (MSI)
- San Jose State U. (MSI)
- Stanford U.
- U. of California, Berkeley
- U. of California, Merced (MSI)
- U. of California, Davis (MSI)
- U. of California, Riverside (MSI)
- U. of California, San Francisco
- U. of California, Santa Cruz (MSI)
- U. of California-Irvine (MSI)
- U. of California-Los Angeles
- U. of California-San Diego
- U. of California-Santa Barbara (MSI)
- U. of Southern California

Hawaii (EPSCoR)

- U. of Hawaii (MSI)

Guam (EPSCoR)

- U. of Guam (MSI)

Oregon

- Oregon State U.
- Willamette U.

Washington

- U. of Washington
- Washington State U.

Montana (EPSCoR)

- Montana State U-Bozeman

Nevada (EPSCoR)

- U. of Nevada, Reno

Wyoming (EPSCoR)

- U. of Wyoming

Colorado

- Colorado School of Mines
- NCAR
- U. of Colorado at Boulder
- U. of Colorado Denver

Utah

- Brigham Young U.

New Mexico (EPSCoR)

- U. of New Mexico

Arizona

- U. of Arizona

**Minority Serving Institutions
NSF EPSCoR States**

Michigan

- Michigan State U.
- U. of Michigan

Wisconsin

- U. of Wisconsin-Madison

Missouri

- Missouri Western State U
- Truman State U
- U. of Missouri System
- Washington U. in St. Louis

Iowa (EPSCoR)

- Iowa State U.

Kansas (EPSCoR)

- Kansas State U.
- U. of Kansas
- Wichita State U

Nebraska (EPSCoR)

- U. of Nebraska at Kearney
- U. of Nebraska-Lincoln

Illinois

- Northwestern University
- U. of Chicago
- **U. of Illinois at Chicago (MSI)**

Indiana

- Purdue U. Main Campus
- U. of Notre Dame

Kentucky (EPSCoR)

- U. of Kentucky

North Dakota (EPSCoR)

- North Dakota State U.

Ohio

- Case Western Reserve U.
- Kent State University
- U. of Akron
- U. of Cincinnati Main Campus

**Minority Serving Institutions
NSF EPSCoR States**

Tennessee

- Vanderbilt U.

South Dakota (EPSCoR)

- Black Hills State U.
- South Dakota School of Mines & Tech.
- South Dakota State U.
- U. of South Dakota

Oklahoma (EPSCoR)

- Oklahoma State U. System
- Southwestern Oklahoma State U.
- U. of Oklahoma

Arkansas (EPSCoR)

- U. of Arkansas
- U. of Arkansas for Medical Sciences

Louisiana (EPSCoR)

- Louisiana Tech U.
- Southeastern Louisiana U.

Texas

- Southern Methodist U.
- U. of Texas at Austin
- U. of Texas at Dallas

Maine (EPSCoR)

- Colby College

New York

- Columbia U.
- New York U.
- Manhattan College
- Rensselaer Polytechnic Institute
- Rochester Institute of Technology
- Syracuse U.
- The State U. of New York at Buffalo

Connecticut

- U. of Connecticut
- Yale U.

Delaware (EPSCoR)

- U. of Delaware

Massachusetts

- MIT
- Northeastern U.

Rhode Island (EPSCoR)

- Brown U.
- U. of Rhode Island

Pennsylvania

- Carnegie Mellon U.

New Jersey

- Rutgers U.

Maryland

- Johns Hopkins U.
- Morgan State U. (MSI)

West Virginia (EPSCoR)

- West Virginia U.

Virginia

- Virginia Commonwealth U

District of Columbia

- American U.

North Carolina

- N. Carolina Agricultural & Technical State U. (MSI)
- N. Carolina State U
- U. of North Carolina at Chapel Hill

South Carolina (EPSCoR)

- Clemson U.

