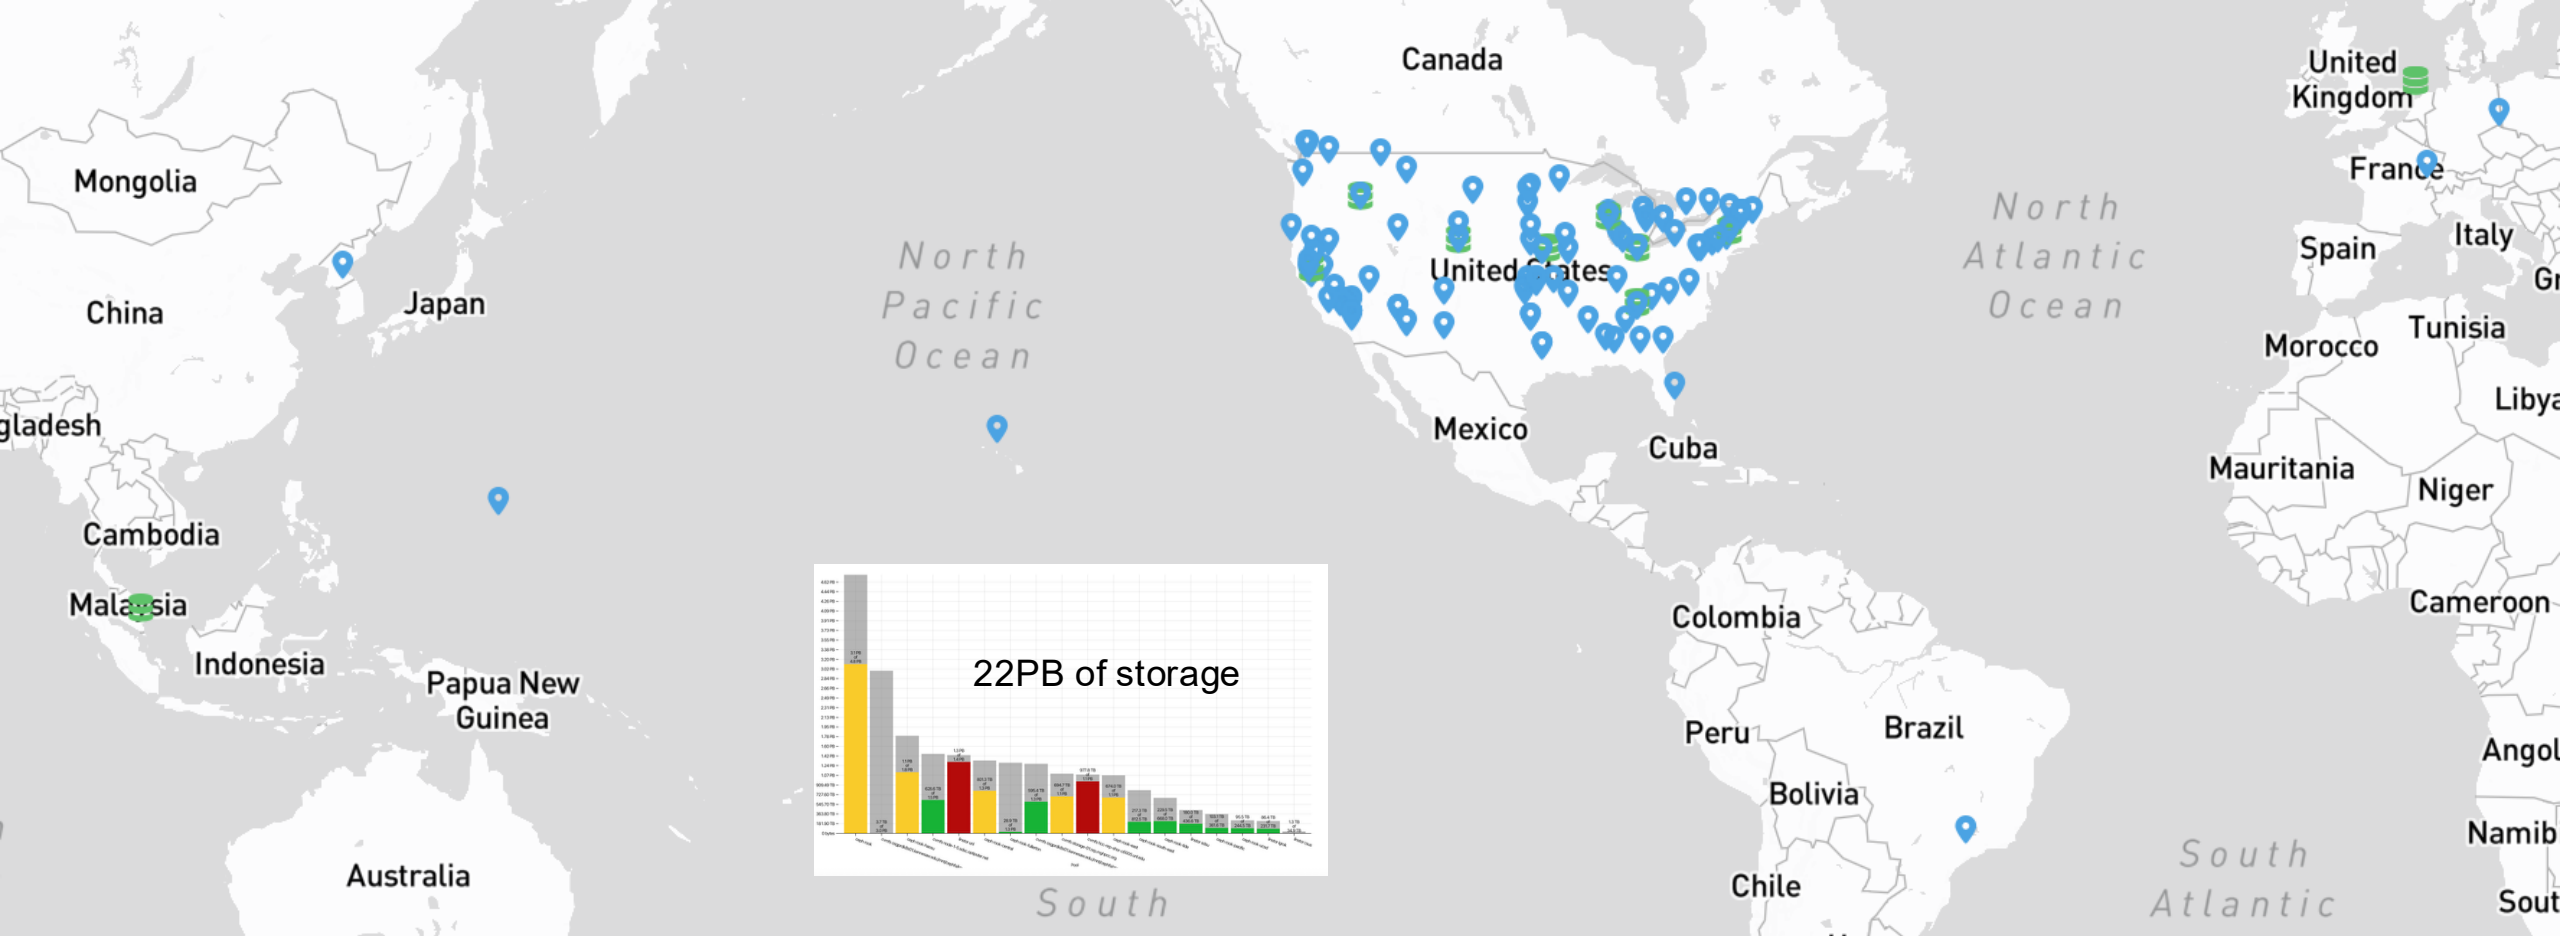



# State of NRP

**Frank Wuerthwein**  
**Director**  
**San Diego Supercomputer Center**

**National Research Platform,  
a community owned cyberinfrastructure for  
Research, Education, and public-private partnerships**




SITES  
**129**  
 Hosting NRP nodes




NODES  
**501**  
 Registered in Kubernetes




GPUS  
**1.5K**  
 Across all nodes



CPU CORES  
**32.3K**  
 Across all nodes



OSDF NODES  
**31**  
 Nodes serving OSDF cache  
[View OSDF Dashboard →](#)



**NRP is globally distributed, with the dominant deployments in the continental USA**



The NRP Team:



UNIVERSITY of NEBRASKA  
LINCOLN



© Mapbox © OpenStreetMap Improve this map

# NRP...

## ... enabling colleges to own AI infrastructure ...

- **Your REN provides network planning, engineering, operations.**
- **We run your hardware from IPMI up**
  - OS maintenance, security monitoring, ...
- **Researchers, Educators, and students see a global scale Kubernetes cluster**
  - While the cluster is shared, we restrict use of individual hardware to owners when necessary
  - Documentation and training for the community
- **Lot's of software to reuse**
  - E.g., we operate JupyterHub for the community & show people how to deploy & customize their own
    - 151 community-run JupyterHubs now on NRP
- **Maintain lots of topical Chat channels for the community to learn from each other and interact with each other**
