



iD Lab

Interactive Discovery Laboratory

AN INTERACTIVE LABORATORY AT NATIONAL SCALE

NSF IDSS | Award #2609583

PI: Amit Chourasia, UCLA

Co-PI: Amitava Majumdar (SDSC), Brett Bode (NCSA), Joe Stubbs (TACC), Phil Blood (PSC)

Site lead: Amiya Maji (Purdue)

Project Manager: Eleanor Koehl (UCLA)

UCLA | NCSA | PSC | Purdue | SDSC | TACC



ORIGIN

Why OneSciencePlace Exists?

OneSciencePlace (OSP) did not begin as a product. It began with a problem we observed and an NSF recommendation to solve it.

01

THE FINDING

The Science Gateways Community Institute (SGCI), NSF-funded from 2017 to 2022, worked with dozens of gateway projects. A pattern recurred: domain-specific gateways delivered real scientific impact, but were costly to maintain, slow to evolve, and difficult to generalize.

02

THE RECOMMENDATION

Based on those findings, NSF recommended that SGCI design a modern, modular framework rather than continue building one-off gateways.

03

THE RESPONSE

OneSciencePlace was conceived in direct response. Not as an extension of any earlier system, but as a synthesis of what each had taught.

A clean start: OSP was designed around open standards, modular integration, and user-space deployment.



25 YEARS OF NSF-FUNDED CYBERINFRASTRUCTURE

A synthesis of what came before

OSP is built by the people who built and operated the systems that came before it. Each contributed lessons that shaped the platform.

HUBzero and nanoHUB

Content-centered design and community building

CIPRES

Complex applications with easier interfaces, backend performance, and strong user support

Apache Airavata

Hosted platform model, workflow execution, and distributed orchestration

Tapis

API-driven compute and data integration. Tapis is the backend for OSP

SeedMeLab

FAIR-ready data sharing and metadata. SeedMeLab is now part of OSP



Tapis and SeedMeLab are part of OSP

Tapis is the backend for compute and data integration, and SeedMeLab's data sharing and metadata system forms the foundation of OSP's data module.



DESIGN PRINCIPLES

Four principles shaped the architecture. They reflect what the team learned from decades of building and operating cyberinfrastructure.



Composable

Each deployment adopts only the components it needs: apps, data, publishing, identity, content management. Streamlined for the project at hand, expandable as needs grow.



Content-centric

Applications, systems, publications, data, and documentation are all searchable, manageable content in one platform, not separate systems with their own tooling.



For users at any skill level

User experience drives design choices. A no-code form builder, a browser-first interface, and contextual documentation make research computing approachable.



Built for longevity

Technology choices favor mature open-source components, with the goal that the platform stays maintainable and extensible for at least a decade.

The mission: Powerful computing should be easy for everyone. OneSciencePlace removes the engineering barriers to delivering science gateways, HPC portals, and data repositories, so people can focus on science, education, and learning.



DESIGNED FOR BROADER REACH

Meet stakeholders where they are

Scalability is usually argued in one direction, toward the national scale. OSP is designed to scale in both, and the same platform serves each level.



Individual labs and PIs

A focused home for a group's applications, data, and students



Educators and courses

Pre-configured applications and single sign-on for a class



Campus and National portals

Integrates in user space, no privileged access required



Science gateways

One or many applications across local and national resources



Repositories

Publish applications, data, and results as citable objects



iD Lab is OSP's national-scale deployment: one tenant, five sites, two clouds. The new capabilities, hardening, applications, and integration patterns iD Lab funds return to the same platform that serves a single lab or a single classroom, so this investment reaches well beyond the project.



FOUR YEARS OF BUILDING AND TESTING

Demonstrating the Platform



Quakeworx

OSP's first gateway and its primary science driver, built with the Statewide California Earthquake Center, UIUC, and Scripps Institute of Oceanography. Interactive and batch earthquake simulation across HPC and cloud, FAIR publication, and federated sign-on. Moving from early user phase to production.



CIPRES

A well-regarded phylogenetics gateway, since concluded. Over two dozen CIPRES applications were ported to OSP, validating the platform's ability to host complex, parameter-rich workflows.

16 presentations

Talks and posters

6 Tutorials, 1 Course

Hands-on tutorials and course for over 250 attendees

30 publications

14 papers, 16 preprints, in-press, and submitted



WHAT USERS ACTUALLY SEE

One click from browser to running application

The screenshot displays the 'Curated Listing' page in the Quakeworx interface. On the left is a sidebar with navigation options: 'My jobs', 'Create' (with sub-items: App, Event, Page, Publication, Sys credentials), 'Apps' (with sub-items: Curated apps, All apps, My apps), 'Publications' (with sub-items: My publications, All publications), 'Help', and 'Support'. The main content area is titled 'Home / Curated Listing' and 'Curated Listing'. It features three application cards:

