Cloud Computing and Information Policy:

Computing in a Policy Cloud?

Forthcoming in the *Journal of Information Technology and Politics, 5*(3).

Paul T. Jaeger

University of Maryland

Jimmy Lin

University of Maryland

Justin M. Grimes

University of Maryland

Abstract

Cloud computing is a computing platform that resides in a large data center and is able to dynamically provide servers the ability to address a wide range of needs, ranging from scientific research to e-commerce. The provision of computing resources as if it were a utility such as electricity, while potentially revolutionary as a computing service, presents many major problems of information policy, including issues of privacy, security, reliability, access, and regulation. This paper explores the nature and potential of cloud computing, the policy issues raised, and research questions related to cloud computing and policy. Ultimately, the policy issues raised by cloud computing are examined as a part of larger issues of public policy attempting to respond to rapid technological evolution.

Keywords: cloud computing, information policy, rechnology policy, grid computing, security, privacy, reliability

Introduction

Cloud computing refers to a computing platform that is able to dynamically provide, configure, and reconfigure servers to address a wide range of needs, ranging from scientific research to e-commerce. While cloud computing is expanding rapidly as service used by a great many individuals and organizations internationally, policy issues related to cloud computing are not being widely discussed or considered. As this paper will demonstrate, there are a wide range of policy issues related to cloud computing that merit considerable attention as cloud computing develops into a widely used commercial enterprise; yet there has thus far been a lack of policy-making or court cases related to cloud computing. The objective of this paper is to introduce the policy concerns, research areas, and potential solutions related to cloud computing that will likely be the focus of discussion and deliberation in coming years. If these problems are considered during the developmental stages of cloud computing, perhaps they can be addressed before the consequences of nonaction are too significant.

Typically, the cloud computing infrastructure resides in a large data center and is managed by a third party, who provides computing resources as if it were a utility such as electricity—accessible by anyone, anywhere with an Internet connection. For the "cloud provider," this consolidation of computing resources yields many benefits deriving from centralized management and economies of scale; for the "cloud user,"[1] the ability to gain rapid access to computing capacity not only reduces overall cost, but also lowers the barrier to entry for many processing-intensive activities, since it eliminates the need for upfront capital investment and the necessity of maintaining dedicated infrastructure. Through cloud computing,

[1] Throughout this article, "user" refers generically to organizations or individuals that exploit cloud computing technology. While these types of users may have some different interests and perhaps different concerns based on scale of data involved, generally a corporate and an individual user are going to have the same policy-related expectations and concerns related to cloud computing.

users transfer the burden of system management and data protection (e.g., in event of system crash or physical theft) over to the cloud provider. In addition, cloud computing provides a potential avenue by which users of handheld devices could have access to computing services. In essence, users of cloud computing "outsource" their data processing needs to a third-party. These ideas are certainly not new, as cloud computing evolved out of earlier technologies for distributed processing, such as "grid computing."

Far more than a theoretical construct, cloud computing technology has reached the commercialization phase of development. A range of cloud providers already offer a variety of services, with users employing clouds for massive database management, data mining, and the deployment of Web services, among other activities (Baker, 2007). Potential specific uses of cloud computing range from using clouds to process huge amounts of data to solve incredibly complicated scientific problems to using clouds to manage and provide access to medical records (Hand, 2007). Commercial and individual cloud computing services are already available from Amazon, Yahoo, Salesforce, desktop Two, Zimdesk, and Sun Secure Global Desktop, while Google's efforts in cloud computing have attracted a great deal of interest (Delaney & Vara, 2007; Gilder, 2007; Ma, 2007; Naone, 2007). Another major effort is an academic-industrial collaboration spearheaded by Google and IBM, in conjunction with six major research universities in the United States, whereby the companies are providing faculty and students with access to clouds for research and education (see: http://googleblog.blogspot.com/2007/10/let-thousand-servers-bloom.html and http://www-03.ibm.com/press/us/en/pressrelease/22414.wss).

In spite of its promise and potential, cloud computing sits at the difficult intersection of new computing concepts and information policy. Not only does cloud computing raise major issues regarding privacy, security, anonymity, telecommunications capacity, liability, reliability, and government surveillance, relevant existing laws do not appear to be applicable to this new idea. This situation is indicative of a growing problem in which technology so far outpaces information policy that the developers and users of an important new technology create, implement, and use it, hoping that the law will ultimately catch up to their activities. While Princeton University's Center for Information Technology Policy held a two-day workshop entitled "Computing in the Cloud" (http://citp.princeton.edu/cloud-workshop/) in January 2008 to discuss some broad policy issues related to cloud computing, this is an area that merits considerable attention.

This paper focuses on a range of policy aspects of cloud computing—specific issues raised by gaps in current laws and regulations. In the case of cloud computing, technological innovation, commercial interest, and consumer interest are all fast outpacing current information policy. The primary goal of this article is to raise awareness of these issues at the intersection of computing and policy—although we propose solutions to the extent possible, ultimate resolution of any specific issue is beyond the scope of this article.

This paper will first discuss the nature and origins of cloud computing and key technical characteristics and benefits to users. Current examples of and initiatives in cloud computing are next examined. The policy issues related to cloud computing are then discussed, followed by the policy gaps raised by cloud computing. Finally, the paper describes how issues of cloud computing and policy might be reconciled to facilitate the development of cloud computing as a beneficial development for individual, corporate, and governmental computer users.

What is Cloud Computing?

Like most technologies, cloud computing evolved from a need. The tremendous growth of the Web over the last decade has given rise to a new class of "Web-scale" problems—

challenges such as supporting thousands of concurrent e-commercial transactions or millions of search queries a day. The natural response of technology companies has been to build increasingly large data centers to handle the ever-growing load; these data centers consolidate a great numbers of servers (hundreds, if not thousands) with associated infrastructure for storage, networking, cooling, etc. Over the years, technology companies, especially Internet companies such as Google, Amazon, eBay, or Yahoo, have acquired a tremendous amount of expertise in operating these large data centers. This "know how" extends beyond physical infrastructure to include experience with process management and other intangibles. Cloud computing represents a commercialization of this combined solution.

