Is Security Lost in the Clouds? (*)

Marjory S. BLUMENTHAL
Georgetown University, Washington, DC

Abstract: "The cloud" can apply to different kinds of services (typically differentiated as platform-as-a-service, infrastructure-as-a-service, and software-as-a-service), and it is the subject of rampant hype about its benefits. This paper draws on extensive readings from the literature (technical, business, and policy) and consultations with a wide range of experts over the past two years. Intended to provide a counter to the cheerleading and a framework for more balanced consideration of public cloud services, in particular, it begins with an exercise in accentuating the negative. In particular, it lays out various ways in which the cloud might be seen as a new platform for malice. The paper enumerates key issues, including kinds and sources of risk (vulnerabilities and threats) associated with providers and/or users and implications for trustworthiness in cloud contexts, as well as the prospects for new technology to counteract apparent sources of risk. It addresses different cloud contexts, and it argues for leveraging cloud concerns to rethink fundamental issues about the nature, handling, and protection of data (which may be stored or processed in the cloud - or not).

Key words: Cloud, privacy, security, cybersecurity, data.

Google's January 2010 news of apparent attacks from China on its Gmail service was a comparatively public alert to the need to rebalance popular thinking about the merits of cloud computing. Touted by its advocates as the next big thing in computing, if not the next incarnation of the Internet, the cloud has a combined market potential that is huge. The market embraces different kinds and extents of cloud service, generally differentiated as (a) the wholly do-it-yourself Infrastructure as a Service (IaaS, such as Amazon Elastic Compute Cloud), (b) the middle ground of Platform as a Service (PaaS, such as Microsoft Azure), and (c) a specific application leveraging the cloud-provider's platform and infrastructure (Software as a Service or SaaS, such as Gmail or

(*) Acknowledgments: This work was supported by a grant from the U.S. Office of Naval Research grant number N000140910037. The ideas were presented at the Telecommunications Policy Research Conference (TPRC), October 2, 2010, Arlington, VA. The author appreciates comments on earlier drafts from Fred B. Schneider, Jim Waldo, Micah Sherr, Jim Fenton, and Scott Charney.

Salesforce.com). [1] Cloud services leverage the technology of virtualization, the use of software to divide up capacity on computing hardware into virtual machines (VMs) associated with specific customers and their data and/or processes. Much of the technology is not new, but the business models are.

Security concerns emerged early for public cloud offerings, which dominate exposure for the general public and for at least smaller enterprise users. [2] By 2008, commercial consortia such as the Cloud Security Alliance (CSA) and conferences for researchers and practitioners were discussing security for the cloud - although implementation of new cloud security ideas lags, [3] at best. This paper responds to these trends by (1) considering the alter ego of the cloud as a platform for malice, and (2) arguing for more syst ematic rethinking about how we handle information.

■ The cloud as a new platform for malice

To balance the hype about the cloud and its benefits, it is a useful thought exercise to consider how we might characterize the cloud as a platform for malice. The negative potential of the cloud spans a range of threats to systems and users.

Perhaps least obvious is the range of concerns associated with the provider: Cloud service providers effectively have access to growing amounts of data and processes. They also have ways of avoiding risk, depending on the type of cloud: users have more control and bear more risk with IaaS offerings than with PaaS or SaaS ones. These two terms-of-service [4] excerpts illustrate how dominant public cloud providers expect their users to bear risks:

---

[1] The cloud is available in various forms (generally described as infrastructure-, platform- or software-as-a-service) offered by different kinds of providers. The National Institute of Standards and Technology (NIST) strove to capture that scope, but cloud definition remains unsettled. See: http://csrc.nist.gov/groups/SNS/cloud-computing/.

[2] Larger enterprises are more likely to be able to afford a private cloud solution, with greater control corresponding to private ownership.

[3] The ideas emerging from research have been characterized, not unreasonably, as "academic" (as opposed to practical) (CACHIN, 2009).

[4] A practical comparison of cloud offerings and the nature of terms of service can be found in WAYNER (2008).

> "Google AppEngine: 5.5. You agree that Google has no responsibility or liability for the deletion or failure to store any Content and other communications maintained or transmitted through use of the Service. You further acknowledge that you are solely responsible for securing and backing up your Application and any Content."
>
> "Amazon Web Services: 7.2. Security. We strive to keep Your Content secure, but cannot guarantee that we will be successful at doing so, given the nature of the Internet. [...] We will have no liability to you for any unauthorized access or use, corruption, deletion, destruction or loss of any of Your Content or Applications."

Although legal (including contractual) mechanisms are an important vehicle for protecting users, the appropriate balance of interests between providers and users is likely to take time to emerge, given the relative newness of cloud offerings and the relatively rapid development ongoing in the marketplace. Some of the balancing will arise from the ways that the inherent principal-agent problems get worked out (FRIEDMAN & WEST, 2010). But as the termination of Amazon service to WikiLeaks illustrates, many factors - including some exogenous to the user-provider relationship - can be at play, and a provider can act quickly to protect its own interests (FOWLER, 2010). The fact that technology for auditing what goes on in a cloud remains immature at best adds to the handicap burdening the user.