Georgia

- Georgia Institute of Technology
- Georgia State U. (MSI)

Florida

- Florida A&M U. (HBCU/MSI)
- Florida International U. (MSI)
- Florida State U.
- U. of Central Florida (MSI)
- U. of Miami

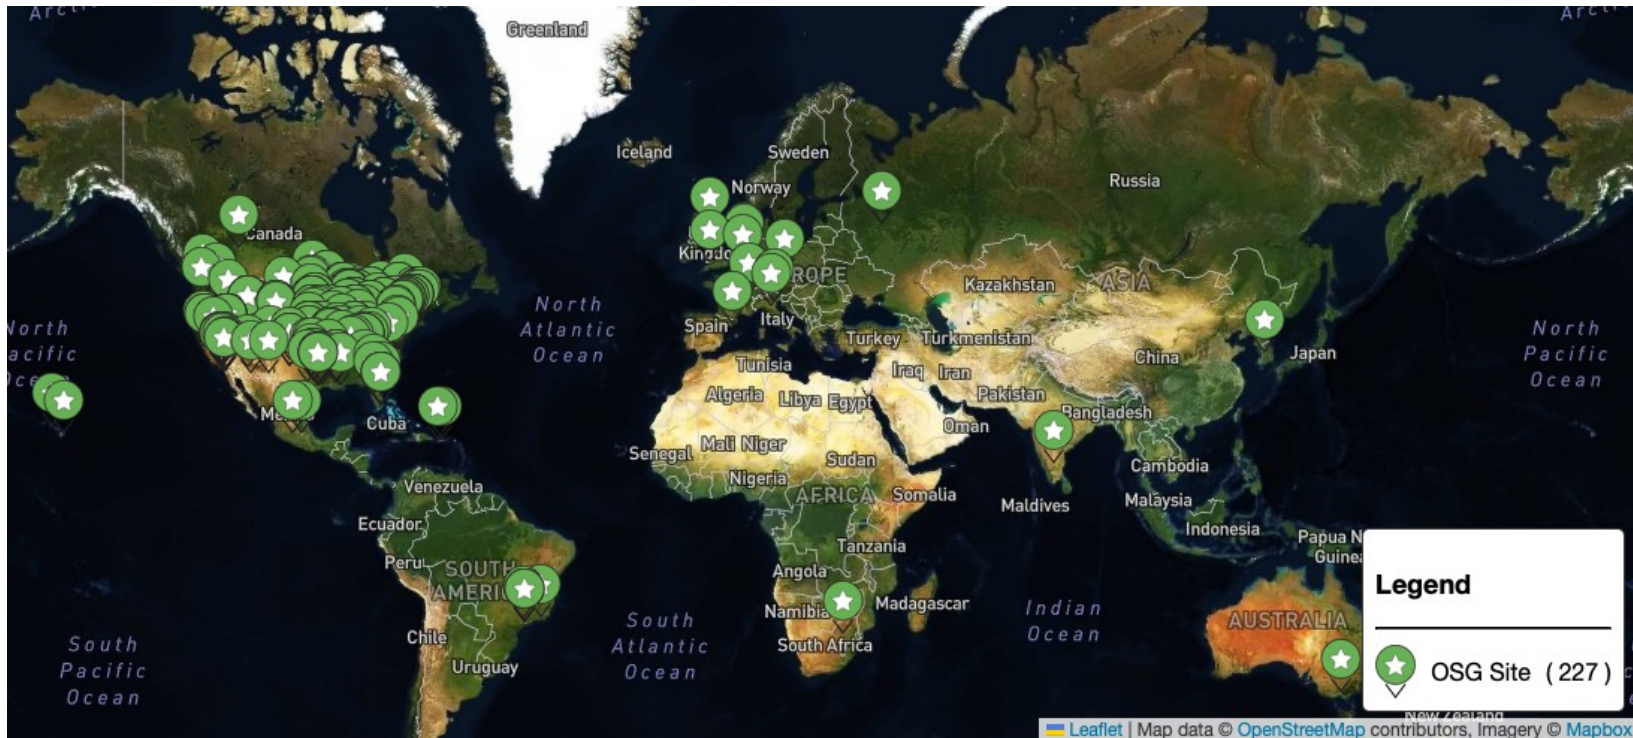
Puerto Rico (EPSCoR)

- Universidad de Puerto Rico (MSI)

**Minority Serving Institutions
NSF EPSCoR States**

Open Science Grid

A Federation of Compute and Storage Clusters
(founded in 2005)



