

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

MAGIC/NITRD July 1, 2020



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

W UNIVERSITY of WASHINGTON

Ed Lazowska, Co-PI

Berkeley

David Culler, Co-PI



NSF Award: 1925001 Deep Medhi, PO



SDSC SAN DIEGO SUPERCOMPUTER CENTER

What problem are you trying to solve?

For the NSF

- Create an entity that will facilitate the use of cloud for CISE-funded PIs
- Cloud resources must be IDC free
- For CISE researchers and educators
 - Remove the "friction" associated with cloud access and account management; hide cloud contract negotiation
 - Help educate them about cloud with appropriate documentation
 - Help with cloud resource request part of their proposal to NSF
- For cloud providers
 - New users, grow the market
 - Get the NSF to pay them, not the other way around





CloudBank seeks to address <u>pain points</u> 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 <mark>opaque</mark> as to what is going on and I don't have the <u>time</u> 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 <mark>time</mark> 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.

* Input: CS researchers UCSD, UCB, UW (time) (hassle)



CloudBank will address these pain points through a set of managed services

- Cloud solution consulting, cloud training, proposal assistance
- Onboarding, account setup, user management, off-boarding
- Tools for PIs to grant permissions to research group members and students
- Unified reporting across cloud providers
- Bundling multiple small requests into a bulk request to cloud providers, providing leverage to secure discounts
- Pass along savings to researchers





CloudBank workflow







CloudBank is focused initially on specific CISE solicitations

- Important Scope: Only specific NSF solicitations are eligible for cloud resources through CloudBank. The solicitation states eligibility explicitly.
- There have been 2 such solicitations so far:
 - Smart and Connected Communities (Closed)
 - Cyber-Physical Systems (Open: Submission window June 8 22, 2020)
- PIs request cloud resources in a Supplementary Document, not the formal budget. Cloud funding does not pass through the awardee institution.
- 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., they are treated like equipment.





CloudBank project is organized into 3 main efforts



SDSC SAN DIEGO SUPERCOMPUTER CENTER



CloudBank project team

- UC San Diego/SDSC: Mike Norman (PI), Vince Kellen (Co-PI), Shava Smallen (Co-PI), Brian DeMeulle, Declan Fleming, Paul Hoover, Kate Kaya, Nick Marangella, Sonia Nayak, Amin Qazi, Shawn Strande, Mary Thomas, Nicole Wolter, Choonhan Youn, Frank Wuerthwein
- University of Washington: Ed Lazowska (Co-PI), Rob Fatland, Sarah Stone, Amanda Tan
- UC Berkeley: David Culler (Co-PI), James Percy
- Strategic Blue: James Mitchell, James Smith, Steve Old, Frank Contrepois
- External Advisory Board: Dana Brunson (Chair), Maytal Dahan, Dennis Gannon, Tracy Camp, Tracy Futhey, Michael Franklin, Jason Liu
- Cloud providers: AWS, GCP, Azure. Interest from IBM and Oracle
- National Science Foundation: Deep Medhi





CloudBank User Portal comprises the bulk of year 1 work

Feature	Status
Design and data structures to support all relevant functions and workflows	Complete
User account management	Complete
Automated cloud account creation on provider platforms	Completed for AWS and GCP. Azure in progress. Discussions underway to add IBM.
Federated login to portal, and providers using CILogon	Complete
Integrated usage monitoring with Nutanix Beam	Completed for AWS, and Azure. GCP planned for September (Nutanix resources are gating)
Scripts for usage monitoring and threshold alerts	In testing
NSF access for award processing and viewing	In progress.
Framework for documentation and training resources	Complete. Content and resources being added regularly.
Early user testing	Underway, and expanding.

