



Managed Services to Simplify Cloud Access for Computer Science Research and Education

UC San Diego

Mike Norman, PI
Shava Smallen, Co-PI
Vince Kellen, Co-PI

W
UNIVERSITY *of*
WASHINGTON

Ed Lazowska, Co-PI

Berkeley
UNIVERSITY OF CALIFORNIA

Kathy Yelick, Co-PI



NSF Award: 1925001
Deep Medhi, PO

Outline

- CloudBank Overview- *Shava Smallen (SDSC/UCSD), Co-PI*
- CloudBank for Researchers – *Naomi Alterman (UW)*
- CloudBank for the Classroom – *Kathy Yelick (UCB), Co-PI*
- Q&A

CloudBank seeks to address pain points of researchers and educators using public cloud

“Setting up your first experiment in the cloud is too time-consuming, and different for each cloud platform.”

“I need documentation specific to my research.”

“Provisioning all students in a class with their own cloud credits and dealing with their cost overruns is a major pain. Financial losses can be substantial if students are given free reign.”

“There is no simple way to run multiple experiments with software already installed in one instance.”

“Research credits are difficult for us to keep track of.”

“We need billing help. It’s opaque as to what is going on and I don’t have the time to wade through this.”

“Cloud is just not designed for typical use case of a small research group.”


“We need an expedited process for setting up experiments very quickly. Often we use cloud resources when we’re on a deadline so quick setup is critical.”

“I just don’t have enough time to learn all this. Can you give me a jump start?”

“It’s a Hobson’s choice between giving everyone root access or losing my life in fine-tuning access.”

“Cloud is a serious learning curve. Just too much documentation, and of the wrong type. Students get lost.”

<http://www.cloudbank.org>




TRAINING - ABOUT - FAQ - HELP - LOG IN

Managed Services to Simplify Cloud Access
for Computer Science Research and Education


FEATURES

*For an overview of CloudBank and key features of the portal, please see [our recent webinar](#) and [write-up in Science](#).
For an introduction to using public clouds, please see our [Cloud 101 videos](#).*




User Portal

Researchers can manage their cloud resources and monitor their usage through the CloudBank user portal.



Financial Operations

CloudBank simplifies accounting processes and helps optimize resource usage for researchers.



Education & Training

CloudBank offers classroom tools, training workshops and videos, curated training materials, and individual support to help researchers get started.

