Enabling Computer and Information Science and Engineering Research and Education in the Cloud

Cloud computing has the potential to transform both research and education in the CISE (computer and information science and engineering) community. The CISE directorate of the National Science Foundation convened a workshop on January 8-9, 2018, to bring together representatives from academia, industry, and government to discuss ways to enable CISE research and education to most effectively use the cloud. The workshop agenda and list of attendees appear as appendices at the end of the report.

Introduction

Cloud computing has revolutionized the way software and hardware resources are acquired and used in every sector. Every company in every sector now looks to the cloud as the means for storing and processing their data and as the means for running their applications. Cloud providers stand up data centers running state-of-the-art processors (e.g., GPUs and FPGAs), storage, and networking, and state-of-the-art services (e.g., machine learning algorithms and models). These resources benefit customers of cloud providers. As more and more companies make their internal processes and external businesses increasingly data-driven, the demand for cloud capability will continue to grow.

But what about the CISE (Computer and Information Science and Engineering) community of researchers and educators? The CISE community, and the academic community more broadly has been slow to move to the cloud, and thus is not accruing the intellectual and financial benefits of using the cloud that other sectors recognize. Academia is falling behind.

The purpose of this workshop was to explore the idea of an *academic cloud,* a cloud that provides the set of services and capabilities that serve the unique needs, workloads, and users of the CISE community. Given the resources and talent needed to build and maintain a professional cloud, our expectation is to look to commercial cloud providers as a way to provide the infrastructure that would be the basis of an academic cloud. However, since commercial cloud providers have not addressed the academic sector directly, there is a gap between the services they provide and the requirements academia has in using the cloud for research and education. A major goal of this workshop was to understand this gap and how to best address it. To gain a qualitative understanding of the gap, the workshop collected ideas, insights, and experiences from a collection of experts across a range of constituencies.

The mission of an academic enterprise is to do research to advance knowledge and to educate the next generation. Current cloud services are tailored for enterprise customers, which have different kinds of users, usage patterns, and needs. This workshop initiated a dialogue between the CISE community and cloud providers to explore what the barriers to cloud adoption by academia are and to argue the unique characteristics of academic users. Government agencies, which funds most of academic research, have a vested interest in ensuring their investments, especially in computing infrastructure, have the most impact. Moreover, they can serve both as a convener and as a neutral party in the dialogue between academia and industry. Hence, this workshop had representatives from all three parties: academia, industry, and government. The workshop, sponsored by the NSF CISE directorate, took place January 8-9, 2018, in Alexandria, Virginia.

Why the cloud?

Cloud platforms provide on-demand, elastic, and self-serve access to resources at scale and are thus capable of supporting large-scale and/or big data computing. They provide access to contemporary hardware and advanced software stacks, with users “riding the technology curve” as new technologies are made available in the cloud. The on-demand nature of access to resources can provide a “fast path” to computing—acquiring cloud resources is much faster than buying, installing, and operating on-premise, localized hardware. Cloud platforms can also support important new modalities of data, such as streaming data and real-time analytics on such data. CISE researchers and educators can leverage modern cloud platforms to accelerate and improve their research and teaching, instead of building and deploying dedicated local infrastructure. (And, beyond CISE, an academic cloud can benefit other academic disciplines.)

However, CISE researchers and educators face various challenges in using the cloud effectively and efficiently. These barriers include costs (including concerns about run-away costs) compared to in-house resources, mismatches between cloud pricing models and the cycles of academic research (measured in terms of the length of government-funded grants) and education (measured in terms of academic semesters), the need for users to navigate a bewildering array of cloud offerings, training and user support, ensuring ongoing access to software and data, and more. Overcoming these and other barriers will require concerted and coordinated effort by several important stakeholders, including cloud providers, academic institutions (including faculty as well as campus CIOs), and government funding agencies. In this report, we outline the main recommendations from this workshop:

1. **Articulate the case for academic institutions to use the cloud**: We recommend preparing a detailed write-up on the rationale for academic institutions to make more extensive use of the cloud, including key steps that campuses can take to move in that direction. The document should include a detailed cost comparison between cloud and on-campus resources, for an “apples to apples” comparison that can help CIOs and university administrators to make informed decisions about how to best support CISE research and education on their campuses.