  - Chat channels monitored by professionals, community is encouraged to interact & help each other
  - Starting to explore AI chatbots trained on the chats
- **LLM as a service**
  - We run most popular LLMs as services for all
  - ... and provide popular APIs to upload any AI model from Hugging Face onto our resources. ← **More on LLM on Thursday in talk by Mohammad**
- **AI workflows via NDP**
  - Reusable workflows that incorporate curated data & AI into classroom exercises ← **More on NDP on Thursday in talk by Ilkay**

# Institutional Resource Commitments to NRP

- 70 colleges
  - **34 Non-R1 Institutions host hardware in NRP**
    - Community Colleges:
      - Brookdale Community College, NJ
      - Los Angeles Community College District, CA
      - San Diego Community College District, CA
    - 36 R1 Institutions
  - **14 Research and Education Networks**
  - 3 National Laboratories
  - 1 Museum

-----  
**88 Institutions in total**  
**129 Geographic locations**

**NRP hardware is owned  
by the NRP community**

**Any institution can own hardware in NRP**

**We do all the operations & support,  
and can even find you a data center  
to place your hardware in.**

# NRP User Community

- NRP had 5,833 users in the last 12 months
  - 5,275 Jupyter users across 151 JupyterHubs
  - 2,851 Kubernetes users
  - 73% of users from **54 R1**
  - 25% of users from **97 non-R1 US colleges**
  - 2% of users from national labs/RENs/Museum