The tremendous amount of information available in electronic format today has translated into a proliferation of data- and processing-intensive problems for a wide variety of organizations and even individuals, in the context of the Web and beyond. For example, genomics research involves huge volumes of sequence data; financial companies maintain mountains of information about clients; even the serious hobbyist may have more video footage than can be reasonably processed by available machines. Common to all these scenarios is the need for large amounts of processing power. Prior to cloud computing, acquiring such resources was an expensive proposition. Upfront capital investment in purchasing the computers themselves is only the initial step; significant resources must then be devoted to maintain the infrastructure. In many cases, users (especially smaller companies, non-profit organizations, and academic research groups) are unable or unwilling to make this investment.

The convergence of need and solution has produced the current conception of cloud computing, which promises to benefit all parties involved. Cloud providers gain an additional source of revenue and are able to commercialize their expertise in managing large data centers.

Overall cost as measured on a capacity-basis is reduced due to consolidation, and capital investment in physical infrastructure is amortized across many customers. Cloud users no longer have to worry about purchasing, configuring, administering, and maintaining their own computing infrastructure, which allows them to focus on their core competencies. This paradigm has also been referred to as “utility computing,” in which computing capacity is treated like any other metered utility service—one pays only for what one uses.

Currently, the best-known example of commercial cloud computing is Amazon’s Elastic Compute Cloud (EC2) (http:// aws.amazon.com/ec2), which allows customers to “rent” compute cycles in Amazon’s data center. Typically, this service is used in conjunction with Amazon’s Simple Storage Service (S3) (http://aws.amazon.com/s3), which provides data storage services. For S3, costs are straightforwardly computed in terms of disk storage used on a monthly basis and additional charges for data transfer. This is attractive for users, since they only pay for space they use, with additional capacity available on demand. For EC2, users are charged in terms of instance-hours; one instance-hour can be intuitively understood as the data processing capabilities of a particular computational unit for one hour. Such a pricing model is attractive for the same reasons as S3: Costs scale predictably with use and no resources are spent on idle processors. A number of startup companies use Amazon’s services with regularity, and even established companies have found these services to be useful. For example, the *New York Times* recently employed EC2 and S3 to process 11 million scanned articles 1851-1980 from the newspaper’s archives into PDF’s accessible by readers (http://open.blogs.nytimes.com/2007/11/01/self-service-prorated-super-computing-fun/). For applications where computing demand is uneven or, in the case of the *New York Times,* nonreoccurring, cloud computing provides an efficient, cost-effective solution to user needs.

The activity and interest in cloud computing is by no means limited to the commercial sphere. Developing applications and algorithms that run on hundreds if not thousands of processors is a daunting challenge, a task most computer science graduates today are ill-equipped to handle. Whereas the present computer science curriculum is mostly focused on sequential processing (e.g., executing one instruction after another), cloud computing requires the programmer to reason about parallel processing, where many operations are executed concurrently, and distributed processing, where operations are executed in different processors.

To address this growing need for expertise, in October 2007, Google and IBM jointly announced the academic cloud computing initiative with six U.S. research universities: Carnegie Mellon University, Massachusetts Institute of Technology, Stanford University, the University of California at Berkeley, the University of Maryland, and the University of Washington. The second author of this article is the lead faculty at the University of Maryland. As part of this initiative, IBM and Google have dedicated a large cluster of several hundred computers for use by faculty and students at the participating institutions. By making these resources available, the companies hope to encourage faculty adoption of cloud computing in their research and also integration of the technology into the classroom. These investments are seen as necessary steps to sustain the growth of cloud computing as an emerging paradigm.

In general, there are two primary ways in which cloud clusters can be used. In one mode, the cloud cluster simply hosts a user’s application, which is typically provided as a Web service accessible to anyone with an Internet connection. For example, the cloud provider can simply take over the task of maintaining and running a company’s inventory database or transaction processing system. In the consumer realm, services such as Google Maps, Gmail, and YouTube can already be thought of as “cloud applications.” The second mode may be thought of as “batch processing," where the user transfers a large amount of data over to the cloud cluster along with associated application codes for manipulating the data. The cloud cluster executes the application code, and the results are returned to the user. Note that in both usage scenarios, the user's data and applications reside (at least for some time) on the cloud cluster, which is owned and maintained by a third party. This characteristic of cloud computing is at the root of many challenging policy issues.

Cloud Computing and Issues of Information Policy

Cloud computing raises a range of important policy issues, which include issues of privacy, security, anonymity, telecommunications capacity, government surveillance, reliability, and liability, among others. This section of the paper introduces and examines these issues individually. While some of the trade press and popular media accounts of cloud computing have raised potential issues of privacy and intellectual property (i.e., Delaney & Vara, 2007; Ma, 2007), the range of policy issues raised by cloud computing merits significant consideration.

A productive approach to begin analysis of the information policy issues related to cloud computing is to consider user expectations. At a minimum, users will likely expect that a cloud will provide:

- **Reliability and Liability.** Users will expect the cloud to be a reliable resource, especially if a cloud provider takes over the task of running "mission-critical" applications and will expect clear delineation of liability if serious problems occur.
- **Security, privacy, and anonymity.** Users will expect that the cloud provider will prevent unauthorized access to both data and code, and that sensitive data will remain private. Users will also expect that the cloud provider, other third parties, and

governments will not monitor their activities. The only possible exception may be for cloud providers, who may need to selectively monitor usage for quality control purposes.

- **Access and usage restrictions.** Users will expect to be able to access and use the cloud where and when they wish without hindrance from the cloud provider or third parties, while their intellectual property rights are upheld.

Each of these interrelated issues will be considered in terms of its importance, what realistic expectations users might have, and the policy implications. While there are other policy issues beyond those mentioned above, we consider those policy considerations central to the successful development of cloud computing.

Reliability and Liability

Users have the expectation that services provided through the cloud will be reliable, as one of the key concepts behind this computing paradigm is the transfer of data center management, along with its associated risks, to a third party. Yet, reliability raises some very significant questions.