2. **Articulate the “business case” for cloud providers to support academic users**: We recommend preparing a detailed write-up on the “business case” for cloud providers to

invest in better supporting academic research and education, including key steps that cloud providers can take in that direction. The document should include a quantitative treatment of the size of the academic market, as well as examples of “success stories” of academic use of cloud computing. It needs to make a clear return-on-investment case for cloud providers. It needs to be clear about why the academic community is different from a typical enterprise customer.

3. **Remove artificial costs that make cloud computing less attractive**: All stakeholders should remove artificial incentives for buying and using local equipment. Academic institutions can expose the true costs of owning and operating dedicated equipment, work toward waiving overhead charges on cloud usage, form bulk agreements with cloud providers, and include cloud credits in faculty start-up packages. Cloud providers can create new pricing structures such as up-front payments and long-term storage, as well as better tools for hierarchical allocation of credits and dashboards for managing cloud usage. Funding agencies such as NSF can work toward removing overhead charges on cloud usage, enabling cloud credits to outlive individual grants, acquiring cloud credits in bulk for academic users, and subsidizing data storage for data sets that are not cost-effective for cloud providers to store for free but are important in academic research.

4. **Create support structures for academics transitioning to the cloud**: The stakeholders should simplify the transition to the cloud for academic researchers and educators. This includes funding to support early adopters in different subfields as they transition their research or teaching to the cloud, to help create successful use cases that other faculty can follow. In addition, cloud providers and academics can work together to identify a base set of requirements for cloud offerings, as well as guidance for selecting from a dizzying array of service offerings and matching cloud offerings with privacy and regulatory constraints on sensitive data. Also, cloud providers, funding agencies, and academic institutions should invest in creating a cadre of experts for using the cloud and creating research and educational software for the cloud.

5. **Form a central entity to serve as a nexus between multiple cloud providers on one side and multiple academic institutions on the other:** A central entity can serve as a single point of contact for the cloud providers and a unified voice for CISE researchers and educators, to further lower the barriers to academic use of the cloud. This entity can dig deeper into the topics outlined in the report, and serve as a convener and trusted third party in further efforts to bring the key stakeholders together. This central entity can help identify base requirements for academic cloud use, arrange for bulk purchases of cloud credits, draft templates for agreements between cloud providers and academic institutions, support creation of a cadre of professional staff to support academic cloud usage, identify and support software tools and data sets unique to academic users, and ensure that the large and diverse set of academic institutions are supported.

We elaborate on the rationale for these recommendations in the main body of the report.

The Case for Cloud Computing for CISE Research and Education

Cloud computing offers unprecedented access to storage, computation, and network resources, as well as large datasets and software platforms. In this section, we briefly summarize the reasons why academic institutions and cloud providers should work together to lower the barriers for academic use of the cloud, as an important step toward a more in-depth treatment of the subject (see Recommendations #1 and #2 above).

Why Academic Institutions Should Embrace Use of the Cloud

Academic institutions are starting to move toward cloud computing, for several key reasons:

- **Lower cost**: Cloud providers can achieve economies of scale beyond what individual academic institutions, let alone individual research projects or groups, can. These benefits come from buying equipment in bulk, greater levels of automation, and statistical multiplexing of demand from many customers. Academic users can enjoy additional discounts through academia consortia (e.g., Internet2, EDUCAUSE).
- **Lower energy consumption**: Cloud providers invest in energy-efficient equipment, and often build data centers in geographic locations with access to inexpensive or environmentally friendly sources of power and cooling (e.g., hydroelectric power, solar power, open-air cooling, etc.).
- **Zero space usage:** Cloud computing obviates the need for taking up valuable space on campus for local data centers.
- **Better security/privacy**: The data centers operated by large cloud providers offer good physical security, as well as heavy use of automation to prevent and respond to cyber attacks (e.g., timely application of software patches, use of firewalls and intrusion detection systems, etc.). Cloud providers have professional staff to protect the servers and networks and are always updating their systems with the latest security patches.
- **Better availability:** Cloud services commonly replicate data and include sophisticated mechanisms to recover from failures and maintain operations in spite of failures. As a result, they offer better availability guarantees than single-PI clusters, where disk, server, and network failures are visible and disruptive.
- **Shared access to data sets**: Having data in the cloud simplifies sharing of data across different institutions, and avoids the bandwidth costs of downloading data locally to each campus. In addition, some cloud providers offer a wide variety of public data sets to their customers.
- **Shared access to software/tools**: Cloud providers offer a variety of software, from virtual machine images to machine-learning libraries, that researchers and educators can use. This avoids the substantial cost of installing, let alone creating, the software locally at each school.