**90% of NRP users use JupyterHubs**

**There are more non-R1 than R1 colleges among the NRP user community**

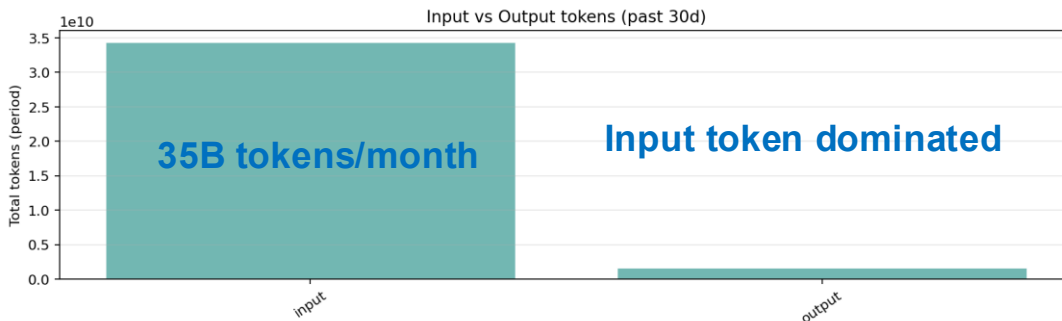
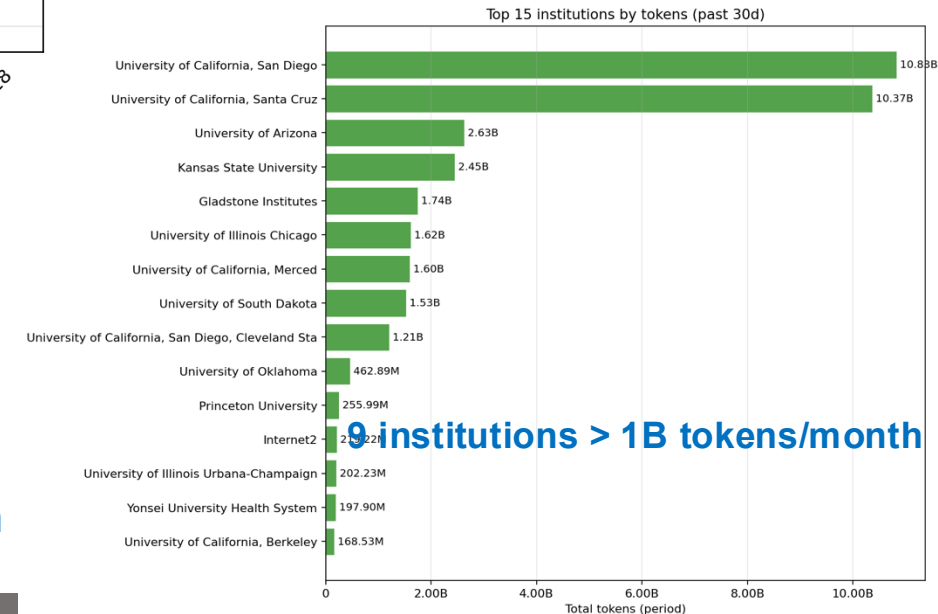
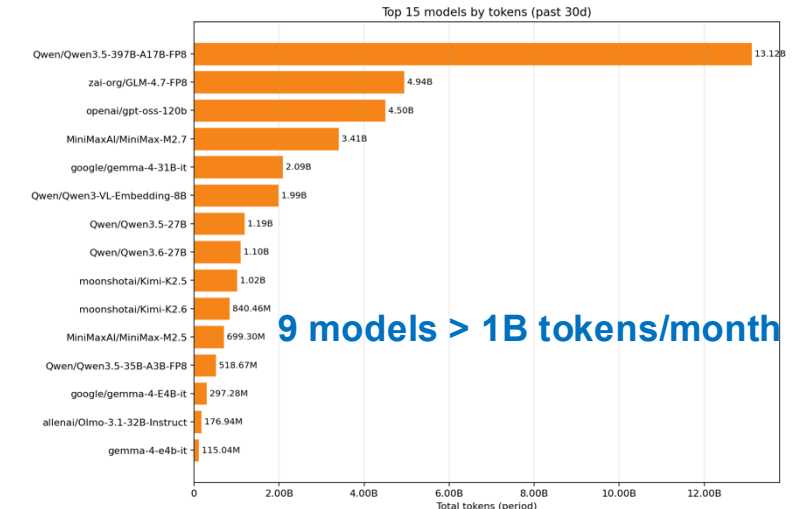
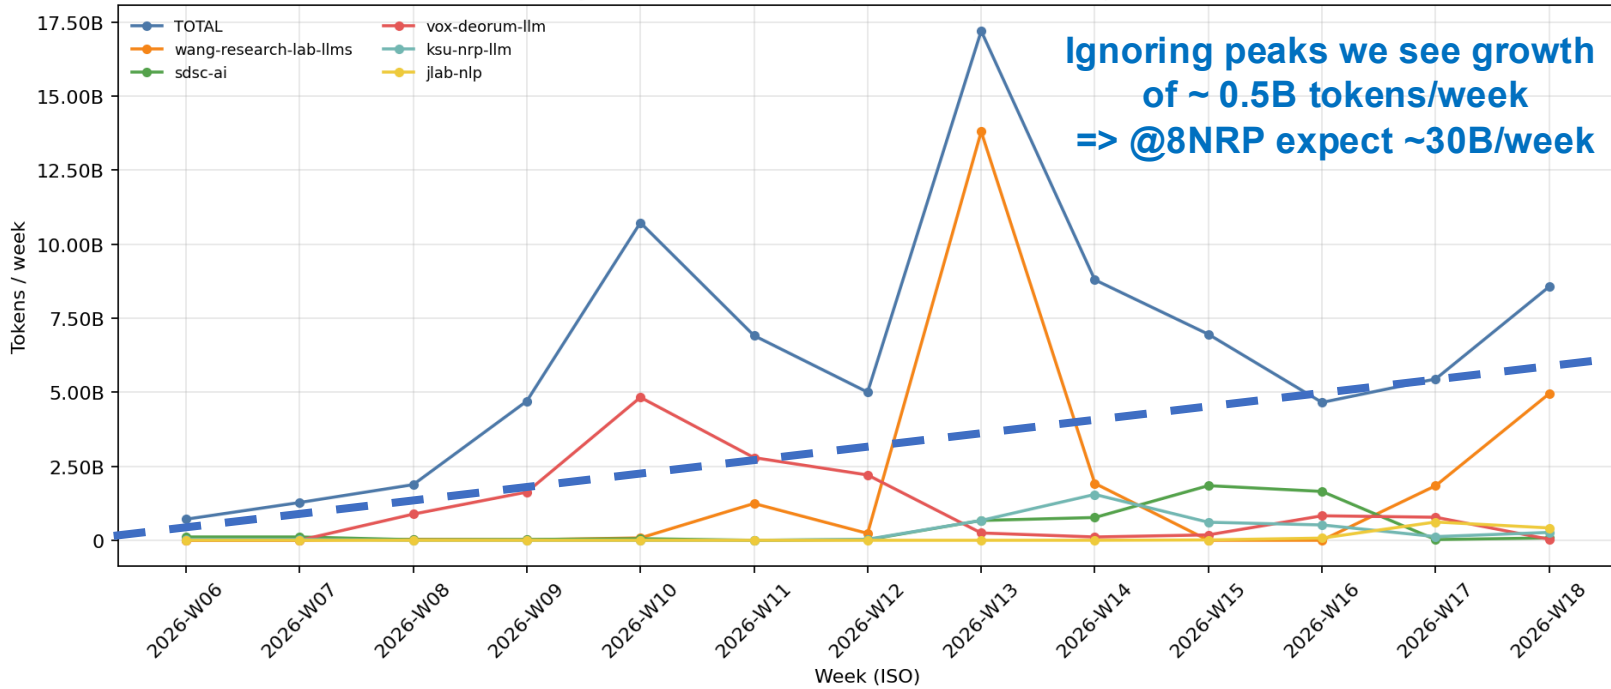
- Since February 2026, we measure LLM as a service use on NRP:
- NRP served LLM services to 751 users across 276 projects this year

**8.4% of users use LLM**

**16.7% of all project use LLM**

# AI Inference as a Service on NRP

Weekly token volume - cluster total + top 5 teams (past 120d)



**751 users**  
**276 projects**  
**35 institutions**

**14 projects each used > 1B tokens/month**

# Let's put NRP in perspective ...

- NSF ACCESS program has 12,000 users per year
- NERSC serves 11,000 researchers per year
- NAIRR has 710 research projects supported since 1/2024
- NRP had 5,833 users across more than 8,000 projects

**NRP user community is about  
1/2 the size of the NSF ACCESS community**

NSF ACCESS is the primary NSF allocation mechanism for all of its funded production systems.

