Best Practices on Using the Cloud for Computing Research

CRA-I
Computing Research Association
Industry

Best Practices on Using the Cloud for Computing Research

Workshop Organizers: Fatma Özcan (Google), Vivek Sarkar (Georgia Tech)

Helen Wright (CRA-Industry Senior Program Associate)

June 2022

CRA-I

Computing Research Association

Industry

1 Executive Summary

Cloud computing has emerged as critical infrastructure for most/all enterprises worldwide, and has also become essential for computing research, both for advancing the frontiers of computing and for advancing research in computing-reliant disciplines. It has introduced a revolutionary model in which computing resources can be accessed as a utility to store, manage, process, and share data thereby enabling novel workloads and efficiencies of scale; this is in contrast to the traditional model of enterprises relying entirely on acquiring and maintaining in-house local servers and personal computers.

The Computing Research Association's newest committee, Computing Research Association-Industry (CRA-Industry), held its first workshop on Best Practices on Using the Cloud for Computing Research in March 2022. This workshop built on a very successful virtual roundtable discussion in September 2021 roundtable event[1], with 5 panelists and over 50 members of the community in attendance. The goal of the workshop was to build on the momentum from the roundtable discussion, and identify best practices on using the cloud for computing research in three different focus areas: research, education, and collaboration.

This report introduces background on cloud computing, discusses challenges, opportunities, and synergies for the focus areas, and summarizes the following recommendations from the workshop:

- Encourage research that focuses on extensions to cloud computing including: hybrid cloud, multi-cloud, and local clouds, research that extends the reach of the cloud to "the edge" and vice-versa, and research that exploits the geospatial distribution of cloud resources.
- Continue to educate and inform the broad computer science community and various related government agencies about the definition of cloud, the inherent properties and tradeoffs in cloud computing, and the importance of cloud access for research.
- Develop a broad ecosystem that encourages collaboration, community and access to computing.

Achieving the above recommendations will require close partnership across industry, academia and government, and CRA-Industry is well positioned to convene discussions and help enable this partnership. Cloud computing is currently at an important juncture in the evolution of its technology innovation. Industry has established cloud computing broadly due to ease of application development and powerful capabilities. This has created gaps and significant opportunities for academia to evolve both curriculum and research approaches to contribute to advances in cloud computing.

2 Introduction

Cloud computing has made a profound impact on the computing research community and the field of computing as a whole. It has introduced a revolutionary model in which computing resources can be accessed as a utility to store, manage, and process data, rather than acquiring and maintaining local servers and personal computers. The use of the cloud also makes it possible to share data at a greater rate than ever before.

The Computing Research Association's newest committee, Computing Research Association-Industry, held its first workshop on Best Practices on Using the Cloud for Computing Research in March 2022. It brought together 30 participants from industry, academia, and government in a hybrid format with on-site participants meeting in Washington, DC. This workshop built on a very successful September 2021 roundtable event[2], with 5 panelists and over 50 members of the community in attendance. The goal of the workshop was to build on the momentum from the roundtable discussion, and identify best practices on using the

---

[1] https://cra.org/industry/events/virtual-roundtable-on-best-practices-on-using-the-cloud-for-computing-research/

[2] https://cra.org/industry/events/virtual-roundtable-on-best-practices-on-using-the-cloud-for-computing-research/

cloud for computing research in three different focus areas: research, education, and collaboration. This report summarizes the findings and recommendations from the workshop, and aspires to be a resource that can help motivate future activities in this area across industry, academia, and government.

3 Background: the Cloud Revolution

There is a long history of successful efforts by the National Science Foundation (NSF) and other federal agencies to provide national access to computing resources to wide and diverse sections of academia. However, cloud computing is far more than just remote access to a large collection of efficiently managed computing resources. At every level, cloud computing provides capabilities in the form of systems, networking, and software solutions that are far beyond what is available in computing research environments, arguably in even the most advanced on-premise computing facilities, and yet these capabilities are now mainstream and accessible worldwide. There are many factors to consider in the cost trade-offs between on-premise vs. cloud computing; while there may be instances when cloud-based approaches are more costly with respect to compute services, on-premise computing often includes additional costs related to staffing and under-utilized capacity.