- **Expandable library of resources:** Academics can add new data sets, software/tools, machine-learned models, course materials, etc. to a library, accessible by all, for research and education.
- **Reproducible research**: Running experiments, such as analyses of large data sets, in the cloud simplifies the process of reproducing research results. Other researchers can more easily replicate an analysis simply by running the same software over the same data, and can then build on top of the work of others.
- **Training students**: Many of the graduates of universities enter a work world where cloud computing is increasingly the norm. Exposing these students to the cloud as part of their undergraduate or graduate education can better prepare them for the jobs they pursue after graduation.

The economies of scale that drive cloud computing offer particular advantages for schools of smaller size or more limited financial resources, providing an effective way to improve access.

That said, academic institutions do have legitimate reasons for maintaining local resources, particularly when:

- The research computing task is adequately managed by a local computer (e.g., laptop or desktop);
- The research computing task involves running some computation non-stop, fully utilizing a cluster for long periods of time (e.g., large-scale simulations). In those scenarios, local compute clusters may remain a cost-effective solution.
- Data sets are created *locally* and used for short durations of time before being replaced (where uploading data to the cloud may be time-consuming or expensive compared with the useful lifetime of the data) or downloaded often (where cloud “egress charges” may overwhelm the cost benefits of using the cloud);
- Experiments require *real-time* responsiveness (e.g., interactive visualization of large data sets or real-time control of experimental infrastructure);
- Teaching benefits from students having hands-on access to equipment (e.g., hardware and operating system labs);
- The research requires more specialized equipment (e.g., the very latest GPUs) or greater flexibility (e.g., systems research that requires modifying low-level hardware or software) than traditional cloud providers offer;
- The research or educational activity requires the use of academic licenses for software that would otherwise incur extra fees to run in the cloud;
- A campus already has significant computation and storage resources that would otherwise sit idle; in these scenarios, it makes sense to use those resources until they “age out” and only consider moving to the cloud at that point; and
- Researchers are unable or unwilling to spend any time learning about the new cloud computing technology, e.g., because the institution lacks support for researchers to move to the cloud.

The goal should be to achieve the considerable benefits of the cloud where appropriate, while using local resources as needed to best meet the research and educational needs. Today, local campus resources are often preferred over cloud computing due to other, artificial barriers, as we discuss in more detail in later sections.

Recommendation #1: We recommend preparing a more detailed write-up on the rationale for academic institutions to make more extensive use of the cloud, including key steps that campuses can take to move in that direction. The document should include a detailed cost comparison between cloud and on-campus resources, for an “apples to apples” comparison that can help campus CIOs and university administrators to make informed decisions. (The CloudMaven site at http://cloudmaven.org is a useful starting point.) Possible groups that could prepare such a document include the Coalition for Academic Scientific Computation (CASC), the Campus Research Computing Consortium (CARCC), and the Computing Research Association (CRA).

Why Cloud Providers Should Embrace Academic Users

On the surface, the academic market may seem relatively small compared to the large enterprises that increasingly rely on the cloud. Revenue aside there are many reasons for cloud providers to focus on academic researchers and educators as customers:

- **Academic mindshare**: Academic institutions train the students who go on to positions in every sector of the economy, and these graduates take what cloud expertise they gain with them. Technology tracks aside, academic institutions also train future leaders who make decisions about what technologies and services their organizations should use.
- **Early visibility into future requirements**: Academic users, especially from the CISE community, are harbingers of things to come. Researchers and educators across all fields look to CISE faculty for guidance on what technologies and services they should use, and academic applications may place requirements on cloud computing (e.g., use of GPUs/TPUs, billing models, etc.) that commercial customers will request down the road.
- **Innovations in cloud computing**: CISE researchers create innovations that can make cloud computing better, whether through new data-center network architectures, better operating systems, or machine-learning libraries. Engaging with academic researchers can accelerate this research and lead to better cloud offerings in the future.
- **Moderately large market**: The academic market in general, and the CISE community in particular, is, actually, sizeable. The United States has more than 4500 institutions of higher education, including more than 3000 four-year institutions and 1700 two-year institutions. About 20 million students attend American colleges and universities at any given time. In addition, a number of academic research projects make substantial use of computation, storage, and network resources.