What if a cloud provider experiences technical problems that result in an organization’s mission-critical applications becoming unavailable? This exact scenario took place on February 15, 2008, when Amazon’s S3 experienced a service outage that lasted for approximately two hours. As many organizations, particularly startup companies, have come to rely on the S3 service, this disruption raised significant awareness regarding the perils of relying on a third party to serve mission-critical needs. Although service disruptions will become increasingly rare as the technology matures, 100% reliability will never be possible. Thus, who bears the risk? Would it be the cloud user, who simply accepts service disruptions as a normal “cost of business” (in the same way that a data center under one’s own control may not be completely reliable either)? Would it be the cloud provider, who might be contractually obligated to guarantee a particular quality of service (e.g., 99.99% uptime) and thus legally liable for lost revenue, productivity, etc.? Theoretically, one could even imagine third parties (e.g., insurance companies) assuming such risk for a premium. Whereas service disruptions are transient, there are even more serious issues associated with data integrity. What if a user's data were lost or corrupted? Once again, such risks could be borne by the user, the provider, or a third party.

In addition to service disruptions and data integrity, correctness of results as generated by the cloud is another facet of reliability. In many cases, the processing-intensive nature of the tasks will make it impractical to validate results independently. Consider the hypothetical example of a financial company running large-scale simulations on the cloud, and then acting on the results (perhaps on behalf of its clients). Unbeknownst to them, flaws in the infrastructure resulted in corrupted data, and hence incorrect results. The outcome might be investments that later prove to be disastrous. Who should bear the liability in this case?

The issue of reliability is not only a technological problem, however, as the adequacy of telecommunications policy to address issues regarding cloud computing needs to be carefully considered, since it challenges fundamental assumptions of telecommunications policy. Internet access and services fall under the umbrella of telecommunications regulation, so cloud computing provision would presumably as well. However, telecommunications laws are based on the assumption that the purpose of a network is to ship "bits" in the form of voice or data transmissions from point A to point B. However, the notion of cloud computing defies these expectations in at least three ways: First, it is not only data, but also instructions for processing that data, that are transmitted; second, the transmitted data are frequently modified as a result of computations occurring in the cloud; and third, there may not be a notion of a "recipient" in the usual sense (results of the computations may be transmitted back to the sender, or simply stored in the cloud itself).

Further, telecommunications policy is not based around ideas of providing a legal compensation framework for lost transmission content. If a telephone call is dropped, the parties talking are expected to call back, resume their conversation, and fill in the missing details. If a company drops too many calls, policy relies on market forces to drive the company out of business. However, cloud computing greatly amplifies the risk associated with data loss and service disruptions. Depending on the nature of the loss, it may be one blow from which a company can never recover. (Due to the fluid nature of Web-based transactions, for example, a major service outage may drive customers to a competitor, from which they never return.) To what extent are cloud computing providers liable for their services, is a question that will need to be addressed. Liability and potential litigation is a growing concern of cloud computing providers, who function in a currently gray area. Unnecessary litigation would stifle innovation, as no technical system is infallible, and 100% uptimes and services simply cannot be guaranteed.

Utilities in both North America and the EU have dealt with similar problems of privacy protection, reliability, and liability; so it seems that telecommunications law may serve as basis, or at least a model, for regulating clouds. This analogy, however, raises the question of how telecommunications law would view a cloud. Would it be considered an ISP? A telecommunications provider? A common carrier? Or, would a new definition under telecommunications law be needed? The closest parallel might be viewing cloud computing as cooperating interconnected networks, which were envisioned in the Telecommunications Act of 1996 as part of universal access. The status accorded to cloud computing under the law would impact the rules applied to cloud computing, the expectations of users for the reliability of cloud provision, and any potential actions available to users beyond market forces.

The ongoing debate of network neutrality may also have a considerable impact on the development and progress of cloud computing. Cloud computing is highly dependent on a consistent and stable Internet platform. If network neutrality is not guaranteed, the telecommunications service providers that control the underlining network connection would have the ability to limit a cloud provider' service through pricing and distribution structures. Without network neutrality, under a differential pricing model, telecommunications service providers could effectively charge cloud providers more, absorbing any potential profits (Odlyzko, 2008). Telecommunications service providers could even become themselves cloud providers, giving preferential treatment to their own services.

For cloud computing to effectively function, however, cloud providers need to be able to manage their systems to schedule and prioritize tasks and data. As noted by Odlyzko (2008), cloud computing is merely an "extreme form of vertical integration, just carried out by other companies than the telecommunications service providers, and at higher levels of the protocol stack" (p. 16). Therefore, even with established network neutrality, cloud providers could become limited in performing necessary optimization functions. In essence, network neutrality could be a double edge sword for cloud computing. Without network neutrality, telecommunications service providers could cutoff or hamper the connection of cloud providers; with network neutrality, cloud providers may be subject to the same regulations as telecommunications service providers meaning they would not be able to regulate and optimize their own services.

Security, Privacy, and Anonymity

After reliability, perhaps the most pressing concerns for cloud users will be security, privacy, and anonymity. Clearly, the levels of privacy and anonymity available to the user of a cloud will be lower than the user of desktop in many cases (Delaney & Vara, 2007). To protect the privacy of cloud users, care must be taken to guard both users' data and applications for manipulating that data. A few examples: Corporations may be concerned about the security of client data and proprietary algorithms; researchers may be concerned about pre-mature release of new data or discoveries; individuals may be concerned about personally sensitive information. Yet, as unauthorized releases of sensitive information by corporations and governments over the past several years have demonstrated, the electronic environment provides innumerable scenarios for the unwanted release of information.

Since the physical infrastructure in a cloud computing environment is shared among a number of users, the fate of sensitive information (e.g., personally identifiable information, medical records, trade secrets) is a tremendous question. For example, if a bank stores records of customers (including, for example, social security numbers) in a cloud, what guarantees can be made about the fate of such information? Even if an organization takes all the necessary steps to protect sensitive information (e.g., encrypting the records), are such efforts sufficient? Depending on the exact type of records, will regulatory obligations be met? In some cases, the proprietary algorithms for data processing are far more valuable than the data itself. Could a particular user receive reassurances with respect to applications that run on the cloud?