Just considering the cloud's computing capabilities, all the public cloud vendors provide state-of-art server platforms with 100s of cores, terabytes of RAM, 100 - 400 gbps NICs, and leading edge SSDs. Extensive concurrency is the norm, not the exception, and high degree of scale-out and scale-up server configurations are routine. Memory hierarchies are extremely sophisticated and are broadening into more general fabrics, with chiplets, shared last-level caches, direct-IO, and increasingly disaggregated memory tiers. Interconnects are fast and wide, with rapid adoption of each PCIe generation, soon to be augmented with CXL (Compute Express Link). Accelerators, including GPUs and TPUs for massive machine learning, smartNICs, encryption, authentication, encoding/decoding and more are absolutely commonplace, along with modern security practices from the moment the BIOS boots. And, a layer of virtualization lies over all this, providing isolation and elasticity, while also introducing its own influences. Operating at immense scale, new generations of resources are deployed at a faster than usual pace, with the design process reaching out years ahead of what is available on the market from chip and technology vendors, and the provisioning and maintenance processes carrying numerous past generations forward till they eventually phase out.

While a cloud computing node today is scarcely recognizable from computer science textbooks, it is but a building block. All public cloud vendors scale this building block in the millions, with massive clusters of computers carefully engineered into data centers of enormous scale, and these are replicated globally to form dozens of regions, extended with hundreds of PoPs and edge sites, interconnected with global communications backbones each on the scale of the public internet. Clouds connect at all tiers and have vast peering capabilities. Orchestration of large distributed collections of microservices is the norm, as is attending to availability, reliability zones, load balancing, replication, and demand management.

Along with the consolidation of these technological advances in cloud architectures, the entire ecosystem of software gravitates towards them. In addition to the run-time environments, VMs, storage, communications and orchestration offered by the cloud vendors, we find foundations of ML frameworks, containers, language infrastructure, data analytics, and more. Cloud infrastructure stacks support automation, streaming, DevOps, diagnosis and debugging, transactional database services, as well as entire ecosystems around monitoring, dependability, security, compliance, data governance.

This massive, distributed, global scale, rapidly evolving infrastructure and ecosystem, engineered to serve billions, is simply how computing is routinely done today, not the model of LANs of laptops and desktops connected to a departmental server that dominated enterprise computing in late 20th century and still persists in academia for research and teaching.

The impact of cloud hardware and networking advances on the software ecosystem called out above is deep and lasting, and requires us to rethink the foundations of many disciplines that we teach and research in our universities, including database systems, operating systems, and programming languages. In fact, there are important HCI challenges and opportunities to be resolved for cloud environments. Quite simply, university curricula are out of date because they have not changed to reflect how cloud services are being built, deployed and managed by the cloud vendors and software vendors (who think cloud-first when building applications); nor do these curricula help potential users of these services understand the principles and complexity underlying cloud computing, including how to take advantage of its unique capabilities, e.g., elasticity, while being aware of and working around its weaknesses, e.g., impact of tiered data access and the immaturity of cost controls and training.

As an example, consider database management systems, which have been taught at both undergraduate and graduate levels for decades based on the classic relational model and how it is implemented in traditional SQL databases. Thanks to the cloud, we now have entirely different storage stacks for both operational and analytic databases; there have been radical changes in how we approach database design and concurrency control; and indeed, there are new families of database systems, e.g., geo-replicated databases, that couldn't exist without cloud infrastructure. Likewise, the cloud has revolutionized how we think of data governance–nowadays, a global catalog of all data across filesystems, multiple clouds, and numerous instances of diverse database systems is considered a core foundation. Contrast this to a traditional world where each relational DBMS managed its own catalog, while catalogs spanning multiple database systems were an after-thought. Finally, from the perspective of database developers and users, the principles of database design and tuning need to be re-thought from the ground up given the new norms of unlimited and elastic storage; separation of storage from unlimited and elastic compute; and the proliferation of specialized data engines for transactional uses to data science to warehouse analytics that all seek to address the same underlying data. Traditional curricula clearly need to be revamped significantly, and likewise, there is significant opportunity for the research community to shape cloud services of the future.