# Digital Assets in the Classroom

# The beginning of a revolution in education

- Compute, Data & AI redefine what's possible to teach in a classroom
    - Programming at the fingertips of every student
    - Simulations, data analytics, and visualizations on any topic
- ⇒ **Students transition from consuming to experiencing education**
- AI makes personalized tutoring before, during, and after the classroom session possible.
- ⇒ **More students understand more as a result of personalized attention, even in classes with hundreds or thousands of students**

# Edutainment ... the norm in College today



The Educator needs to be an entertainer to retain the attention of the class.

**Students are consumers of power point presentations, or worse, endless math scribbled on blackboards.**

**Majority of humans learn better when they do rather than watch.**

Jupyter notebooks allow interactive engagement with the curriculum, in realtime in the classroom.

**Key: With Jupyter, Students *do* instead of *watching* the instructor perform.**

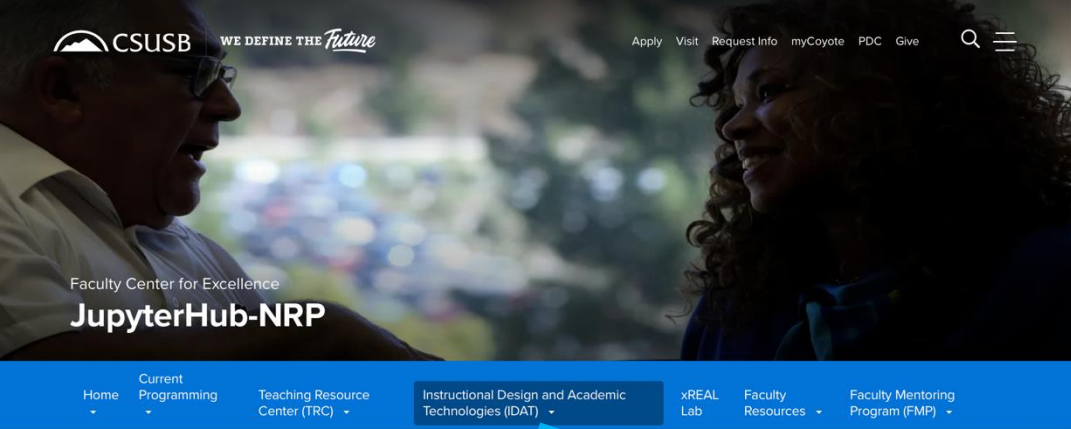
# Jupyter on NRP

- 90% of the NRP user community uses Jupyter
- We operate 151 JupyterHubs
  - A couple we run for the community
  - The other 149 Hubs are run by individual institutions for their communities
    - We provide a template ...
    - ... and some instructions for deploying and operating

=> [See tutorial yesterday afternoon](#)

2:00 - 3:00 PM	Setting Up Custom JupyterHubs for Classroom and Research <i>Mohammad Firas Sada, UCSD/SDSC</i> <i>Daniel Diaz, UCSD/SDSC</i>
----------------	--

**We want to support your institutional brand on your JupyterHub**



# Faculty Center for Excellence JupyterHub-NRP

### CSUSB JupyterHub

- Accessible to anyone with a CoyoteID
- Each user is provided with 50 GB of dedicated storage, expandable upon request
- Resources include large RAM and storage and multiple CPUs and GPUs
- Pre-installed software includes Linux desktop, Python, Matlab, R, Scipy, Pyspark, Tensorflow, a data science package, and SageMath
- Other software can be installed upon request
- Can be used for research and classes

### CSUSB JupyterHub for faculty research

- Dedicated to faculty research
- Contact Dung Vu to access this server

### National Research Platform

- National Research Platform (NRP) Nautilus Cluster

### Contact Us

**Faculty Center for Excellence**

✉ fce@csusb.edu  
☎ (909) 537-7684  
📍 PL-4005

**Office Hours**  
Monday - Friday: 8:00 am-5:00 pm  
Saturday - Sunday: Closed

FCE Reservations

[Canvas or Instructional Design Support](#) →

## Titan Supercomputing Center (TSC)



**CSUF-Titans** : For students research

**CSUF-Classroom** : Classroom use only

**CSUF-Research** : For faculty, staff research

### Nautilus Data Storage Limits

⚠ Be advised, data may be DELETED in accordance with the Acceptable Use Policy. Nautilus environments is not intended as long term storage solutions.

### Usage Policy

The JupyterHub environment provided by the **National Research Platform (NRP)** is a shared, research-focused computing platform. Users must follow strict guidelines for appropriate use:

**Do NOT store protected or sensitive data**, including: HIPAA (Health Information), PID (Personal Identifiable Data), FERPA-regulated data, and FERPA-protected student records.

## Tutorials

- [National Research Platform](#)
- [Nautilus Documentation](#)
- [Step-by-Step: How to Get Started with JupyterHub](#)
- [Nautilus JupyterHub - Image & Resource Guide](#)
- [How to Save Your Work from Jupyter Notebook/JupyterLite](#)
- [Releasing Resources on JupyterHub Services](#)
- [How to Use the Custom Image in JupyterHub](#)
- [How to Start R in JupyterHub](#)
- [How to Upload Your Python File into Jupyter Notebook/JupyterLite](#)
- [Track Memory and Storage Usage in JupyterHub](#)
- [How to Get an LLM Token and Use It in JupyterHub](#)
- [How to Acknowledge Support from NRP/Nautilus in Research Papers](#)

Click on “Instructional Design” and you see:

Home | Current Programming | Teaching Resource Center (TRC) | **Instructional Design and Academic Technologies (IDAT)** | xREAL Lab | Faculty Resources

- Next Generation Smart Classrooms
- Make an appointment with an Instructional Designer
- Affordable Learning Solutions
- High Performance Computing
  - Awards
  - HPC Trainings
- Canvas Resources
- Online and Hybrid Teaching Resources
- Teaching with Generative AI
  - GenAI Resources for Students
- Quality Online Teaching Program

# The Emergence of AI Tutors

- In Fall 2025, UC San Diego rolled out multiple AI tutors across a wide range of courses.
- They all have in common:
  - Attention to individual students even in large classes
  - Interactivity to change from student as a consumer to students as an actively engaged participant
  - Tutoring is available 24/7
- **We expect that most large lecture classes will offer personalized AI tutors before, during, and after the classroom in the future.**

See next talk on ASPIRE Math for an example AI Tutor that is measured to have a dramatic impact on education outcomes.

# Key Messages for the future of Higher Ed

- **Humans learn better by doing than watching**
  - Jupyter brings the doing back to large classes
- **Digital Assets in the classroom redefine what's possible**
  - AI makes complex programming possible for novice programmers
  - Large scale data processing as classroom exercises is now possible
- **More students learn more with personalized tutoring**
  - AI makes personalized tutoring before, during, and after the classroom possible even for large lecture classes.