Wesabe and Mint are two early examples of what the future may hold for sensitive data and cloud computing. Both of these Web-based applications are money management software packages akin to the popular software Quicken or Microsoft Money. Each Web application functions by encouraging users to provide all their financial information to their service, which is then aggregated and systematically analyzed. These services push the envelope of privacy, security, and confidentially of sensitive data, yet they could be forerunners of the types of cloud services to come.

While there are basic issues that may be inherent in the nature of the cloud, such as a cloud provider needing to profile users to ensure service quality, users will not likely want their actual content monitored. Many companies already provide contextualized ads based on keywords, Web sites viewed, and other forms of automated learning of users' activities and content. Will users accept those same strategies in a cloud? It appears that users are willing to accept these strategies in e-mail systems in exchange for increased convenience, storage capacity, and searching ability, if the success of Google's Gmail is any indication. Perhaps user attitudes will be the same for cloud computing. However, corporate users may be more concerned about monitoring of information, even for ad placements, than the individual users of Gmail accounts.

Furthermore, cloud computing could also easily open up the ability of third parties to monitor content in a similar fashion. Users of all types may prove less comfortable with both providers and third parties monitoring or possibly using their information. One Microsoft Vice President working on their cloud computing efforts has suggested "you may win a Nobel Prize by analyzing data assembled by someone else" (quoted in Baker, 2007, ¶ 35). Such statements will likely have a chilling effect on the enthusiasm of researchers for cloud computing who wish to keep control of their own data and have no one else looking at it.

To ensure the growth and adoption of cloud computing, it will be necessary to find technological and policy solutions for ensuring privacy (in some form) and assuring information security. In these circumstances, ensuring anonymity will not be sufficient. Solutions have been developed to ensure user anonymity on peer-to-peer networks, and these technologies may transfer to the cloud concept (Singh, Gedik, & Liu, 2006). However, while anonymity of users' activities will clearly be a central aspect of protecting user privacy, much of the information flowing through the cloud will not only have to be protected in terms of who it belongs to, but also what it is. In the United States, since there is no universal standard for privacy protection, it may be a significant issue to determine what type of privacy is guaranteed in cloud computing, if any, or it will be left up to individual cloud providers to decide. Furthermore, cloud computing is a global service, crossing multiple governments and their differing sets of regulations and servicing users across the world; it will also have to account for the privacy concerns of different cultures and the privacy laws of numerous countries. Even between European Union member nations and the United States, there are significant differences in the definitions of privacy and variations in types of privacy protection available (Sunosky, 2000). While the World Summit for the Information Society (http://www.itu.int/wsis/index.html) has spearheaded a number of efforts to use cloud computing to encourage collaboration and reduce deficits of scientific knowledge in certain regions of the world, intergovernmental collaboration on cloud computing standards seems not to have been explored yet.

Such differences may not only confuse users, but may become more problematic if users are using multiple cloud providers, with or without their knowledge. One can imagine in the future a network of cloud providers, and an efficient mechanism for "routing" computing resources to where they are needed, much in the same way the power grid is managed today. Say a user connects to provider A, but provider A doesn't have enough "compute cycles." To guard against exactly such a scenario, provider A has contracted to purchase computing resources from provider B for limited durations. Thus, provider A proceeds to transfer the user's data and code over to provider B, without the user's explicit knowledge.

This type of scenario also raises another issue of privacy and security: the establishment of controls on what cloud providers can and cannot do with users' data. Cloud providers have legitimate reasons for monitoring use of the resource, which is necessary to perform routine maintenance, to more effectively balance load across servers, to optimize for certain usage profiles, etc. However, at what point does information gathering become nefarious data mining? While clearly-stated usage agreements might address this problem, such a solution may prove to be overly reliant on industry to make these decisions, given the potential revenue that could come from, for example, advertising targeted to cloud users, as discussed above.

Yet another related issue concerns the cloud provider's obligation (if any) to maintain the anonymity of its users and serve as a neutral provider of computing capacity. Would it represent a conflict of interest for a cloud provider to supply services to two competing companies? In this case, technical solutions may prevent unauthorized access of data. Nevertheless, there are no clear guidelines with respect to the obligations of the cloud provider in these cases.

Perhaps more complicated is the fact that cloud computing also highlights the difficulty of protecting information stored on remote systems from government investigations. Search warrants for off-line computers are much harder to get than for online databases, and providers would likely have no requirement of informing users of such a search warrant (Ma, 2007). Since 2001, the United States government has engaged in amazingly comprehensive gathering, surveillance, and analysis of electronic information through a mix of laws, executive orders, and secret programs (Jaeger, 2007, in press; Jaeger & Burnett, 2005). When a subject of interest is identified, the more targeted information collection may extend well beyond that subject. For example, Verizon acknowledged to Congress that the government requested data about not only specific subjects, but also about individuals several generations out from a subject—that is, an individual who phoned a second person, who phoned a third person, who phoned the target of the investigation. would still have all of his or her records collected even if she or he has never heard of the target of the investigation, who was three generations of phone contact away (MSNBC, 2007).

The USA PATRIOT Act, the Homeland Security Act, and other related security legislation, coupled with sophisticated electronic information-gathering technologies, make it possible for the government to gain access to electronic information in virtually any context. With personal, corporate, and even secret data and code flowing through a cloud computing network, the concept of generations of data being gathered by the government in an investigation become particularly problematic. As such, the current government approach to surveillance and security hinders trust and use of cloud computing for fear that innocent but sensitive data or code might become snared in an investigation. This lingering mistrust and fear of governmental snooping is already having a negative backlash on certain Google services (Avery, 2008).

For all of these significant concerns, it is clear that having the cloud provider define a usage policy isn’t enough to address privacy and security. One obvious solution is for users to encrypt whatever data resides in the cloud. Although this solution may suffice for the storage of data, it may prove technically cumbersome to implement algorithms that could process the data in an encrypted form. Decryption on-the-fly prior to processing is not a solution, since it creates a point of vulnerability in the cloud. Furthermore, for competitive intelligence purposes, it may not be necessary to gain knowledge of the exact processing being performed by a user: simple observations of *how much* and *for how long* may provide valuable clues themselves.