Cloud computing enables important, novel workloads and is playing an increasingly important role in research, both as the subject of computer science exploration and as an instrument to advance research across multiple compute-reliant disciplines. Private clouds, an evolution of the information technology (IT) data centers, offer centralized services to a specific community such as an enterprise or a university. A hybrid cloud is one in which applications are running in a combination of different environments. Commercial clouds are attractive to users because they offer elastic compute, networking, and storage resources at scale while outsourcing the management of those resources. Commercial clouds share resources across a broad set of users, and open new research frontiers while enabling previously unimaginable workloads. The cloud has introduced new practices in software engineering and focused on specific elements of the computer science curriculum that deserve increased attention. In combination, these factors imply that commercial cloud access is a gating factor for relevance and impact of computing research across both computing and a broad range of computing-reliant disciplines.

4 Cloud+Research: Challenges, Opportunities, Synergies

Cloud computing creates new challenges, opportunities, and synergies for computing research. The convenience of outsourced IT operations comes with a consumer-producer tension that is common to all shared research instruments, insofar as the operators limit access and usage models for the shared infrastructure so as to ensure security, availability, and maintainability for large numbers of users. The tension is exacerbated by the "vendor lock" phenomenon in which it can be challenging to use cloud resources from different vendors interchangeably. This tension is not easy to resolve, in part because the many subfields of computer science leverage and focus on different aspects of the cloud. For example, HPC researchers need to leverage compute (whether general-purpose or specialized) at scale; networking researchers leverage distribution and access to fiber and edge resources; and storage and database research builds on yet other advantages. Underpinning the tension is the fact that commercial clouds, sometimes called "hyperscale" clouds, invest billions in infrastructure and operations that cannot be rivaled by any research institution. Given the diverse focus of various subfields, it is important to distinguish the researchers who need access to "bare metal" and those who do not.

Because cloud computing is a necessary, fertile, and relatively untapped ground for computer science research, addressing these challenges is a key opportunity for the community. It is especially important to support research opportunities that the cloud/IT industry may not naturally prioritize: ideas that disrupt the cloud itself, ideas that radically reduce costs for research institutions, and ideas that increase research access to otherwise proprietary infrastructure. Longer-term opportunities include advances and abstractions that enable cost savings and specialized support in hybrid clouds, multi-clouds, and local clouds; research that extends the reach of the cloud to "the edge" and vice-versa; research that exploits and restricts the geospatial distribution of cloud resources. These advances will likely consider a combination of networking, storage, and compute across all layers of the stack; for example, data analytics of large, common data sets would today require each customer of a commercial cloud to replicate and conduct secure experiments on private copies of a scientific dataset, while the right abstractions would support optimizations that share (without compromising security or privacy) intermediate results, locally cached datasets, and optimized movement/access across policy-based access to cloud networking resources. These kinds of advanced methods could democratize large-scale scientific research. Finally, to the extent regulation and policy can represent both constraints and tools in the context of cloud computing (e.g., GDPR, security compliance, sovereign cloud), it may be important to consider multidisciplinary collaborations that can steer the cloud toward equitable access by both researchers and cloud users.