**Colleges that fail to adopt these methods will go bankrupt**

# UC San Diego is proving that these methods can be deployed at scale

1/3 of all graduate and 1/4 of all undergraduate students have taken at least one class with digital assets in the classroom in AY25

UC San Diego educated ~8,000 Graduate Students and ~35,000 undergraduates in AY25

# UC San Diego Digital Assets in the Classroom

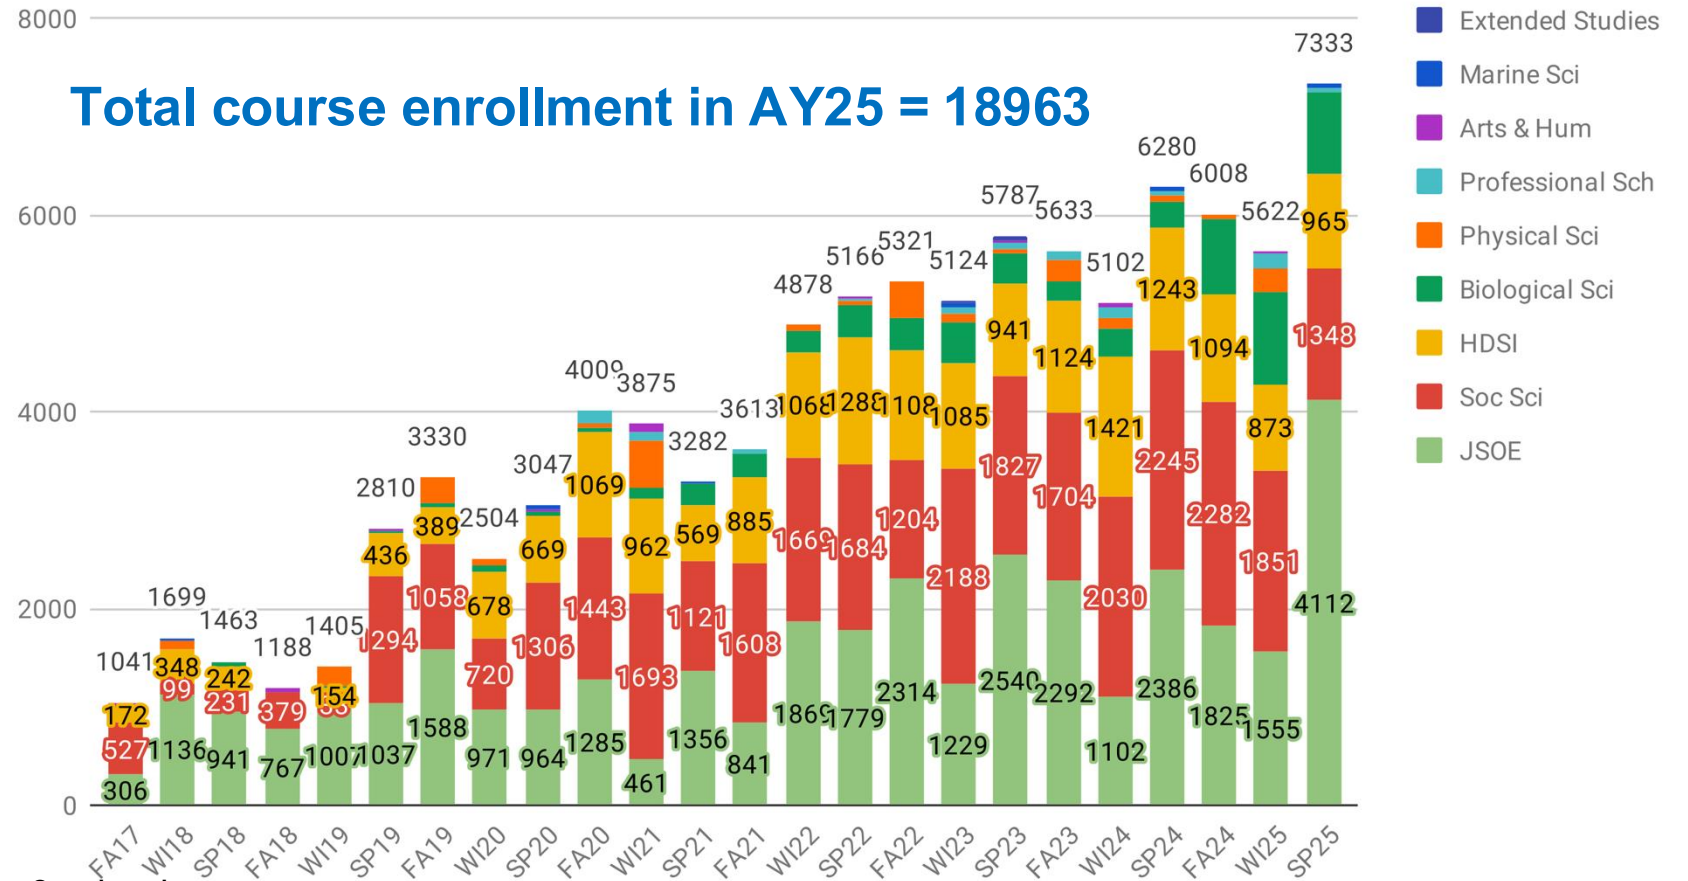
24% of all UCSD undergraduates used compute & data infrastructure in AY25 as part of their classroom education

Total in AY25	11092
UG Students	8538
Grad/Prof Students	2554

Some students take multiple courses/AY

34,955 undergraduates enrolled in Fall 2024

#Students by School/Division, Quarter



In AY25: 132 courses, 81 instructors, across 8 schools

# 1523 Students in Machine Learning for Music (SP25)

CSE 153 and 253 (and 153R / 253R) are **undergraduate and graduate courses devoted to the application of machine learning to understand and generate music**. After taking this course, students will be able to understand and manipulate data (and data structures) for music representation; students will be able to build predictive models and pipelines for music information retrieval; and students will be able to algorithmically synthesize music, culminating in a project demonstrating their creative work.

No previous background in machine learning or music is required, but all participants should be comfortable with programming (all example code will be in Python), and with basic optimization and linear algebra.

module	week (approx)	slides	workbook
		course introduction	
1	1-2	data structures for music and data ingestion	workbook 1
2	3-4	predictive pipelines for music & MIR	workbook 2
3	5-6	symbolic music generation	workbook 3
4	7-8	continuous-domain music generation	workbook 4
5	9	other topics in Music ML	
	10	musical performances!	