AN DIEGO

PERCOMPUTER CENTER

CloudBank Using the CloudBank Catalog: Major categories of services are highlighted below. Click on any category (e.g., Compute) and it will expand to show a table with types of services represented in the rows and public clouds in the columns. Hover over the row header in the leftmost column o see an expanded definition of the type of service. Compute \odot 0 Storage 0 = Network loudBank Catalog Search Cloud Tools Manage Y III Database Services Explore available Public Clouds and do an at a View and search all Cloud tools by category or Manage you plance comparison of features. vendor billing acc HI Big Data & Analytics Machine Learning 50° Machine learning (ML) is the scientific study of algorithms and statistical models that computer systems use to perform a specific task CloudBank without using explicit instructions, relying on patterns and inference instead. It is seen as a subset of artificial intelligence. [Source: DASHBOARD - TR Wikinedia **Manage Funds Google Compute** Amazon Web Services Microsoft Azure IBM Cloud Platform are Speech to Text , Azure Speech Amazon Transcribe Soogle Speech-to-Tex Text to Speech Azure Computer Vision , Azure Displaying 1 - 12 of 12 funds Visior Amazon Bekognitic Google Vision A Face Filter by PI or Co-\$1,000.00

Portal screenshots

Key technology and collaboration components



UC San Diego

The CloudBank portal provides project-developed and provider content to assist researchers. Tiered user support augments this. CloudBank User Support

Cloud Resource Catalog

payment bar [Source: Del	sed on usage. Cloud computing servic	es range from full applications and develope	ment platforms, to servers, storage	a, and virtual desktops
	Amazon Web Services	Google Compute Platform	Microsoft Azure	IBM Cloud
IAAS	Amazon Elastic Compute Cloud	Google Compute Engine	Azure Virtual Machines	
General Purpose	٥	٥	0	
Compute Optimized	٥	٥	•	
Homes				



Video Tutorials

Think Cloudily



Provider Tools

lisplaying 1 - 30 of 170 tools				
LOGO	TITLE	CATEGORY		
Amazon Aurora	Amazon Aurora	Relational Database		
Enazee clustront	Amazon CloudFront	Content Delivery Network		
Amazon Cloudfinitch	Amazon CloudWatch	Monitoring, Logging		
Amazon Comprehend	Amazon Comprehend	Natural Language Processing		







Substantial progress has been made. On track for a transition to production operations on <u>August 1, 2020</u>

- Majority of portal work complete
- Public cloud agreements in place
- Education and training material development continuing
- Provider engagement on support and training
- Early user testing underway
- Community outreach and input
- Advisory Board has been formed and has met

Year 1: Development 8/1/19 – 7/31/20

- User portal design and development
- Implement stakeholder workflows
- Cloud provider contract negotiations
- Integrate usage and account monitoring
- Early user testing
- Outreach
- Advisory Board



Years 2-5: Production Operations 8/1/20 – 7/31/24

- Operational support for NSF solicitations
- PI and user training and support
- Feature enhancements of portal
- Ongoing bugfix
- Annual cloud workshop
- Expansion to other providers
- Track project progress via KPIs





Future CloudBank Center of Excellence is an opportunity for partnering

- Best practices and experiences in using public cloud
- Cross-promotion of resources, programs, and expertise
- Annual workshop, training, and other events
- Participations in annual meetings like PEARC
- Leverage existing programs like XSEDE, Campus Champions, SGCI
- Cloud provider workshops, webinars, training
- And more...





For more information

Go to www.cloudbank.org







"Any opinions, findings, conclusions or recommendations expressed in this material are those of the author(s) and do not necessarily reflect the views of the Networking and Information Technology Research and Development Program."

The Networking and Information Technology Research and Development (NITRD) Program

Mailing Address: NCO/NITRD, 2415 Eisenhower Avenue, Alexandria, VA 22314

Physical Address: 490 L'Enfant Plaza SW, Suite 8001, Washington, DC 20024, USA Tel: 202-459-9674, Fax: 202-459-9673, Email: <u>nco@nitrd.gov</u>, Website: <u>https://www.nitrd.gov</u>