Complementarity in Implementation of “Bring Your Own Resource” model

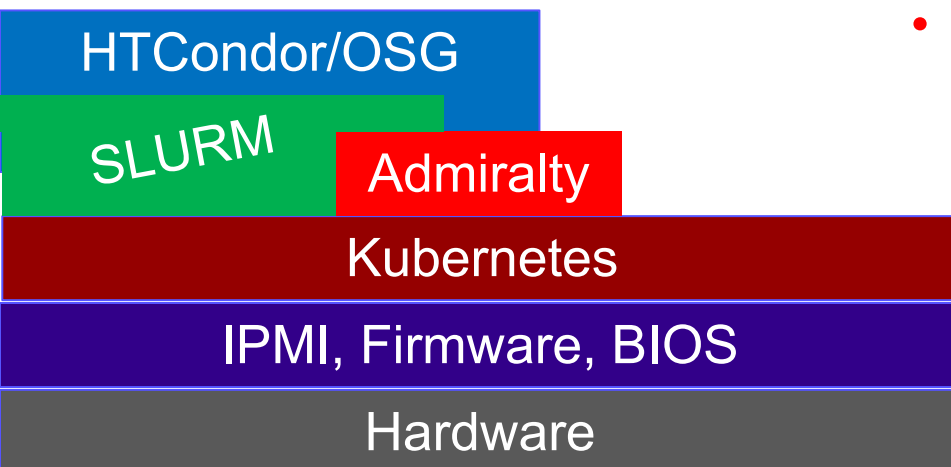
OSG/PATh focused on campus cluster integration.

NRP focused on individual node integration instead of clusters.

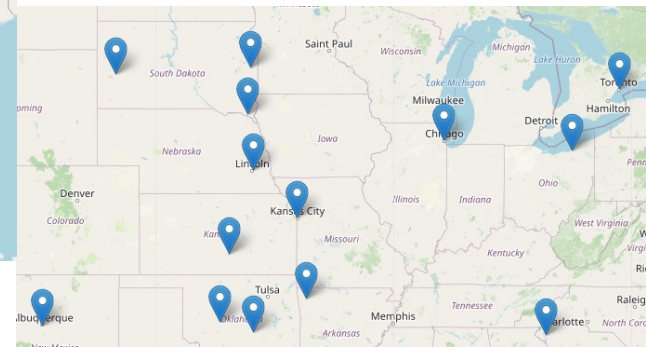
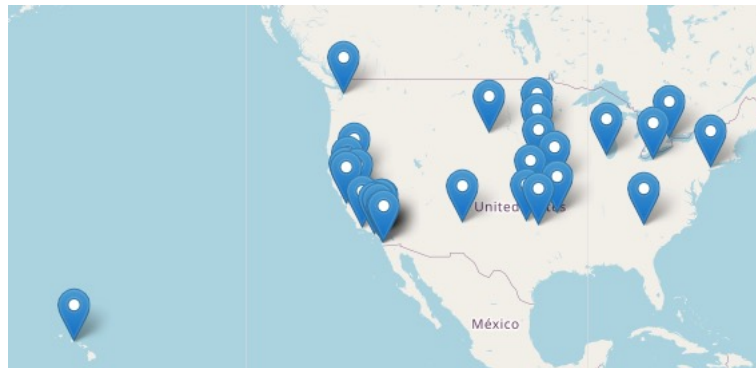

NRP operates at all layers of the stack, from IPMI up

- **IPMI reduces TCO and lowers threshold to entry**
- Kubernetes allows service deployments
 - Also the natural layer for application container deployment
- Admiralty allows K8S federation with folks who want control
 - Including cloud integration to access TPUs & other cloud only architectures
- HTCondor allows NRP to show up as site in OSG

- **Under-resourced institutions**
- **Network providers and their POPs**
- **CS & ECE faculty specialized on:**
 - **AI/ML => gaming GPUs**
 - **systems R&D**



All of these find it difficult to justify staff to support all layers




NRP operates at all layers of the stack, from IPMI up

- IPMI reduces TCO and lowers threshold to entry
- **Kubernetes allows service deployments**
 - Also the natural layer for application container deployment
- Admiralty allows K8S federation with folks who want control
 - Including cloud integration to access TPUs & other cloud only architectures
- HTCondor allows NRP to show up as site in OSG

NRP is unique in its support of global service deployments

- **Open Science Data Federation**
 - Origins & Caches in US, EU, Asia
- **Protein Data Bank**
 - (Future) Replicas in EU & Asia



Supporting Nautilus for the next decade

Nautilus = K8S infrastructure of PRP for the last 5+ years

Nautilus = K8S of NRP for the next 10 years

- NSF supports via the Cat-II program novel systems ideas.
 - 3 year “testbed” phase
 - The PI owns the resource, and has (some) freedom regarding who uses it.
 - No requirements for making it available via any specific allocation mechanism.
 - It is expected that not all features work on day 1.
 - 3 years of experimentation & development of features
 - 2 year “allocation” phase
 - The resource is made available via an NSF supported allocation mechanism.
 - **The solicitation mentions the possibility of an additional 5 year renewal without re-competition if system is successful.**
- **We decided that this is an ideal program to try and secure NRP core operations funding for the next 10 years**
 - And thus provide the stability necessary for growth of NRP.