Access and Usage Restrictions

The concept of cloud computing also raises questions of proper access and use. Users will want to access and use the cloud on their own terms, yet will also want their intellectual property rights protected. The protection of copyright in types of materials stored on and shared through clouds has received some consideration, and cloud providers believe the ability to trace usage will serve as a means of preventing illegal activity (Delaney & Vara, 2007; Ma, 2007). However, this does not fully consider all of the dimensions of access and usage that are relevant to cloud computing, most significantly the issue of licensing. The unique issues of cloud computing may ultimately illustrate that technology has reached the point that there is serious need to rethink how intellectual property is licensed.

If a user has a license for a particular software product or dataset and uses it in their work, can they still use the licensed product *in the cloud*? Recall that to make effective use of the cloud, both data and applications may need to be transferred over to the cloud provider. Licenses typically forbid redistribution, but whether use of cloud computing could be viewed as redistribution is an interesting question. Data and applications are being sent to a third party, but that third party is not “using” the licensed products in any real sense. And yet, licenses tend to be written to be as inclusive as possible (Carrico & Smalldon, 2004).

A parallel issue of access related to cloud computing exists in the provision of international access to a cloud. Cloud computing means that anyone with an Internet connection can access the cloud, including people in other countries. Licensing and use agreements may be different across national markets and certain products may only be available in certain markets, but the cloud eliminates such differences. This may raise some problems for cloud providers in the types of software and processing capacities available in the cloud. Further, cloud providers may be affected by export control regulations and limitations on sharing of scientific information with certain nations (Jaeger, 2007; Jaeger & Burnett, 2005).

Similarly, cloud users may engage in activities that violate intellectual property rights of others or represent some other forms of illegal activity. Just as Internet Service Providers (ISPs) are liable to limit illegal file-sharing activities on their networks and Web sites must remove materials that they do not own the copyright to, cloud providers may have to place controls on the use of their computing resources. (Note that this would be technically very challenging.) If users engage in criminal activities by employing the processing power of the cloud—breaking government encryption or cracking Web sites—would the provider of the cloud be liable? If so, that would mean cloud providers would have to engage in very close monitoring of cloud activities, which would further enhance issues of privacy and security within the network. Similarly, cloud providers may have to worry about users trying to engage in criminal activities by stealing other users' data, such as corporate espionage.

There may also be issues of public access in the provision of cloud computing. If any computer can be connected to a cloud, will cloud providers be willing to accept public access computers? A large percentage of Americans rely on the public access computers, such as those provided by public libraries, as their sole means of getting online (McClure, Jaeger, & Bertot, 2007). If public access computers are not included by cloud providers, then a large numbers of users will be unable to access the capacities of the clouds, creating new forms of social disparities in technology access. Cloud computing would also offer a means of assisting organizations that provide public computing access, many of which are overwhelmed by the capacity and technical demands of Web 2.0 content and services (Jaeger et al., 2007; McClure et al., 2007).

In spite of all of the policy concerns, cloud computing ultimately has the potential to help bridge certain gaps in access to digital content. By moving computing and storage away from the users, cloud computing reduces the demands and requirements on local hardware that individuals have to purchase. While the gap in broadband access would still need to be addressed, hardware costs can be eliminated or greatly marginalized by cloud computing. Storing and computing in the cloud would be beneficial and more economical for the many people who do not own a computer or who have trouble affording a computer (Carr, 2008).

Computing in a Policy Cloud?

As the selection of issues discussed above demonstrate, there are significant uncertainties about and tensions between public policy and technological capacity in the development and provision of cloud computing. In some cases, the technical solutions available might not be compatible with policy mandates. For example, only so much security can be ever guaranteed to users by providers as a result of government surveillance and data collection activities. While the policy issues raised above clearly demonstrate that there are many potential concerns raised by the unique nature of cloud computing combining aspects of computer, information, and telecommunications issues, it also demonstrates a larger policy problem.

Information policy in the United States, simply put, is continuing to fall further and further behind in policies related to new technology developments and how these developments are being employed. This gap between policy and technology has been noted, as has the increasing speed and distance of the gap as the United States continues to make laws reactively and based on a pre-electronic mentality (i.e., Braman, 2006). However, by raising so many different policy issues, cloud computing demands meditation on the gap between policy and technology.

Many of the problems at the nexus of policy and technology derive from trying to use print-based concepts of policy in an electronic world. Intellectual property problems caused by this gap may be the easiest to understand. For example, the extensions of copyright protection to such incredible lengths—life of the author plus 80 years—create many questions of ownership, and these extensions create significant tensions with the increases in access to information brought about by the Internet and electronic files. The exceptions created to try to address these issues, such as the fair use exemption and the exemptions for use in distance education, only serve to make the issues murkier and leave many information providers and users uncertain of their legal positions (Butler, 2003; Travis, 2006). Orphan works—older works where the copyright owner is untraceable—are virtually unusable, even by archives that own the items (Brito & Dooling, 2006; Carlson, 2005). Libraries struggle mightily with previously much clearer issues of interlibrary loan, electronic resources, and services to distance learners, while universities must determine how to try to provide resources to distance education students (Allner, 2004; Carrico & Smalldon, 2004; Ferullo, 2004; Gasaway, 2000). At the same time, industries, educational institutions, and users struggle with the implications of electronic files and the ability to share files for music, movies, books, and other content formats (Strickland, 2003, 2004).

Perhaps the most striking result of this uncertainty of intellectual property in an electronic world is the Goggle Books project (http://books.google.com/). The goal of this project is to digitize and make searchable every book printed, whether or not it is in copyright, and a number of major libraries have made their collections available for Google to digitize. This legality of the effort has attracted some note from the mainstream press (i.e., Thompson, 2006). However, the legal defenses Google has raised for digitizing and making freely available copyright-protected materials without permission of the copyright holders do not hold up well under basic legal analysis (Hanratty, 2005). Yet, the confused nature of applying print-based conceptions of intellectual property to an electronic world allow for projects like Google Books to plough forward under a cloud of legal uncertainty. And intellectual property questions are but one example of these kinds of uncertainties at the intersection of technology and policy in the United States. It seems that cloud computing is similarly moving forward in uncertain legal territory, which might significantly hamper its legitimacy in the eyes of many users.