Recommendation #2: We recommend preparing a more detailed write-up on the "business case" for cloud providers to invest in better supporting academic research and education, including key steps that cloud providers can take in that direction. The document should include a more quantitative treatment of the size of the academic market, as well as existing examples of "success stories" of academic use of cloud computing. The write-up needs to make a clear return-on-investment case to cloud providers. It needs to be clear about why the academic community is different from a typical enterprise customer.

Overcoming Cloud Barriers for CISE Research and Education

Despite the many benefits of using the cloud, academic research and education often rely on local equipment and software. In this section, we outline the main barriers to academic use of the cloud, as well as recommendations for overcoming these barriers.

Cloud Costs are Artificially High, Relative to Local Infrastructure

While using the cloud should be less expensive than local alternatives, a number of factors can make local computing resources more attractive financially:

- **Artificially low costs for managing equipment:** At many academic institutions, faculty do not directly pay for storing, powering, and cooling the equipment they buy or use. In addition, having graduate students in a research group serve as de facto system administrators may seem relatively inexpensive, compared to professional technical support. We recommend that academic institutions work to expose more of the costs of running local equipment to the faculty, or to offer incentives for cloud usage to compensate for these costs to get closer to an "apples to apples" cost comparison.
- **Local compute resources included in tuition but not cloud resources:** Students taking classes get access to a wide selection of laboratories equipped with desktops or connected to private clusters located on university campuses. In comparison, using a public cloud requires that instructors seek credits for their students. If cloud usage was included in tuition and cloud access streamlined the way access to local compute resources is streamlined, teaching using public clouds would become dramatically easier for faculty.
- **Overhead charges on cloud computing**: Most academic institutions charge "overhead" on grants, to the tune of 50%-60% depending on the school, for using the cloud, whereas equipment purchases are not charged overhead. This can make cloud computing artificially more expensive than buying and using local equipment. We recommend that NSF work toward eliminating overhead charges for cloud computing, and that academic institutions consider waiving these charges as some schools (e.g., University of Washington and University of California at San Diego) have already done, while recognizing that academic institutions will need time to prepare for this transition. Since academic institutions rely on overhead to run their operations, any "loss" from not

being able to charge this overhead has to show up as a “gain” somewhere else when doing the “apples to apple” cost comparison.

- **Ability to retain equipment after a grant**: When researchers buy equipment as part of a grant, they can continue using this equipment after the period of the grant ends. However, they can only pay for cloud usage during the limited period of the grant. We recommend that funding agencies and cloud providers explore funding models that allow faculty to pay up front for cloud usage, and retain any remaining cloud credits beyond the lifetime of a grant.
- **Difficulty managing cost risks and overrun in charges**: Researchers and educators understandably worry about the risk that a research experiment or course assignment might unintentionally overrun the money budgeted for cloud usage. In addition, spot pricing can make cloud usage more expensive at critical times (e.g., just before a conference paper deadline). These issues do not arise with local equipment, as the cost of the equipment is separate from its use. Faculty often have multiple research projects (each with multiple students), and students often take multiple courses (over multiple years of an education), making it hard to manage aggregate resources effectively and to manage multiple identities and multiple accounts. Overruns occur invariably in both classes and research projects, causing significant stress and time commitment in working with cloud providers on resolving those overruns.
- **Limited cost savings for small users**: Individual faculty and students can already use the cloud, but often the cost savings for the cloud are (at best) modest for individual users. Bulk agreements can offer more significant cost savings. We recommend that academic institutions and funding agencies work with cloud providers to acquire cloud resources in bulk, and pass along those benefits to individual faculty and students.
- **Data sets not available in the cloud**: Some data sets important for academic research are not readily available in the cloud yet, in part because they may have limited commercial use. This puts individual researchers in the position of needing to “upload” that data to the cloud, often at their own expense. Instead, cloud providers could move toward a “data is free, but pay for compute” pricing model, or funding agencies or large scientific projects could subsidize support for making large research data sets available in the cloud.

