# Broadening the Reach for Access to Advanced Computing

## Leveraging the Cloud for Research

Forough Ghahramani
NJEdge
forough.ghahramani@njedge.net

John Hicks
Internet2
jhicks@internet2.edu

Barr von Oehsen
Rutgers University
barr.vonoehsen@rutgers.edu

## ABSTRACT

Many smaller, mid-sized and under-resourced campuses, including MSIs, HSIs, HBCUs and EPSCoR institutions, have compelling science research and education activities along with an awareness of the benefits associated with better access to cyberinfrastructure (CI) resources. These schools can benefit greatly from resources and expertise for cloud adoption for research to augment their in-house efforts. The Ecosystem for Research Networking (ERN), formerly the Eastern Regional Network, Broadening the Reach (BTR) working group is addressing this by focusing on learning directly from the institutions on how best to support them. ERN BTR findings and recommendations will be shared based on engagement with the community, including results of workshops and surveys related to challenges and opportunities as institutions are evaluating using the cloud for research and education, as part of the NSF sponsored CC*CRIA: OAC-2018927.

## CCS CONCEPTS

• **General and reference** → Cross-computing tools and techniques; • **Social and professional topics** → Professional topics; Computing education; Computing / technology policy.

## KEYWORDS

cloud computing, research computing, advanced cyberinfrastructure, higher education

**ACM Reference Format:**
Forough Ghahramani, John Hicks, and Barr von Oehsen. 2022. Broadening the Reach for Access to Advanced Computing: Leveraging the Cloud for Research. In *Practice and Experience in Advanced Research Computing (PEARC '22), July 10–14, 2022, Boston, MA, USA.* ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3491418.3535143

## 1 INTRODUCTION

New technology capabilities, high performance networking and computing, and availability of data are driving emergence of new models of research computing to support research both on-premise and in the cloud. The role of research computing and data in scientific discovery and scholarship across all disciplines presents challenges as well as opportunities. The Ecosystem for Research and Networking (ERN) [4] was formed as a consortium of institutions and regional network providers that agree to simplify and enable multi-campus research collaborations and partnerships between academic institutions of all types and sizes that advance the frontiers of research, pedagogy, and innovation. The ERN adheres to the principles that promote team science and simplifies access to both instruments and data, and supports democratizing research data and instruments with a particular emphasis on broadening the reach to under-resourced and under-represented colleges and universities. The ERN includes the Broadening the Reach (BTR), Architecture and Federation, Policies, Structural Biology, and Materials Discovery Working Groups.

Gaining a better understanding of the advanced computing requirements, and outreach to best support researchers, educators, and upper administration from the nearly 2000 public and private colleges and universities of all shapes and sizes within the Northeast is essential. The majority of these institutions are small to medium sized non-R1 colleges and universities, including MSIs, HSIs, and HBCUs. The ERN BTR working group is representative of the diverseness of these communities and focuses on how the ERN can be more inclusive, and through a deeper understanding to enable the ERN to have the broadest impact across multiple research disciplines, pedagogical approaches, university and college stakeholders, and organizations within the region and beyond.

The ERN BTR has engaged with the community since August 2020 to a) identify the needs, b) raise awareness to existing regional and national resources, and funding opportunities, c) identify existing best practices and models. Results of this engagement has provided a baseline of the issues most important to the community from the perspectives of researchers, faculty, research computing staff, central-IT, and campus leadership, including CIO's, vice presidents for research, and provosts. The ERN BTR has organized workshops and surveys focused on the topic areas of importance to the community including: Democratizing Campus CI, Research Collaborations, Workforce Development and Expertise, Trends in Research Computing and Data.

Addressing the topic of challenges and opportunities in leveraging the cloud, the December 2021 virtual workshop, "Leveraging the Cloud for Research," brought together researchers, faculty, research computing and central IT professionals, as well as other institutional stakeholders, including students and administrators, to share information and facilitate discussions about successful implementation of the cloud for research from the researcher and institutional perspectives. The workshop provided a general introduction to the Cloud; challenges and opportunities in leveraging the cloud in research projects, featuring NSF-Funded Exploring Clouds for Acceleration of Science (E-CAS) projects; the institutional support and onboarding for research in the cloud; funding


*PEARC '22, July 10–14, 2022, Boston, MA, USA*

ACM ISBN 978-1-4503-9161-0/22/07…$15.00
https://doi.org/10.1145/3491418.3535143

onboarding for research in the cloud; and available training, tools and resources for cloud implementation. Participants represented a range of R1's through smaller, mid-sized and under-resourced campuses, including MSIs, HSIs, HBCUs, and EPSCoR institutions.