- autoTandem**: Marked as 'Beta' and 'Updated'. Description: 'Tandem simulates sequences of earthquakes and aseismic slip.' Options: 'Batch', 'Executable', 'Cluster'. System: 'Expansive service'. Version: 'main_17c42dc9ae0ec519d cc1b5732681b2e4054666f1'. Buttons: 'View', 'Launch'.
- Broadband Platform**: Marked as 'Dev'. Description: 'Broadband Platform - can be used to generate 0-20+Hz seismograms for historical and scenario earthquakes.' Options: 'Batch', 'Executable', 'Cluster'. System: 'Expansive service'. Version: '0.1.0'. Buttons: 'View', 'Launch'.
- HFQsim**: Marked as 'Stable'. Description: 'The app simulates fault-and-earthquake behaviour by modelling the interplay between seismic slip and aseismic deformation in a unified framework.' Options: 'Batch', 'Executable', 'Cluster'. System: 'Expansive service'. Version: '1.0.0'. Buttons: 'View', 'Launch'.

The curated catalog running today in Quakeworx. iDLab deploys the same catalog across five centers.

THE SHIFT

Research progresses interactively


Researchers develop ideas through rapid, interactive exploration of data, often in environments like Jupyter. These workflows are becoming increasingly data-intensive and computationally demanding.

The challenge is no longer building powerful infrastructure. It is making that infrastructure usable by far more researchers, educators, and students.



Who we are not reaching yet

Research computing serves a spectrum of users; however, the most prevalent access model is the command line.



The user spectrum across access modalities



Hard-won lesson

We are not our target user



The opportunity

Available capabilities are not the constraint. Access is.
Moving down this pyramid is how national cyberinfrastructure reaches the people it was built to serve.



iDLab complements, not replaces


The national investment is significant, and it works.
NAIRR and ACCESS resources are the instruments; iDLab assembles them into one interactive laboratory.



WHAT IDLAB DELIVERS


iD Lab: one interactive laboratory

iD Lab transforms the site's HPC, AI, data, and cloud capabilities, bringing them together as a single interactive laboratory.





A production architecture, already built and operating



iDLab runs as one tenant of OneSciencePlace.

The platform, orchestration, and identity trust plane are inherited, improved, and strengthened.



WHO THE PLATFORM SERVES

Research and education across disciplines

iDLab serves data-intensive and AI-enabled research across disciplines and provides ready-to-use environments for teaching.

Research domains



Natural hazards

Earthquake simulation and seismic hazard modeling



Spatial biology

Tissue and cellular imaging at scale



Neuroscience

AI-driven analysis of brain imaging and signals



Geospatial analytics

Mapping and spatial AI



Physics and engineering

Simulation and modeling workflows



Cryo-EM and agriculture

Structure determination and crop science



Education is a first-class use case, not an afterthought

Pre-configured applications and single sign-on for courses, workshops, and training, so students reach advanced computing without local infrastructure.

Package → Deploy → Onboard → Validate → Refine



HOW THE PROJECT RUNS

Timeline and governance

Three-year arc

Year 1

Build and advance in phases

Deploy, integrate, and reach production quality, with early users on the platform by year-end

Years 2 to 3

Operate

National operations at scale: onboarding, support, new applications, new capabilities, quarterly public reporting

Governance



Site leads

Each center runs its own systems; together they settle what must be the same everywhere.



Working groups

Four standing groups: platform, Tapis, site systems, and outreach.



External advisory board

Four to six members; annual meeting with written updates every six months



HOW WE WILL KNOW IT WORKED

Success: readiness, adoption, national impact

The project is executed in phases. Each year has defined milestones, with progress measured monthly and reported publicly each quarter.

YEAR 1

Readiness

From pilot to production

Production-quality operations, identity across all sites, and applications deployed. Capability is delivered through rolling updates rather than a single gated release.

- Identity and applications operational at all sites
- Early users working on the platform
- Security posture documented and reviewed

YEAR 2

Adoption

Communities onboard and grow

Communities onboard through their own allocations, applications, and courses. Support and reliability scale with them.

- Communities onboarding, bringing their own applications, courses and if needed allocations
- Application catalog expanding across disciplines
- Support and reliability scaling with demand

YEAR 3

National impact

A routine path to national computing

Interactive access becomes a normal way to use national infrastructure, for researchers and students who could not before.

- Sustained national operation across all sites
- Interactive access as a normal route to national resources
- A platform other projects and courses build on

Capability ships through rolling updates, so progress is continuous rather than gated on a single milestone.

Many systems.

One interactive laboratory.

*National cyberinfrastructure already exists. iDLab makes it interactive, accessible, and usable at national scale.
Unlock, accelerate, and amplify science, teaching, and learning.*

iDLab welcomes users, stakeholders, institutions, resource providers, and all current and future IDSS projects.



Project: idlab.onscienceplace.org

NSF IDSS | Award #2609583 | UCLA, NCSA, PSC, Purdue, SDSC, TACC

Platform: onscienceplace.org