These concerns motivate a number of related ideas for avoiding artificial disincentives for using the cloud.

Recommendation #3: We recommend that all stakeholders work to remove artificial incentives for buying and using local equipment. Academic institutions can expose the true costs of owning and operating dedicated equipment, work toward waiving overhead charges on cloud usage, form bulk agreements with cloud providers, and include cloud credits in faculty start-up packages. Cloud providers can create new pricing structures such as up-front payments and long-term storage, as well as better tools for hierarchical allocation of credits and dashboards for managing cloud usage. Funding agencies such as NSF can work toward removing overhead charges on cloud usage, enabling cloud credits to outlive individual grants, acquiring cloud credits in bulk for academic users, and subsidizing data storage for data sets that are not cost-effective for cloud providers to store for free (e.g., from climate science, astronomy, or medicine) that are important in academic research.

Transitioning to and Using the Cloud is Challenging

Even when using the cloud is more attractive than relying on local infrastructure, faculty often avoid shifting their research and teaching to the cloud, for several reasons:

- **Challenges for early adopters**: As in any enterprise, early adopters of the cloud face additional challenges. Junior faculty are understandably loathe to take risks on new platforms during their crucial pre-tenure years, and senior faculty may have significant inertia due to existing equipment investments and local expertise. Cloud providers can lower the barriers by publishing detailed “use cases” of academic use of the cloud. Funding agencies and cloud providers can work together in a government-industry partnership to offer grants for early adopters to experiment with moving their research or teaching to the cloud or even support “virtual centers” in particular research areas. Such programs were offered in the past (e.g., with Google+IBM in 2008 and with Microsoft in 2009) and recently through the NSF Big Data Hub program. Fortunately, now, the commercial cloud is much more mature, increasing the likelihood of success. Still, CISE can increase the impact of support for early adopters by requiring awardees to have a plan for spreading best practices so other academics can learn from their experiences.
- **Large and diverse set of cloud offerings**: Each cloud provider offers a dizzying array of offerings, making it hard for academic users to know whether and which options to pick. In some cases, the offerings may lack sufficient capabilities for hierarchical allocation and management of resources (see previous section). Cloud providers and academics can work together to identify a minimum baseline of common services, as well as contracts and memorandums of understanding (MOUs). CISE researchers can also conduct research on novel techniques for automatically identifying the right mix of resources (e.g., processing, memory, storage) for a particular computational task.
- **Lack of local expertise**: Academic institutions often lack local expertise to train and support faculty and students who want to make greater use of the cloud. At some schools, the campus CIO office has become a place for “cloud czars” (who help local users select cloud offerings and make the most effective use of the cloud) or research software engineers (who help researchers create software that can run at scale in the cloud). Cloud providers could help in training these local staff and facilitating their professional development, and funding agencies can create new structures to “federate” these kinds of staff to ensure a broad range of academic institutions have support. Cloud providers and funding agencies can offer jointly funded programs for traineeships, where a recipient would become an expert on a cloud and bring back that knowledge to his/her campus and be a go-to expert on that cloud.
- **Lack of cloud usage support:** Today, when faculty seek to use local compute resources, they get plenty of support in managing the purchase, set-up, configuration,

and maintenance of that infrastructure. Similarly, for teaching, local resources are managed by IT professionals, allowing faculty to focus on teaching the material. In contrast, when moving to the cloud, faculty are often completely on their own: They have to manage account set-ups, consolidated billing, cost and usage monitoring, cloud credits, cost overruns, etc. This is a significant burden both when using cloud resources for research and for teaching. Funding agencies and cloud providers should consider offering grants to institutions that seek to overhaul their IT teams to provide cloud support at the level of local infrastructure support, and campuses can prioritize the hiring cloud-savvy employees into their IT teams.