In the short term, the computing research community needs to educate faculty, students, and agencies about how the cloud abstraction should be defined and understood in a vendor-independent manner, its inherent properties and tradeoffs, the importance of cloud access for research, and the evident equity issues given the current state of access. These aspects of the computing research ecosystem are critical to consider while thinking about how to support the next advances in computing, given that it is economically infeasible to replicate "hyperscale" capabilities for research purposes. Today, the NSF is supporting research access by subsidizing the use of commercial clouds through the CloudBank accounting interface, but CloudBank itself could potentially grow to serve as a testbed and a platform for new abstraction layers across and between clouds, and commercial cloud providers could potentially offer specialized research interfaces to CloudBank to open the range of supportable experiments. Commercial providers might also be inspired to support economic experimentation, for example, by offering flexibility in pricing models related to data egress. Lessons learned from prior shared research infrastructure (e.g., PlanetLab, GENI) may apply – in particular, designing a computing research testbed itself is in fact almost synonymous with conducting some of the desired research.

5 Cloud+Collaboration: Challenges, Opportunities, Synergies

Cutting edge research frequently requires large datasets that are difficult and expensive to replicate to multiple on-premise locations. Many on-premise locations simply don't have the local capacity to store the desired datasets. The cloud can provide both the necessary capacity and a "single instance" location where many parties can both independently and collectively process large (shared) datasets.

Aside from enabling the cost of storing datasets to be amortized over many consumers of those datasets by avoiding replication and the associated data transfer copy costs, clouds also offer the potential to amortize computational costs by various means. These include techniques such as caching the results of common analysis (sub)queries across multiple users, intelligent caching of data on computational hosts across multiple users, and batching and ordering of computational tasks so as to most efficiently use shared resources. Similar opportunities may exist for training multiple ML models on such datasets.

Certain impediments must be overcome in order to enable this kind of collaborative sharing of datasets in the cloud. Shared use requires potentially fine-grained tracking of usage to enable fair allocation of costs among users. This, in turn, requires detailed provenance tracking for all derived data — i.e. data generated through computational/analytic means from other data. Shared must be enabled in a manner that still preserves privacy and security requirements for the data.

It is also important to consider synergies between industry and academia in computing research. Many of the synergies are obvious, for example, academic research can offload industry research and in some cases may be able to anticipate industry needs and opportunities. However, industry can and does move more quickly on opportunities that are clearly recognized, so it is important to ensure that computer science researchers formulate problems that will not be mooted in the short term. Another synergy is between research and customer use cases – there are times when academia considers workloads, use cases, and challenges that foreshadow commercial customer uses; in this regard cloud providers are highly motivated to support academic researchers as cloud users. A third synergy comes in the area of curriculum. Cloud providers are motivated to encourage the development of an educated workforce; meanwhile academia is motivated to ensure that it provides a relevant modern curriculum. Traps and pitfalls exist in the distinction between products and principles, but to the extent cloud computing emphasizes challenges that have evolved from prior generations, it would be worthwhile to collaborate on a curriculum that incorporates cloud computing concepts such as consistency, privacy, devops, and distributed systems at scale. Finally, in terms of synergies in collaboration, it is helpful that open source projects form the basis of significant fractions of the cloud computing stack and can form as a basis of synergies in research collaboration.

6 Cloud+Education: Challenges, Opportunities, Synergies

The impact of cloud hardware and networking advances on the software ecosystem in the 21st century is deep and lasting, and requires us to rethink how the foundations of computing is taught in our universities, starting with the first lecture that students hear on “What is a computer?” Courses that will be impacted deeply include those on database systems, operating systems, and programming languages, since they need to explain the principles of cloud computing and its unique capabilities, e.g., elasticity, as well as its limitations, e.g., impact of tiered data access and the immaturity of cost management. Further, the cloud abstraction should be defined and understood in a vendor-independent manner, even though course projects and practicums will be hosted on vendor-specific cloud platforms.

In a nutshell, the cloud fundamentally changes the way we need to think about computing and its existence requires systemic changes to computer science education at the high school and university level. There are a number of obvious benefits to using the cloud in the classroom. These include the ability to provide uniformity and equity of experience for students, the access to compute, data, and specialized resources that may not be available locally or through the education institution, and the ability to collaborate among students.