# California as an example for the importance of community colleges

# CA Master Plan for Social Mobility

- **California Community College (CCC) system dominates college student population in CA**
  - 2.2 Million students across 116 community colleges => 2-year colleges
  - 0.48 Million students across 23 CSUs => 4-year colleges
  - 0.44 Million students across 10 UCs => PhD granting research universities
- **3M students enrolled in public colleges in CA => 15% of all US college students**
- **CCC transfer students to CSU & UC drive social mobility in CA**
  - 30% of UC incoming class are transfer students from CCC
  - 50% of CSU incoming class are transfer students from CCC
- **California has strong path from K-12 to CCC to UC/CSU**
  - Public high schools in CA often offer advanced courses with their local CCC

**As CSUs and UCs deploy digital assets in the classroom  
CCCs must have the same capabilities to prepare transfer students**

# NDP is curating curriculum for use on NRP

NATIONAL DATA PLATFORM

My Dashboard

← Go BACK

## NDPDATA8 : Data 8

Institution UC Berkeley

Academic Term Spring 2026

### Modules

DATA 8 - Week 1  
UC Berkeley  
DATA 8: Week 1  
**To launch your workspace:**  
Open JupyterHub and keep all default settings.  
View →

DATA 8 - Week 2  
UC Berkeley  
DATA 8: Week 2  
**To launch your workspace:**  
Open JupyterHub and keep all default settings.  
View →

DATA 8 - Week 3  
DATA 8: Week 3  
**To launch your workspace:**  
Open JupyterHub and keep all default settings.  
View →

DATA 8 - Week 4  
DATA 8: Week 4  
**To launch your workspace:**  
Open JupyterHub and keep all default settings.  
View →

DATA 8 - Week 5  
DATA 8: Week 5  
**To launch your workspace:**  
Open JupyterHub and keep all default settings.  
View →

DATA 8 - Week 6  
DATA 8: Week 6  
**To launch your workspace:**  
Open JupyterHub and keep all default settings.  
View →

Report Issue/Feedback

**Data8 is the intro to Data Science curriculum** created at Berkeley, and adopted at UC Merced & UC San Diego.

This curriculum is offered to all community colleges in CA by Eric, Sean et al. => [See Thursday talk by Sean](#)

11:40 AM – 12:00 PM

**Successes and Challenges in Bringing data8 to the Community Colleges**  
*Sean Michael Morris, UC Berkeley*

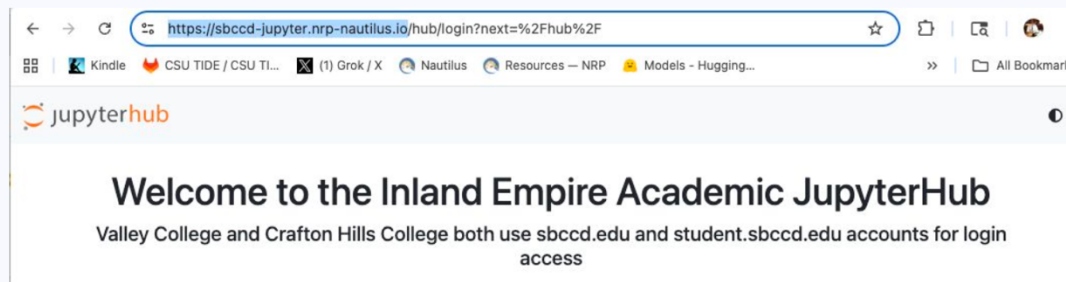
**Goal:**  
**Make transition from CCC to UC Berkeley/Merced/San Diego data science major more seamless.**

Data Science is just the first major to establish proof of concept. In theory, we can do this for a wide range of course in many majors.

22

# CSU San Bernardino's CCC Engagement

## Hands-On and User-Friendly JupyterHub Environment



### Images:

- NRP Minimal Python
- NRP Stack PyTorch2
- CSUSB R Studio Server
- CSUSB Tensorflow Deep Learning & Data Science
- CSUSB Data Science Python, and other scientific computing libraries for data science and machine learning
- CSUSB Stack SageMath SageMath, LaTeX, and AI Tools
- QUAY Julia Julia and core packages for numerical computing.
- QUAY Scipy Python, SciPy, and other scientific computing libraries for data science and machine learning

Architecture

amd64

*Important: To avoid unnecessary resource usage, please stop your server when it's no longer needed or before launching a different stack. Navigate to: File → Hub Control Panel → Stop Server*

Start

## Collaborating Community Colleges

- ❖ San Bernardino Valley Community College [14,500 stds]
- ❖ Crafton Hills Community College [6,900 stds]  
<https://sbccd-jupyter.nrp-nautilus.io>
  - Faculty members have been experiencing the Jupyterhub.
  - Some are interested in using Jupyterhub for classes.
  
- ❖ Chaffey Community College [22,500 stds]  
<https://chaffey-jupyter.nrp-nautilus.io/>
  - An AI degree has been approved - Working on building Curriculums

# Let's play with some numbers ...

- Roughly 20M students attend college in the USA
- $\frac{1}{4}$  of UC San Diego undergraduates have used data & compute in their classroom in AY25

⇒ **There are ~5 Million college students in the USA that need data, compute & AI resources for their classroom education annually.**

# More numbers games ...

- 20M students attend college in the USA
  - 35k undergraduates + 8k graduate students at UC San Diego enrolled in 132 courses that requires an infrastructure of roughly 100 GPUs
    - $20M / (35k + 8k) = 465$ 
      - ... let's just blindly scale by that to get an estimate for CA
- ⇒ **60,000 courses in USA need digital assets in the classroom**
- ⇒ **45,000 GPUs needed for those classes nationwide**

Probably an overestimate because not all courses need GPUs.  
Compute needs are generally larger for advanced courses than intro courses.