A pre-workshop survey was designed to help identify perceived challenges and opportunities for leveraging the cloud for research at participating institutions, and to use the knowledge from the survey to include relevant topics in the workshop.

The next three sections of this paper focus on the ERN BTR cloud workshop findings, recommendations, and ERN BTR plans for the future.

## 2 LEVERAGING THE CLOUD: CHALLENGES AND OPPORTUNITIES

Advances in cloud computing allow almost unlimited access to high-end computing resources for researchers at every type of institution, creating a more level playing field for experimentation than has ever existed before [3].Traditional high-performance computing (HPC) presents challenges for researchers at most campuses. The high cost of acquisition is a major barrier. Most academic and IT departments lack the specialized expertise to support such equipment, and in some cases, they may not even have an appropriate place to house even a small cluster. Researchers spend valuable time specifying, configuring, and operating HPC, and unless HPC is the individual's research area of focus, that time detracts from valuable research time. Large research universities may use graduate students to do some of this work, this is often not the most productive use of their time either. HPC is often funded by one-time research grants and quickly becomes obsolete. Researchers look to institutional support to keep them updated, support that is often not available. Without technical and administrative support, in-house technology can be underutilized. HPC in one department might be used by researchers in another department, if they know it exists, if they have documentation and assistance, and if a scheduling process is available to support equipment sharing. In the absence of this infrastructure, valuable equipment in one department or college often sits idle while another department struggles to obtain funding for its own HPC.

Until recently, the only researchers who had access to large-scale computing were those with access to supercomputers or large computing clusters. With the power of the cloud, the tools that scientists and engineers work with every day can become infinitely more powerful, accessing and manipulating more data and applying more complex calculations. This ability to provide easy access to data and computation at an arbitrary scale has the power to democratize research by making computing power available for the vast majority of researchers [2].

Virtually all research has become data-centric. Scientists and academics often must access and share large datasets and collaborate with other researchers in many disparate locations. These changes in how we go about scientific discovery are amplified by the growing trend toward interdisciplinary research. To truly address increasingly complex global issues such as climate change, genetic diversity, particle physics, study of the brain, and personalized medicine, researchers will be expected to develop ever more complex simulations and models. To do so, they will need to mine, search, and analyze huge sets of data in nearly real time and collaborate across disciplines at unprecedented rates. The ability to extract insight based on deep analysis of data and to collaborate is creating transformative change in research.

To address the massive challenge of extracting knowledge from vast data collections, researchers have turned to computing in the cloud. Variations exist for the definition of cloud concept. For purposes of this paper the NIST definition is used, cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction [6].

Complex large-scale simulations in cloud computing environments have great advantage. The cloud can scale and provision resources quickly as needed, run very large simulations that use many thousands of cores, while providing the ability to use pre-built models and functions without technical knowledge. Modern particle physics relies on powerful, large scale computing resources, with increasing emphasis on co-processors, such as GPUs. For most research groups, acquiring these resources is often expensive and inefficient. Based on the NSF funded, Exploring Clouds for Acceleration of Science (ECAS) program, "Particle Physics at the Largest Scale in Google Cloud" project, the cloud provides flexible, customizable, and scalable access to these resources. And, it can be used as an R&D testbed for groups that may be planning to purchase larger-scale, on-premises resources. Similarly, through the NSF ECAS "Deciphering the brain's neural code through large-scale simulations of cortical circuits" project, researchers were awarded significant cloud credits, and are able to develop AI methods along with improved brain modeling tools resulting in large scale models of areas of the brain that would have taken years. For these researchers, the cloud provides access to unlimited number of nodes, with no queues, and the ability to scale up resources as necessary. A fully customized SLURM Cluster is deployed, and kubernetes is used to deploy the tool for the wider community. ECAS accelerated the research and technology. Given funding support is a major barrier, more programs such as NSF ECAS are needed to help accelerate cloud adoption for research.