Understanding how all the familiar computing systems work today (e.g., mobile apps, social networks, etc.) requires understanding fundamentals of distributed computing and knowledge of cloud computing. These concepts historically are introduced in relatively few upper division university computing courses related to advanced operating systems concepts. It is necessary to rethink how the most familiar and basic computing experiences (like sharing a photo with friends from a mobile device) are explained to students early in the computing curriculum (high school and freshman level) in terms of the cloud. Many of the upper division systems courses have to be revisited to account for how cloud computing has impacted them. For example, explaining database systems, operating systems, software engineering, etc.

Resources for research projects increasingly require the cloud to leverage cloud resources (computing, data, etc.), economies of scale, opportunities for sharing results. The widespread use of AI technology in computing, for example, relies heavily on understanding cloud computing both to understand where AI models come from (how they are trained) but also how they can be used across mobile and edge devices (where inference happens in the cloud). Many current and future job and business opportunities require knowledge of the cloud to make informed decisions, but students graduating today from computer science programs don’t necessarily have cloud skills and knowledge required to be effective and ready to hit the ground running when they enter the workforce. It is clear that cloud provider companies benefit from educators providing students with a basic understanding of cloud technology and how to leverage it.

There are a number of challenges when it comes to educating and using the cloud in the classroom. For example, the cost of using the cloud for classrooms/teaching can vary considerably depending on resources required (< $100 per student a semester to over $200-500 per student a semester) , e.g. storage, virtual machines, GPUs, etc. There is a need to understand which functionalities and resources are affordable for teaching/education purposes and a need to have a way of managing costs. Another challenge is converting a course to incorporate cloud requires a dedicated effort upfront by faculty and TAs to develop. There is a high barrier to entry for faculty/students that do not have experience with cloud and sustaining a class after it has moved to cloud needs continued investment including a dedicated cloud budget or credits. Finally, the use of the cloud may be at odds with current campus IT infrastructure and staffing investments.

Academia should recreate an incentive structure that connects companies (subject matter experts) with academics to co-create educational materials that embed an understanding of cloud computing at all levels of computing education. They should encourage a modular approach to curriculum that can allow sharing and composition of materials more freely. Materials should be made widely available and include “quick-start” on-ramps for educators that want to start using the cloud in their courses without any prior experience. Finally, there should be incentives for educators to both teach about the cloud in their courses, increase their understanding of how cloud computing can benefit their instructional mission, and utilize the cloud for teaching when appropriate.

Industry can help as well by supporting new course development and encouraging academic institutions to teach cloud classes by funding cloud for student offers and providing discounts. They can encourage participation by subject matter experts, potentially their own employees, in course development and guest lectures. Companies could set up an incentive program for their employees to “give back” to their alma maters and teach special classes. Industry could share their training content and provide tutorials on cloud and developing cloud workflows.

Government could play a key role in developing cloud-based academic curriculums by providing funding. A potential long term, sustained funding commitment (multi-year) for credits or cloud teaching, could make a huge impact. It is critical that funding agencies provide grants to support innovative approaches to enhancing education through the use of cloud. More programs like the Improving Undergraduate STEM Education: Computing in Undergraduate Education (IUSE: CUE) solicitation (NSF 22-588)[3] are essential are departments wanting to revamp their curricula to better account for cloud programming and cloud systems topics.

A key synergy comes in the area of curriculum. Cloud providers are motivated to encourage the development of an educated workforce; meanwhile academia is motivated to ensure that it provides a relevant modern curriculum. Traps and pitfalls exist in the distinction between products and principles, but to the extent cloud computing emphasizes challenges that have evolved from prior generations, it would be worthwhile to collaborate on a curriculum that incorporates cloud computing concepts such as consistency, privacy, devops, and distributed systems at scale.

Finally, the Computing Research Association and other organizations can help by creating regular opportunities for educators, researchers, and cloud computing providers to share perspectives and experiences to benefit the teaching mission through workshops and events like the one that led to this report. They can regularly survey computing professionals on 1) how they use the cloud, 2) what impact it has on their profession, 3) perspectives on what additional information about the cloud should be taught, 4) what weaknesses they see in how the current curriculum teaches students about cloud computing.