- **Uncertainty about privacy restrictions**: In some cases, academic researchers and educators use local computation and storage resources out of concerns about data privacy, whether concern about student information, about personally identifying information in research data sets, or about patient health records in a medical school. Sometimes these restrictions about maintaining data on site come from the funding agencies themselves. In practice, cloud providers have service offerings that comply with a variety of different regulatory frameworks, e.g., HIPAA or EU GDPR. Cloud providers, funding agencies, and academic institutions could work together to “match” the privacy requirements with the associated cloud offerings. Seed funding for some initial “use cases,” where specific research projects navigate these concerns while migrating to the cloud could lower the barrier for future research projects and also identify any additional barriers that cloud providers and funding agencies need to overcome.

These concerns motivate a number of related ideas for making it easier for academic research and education to transition to the cloud.

Recommendation #4: We recommend that all stakeholders work to simplify the transition to the cloud for academic researchers and educators. CISE can support an effort to collect and analyze CISE research and education usage data and use cases, and track progress in supporting these use cases overtime. In addition, CISE and the cloud providers can support early adopters in different subfields as they transition their research or teaching to the cloud, to help create more successful use cases that other faculty can follow. In addition, cloud providers and academics can conduct a study to identify a base set of requirements for cloud offerings, as well as guidance for selecting from a vast array of service offerings and matching cloud offerings with privacy and regulatory constraints on sensitive data. Also, cloud providers, funding agencies, and academic institutions should invest in creating a cadre of experts for using the cloud and creating research and educational software for the cloud.

A Nexus for CISE Cloud Adoption

Many of the barriers to and recommendations for academic cloud usage warrant further discussion and deeper investigation, beyond what any single workshop report can address. Also, individual cloud providers working with individual academics (or even individual academic institutions) could lead to a proliferation of a large and uncoordinated mix of policies, contracts, pricing models, support staff, training programs, service offerings, data sets, and more. While letting "a thousand flowers bloom" clearly offers some advantages, even bottom-up experimentation with new ideas benefits from having a clearinghouse that can codify the lessons learned and identify best practices. As such, we believe there is value in having some sort of central entity that can serve as a convener, or trusted party, in supporting the relationship between academic institutions and cloud providers.

The central entity would serve as a *nexus* between multiple cloud providers on one side and multiple academic institutions on the other side. It could serve as the single point of contact for the different cloud providers and a single voice and representative for academic institutions. This model has the advantage that both new cloud providers and new academic institutions could join the "academic cloud" over time.

Such a central entity could do the work of recommendations 1-4, as well as other tasks, including:

- Complete some of the recommended studies (e.g., making the case for academic research and education to move to the cloud);
- Identify the common base requirements of cloud offerings for academic research and education;
- Arrange bulk purchase and allocation of cloud credits;
- Draft templates for agreements between cloud providers and academic institutions;
- Draft templates for data management plans for researchers who use the cloud in their research projects;
- Identify and support software tools and data sets particularly relevant to academic user;
- Support creation of and training for a cadre of professionals to support academic researchers and educators in using the cloud;
- Ensure the full range of academic institutions (e.g., research universities, liberal arts colleges, public and private institutions, community colleges, minority-serving institutions, etc.) are well served;
- Stand up "virtual centers" each with a disciplinary focus and an emphasis on the unique opportunities and challenges for using the cloud (e.g., equipment, data sets, software), in partnership with existing professional societies and organizations; and
- Serve as a unified voice for academic CISE researchers and educators with regards to the cloud.

Recommendation #5: Form a central entity to serve as a nexus between multiple cloud providers on one side and multiple academic institutions on the other. It can serve as a single point of contact for the cloud providers and as a unified voice for CISE researchers and educators, to further the mission of lowering barriers to academic use of the cloud. This entity can dig deeper into the topics outlined in the report, and serve as a convener and trusted third party in further efforts to bring the key stakeholders together. This central entity can help identify base requirements for academic cloud use, arrange for bulk purchases of cloud credits, draft templates for agreements between cloud providers and academic institutions, support creation of a cadre of professional staff to support academic cloud usage, identify and support software tools and data-sets unique to academic users, and ensure that the large and diverse set of academic institutions are supported.

Conclusions

Cloud computing offers a great opportunity for academic researchers and educators to perform research and teach students in new, innovative, and cost-effective ways. Yet, academic users face a number of barriers to using the cloud. We believe that the recommendations we have outlined offer a productive path forward, to bring the main stakeholders --- academic institutions, cloud providers, and NSF CISE --- together to address these challenges effectively, together.