Significant challenges to widespread adoption of cloud for research exist. Cloud computing for research requires a different model for research support, just as it does for business information technology. Instead of one-time capital investments, cloud-based HPC requires ongoing financial support. Compute optimized, placement groups, bulk API for reduced latency are all great but also generate huge bills. With the cloud, you continue to get charged, however, the problems associated with supporting aging, obsolete equipment is eliminated. A lack of experience among researchers and IT support staff creates a fear of on-going costs. Cloud providers offer sophisticated tools for estimating, controlling, and reducing the costs of cloud usage, and provide training and simplified interfaces. The volume of users and the amount of standardization that cloud providers have developed, along with available tools have created a more common HPC infrastructure that enables a quick ramp-up for accelerating research productivity.

Network bandwidth poses another issue. Moving large data sets on commodity networks, or even on regional research and education networks, simply doesn't work well for hundreds of terabytes or petabytes of data, which is the scale required by researchers in many fields. Often researchers resort to shipping hard drives. Cloud services have been developed to support the process. The steps in the network path, including from the lab to the campus border, from the campus to the ISP or regional research and education network, and from the network provider to the cloud provider, can each pose a challenge.

As is the case with any use of cloud services, information security considerations play a role. Intellectual property protection, personally identifiable information, and export controls can all present issues, requiring appropriate review and analysis. The responsible use of cloud resources often requires collaboration with campus information security personnel. Just as with traditional enterprise computing, most research can be moved to a cloud provider if the appropriate protocols are in place, but some may benefit from being kept in the local environment. There may be good practical and financial reasons to stick with on-premise resources.

A promising strategy today is a hybrid environment that takes advantage of both local resources and cloud resources as appropriate. The tools such as Jupyter Notebook, Docker, and Kubernetes, make executing the same code in the local environment and in the cloud relatively easy. Researchers can develop and test their code using owned equipment, and then if they need to access more processor cores and memory, they can run their final experiments in the cloud once they're confident that they understand how it will work, thereby reducing the chance of incurring cloud fees for an unsuccessful run.

Most researchers use the cloud on their own with no centralized approach. Some institutions are working proactively to organize the approach for utilizing the cloud across the institution. Large scalability requires institutional backing and negotiation on restrictions, and proper elasticity and cost management.

From an institutional perspective, universities are engaged in various phases to adopt cloud computing service model. There is recognition of the value to the migration to the cloud in terms of online education, research, economic crisis, globalization, and high and constantly changing requirements [1]. Several key opportunities to be achieved include cost savings, flexibility and scalability, security and improved experiences. Universities need to be aware of cloud computing technology to be able to choose the right cloud model that is most appropriate for existing and future needs. The journey to cloud is not without storms. Challenges exist as institutions are adopting and navigating the cloud. The key barriers identified to this transition include budget and financial, cybersecurity management, procurement, integration complexity, and a skilled workforce.

Many institutions have defined a cloud strategy to best tackle their digital evolution. They are changing their business operating models to be more agile, innovative, and adaptive through cloud tools and services. While adopting the cloud may offer immediate value, it also introduces its obstacles and operational complexity. New tools, workflows, and resulting data streams can slow organizations down. There is also the inherent complexity of multi-cloud environments, which have become the norm as organizations adopt cloud-native tools and integrate them with their public cloud and on-premises environments. Public cloud is a commercial relationship that includes a complex ecosystem of dealing with individual stakeholders from funding agencies, to legal, and business entities. Ultimately, failure to adopt cloud with the right strategies will lead to failed digital transformation projects.

Different methodologies and results have been used in the adoption of cloud computing. Some of the critical success factors (CSFs) for the effective implementation of cloud at universities are data security, availability and reliability, customizable service level agreement, network bandwidth, compatibility, technical support, management support, human and resource readiness, complexity, cost flexibility, ease of use, and relative advantage [7]. The cost-effectiveness of cloud computing is one of its most prominent advantages that causes organizations to adopt and use cloud computing.

The cloud market has evolved significantly, the silos between the service categories, software as a service (SaaS), infrastructure as a service (IaaS) providers and platform as a service (PaaS), is blurring. With such a large number of cloud service providers, some institutions have turned to an external partner to help manage their services. Individual institutions can leverage the RENs to build relationships with cloud providers and cloud commercial partner solutions, as well as determining which cloud service provider meets the specific need. Various services are offered for improving network access to the big three public cloud providers. Organizations such as Omnibond and Internet2 offer tools and resources for navigating the cloud.