---

[3] https://www.nsf.gov/pubs/2022/nsf22588/nsf22588.htm

7 Broader Opportunities

7.1 Cloud for All

While addressing the access and education challenges for the field as a whole, we have the distinct opportunity to simultaneously address long standing resource gaps for traditionally underserved institutions. These resource gaps include education and training materials as well as compute resources. The majority of US academic institutions today also lack sufficient education and training materials and compute resources on cloud computing. Common challenges across these institutional contexts include determining how to make room for cloud computing topics in the curricula, which topics are most important to include, what levels of cloud access are necessary to support the curricula, and how to gain those levels of access. By prioritizing traditionally underserved institutions in development of respective solutions, we can not only address historical gaps, but also position ourselves to produce working solutions for majority, minority, well-served, underserved, (all) institutions.

It is likely the case that the level of scaffolding necessary across institutional contexts will vary, and that curricula developed for traditionally underserved institutions will necessarily integrate more scaffolding than those developed for institutions operating at the state of the art in research. However, in terms of educational design, reducing the scaffolding of a well-formed curriculum (a common basis) is a much simpler problem than attaching scaffolding to a relatively advanced curriculum.

Without prioritizing traditionally underserved institutions, we can expect the existing gaps only to grow. Allowing the gaps to persist and grow is detrimental to workforce development, and further translates to losses in innovation and productivity in research and in industry.

One immediate recommendation for moving forward in this direction is to form partnerships with traditionally underserved institutions including but not limited to HBCUs, tribal colleges, and others. One example that could be adopted and adapted to cloud computing is the IBM-HBCU Quantum Center initiative.[4]

Further, Identifying and enabling particular faculty members of these institutions to lead in the cloud for computing research is a critical, concrete short term step. Note that enabling faculty participating (and leadership) must include offsetting the relatively greater cost of their participation. It cannot be assumed that the same level of human resources (e.g. secretarial assistance, students) and financial resources (e.g. travel budgets, laptops and student stipends) as are available at larger U.S. institutions are necessarily available to them. Having identified potential faculty leaders, the next step would be to work with them to identify supplementary support that would enable them to participate actively and sustainably (and lead).

7.2 Social Responsibility

Cloud computing is not unique in raising social responsibility issues: for society as a whole, for computing-dependent researchers, for computer science researchers, and for computer science educators. In some cases, cloud computing is clearly advantageous, for example in the context of environmental impact, virtualization and server consolidation have a clear lifecycle benefit measurable in the industry's carbon footprint; cloud computing workloads likewise have had positive environmental impact, for example, through more efficient mapping and vehicle guidance. But the carbon footprint of the industry is nonetheless growing, and there are research opportunities in further optimizing the environmental impact of the cloud. In areas such as education, the cloud supports democratizing applications (cf. Khan Academy which many credit as an equalizing force in secondary education).

But although many of the cloud's consumer services are subsidized with advertising, the scientific enterprise calls for access that can be costly and (especially in the case of fundamental computer science research) require services and access that

---

[4] https://www.ibm.com/blogs/research/2021/02/ibm-hbcu-quantum-center-expands/

is available only to wealthy institutions and/or well-connected faculty. Attention is needed to ensure the democratization of access through subsidies and initiatives that create opportunities to create and leverage new cloud abstractions. Replicating public cloud capabilities is beyond the economic means of the NSF and government agencies, so a different approach is needed. NSF's CloudBank program has taken initial steps to provide low-cost access to NSF researchers, but is primarily useful for consumers of cloud computing as opposed to cloud computing researchers. A new approach is needed that supports computer science research "within" state-of-the-art public clouds.