NATIONAL RESEARCH PLATFORM

Designed for Growth & Inclusion

HPC/HTC Resource

32 ALVEO FPGAs

A10 288 NVIDIA FP32 GPUs

80GB A100 64 NVIDIA FP64 GPUs

Tbps WAN IO Capabilities

Configurable Low Latency HPC Fabric

Massachusetts Green HPC Center

Data Intensive S&E

Life Sciences

Physical Sciences

Systems Engineering

Disaster Response

Multi-Messenger Astrophysics

U Nebraska, Lincoln

SDSC, UC-San Diego

Distributed Data Infrastructure

National Scale Content Delivery Network

50TB 100Gbps NVMe Caches in 8 locations

4.5PB Distributed Data Origin across 3 Sites

Composable & Scalable Innovation

Open to Campus Resource Integration

Open Community Support Model

Campus-Scale Instrument integration

BYOR & BYOD

Any Data, Anytime, Anywhere

5 year project with \$5M hardware & \$6.45M people

Supports Nautilus, and thus the core NRP infrastructure

Promises to build on “PRP” functionality, and go beyond
NSF Acceptance Review of System passed on March 9th 2023

PI = Wuerthwein; Co-PIs: DeFanti, Rosing, Tatineni, Weitzel

- I1: Innovative network fabric that allows “rack” of hardware to behave like a single “node” connected via PCIe.
- I2: Innovative application libraries to expose FPGAs hardware to science apps at language constructs scientists understand (C, C++ rather than firmware)
- I3: A “Bring Your Own Resource” model that allows campuses nationwide to join their resources to the system.
- I4: Innovative scheduling to support urgent computing, including interactive via Jupyter.
- I5: Innovative Data Infrastructure, including national scale Content Delivery System like YouTube for science.

**I3 & I4 & I5 turn “PRP” into “NRP” and sustains it into the future.
I1 & I2 are totally new.**

- Support regional Ceph storage systems across the USA.
 - Campuses can join individual storage hosts to the Ceph system in their region.
 - All regional storage systems are Origins in OSG Data Federation (OSDF)
 - **Deploy replication system such that researchers can decide what part of their namespace should be in which regional storage.**
- Deploy caches in Internet2 backbone such that no campus nationwide is more than 500 miles from a cache.

NRP data infrastructure model combines best of PRP & OSG

From PRP we take the regional Ceph storage concept
From OSG/PATH we take the data origin & caching concepts

And then we add as a totally new feature:

User controlled replication of partial namespaces across regions.

(We will develop this during 3 year “testbed” phase)

Want Others to build higher level data services on top

Table 3.1 Representative Science and Engineering Use Cases

Application domain	Lead researcher & Institution	Science Driver Themes	NRP Innovations
LIGO	Peter Couvares, LIGO Lab; Erik Katsavounidis, MIT	BGS, UC, AI	I2, I3, I4, I5
IceCube	Benedikt Riedel, UW Madison	BGS, UC, AI	I3, I4
Astronomy (DKIST & Sky Surveys)	Curt Dodds, U. Hawai'i	BGS, AI	I3, I5,
Campus Scale Instrument Facilities	Mark Ellisman, NCMIR; Samara Reck-Peterson, Nikon Imaging Center; Johannes Schoeneberg, Adaptive Optics Lightsheet Microscopy; Kristen Jepsen, Institute for Genomic Medicine; Tami Brown-Brandl, Precision Animal Management	SD, UC, H	I1, I2, I3, I4, I5
Molecular Dynamics	Rommie Amaro, UCSD; Andreas Goetz, SDSC; Jonathan Allen, LLNL	MD, AI, H	I1, I2, I3
Human microbiome	Rob Knight, UCSD	G, AI, H	I1, I2, I3
Genomics & Bioinformatics	Alex Feltus, Clemson	G, AI, H	I3, I4, I5
Fluid Dynamics	Rose Yu, UCSD	AI	I1, I2, I3
Experimental Particle Physics, IAIFI	Phil Harris, MIT	AI, BGS, SD	I1, I2
Computer Vision	Nuno Vasconcelos, UCSD	AI, CV	I3
Computer Graphics	Robert Twomey, UNL	CV, AI	I3
Programmable Storage	Carlos Malzahn, UCSC	SD	I1, I2, I5
AI systems software stack for FPGAs	Hadi Esmaeilzadeh, UCSD	SD	I1, I2
WildFire Analysis & Prediction	Ilkay Altintas, UCSD	UC, AI, CV	I3, I4

Lot's of AI ...
but so much more ...