Availability of expertise for managing the cloud resources is a limitation. While there are training programs for the enterprise, there are limited training options for the research community. Formal, instructor led training programs are also provided by Internet2, NIH STRIDES, and cloud service providers. On-demand courses are available through Coursera, Linkedin Learning, Quicklabs. The cost of these resources are sometimes prohibitive for many smaller institutions.

The pandemic has accelerated what had already been a steady migration of IT workloads to the cloud over the past several years. Research community is relatively new to the cloud environment. The community needs best practices for getting started with the cloud. For example, guidelines and templates are necessary for how institutions are handling policy issues and procedures pertaining to cloud security that includes compliance with HIPAA regulations. Based on institutional policy, grant overhead for cloud services continues to be a challenge. CloudBank provides managed services to simplify cloud access for research and education, and resources are free of indirect costs for a subset of NSF solicitations that are CloudBank-eligible.

The integration complexity is a big challenge organization face when scaling cloud implementation across the enterprise, followed by the diverse skill sets required, service integration and management, data governance and management, regulatory compliance, and lack of standardization for cloud management tools and configuration. Although technology is evolving to address integration challenges, most enterprises want to become multi-cloud but know building the skills for it can be difficult, as they can't afford so many initiatives to ramp up skills of existing personnel. Challenges

exist as campus IT is striking the balance between the support requirements for the enterprise and those of the research community, especially areas such as security policies, identity access management, and using the cloud [5].

Policies are needed to make better use of cloud computing. A framework focused on the use of the cloud service model supported by plans, programs, budgets, and practices, is necessary. The power of the cloud comes from agility and limitless potential, while a goal is to empower users to get the most from what the cloud has to offer, institutions have a responsibility to do that in a safe way. Governance and security controls can be put in place for cloud, and changes could follow an exception process. Billing can be tied to a central billing system, eliminating the need for using credit cards. IT operations can partner with Sponsored Programs office to find grants and to jump start users' cloud experimentation.

In summary, the major themes regarding adoption of the cloud in research include costs; getting comfortable with cloud computing; weighing the advantages and the risks of cloud computing; reconciling cloud computing with data privacy; maintaining trust and anticipating the cloud by creating the conditions for cloud adoption. Lack of familiarity with cloud technology, prompts researchers to adopt an overly-cautious approach. Many tend towards gradual adoption of the technology through hybrid models, keeping their data stored in a local environment and using the cloud for computation. Some try-out cloud technology as a form of gradual adoption. Concerns over cost and security are ongoing, even though pre-configured platform resources have been developed by cloud service providers to facilitate the use of the technology for domain researchers.

## 3 RECOMMENDTIONS

Results of the findings of the ERN Broadening the Reach working group indicate important recommendations to institutions and the ERN and its partner organizations for helping institutions address the challenges identified in leveraging the cloud for research.

Institutional support is necessary for cloud adoption. Campus stakeholders need to be familiar with what the cloud is and what it can offer. From this point of view, it is necessary to be aware of cloud computing technology to be able to choose the right cloud model that is most appropriate in terms of information services and research community that the campus needs. Institutions need to be intentional about architecting their cloud solutions for transformation through investments in architecture and people who understand the systems to optimize their cloud solutions. Workshops and educational programs are necessary to raise awareness of existing best practices, tools, and resources.

***Recommendations to Funding Agencies:***

- Investing in workforce development projects that address the gaps for creating a diverse campus CI, cloud, and research computing expertise with skills that bridge domain science and computation.
- Expand funding support for programs to lower the barriers for cloud adoption and enable and foster regional collaborations to facilitate sharing, education, and outreach related to cloud technologies.
- Provide funding support that encourages universities to carry out research and development studies on cloud computing.

***Recommendations to the* ERN:**

- Facilitate cloud awareness campaigns and training (online and in-person) with accompanying educational material, including case studies for cloud adoption in research.
- Design a framework for leveraging the cloud that includes plans, programs, policies, guidelines, and best practices to address and overcome the major challenges in cloud adoption, while balancing the requirements of the enterprise and the research community.
- Leverage the partnership with the RENs to understand the nuances of workflow through the cloud
- Engage with the broader CI community, and advocate for a common definition and shared language of what is meant by cloud.