Cloud computing has implications for equitable research. Moving curriculum of any kind to the cloud can be democratizing (cf. Khan Academy) but requires effort and financing. Providing curriculum about modern cloud practices is advantageous to those with the finances to access the cloud, but the computer science curriculum has not yet caught up to modern cloud practices and needs to be refreshed. That same curriculum is likely to require substantial cloud access by students both as users and as researchers; it needs to be made available in a variation of "universal access" by students.

Finally, researchers in many disciplines are likely to imagine new applications and workloads of the cloud. To avoid the unfortunate missteps that have been seen in certain communities (e.g., computer vision / face recognition, ML / sentencing guidelines) it is essential that diverse stakeholders be at the table in both research problem formulation and execution.

8 Recommendations

There were a number of potential opportunities discussed by the workshop participants, and we summarize below the key recommendations that were selected to implement the best practices detailed above.

8.1 Research

Need to explore research directions in which the cloud becomes the foundation of all computing. This includes addressing key challenges for the future such as post-Moore computing, scalability, and security & privacy. Potential directions include advances and extensions to cloud computing such as hybrid cloud, multi-cloud, secure clouds, local clouds; research that extends the reach of the cloud to "the edge" and vice-versa; and research that exploits the geospatial distribution of cloud resources. This can be done by contributing funding in support of research solicitations from various government agencies with multi-year goals and clear objectives for sustained growth.

8.2 Education

Need to continue to educate and inform the broad computer science community and various related government agencies about the definition of cloud, the inherent properties and tradeoffs in cloud computing, and the importance of cloud access for research. This can be done by encouraging academic institutions to incorporate more cloud-related content in their classes. Incentives can also be created that connect companies (subject matter experts) with academics to co-create educational materials that embed an understanding of cloud computing at all levels of computing education. In addition, sustained funding commitment for credits for cloud teaching could make a huge impact on the future workforce engaged in computing research and development.

8.3 Ecosystem

Need to develop a broad ecosystem that encourages collaboration, community and access to computing. This can be done by using the cloud to share datasets and applying them to reproduce computations launched by multiple collaborating users. Organizations, like CRA-Industry, can create regular opportunities for educators, researchers, and cloud computing providers to share perspectives and experiences through workshops and events. These workshops can be similar to the one that created this report. Finally, this ecosystem can be developed by supporting equitable research and education opportunities to provide uniform experiences for all students and researchers across a range of institutions whether their access is on-site, remote, or hybrid.

9 Conclusions

Achieving the above recommendations and goals will require close partnership across industry, academia and government. For example, industry can contribute funding for cloud-related research and course development, cloud computing credits, and subject matter expertise; academia can invest more faculty and student resources in cloud-related research and education activities; and government agencies and labs can launch new programs to increase the level of attention being paid to advance the foundations of future cloud computing and its transition to practice.

In general, cloud computing is at an important juncture in the evolution of technology innovation. Industry has established cloud computing broadly due to ease of application development and powerful capabilities. This has created gaps and significant opportunities for academia especially in the field of computer science to evolve both curriculum and research approaches to contribute to advances in cloud computing. The need for curriculum development is critical in order to train a cloud ready workforce across all sectors that are technology dependent. The need for diversification of computer science research approaches to deeply incorporate cloud research is important to advance the science behind cloud computing in a manner that is free of commercial interests. Industry bears a responsibility to partner with the academic computer science community to ensure access to cloud infrastructure such that this research is possible. Industry also bears a responsibility to freely share training material that can be leveraged for developing cloud based curricula.

In all of this, it is important to recognize the inequities and tiered availability to vast cloud computing resources and ensure participation and access to resources for traditionally underserved populations such as those who are enrolled at HBCUs, HCIs and other MSIs.