Currently, the decisions related to issues of information policy raised by cloud computing are being made entirely by cloud providers. Amazon's EC2 service presents an excellent example of this. In January 2008, Amazon Web Services was storing 14 billion units of data, varying in size from a couple of bytes to 5 gigabytes, and handling 30,000 requests to its database per second (Hardy, 2008). A major reason for these impressive numbers is that Amazon's services present enormous cost-savings for users in many cases (Hardy, 2008). However, many of the issues discussed above are entirely absent from these agreements or absent in terms of the user.

The Customer Agreement (http://www.amazon.com/AWS-License-home-page-Money/b/ref=sc_fe_c_0_201590011_10?ie=UTF8&node=3440661&no=201590011&me=A36L942TSJ2AJA#10) focuses solely on the business transaction elements of the service, on protecting Amazon's intellectual property, and on releasing Amazon from responsibility regarding any of the concerns raised above. If you are concerned about the security of your information, Section 7.2 unreassuringly states "you acknowledge that you bear sole responsibility for adequate security, protection and backup of Your Content." If you are concerned about the privacy of your information, your only protection is the standard Privacy

Notice (http://www.amazon.com/gp/help/customer/display.html/002-0131023-8675278?ie=UTF8&nodeId=468496) that applies to all interaction with Amazon.com, not taking into account any of the unique aspects of cloud computing.

As such, the businesses providing cloud computing services are establishing corporate polices that protect the providers, while many users are unaware of the potential policy implications and the federal government remains silent on the issues related to cloud computing. Ultimately, however, the government may choose not to act and allow the concerns to be settled by the free market. Perhaps the market may develop a cloud provider that charges more for stronger guarantees of privacy, security, and reliability. As things stand, users of cloud computing services may benefit from the cost-savings but may be surprised by important issues not addressed in relation to their cloud computing activities.

In the end, trust is a key reason that both cloud providers and cloud users would want to have these policy issues clarified and settled. Exactly one half of the respondents to a 2007 Pew Research Center study agreed with the statement “You can’t be too careful in dealing with people” (p. 2). People often find it harder to trust online services than off-line services. As examples, people perceive greater potential risk in shopping online, in engaging in social activities, and in participating in online political activities than they do in performing similar activities offline (Best, Kreuger, & Ladewig, 2005). The distrust of online services is even negatively affecting the level of trust accorded to organizations than have long been respected as trustworthy, such as fears about electronic records that are undermining the long-held trust of public libraries (Jaeger & Fleischmann, 2007).

If highly personal information were being stored in a cloud, the importance of rust grows exponentially. The Health Insurance Portability and Accountability Act (HIPAA) of 1996 created very intricate standards to protect the privacy of patient information in medical records (Wun & Dym, 2008). Yet, currently, users would have no idea how highly sensitive information would be handled by the cloud provider. As such, if users do not feel that they can trust cloud providers to keep their information secure, private, protected, and reliably accessible, then many users will opt not to use cloud computing. It is therefore imperative for cloud providers—in interest of the users and of the profitability of their systems—to embrace the development of policies for cloud computing.

Conclusion: Bridging Opportunities and Policies in Cloud Computing

Cloud computing is likely to present many policy questions and raise many issues as it becomes more commonplace. However, it will also likely serve to further highlight the ever-widening gaps between the capacity of technology and the focus of policy. Though no technological development has yet to force policy-makers in the United States to begin to conceive of information policy as proactive rather than reactive and to move away from pre-electronic bases for concepts of policy, it is worth considering whether cloud computing might serve as such an instigator.

Thomas Watson, Sr., former chairman of IBM, was once allegedly misquoted to have suggested in 1943 that the total market for computers around the world would not exceed five (Carr, 2008). Although this quote was never verified, it is often lampooned and propagated by the tech savvy as a silly shortsighted prediction of the future. However, with the advent of cloud computing, it might ironically become an eerily accurate prediction of the future. The history of computing is becoming cyclic. Users once all connected to a central mainframe to do their computing, only to later have that paradigm eventually shift towards desktop computing, with each user having their own computer. Cloud computing completes this cycle, as computing returns toward a centralized source.

Given the factors of economies of scale, first mover advantage, network effect, and path dependency, the future may only be able to support a few massive cloud computing providers. Those first to develop platforms will have the first mover advantage that would later be further supported by the network effect. Latecomers would have difficultly overcoming the network effect and smaller cloud providers could be unsuccessful in competing due to the advantage of cost given by the economies to scale. These cumulative dynamics would result in the future of Watson's erroneous quotation, leading towards the potential monopoly or duopoly of all computing. Moreover, due to issues of vertical integration, cloud providers could have the potential to achieve a vertical monopoly (Odlyzko, 2008).

As it simultaneously offers potential benefits to such a wide range of users and raises so many different issues of information policy, cloud computing might not only serve as a means of technological but also policy advancement. The gaps between policies and technological realities are becoming so significant in some cases that arguments can be made that information policies may have to be completely rethought (Travis, 2006). This situation is further confounded by the number of policy decisions left to the marketplace in the United States that are more heavily regulated through policy in other nations. The protection of personally identifiable information provides such an example—there are enormous differences between the minimal regulation of the United States and the intricate protection structures of the European Union (Sunosky, 2000).

In approaching cloud computing from a policy-making standpoint, the regulation of e-government may serve instructive. While e-government is clearly a government enterprise and cloud computing a nongovernmental one, e-government is an online technology that developed quickly and needed guiding policies after development was well under way. A series of policies, regulations, and the E-government Act of 2002, however, gave shape to development and regulation of e-government. In a similar fashion, the development of policies, regulations, or even a law to cover cloud computing issues would be extremely helpful in sorting out the concerns and uncertainties currently related to cloud computing for providers and for individual and organizational users. With the current lack of policies or court decisions about cloud computing, however, the lack of guidance presents an impediment to the development of the potential of and to user adoption of cloud computing.