***Recommendations to Institutions:***

- Support participation of the campus community in conferences, and seminars on cloud computing. Organize information sharing through seminars and lunch and learns on cloud computing.
- Leverage the RENs and their relationships with cloud providers and commercial partner solutions.
- Establish and document institutional policies for the cloud implementation.
- Provide cloud migration consultation services to the research community to help overcome obstacles and make their path more about science and less about purchase orders.

## 4 OPPORTUNITIES AHEAD

The cloud has the potential to allow almost unlimited access to high-end computing resources for researchers at every type of institution, creating a more level playing field for experimentation than has ever existed before. While there are still challenges to using the cloud for research, the advantages in many cases are so compelling that research will continue to migrate in that direction. From large research universities to small liberal arts colleges, the research cloud will become a growing part of the research instrumentation portfolio, and campus IT departments will likely have to add cloud facilitation for research to the ever-growing list of capabilities expected.

Cloud computing awareness level, compliance of existing capabilities with cloud structure, the level of usage of cloud applications, and the extent of support provided by the university administration are all factors in adopting the cloud for research. Lowering the barrier for leveraging the cloud in research is complex and requires funding support; campus cultural transformation that promotes coordination, communication, and collaboration among multiple stakeholders; standards and policy requirements; ease of use; expertise; communities for information sharing and collaboration; and education and outreach, including awareness campaign and professional development about existing best practices, resources, and tools.

The knowledge gained from the ERN Broadening the Reach working group's community engagement efforts will be used to

determine how the ERN can support the needs of the institutions. Collaboration and partnership opportunities among the community for current and future ERN projects are in progress including the ERN infrastructure architecture design, building the next generation workforce to support important areas such as cloud adoption, designing workshops and webinars to raise awareness to existing resources and to gain further insights from the community. We welcome suggestions to explore opportunities to leverage the learnings and best practices identified to make cloud resources more widely available to researchers and educators.

## ACKNOWLEDGMENTS

The co-authors thank the ERN Working Group members for workshop program development and execution, the ERN Steering Committee for guidance and support, and the Rutgers University Office of Advanced Research Computing assistance for event management support. Special acknowledgement to John McNutt for his work in drafting the survey. The workshop was funded in part by NSF sponsored CC*CRIA: OAC-2018927.

## REFERENCES

[1] Hakan, Aydin, 2021. A Study of Cloud Computing Adoption in Universities as a Guideline to Cloud Migration. SAGE Open, vol. 11, 3 (July 6, 2021). https://doi.org/10.1177/21582440211030280. [Accessed 15 March 2022]

[2] R. Barga, D. Gannon, & D. Reed, D., (0AD), 2000. The Client and the Cloud: Democratizing Research Computing. IEEE Internet Computing, 15(1), 72–75. https://doi.org/10.1109/MIC.2011.20

[3] Michael Berman, 2019. Research Computing in the Cloud: Leveling the Playing Field. EDUCAUSE Review 54, no. 1 (Winter 2019). Available at https://er.educause.edu/articles/2019/3/research-computing-in-the-cloud-leveling-the-playing-field. [Accessed February 20, 2022]

[4] Ernrp.org. 2022. The Ecosystem for Research Networking. [online] Available at: https:/ernrp.org [Accessed 29 March 2022]

[5] Forough Ghahramani, John Hicks and Barr von Oehsen. 2021. Broadening the Reach for Access to Advanced Cyberinfrastructure: Accelerating Research and Education. In Proceedings of ACM PEARC21 conference (PEARC '21), July 19 -22, 2021. ACM, Boston, MA, USA, 4 pages.

[6] Liu, F., Tong, J., Mao, J., Bohn, R., Messina, J., Badger, L., Leaf, D., 2012. NIST cloud computing reference architecture: Recommendations of the National Institute of Standards and Technology (Special Publication 500-292). Available at https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication500-292.pdf. [Accessed March 01, 2022]

[7] Q. N. Naveed, N. Ahmad, 2019. Critical success factors (CSFs) for cloud-based e-learning. International Journal of Emerging Technologies in Learning (iJET), 14(01), 140–149. https://doi.org/10.3991/ijet.v14i01.9170. [Accessed 05 April 2022]