# Possible path for a college new to NRP

- **Some faculty start exploring NRP**
  - <https://nrp.ai/training/>
  - <https://nrp.ai/documentation/>
  - Zoom based office hour every other week on Tuesdays 10am PST
    - Next office hour on May 12<sup>th</sup> (i.e. next Tuesday)
- **When there is sufficient interest among the faculty, college buys hardware for NRP to operate.**
  - You work with your REN on network connectivity necessary
  - You rack up hardware in your data center
  - NRP team operates hardware from IPMI up

# Digital Assets in the Classroom

## Tuesday

9:00 - 10:30 AM	<b>NRP and Kubernetes for Education and Research</b> <i>Mahidhar Tatineni, UCSD/SDSC</i> <i>Mohammad Firas Sada, UCSD/SDSC</i> <i>Daniel Diaz, UCSD/SDSC</i>
10:30 - 11:00 AM	<b>Break</b>
11:00 - 12:00 PM	<b>National Data Platform Education Hub on NRP</b> <i>Mohammad Firas Sada, UCSD/SDSC</i> <i>Pedro Ramonetti, UCSD/SDSC</i> <i>Katie O'Laughlin, UCSD/SDSC</i>
2:00 - 3:00 PM	<b>Setting Up Custom JupyterHubs for Classroom and Research</b> <i>Mohammad Firas Sada, UCSD/SDSC</i> <i>Daniel Diaz, UCSD/SDSC</i>

## Wednesday

11:40 - 12:15 PM	<b>ASPIRE Math Tutor</b> <i>Sesh Murthy, UCSD</i> <i>Portia Restuccia, Vocareum</i>
<b>AI in Education - Master of Ceremony: Eric Van Dusen, UC Berkeley</b>	
1:30 - 1:50 PM	<b>NRP and the CSUs: Case Studies and Lessons Learned Supporting Teaching and Learning</b> <i>Kyle Krick, SDSU</i> <i>Dung Vu, CSU San Bernardino</i> <i>Dhusdee Chandswang, CSU Fullerton</i>
1:50 - 2:10 PM	<b>Use of Digital Assets in Undergraduate Education at UCSD</b> <i>Sam Lau, UCSD</i>
2:10 - 3:00 PM	<b>Panel on Education on NRP</b> <i>Robert Malouf, SDSU</i> <i>Valeria Barra, SDSU</i> <i>Martin Kim, CSU San Bernardino</i> <i>Rosanna Overholser, Cal Poly Humboldt</i> <i>Rakesh Mahto, CSU Fullerton</i> <i>Kyle Krick, SDSU</i>

## Thursday

11:40 AM - 12:00 PM	<b>Successes and Challenges in Bringing data8 to the Community Colleges</b> <i>Sean Michael Morris, UC Berkeley</i>
---------------------	--

# AI for Agriculture

# Research & Engagement with Industry

- Community Colleges and CSUs have a strong local focus
  - Serve local industries
- **We aspire to create research capacity & engagement between colleges and their local industries via the college's ownership of AI Infrastructure**
- **Example:**
  - Agriculture Technology => Instrumenting Iron Horse Vineyards in Sonoma

# Louis Fox: Minds in the Vines: Goals

1. Create a digital twin of the vineyard;
2. Use mixes of the latest technologies that are open source/non-proprietary;
3. Contribute data to a shared and increasingly large database open to agriculturists, researchers and students — highly curated data — where the value of AI agents will continue to increase as the data store increases;
- 4. Involve community college and university faculty and students;**
- 5 Demonstrate the value of 10G robust broadband to ISPs and the tech industry;
6. See if these open source efforts will eventually create new entrepreneurial enterprises with add on services, tool integration, etc. and,
- 7. Inspire other testbeds in precision agricultural and perhaps among other enterprises.**

# Chardonnay and Pinot Noir Grapes, Mainly for Sparkling Wines

**Vineyard Blocks - 156.3 ac.**

- Chardonnay - 64.3 ac.
- Pinot noir - 91.8 ac.

*(BazzanoMapping)*

# Instrumenting Iron Horse Vineyards



Research-grade Weather, Water Pipe Pressure, Quake Sensors



Tom taking 360 video on e-bike



David, Emily & John Chose Sites for Moisture, Temperature and Conductivity Sensor w/LoRaWAN Transceiver & Solar Charger & CO2 Sensors