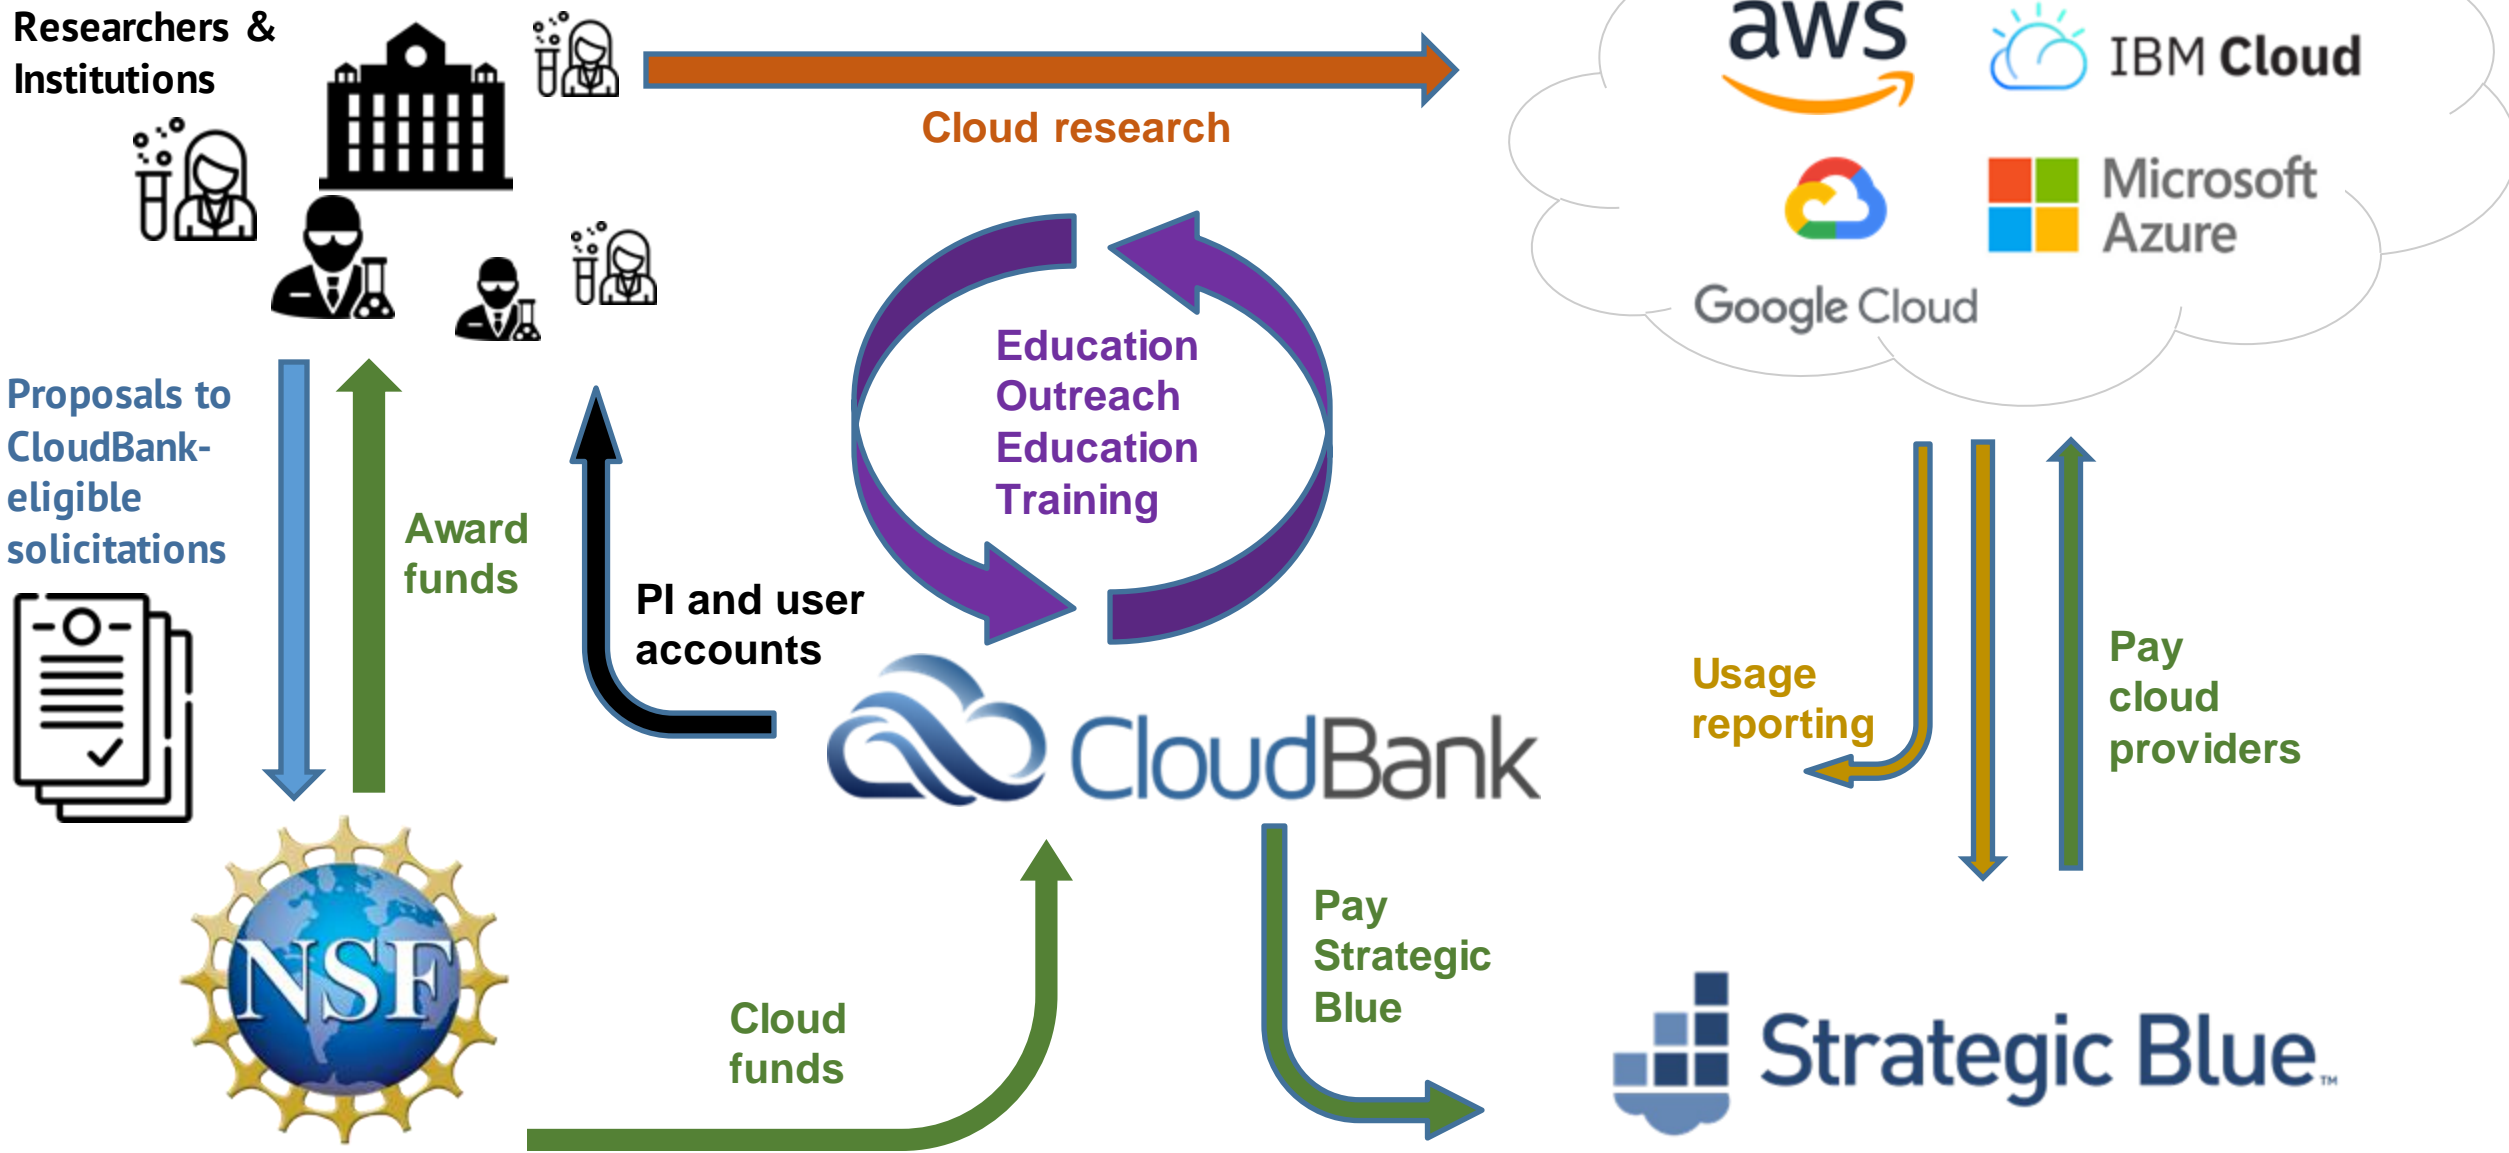
CloudBank is focused on specific NSF solicitations