Based on the discussion herein, there seem to be two approaches that could be taken individually or concurrently: either regulation by a designated federal government agency, such as the Federal Communications Commission, or legislation mandating greater precision in service agreements between the providers and users of cloud computing. Either of these methods would have the goal of establishing accepted cloud computing standards for:

- Basic thresholds for reliability;
- Assignment of liability for loss or other violation of the data;
- Expectations for data security;
- Protections of privacy;
- Any potential expectations for anonymity;
- Access and usage rights; and
- International standardization to promote transborder data flows in clouds.

The specific elements of these standards could be established through regulation or creation of parameters for future service agreements. Whichever approach is taken, this issues will be key elements to address the policy issues and foster user trust in cloud computing.

Aside from direct interventions in terms of public policy, there are other efforts that can be very helpful in the development of some notion of cloud policy. Education is a key first step. The six schools that are a part of the IBM/Google initiative are developing and implementing educational courses on cloud computing. Courses such as these will help to prepare future information professionals and computer scientists to be more cognizant of the issues raised by cloud computing and to be attuned to addressing these issues. The University of Washington is a leader in this regard, having successfully designed and implemented undergraduate computer science courses in 2007 on cloud computing (Kimball, Michels-Slettvet, & Bisciglia, 2008). Ongoing efforts at the University of Maryland include a joint research and education initiative that pairs undergraduate students with Ph.D. students in working on cutting-edge research problems in text processing, such as statistical machine translation and analysis of e-mail archives (Lin, 2008). While both primarily focus on technological issues, at least the latter program includes a discussion of policy issues such as those raised in this paper.

Policy research also clearly needs to bring greater focus on this area. While many computer scientists may be drawn to the interesting technological issues raised by cloud computing and technological futurists will focus on the utopian or dystopian possibilities inherent in the technology (i.e., Carr, 2008), there are a wealth of important social and policy questions raised by cloud computing, such as:

- What expectations for privacy, security, reliability, and anonymity do users of cloud computing have?
- Are there variations in these expectations among individuals, corporate users, academic users, and governmental users?

- Would greater degrees of privacy, security, reliability, and anonymity influence users' decisions about which cloud providers to use?
- Have users even considered issues like privacy, security, reliability, and anonymity?
- Do users have any concerns about protection of their intellectual property in the cloud?
- Do users have any concerns about the monitoring of their activities in the cloud by providers or by the government?
- Do users only trust cloud computing for certain types of functions?
- Can information policies for print-based environments intelligibly translate to cloud computing? If so, how?
- What issues of cloud computing are completely unaddressed in the current policy environment?

These are just a sampling of the types of research questions raised by cloud computing that would not only illuminate this particular technology, but that would also have relevance to greater questions of the relationships between technology and policy in the electronic environment. Data on these questions would also be extremely valuable in crafting a public policy response to the issues raised by cloud computing.

If a particular technology is used widely enough and encounters enough significant policy issues, perhaps it will generate momentum for a broad reassessment of the information-policy making process in the electronic age. While cloud computing is very new, it is worth considering whether and how its development may ultimately impact the large mosaic of policies related to technology. Though it is not difficult to identify potential information policy issues in cloud computing, potential policy-based solutions to these issues are less obvious due to the newness of the technology and to the previously discussed disjunctions between technology and policy in the United States.

The unique nature of cloud computing—and the potential for it to become a truly ubiquitous technology employed by individuals, academic institutions, corporations, and perhaps even government agencies—provides an opportunity to consider essential issues of technology and policy that seem destined to continue to grow in significance as technology continues to evolve. This paper is intended to identify and encourage discussion about the policy issues related to cloud computing. It is hoped that providing serious consideration for and research about these issues as cloud computing is in its developmental stages will allow the policy concerns to be addressed in a timely and satisfactory manner before cloud computing becomes too large or too widely used to regulate effectively.

References

Allner, I. (2004). Copyright and the delivery of library services to distance learners. *Internet Reference Services Quarterly, 9*(3), 179-192.

Avery, S. (2008, March 24). Patriot Act haunts Google service. *Globe and Mail*. Retrieved July 2, 2008, from http://www.theglobeandmail.com/servlet/story/RTGAM.20080324.wrgoogle24/BNStory/Technology/home

Baker, S. (2007, December 14). Google and the wisdom of the clouds. *Business Week*. Retrieved July 2, 2008, from http://www.msnbc.msn.com/id/22261846/

Best, S. J., Kreuger, B. S., & Ladewig, J. (2005). The effect of risk perceptions on online political participatory decisions. *Journal of Information Technology & Politics, 4*(1), 5-17.

Braman, S. (2006). *Change of state: Information, policy, and power*. Cambridge: MIT Press.

Brito, J., & Dooling, B. (2006, March 25) Who's Your Daddy? *Wall Street Journal, A*(9).

Butler, R. P. (2003). Copyright law and organizing the Internet. *Library Trends, 52*(2), 307-317.

Carlson, S. (2005). Whose work is it, anyway? *Chronicle of Higher Education, 51*(47), A33-A35.

Carr, N. (2008). *Big switch: Rewiring the world, from Edison to Google*. New York: Norton.

Carrico, J. C., & Smalldon, K. L. (2004). Licensed to ILL: A beginning guide to negotiating e-resources licenses to permit resource sharing. *Journal of Library Administration,* 40(1-2), 41-54.

Delaney, K. J., & Vara, V. (2007, November 27). Google plans services to store users' data. *Wall Street Journal*. Retrieved July 2, 2008, from http://online.wsj.com/article/SB119612660573504716.html?mod=hps_us_whats_news

Ferullo, D. L. (2004). Major copyright issues in academic libraries: Legal implications of a digital environment. *Journal of Library Administration, 40*(1-2), 23-40.