10 Appendices

10.1 Resources

- Best Practices on Using the Cloud for Computing Research (https://cccblog.org/2021/11/03/best-practices-on-using-the-cloud-for-computing-research/)
- A National Discovery Cloud: Preparing the US for Global Competitiveness in the New Era of 21st Century Digital Transformation (https://cra.org/ccc/wp-content/uploads/sites/2/2021/04/CCC-Whitepaper-National-Discovery-Cloud-2021.pdf)
- SmartFarm: Computing Research for the Next-Generation of Precision Agriculture (https://cra.org/ccc/wp-content/uploads/sites/2/2020/02/Krintz-AAAS20.pdf)
- A National Research Agenda for Intelligent Infrastructure: 2021 Update (https://cra.org/ccc/wp-content/uploads/sites/2/2021/01/A-National-Research-Agenda-for-Intelligent-Infrastructure_-2021-Update-FINAL.pdf)
- NSTC Subcommittee Report: Recommendations for Leveraging Cloud Computing Resources for Federally Funded Artificial Intelligence Research and Development (https://cccblog.org/2020/11/20/nstc-subcommittee-report-recommendations-for-leveraging-cloud-computing-resources-for-federally-funded-artificial-intelligence-research-and-development/)

10.2 Workshop Participants

First Name	Last Name	Affiliation
David	Bader	New Jersey Institute of Technology
Vinamra	Benara	UC Berkeley
Cristian	Borcea	New Jersey Institute of Technology
Elizabeth	Bruce	Microsoft
Ron	Chiang	University of St. Thomas
David	Culler	Google
Marquita	Ellis	IBM Research
Robert	Fatland	eScience Institute, University of Washington
Ian	Foster	Argonne National Laboratory
Bill	Gropp	University of Illinois at Urbana-Champaign
Giulia	Guidi	UC Berkeley
Peter	Harsha	Computing Research Association
Eric	Kolaczyk	Boston University
Chandra	Krintz	UC Santa Barbara
Ed	Lazowska	University of Washington
Dan	Lopresti	Lehigh University
Vani	Mandava	Microsoft
Margaret	Martonosi	NSF
Fatma	Ozcan	Google
Alexander	Pacheco	Lehigh University
Smruti	Padhy	University of Texas at Austin
Manish	Parashar	NSF
Raghu	Ramakrishnan	Microsoft
Lavanya	Ramakrishnan	Lawrence Berkeley National Lab
Chris	Ramming	VMware
Vivek	Sarkar	Georgia Institute of Technology
Ann	Schwartz	Computing Research Association
Gurdip	Singh	NSF

Shava	Smallen	SDSC/UCSD
Divesh	Srivastava	AT&T Research
Ion	Stoica	UC Berkeley
Vaidy	Sunderam	Emory University
Marvin	Theimer	Amazon Web Services
Jeffrey	Vetter	Oak Ridge National Laboratory
Wenwen	Wang	University of Georgia
Crystal	Welliver	University of Washington
Jack	Wells	NVIDIA
Pam	Williams	LMI
Helen	Wright	Computing Research Association
Ben	Zorn	Microsoft

10.3 Acknowledgements

The organizers would like to thank all who have contributed to the Best Practices on Using the Cloud for Computing Research roundtable and workshop. Your input was invaluable to crafting this report for the community. In addition to all the workshop participants listed above, we would like to thank Kishore Ramachandran and Mary Hall.

10.4 Acronyms

BIOS	Basic Input/Output System
CPU	Central Processing Unit
CRA	Computing Research Association
CXL	Compute Express Link
DevOps	Software Development and IT Operations.
GENI	Global Environment for Network Innovations
GPU	Graphical Processing Unit
HBCUs	Historically Black Colleges and Universities
HPC	High Performance Computing
IT	Information Technology
LANs	Local Area Network
ML	Machine Learning
NSF	National Science Foundation
PCIe	Peripheral Component Interconnect Express
PoPs	Post Office Protocol
RAM	Random Access Memory
smartNICs	Smart Network Interface Card
SSD	Solid-State Drive
STEM	Science, Technology, Engineering, and Mathematics
TPU	Tensor Processing Unit
VMs	Virtual Memory System

NOTES

NOTES

NOTES

CRA-I
Computing Research Association
Industry

1828 L Street, NW, Suite 800
Washington, DC 20036
www.cra.org/industry/