- **IMPORTANT SCOPE:** Only specific NSF solicitations are eligible for cloud resources through CloudBank. The solicitation will state eligibility explicitly. There have been 6 such solicitations so far:
 - **NSF 21-533: Computer and Information Science and Engineering Minority-Serving Institutions Research Expansion Program (CISE-MSI), due April 15.**
 - NSF 21-522: Designing Materials to Revolutionize and Engineer our Future (DMREF): CLOSED
 - NSF 20-563: Cyber-Physical Systems (CPS): CLOSED
 - N
 - N
 - N
- Pls req
does n
- The total of your formal budget request and the cloud request must not exceed the total allowable request for the solicitation.
- Cloud resources are free of indirect costs. I.e., are treated like equipment

Without CloudBank, institutions would add cloud funding to their budget and many would pay indirect costs.



CloudBank workflow



Other CloudBank “Doors”



Discretionary/Community awards

Some funds are available to researchers and educators to “kick the tires”

Enterprise

For institutes that would like to adopt CloudBank’s model

Email help@cloudbank.org if interested

CloudBank for Researchers

Naomi Alterman (UW)

CloudBank for Researchers

- Cloud platforms help with problems of scale
 - Store more
 - Compute more
 - Compute faster
- The usual story: “I don’t have the resources to do what I need, maybe moving my work into the cloud can help”
 - But this move is a very non-trivial problem
 - Can feel like a risky leap into the dark
- CloudBank **helps you evaluate whether the leap is worth it**, and then **curates the resources you need to hit the ground running**
 - So you don’t have to learn everything about everything just to decide if it’s even worth your time

Your support team

Platform

How do I give my lab members access?
How do the bills get paid?
How do I track my spending?

UC San Diego

Design

What cloud services are useful for me?
How do I estimate costs for my project?
Who do I reach out to for technical help?

W
UNIVERSITY of
WASHINGTON

Berkeley
UNIVERSITY OF CALIFORNIA


Implementation

What is the best way to get bang for my buck?
I've hit a roadblock using this service, how do I fix what's broken?



When you send a message to help@cloudbank.org, you're getting the support and expertise of many organizations

Success stories



Folding petabytes of proteins in parallel
Hosseinzadeh et. al.

<https://www.cloudbank.org/training/use-case/htc-rosetta>

Success stories

Tracking transit quality at unprecedented scale

Frias-Martinez, et. al.

<https://www.sciencemag.org/news/2021/01/us-law-sets-stage-boost-artificial-intelligence-research>

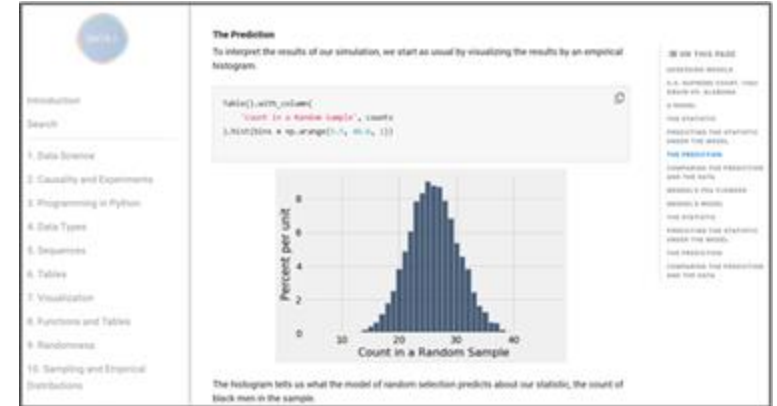


CloudBank for the Classroom

Kathy Yelick (UCB)

Berkeley Data Stack: Education Using the cloud

- Scalable teaching both on campus and remote
- Popular Data Science curriculum -- fastest growing major at UC Berkeley
- Modern education: online textbooks, videos, interactive assignments, etc.