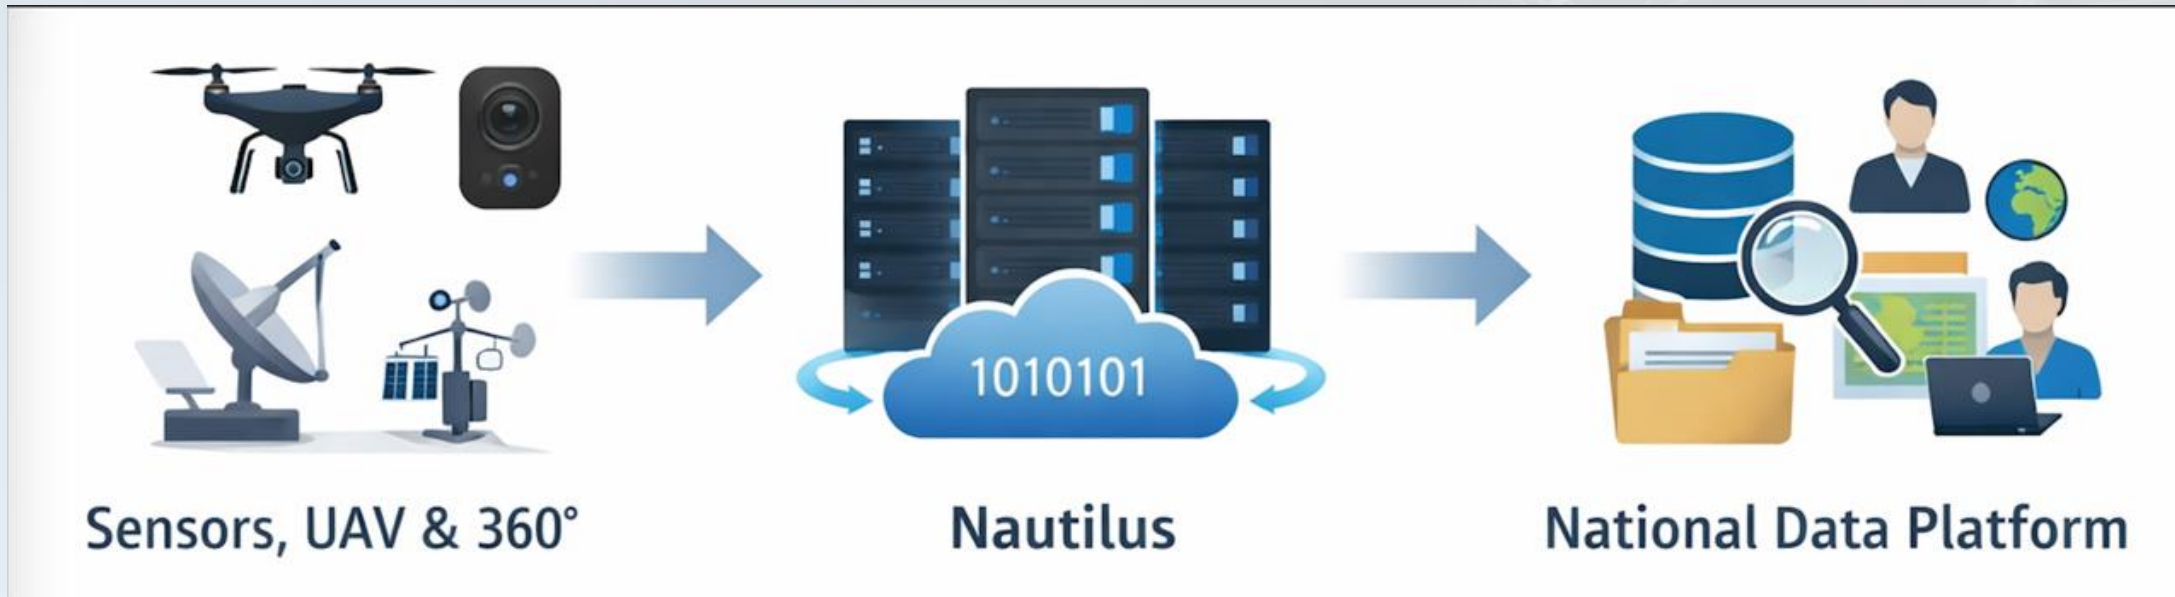


adding cameras for wildfire detection



# End-to-End Research Pipeline

- Capture → Nautilus → NDP
- Not just video — reusable data
- Supports SSU faculty and students + other campuses' research
- Enables scalable collaboration



# Bigger Picture

- The NSF agreed that NRP can provide services to industry small and large as long as ...
  - The total consumed GPU time is less than 10% of the NSF hardware in PNRP
    - $10\% \times (288+64) \times 24 \times 30 \text{ hours/month} = 25,344 \text{ GPU-hours/month}$
  - The industry customers pay for the services if the total exceeds the above, and we use that funding to sustain NRP for the community at large.

**We therefore encourage the NRP community to think creatively about possible joint ventures with industry on NRP.**

See example from Fresno on F3 Innovate this afternoon

# AI for Agriculture

## Tuesday

1:00 - 2:00 PM	<b>Tutorial on Drone Usage and IoT</b> <i>John Graham, UCSD/SDSC</i>
2:00 - 3:00 PM	<b>Tutorial Using ICICLE Digital Agriculture Pipeline</b> <i>Hari Subramoni, Ohio State University</i>
3:30 - 5:00 PM	<b>AgTech for the Curriculum Discussion</b> <i>Frank Wuerthwein, UCSD/SDSC</i>

## Wednesday

AI for Agriculture- Master of Ceremony: Noely Gonzalez-Maldonado, USDA-ARS	
3:30 - 3:50 PM	<b>AI in Agriculture</b> <i>Ryan Dinubilo, F3 Innovate</i>
3:50 - 4:10 PM	<b>AI Institute</b> <i>Hari Subramoni, Ohio State University</i>
4:10 - 4:30 PM	<b>Agricultural Mobile Robots for Automated Field Data Collection and Analysis</b> <i>Konstantin Karydis, UC Riverside</i>
4:30 - 5:00 PM	<b>Panel Discussion</b> <i>Ryan Dinubilo, F3 Innovate</i> <i>Hari Subramoni, Ohio State University</i> <i>Konstantin Karydis, UC Riverside</i>

## Thursday

AI for Agriculture- Master of Ceremony: Melissa Floca, UC San Diego/SDSC	
9:00 - 9:45 AM	<b>Keynote</b> <i>Baskar Ganapathysubramanian, Iowa State</i>
9:45 - 10:10 AM	<b>Iron Horse Vineyards Test Bed</b> <i>John Graham, UCSD/SDSC</i> <i>Tom Defanti, UCSD/SDSC</i>
10:10 - 10:30 AM	<b>CALFIRE Smoke and Fire Modeling for Prescribed Fire in Napa</b> <i>Leticia Lee, UCSD/SDSC</i> <i>Ilkay Altintas, UCSD/SDSC</i>
11:00 AM - 11:40 AM	<b>NAFSI Data Hub</b> <i>Kate Trout, University of Missouri-Columbia</i>

Thank You to our Gold Sponsors

**ciena**®

**IBM**

# Thank You to our Silver & Bronze Sponsors

## Silver Sponsor



## Bronze Sponsor



# Thank you to the team

- SDSC Events Team: Susan, Cindy, Chante
- Program Committee:
- All the speakers ...
- ... and all the attendees

**General Chair - Frank Wuerthwein**

SDSC, UC San Diego

**Kendra Ard**

CSU, Office of the Chancellor

**Tom DeFanti**

*6NRP Program Chair*  
Qualcomm Institute, UC San Diego

**Michael Farley**

San Diego State University

**Melissa Floca**

SDSC, UC San Diego

**James Alexander Hurt**

University of Missouri

**Jen Leasure**

The Quilt

**James Letts**

7NRP Program Chair  
SDSC, UC San Diego

**Larry Smarr**

UC San Diego

**Mahidhar Tatineni**

*Tutorials Chair*  
SDSC, UC San Diego

**Susan Rathbun**

*Local/Logistical Committee*  
SDSC, UC San Diego

**Cindy Wong**

*Local/Logistical Committee*  
SDSC, UC San Diego

**Chante Powell**

*Local/Logistical Committee*  
SDSC, UC San Diego

# Attendance at 7NRP

- 140 people registered representing ...
  - 41 colleges
    - 18 non-R1 colleges
  - 5 Research and Education Networks
  - 4 National laboratories & research institutions
  - Industry
    - 10 Sponsors
    - 4 Not Sponsors
  - 2 Federal Agencies

# Take Home Messages

- **There is a revolution in higher education coming your way**
  - You either jump on board or will go bankrupt
- **NRP allows any college to own their own AI, Compute & Data resources.**
- **Ownership has its privileges:**
  - Educational tools necessary for survival create research capacity to thrive ...
  - ... and opportunities for engagement with local industry

# Acknowledgements

The NRP is supported by NSF grants OAC-1541349, OAC-1826967, OAC-2030508, OAC-1841530, OAC-2005369, OAC-21121167, CISE-1713149, CISE-2100237, CISE-2120019, & OAC-2112167

And by CENIC, Pacific Wave, MREN, GPN, NYSERNet, FLR, NEREN, SunCorridor, OARnet, SCLR, the Albuquerque GigaPoP, and Internet2

As well as a long list of Universities and colleges that host hardware, and share their hardware with the community.