NSF MREFCs

Incl. 4 campus scale instrument facilities

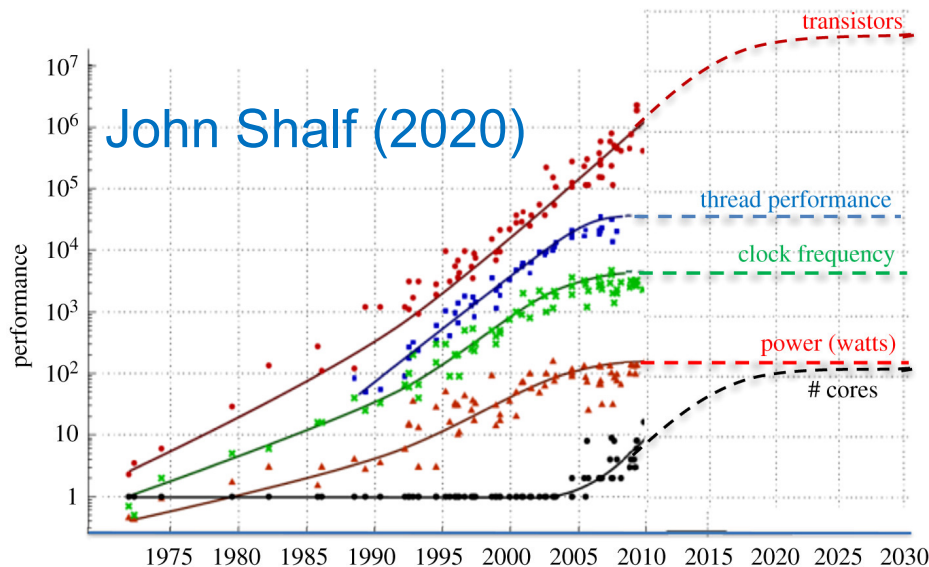
Incl. a very diverse set of sciences and engineering

Key: The NRP innovations column lists those innovations among I1 through I5 listed in Section 2.1 that a given science driver most benefits from.

- **Growth of NRP infrastructure**
 - 1,000++ GPUs end of 2022
 - 50 PB storage end of 2024
 - Growth in diversity of community
 - # and types of campuses and their researchers
- **Introduce new capabilities to NRP**
 - Machine learning at 100TB scale
 - Support Domain Specific Architecture R&D on NRP
 - Expand NRP into Wireless, Edge, IoT
 - Towards “FAIR” on OSDF
- **New Directions initiated by the Community**

- I1: Innovative network fabric allowing “composable hardware”.
- I2: Innovative application libraries allowing “domain optimized architectures” on FPGAs

“end of Moore’s law” motivates new architectures



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

PI, Tajana Rosing

APPLICATION OPTIMIZED ARCHITECTURES
REQUIRED TO KEEP PACE WITH COMPUTE DEMANDS

Mark Papermaster, CTO of AMD

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

- The NSF CC* 2022 program awarded 9 campuses with \$500k storage awards each.
 - We guess this pays for 5PB of storage each.
- Some of these campuses may decide to integrate their CC* storage into the OSDF.
- Some of these campuses have storage from other projects that they may integrate with the OSDF in addition.
- **NSF 23-523 includes \$500k storage solicitation again, Spring & Fall 2023.**

- **PRP ended, and was replaced by NRP**
 - Significant new capabilities via Cat-II system “PNRP”
 - PNRP provides ops effort for Nautilus for the future
 - # of GPUs available double in 2022.
 - new GPUs (A10, 3080, 3090, A100) much more powerful than older GPUs
 - # of FPGAs increase from a few to a few dozen in 2022.
 - # of caches grow by 50% in 22/23
 - => more consistent coverage across USA
 - Data volume served expected to grow substantially in 23/24/25.
 - How much? As yet too hard to predict.
- Hoping to recruit new partners to build **FAIR capabilities on top of OSDF within the next 5 years.**
- Hoping to expand NRP into **sensor networks using 5G & 6G in the next 10 years.**

- 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