Open infrastructure for introductory data science

Provide a shared computing environment

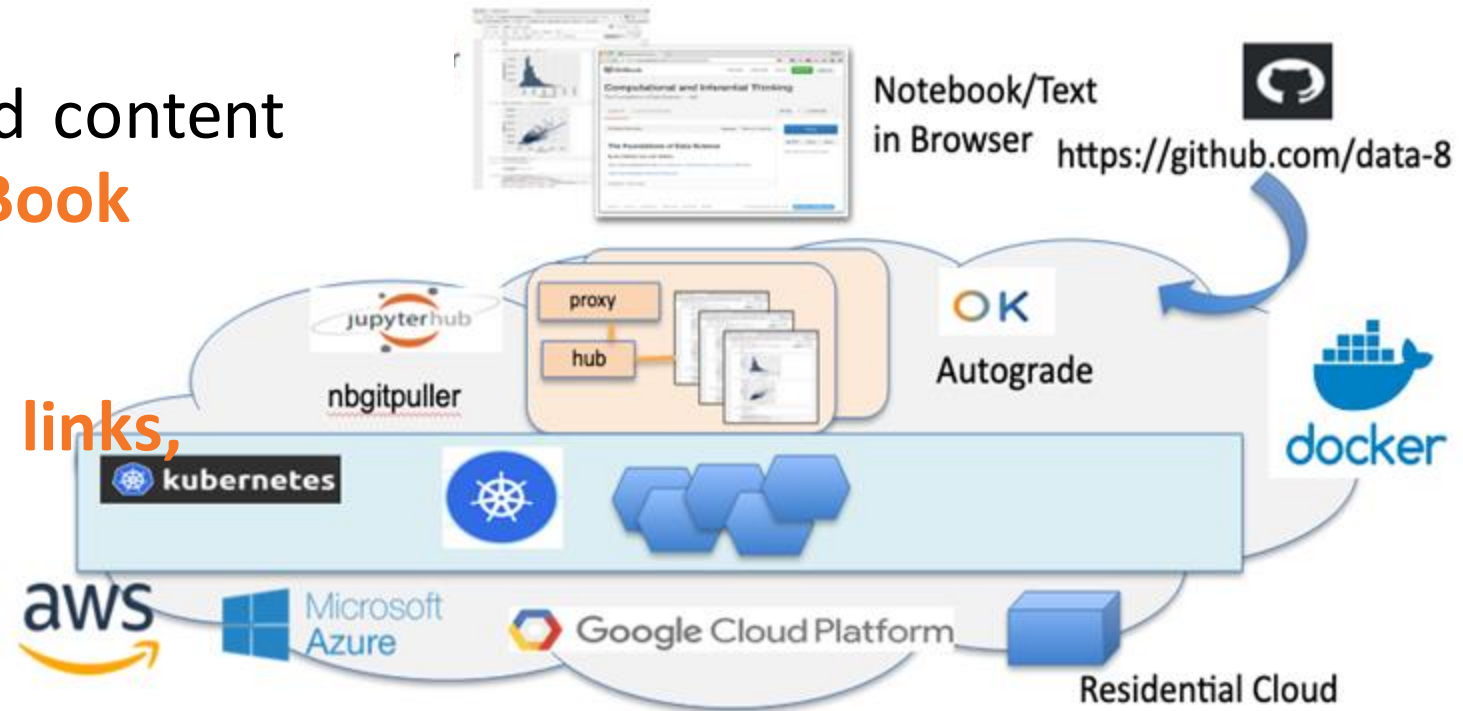
JupyterHub, BinderHub

Guide learners with structured content

Jupyter Notebooks, Jupyter Book

Tools for instruction

Online textbooks, interactive links, autograders, etc.



CloudBank Computing for Data Science Education



1867
HOWARD
UNIVERSITY



Skyline
COLLEGE



Clovis
Community
College

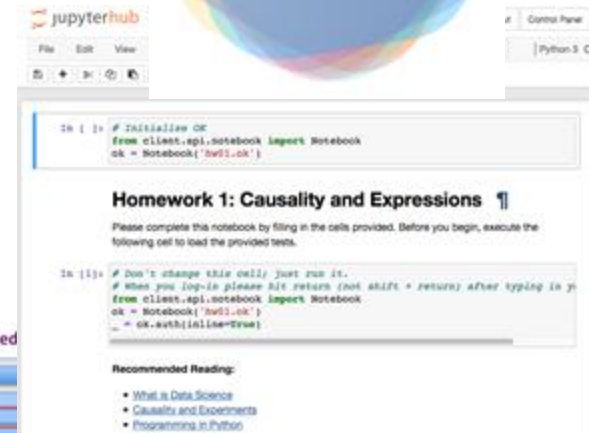
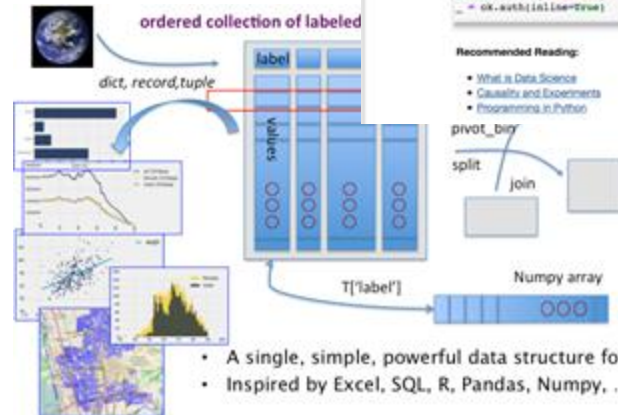
Current

- El Camino
- Howard
- Skyline

Pilots for future semesters

- City College of San Francisco
- Clovis Community College

Part of 2i2c consortium which has several other pilots





Q&A

<https://www.cloudbank.org>

help@cloudbank.org