Looking ahead, much as the CISE community serves as a harbinger of things to come, the academic cloud as proposed, driven by the CISE community now, eventually can and should serve all disciplines on all campuses, supporting the research and education mission of all of higher education.

Appendix A: Workshop Agenda

Monday January 8, 2018

- 10:30am-12:30pm Opening Session
    - James Kurose (NSF): Welcome
    - Magda Balazinska (U. Washington): Introductions and agenda
    - Jeannette Wing (Columbia): Perspectives from NSF, Microsoft, and Academia
    - Jen Rexford (Princeton): Using the Cloud for Academic Research
    - David Culler (Berkeley): Using the Cloud for Education
    - Vani Mandava (Microsoft), Karan Bhatia (Google), and Sanjay Padhi (Amazon): Reflections from Prior NSF Cloud Partners
- 1:45-3:00pm: Short talks
    - Rob Fatland: Cloud and Data in Research Computing
    - Eric Brewer: Maximizing the Cloud
    - Mike Huerta: Getting Cirrus About Cloud Computing at NIH
    - Kathy Yelick: High-Performance Computing and Clouds
    - Ion Stoica: Experience with Using Cloud in Research, Teaching, and Industry
    - David Lifka: Aristotle Cloud Federation: Reducing the time to science
- 3:30-5:00pm: Break-out sessions

    - Research: services needed for research, necessary software/data/people, etc.
    - Education: non-CS students, early-stage CS students, upper-level CS students; training vs. concepts; etc.
    - Business models: unique issues for academics, and appropriate pricing models in response
- 5:00-5:45pm: Report back to the group

Tuesday January 9, 2018

- 9:00am-10:15am: Plenary Discussion
- 10:15am-10:30am: Break
- 10:30am-11:00am: Report back to the group
- 11:00am-12:30pm: Breakouts on draft recommendations
    - Universities (CIOs, faculty)
    - Cloud vendors/providers
- 12:30-1:30pm: Working lunch and discussion, to finalize recommendations
- 1:30-3:00pm: Final Recommendations

Appendix B: Workshop Participants

Attendees

- Ryan Abernathy (Columbia)
- Magdalena Balazinska (University of Washington)
- Chaitanya Baru (NSF)
- Karan Bhatia (Google)
- Gautam Biswas (Vanderbilt)
- Jim Bottum (Internet2)
- Bram Bout (Google)
- Jack Brassil (NSF)
- Eric Brewer (UC Berkeley)
- Elizabeth Bruce (Microsoft)
- Mike Burns (Advanced Clustering Technologies)
- Steven Butschi (Google)
- Ken Calvert (NSF)
- May Casterline (NVIDIA)
- David Culler (UC Berkeley)
- Amol Deshpande (University of Maryland)
- Jay Dominick (Princeton)
- Karina Edmonds (Google)
- Rob Fatland (University of Washington)
- Valerie Florance (NIH)
- Gary Gabriel (Oracle)

- Erwin Gianchandani (NSF)
- Marc Hoit (NCSU)
- Nick Horton (Amherst College)
- Meghan Houghton (NSF)
- Mike Huerta (NIH)
- Zack Ives (University of Pennsylvania)
- Vandana Janeja (NSF)
- Vince Kellen (UCSD)
- Grace Kitzmiller (Amazon)
- Tyler Kloefkorn (NSF)
- James Kurose (NSF)
- David Lifka (Cornell)
- Vani Mandava (Microsoft)
- Sanjay Padhi (Amazon)
- Irene Qualters (NSF)
- Jennifer Rexford (Princeton)
- Majd Sakr (CMU)
- Alicia Salmeron (Google)
- Mehul Shah (Amazon)
- Ion Stoica (UC Berkeley)
- Dan Stone (University of Kentucky)
- J. Barr von Oehsen (Rutgers)
- Mark Wegman (IBM)
- Karen Wetzel (EDUCAUSE)
- Jeannette Wing (Columbia)
- Kathy Yelick (UC Berkeley)

Organizers

- Jennifer Rexford (Princeton)
- Magdalena Balazinska (University of Washington)
- David Culler (UC Berkeley)
- Jeannette Wing (Columbia)