Gasaway, L. N. (2000). Values conflict in the digital environment: Librarians versus copyright holders. *Columbia – VLA Journal of Law & the Arts,* Fall 2000, 115-161.

Gilder, G. (2007). The information factories. *Wired, 14*(10). Retrieved July 2, 2008, from http://www.wired.com/wired/archive/14.10/cloudware_pr.html

Hand, E. (2007, October 24). Head in the clouds. *Nature, 449*, 963.

Hanratty, E. (2005). Google library: Beyond fair use? *Duke Law & Technology Review, 10*. Retrieved July 2, 2008 from http://www.law.duke.edu/journals/dltr/articles/pdf/2005dltr0010.pdf

Hardy, Q. (2008, February 11). The death of hardware. *Forbes*. Retrieved July 2, 2008, from http://www.forbes.com/technology/forbes/2008/0211/036.html

Jaeger, P. T. (in press). The fourth branch of government and the historical legacy of the Bush administration's information policies. *Government Information Quarterly*.

Jaeger, P. T. (2007). Information policy, information access, and democratic participation: The national and international implications of the Bush administration's information politics. *Government Information Quarterly, 24*, 840-859.

Jaeger, P. T., Bertot, J. C., & McClure, C. R. (2007). Public libraries and the Internet 2006: Issues, funding, and challenges. *Public Libraries, 46*(5), 71-78.

Jaeger, P. T., & Burnett, G. (2005). Information access and exchange among small worlds in a democratic society: The role of policy in redefining information behavior in the post-9/11 United States. *Library Quarterly, 75*(4), 464-495.

Jaeger, P. T., & Fleischmann, K. R. (2007). Public libraries, values, trust, and e-government. *Information Technology and Libraries, 26*(4), 35-43.

Kimball, A., Michels-Slettvet, S., & Bisciglia, C. (2008). Cluster computing for Web-scale data processing. *Proceedings of the 39th ACM Technical Symposium on Computer Science Education (SIGCSE 2008), Portland, Oregon* (pp. 116-120).

Lin, J. (2008). Exploring large-data issues in the curriculum: A case study with MapReduce. *Proceedings of the Third Workshop on Issues in Teaching Computational Linguistics at ACL 2008, June 2008, Columbus, Ohio* (pp. 54-61).

Lohr, S. (2007, October 8). Google and IBM join in 'cloud computing' research. *New York Times*. Retrieved July 2, 2008, from http://www.nytimes.com/2007/10/08/technology/08cloud.html

Ma, W. (2007, November 29). Google's Gdrive (and its ad potential) raise privacy concerns. *Popular Mechanics*. Retrieved July 2, 2008, from http://www.popularmechanics.com/technology/industry/4234444.html

McClure, C. R., Jaeger, P. T., & Bertot, J. C. (2007). The looming infrastructure plateau?: Space, funding, connection speed, and the ability of public libraries to meet the demand for free Internet access. *First Monday, 12*(12). Retrieved July 2, 2008, from http://www.uic.edu/htbin/cgiwrap/bin/ojs/index.php/fm/article/view/2017/1907

MSNBC. (2007, October 16). Telecoms won't talk about surveillance. *MSNBC*. Retrieved July 2, 2008, from http://www.msnbc.msn.com/id/21322332

Naone, E. (2007, September 18). Computer in the cloud. *Technology Review*. Retrieved July 2, 2008, from http://www.technologyreview.com/Infotech/19397/?a=f

Odlyzko, A. (2008). *Network neutrality, search neutrality, and the never-ending conflict between efficiency and fairness in markets*. Minneapolis, MN: Digital Technology Center. Retrieved July 2, 2008, from http://www.dtc.umn.edu/~odlyzko/doc/net.neutrality.pdf

Pew Research Center. (2007). *Americans and social trust: Who, where and why*. Washington, DC: Author.

Singh, A., Gedik, B., & Liu, L. (2006). Agyaat: Mutual anonymity over structured P2P networks. *Internet Research, 16*(2), 189-212.

Strickland, L. S. (2003, October/November). Copyright's digital dilemma today: Fair use or unfair constraints? Part 1: The battle over file sharing. *Bulletin of the American Society for Information Science and Technology*, *30*(1), 7-11.

Strickland, L. S. (2004, December/January). Copyright's digital dilemma today: Fair use or unfair constraints? Part 2: The DCMA, the TEACH Act, and e-copying restrictions. *Bulletin of the American Society for Information Science and Technology*, *30*(2), 18-23.

Sunosky, J. T. (2000). Privacy online: A primer on the European Union's Directive and the United States' Safe Harbor privacy principles. *Currents: International Trade Law Journal, 9*, 80-88.

Thompson, B. (2006, August 13). Search me? Google wants to digitize every book, publishers say read the fine print first. *Washington Post*, pp. D1 & D7.

Travis, H. (2006). Building universal digital libraries: An agenda for copyright reform. *Pepperdine Law Review, 33,* 761-833.

Wun, E., & Dym, H. (2008). How to implement a HIPAA compliance plan into a practice. *Dental Clinics of North America*, *52*(3), 669-682.

Author Notes

Paul T. Jaeger

University of Maryland

Paul T. Jaeger, Ph.D., J.D., is an Assistant Professor in the College of Information Studies and is the Director of the Center for Information Policy and Electronic Government at the University of Maryland. His research focuses on the ways in which law and public policy shape access to and use of information. Dr. Jaeger is the author of more than sixty journal articles and book chapters, along with four books.

Jimmy Lin

University of Maryland

Jimmy Lin, Ph.D., is an Assistant Professor in the College of Information Studies at the University of Maryland. His research lies at the intersection of natural language processing and information retrieval. Dr. Lin leads the Google/IBM Academic Cloud Computing Initiative at the University of Maryland.

Justin M. Grimes

University of Maryland

Justin Grimes is a doctoral student in the College of Information Studies at the University of Maryland and a graduate research assistant at the Center for Information Policy and E-Government. His research interests include information policy, e-government, intellectual property, and issues of technology and policy in virtual world communities.

Correspondence concerning this article should be addressed to Paul T. Jaeger at

pjaeger@umd.edu.