CSA and, in more detail, the European Network and Information Security Agency (ENISA) encourage users to assess their tolerance for the risk associated with specific deployment and service alternatives (ENISA, 2009). That guidance is new, abstract, and lengthy, with ENISA's top ten cloud security risks covering a lot of territory. [5] Meanwhile, there has already been at least one case of a provider shutting down after an egregious error caused substantial customer data-loss (KRIGSMAN, 2008), and there is reputational damage to providers even when lost data is recovered, as in the case of the T-Mobile/Sidekick loss of stored personal data resulting from a server failure (WINGFIELD, 2009).

Errors, of course, are only the beginning. Providers can and do go rogue, and history with outsourcing illuminates both the potential problems and ways of coping. What is different with today's clouds from yesterday's timesharing and outsourcing is the intervening growth in criminal exploitation of the Internet. ENISA suggests that the growth in cloud use implies that

---

[5] Loss of governance, lock-in, isolation failure, compliance risks, management interface compromise, data protection, insecure or incomplete data deletion, malicious insiders. See ENISA (2009).

provider employees are increasingly likely to be targeted by criminal gangs (ENISA, 2009). More generally, insider threat may be a particular concern for the cloud, given the growing value of what goes on in the cloud - including the intellectual property associated with both proprietary algorithms and data - and the expectation that providers will try to provide some security. According to ENISA, there is a medium probability of insider abuse of privilege, but a very high impact if it happens.

Provider-based threat may be subtle. For example, many who focus on privacy are troubled by the content-scanning of e-mail by Google in support of its advertising placement, or the analytical tool provided by Twitter for public analysis of data from its service (for which there is less presumption of privacy than for e-mail). Users trust providers like Google, but they know too little about what might be done with their data to judge the real risks, especially when that data endures for long periods of time on the provider's servers.

Industry structure raises indirect concerns: Given that there appear to be significant economies of scale in the provision of cloud services, how concentrated will cloud supply be, and how might that concentration translate into undesirable competitive conduct? For example, observers already remark on high switching costs: the difficulty of moving data to competing providers has led one commentator to characterize cloud computing as the "Hotel California of technology" (ASAY, 2009). There is also the more straightforward concern that a few dominant players may lead to a smaller number of very large data centers that provide economies of scale for the providers but also large targets for attackers. [6] Further, to the extent that, as in other quarters of the information-technology sector, there is a first-mover advantage, one might expect premature commercialization of cloud technology/ies and the possibility of a stream of adjustments if the offering succeeds, a known route to security problems. This has been seen with social media, where incentives aim providers in directions other than user security and providers capitalize on user assumptions of security:

> "[C]onsider a choice before a hypothetical social network: (1) spend time and money securing personal information against unauthorized access by corrupt insiders, or (2) spend time and money exposing personal information to advertisers to increase the value of their ads

[6] Scott Charney points out (personal communication) that a factor mitigating data-aggregation risks may be the corresponding aggregation of expertise, which in cybersecurity remains comparatively scarce.

> [...] [A] social network must allow information sharing in order to be useful. […] [T]his sharing often depends on the assumption of effective access control. […] [S]ocial networks are fun and easy to use, but their access control schemes are tedious and incomprehensible." (ANDERSON & STAJANO, 2009).

Wittingly or unwittingly, cloud providers may enable new ways for malicious users to hide in the cloud. Consider two possibilities:

- Clouds as cutouts or fronts. The rise of "hacking as a service" suggests that clouds may have the same kind of appeal to the malicious as to the conventional user (POULSEN, 2009). Users select providers based on what they have to offer, and the model of certain kinds of ISP supporting the likes of the organization formerly known as the Russian Business Network is not too hard to extrapolate. [7] That prospect raises questions about how the industry is monitored and the interplay of legal and technical mechanisms. Of course, the law itself may be a kind of enabler, as those who focus on digital rights management argue: If copyright holders can invoke the law to scan a cloud for content that violates their rights, what other kinds of scanning might be done, and by whom? For example, some kinds of monitoring of VMs are being developed to enhance security, [8] yet one can wonder about unintended, malign uses as well. After all, the history of filtering technology points to its being put to uses unintended by their developers (notably for surveillance).

- Clouds as havens. Although today cloud infrastructure is concentrated in the United States, there is a general expectation of it spreading in other countries, not least because of the desires of governments for local infrastructure. This presents the prospect, as Stewart Baker once quipped, of "the cloud fleeing the subpoena," or more generally, the cloud providing a haven for those eluding scrutiny of some kind. Cloud technology development has included the ability to move VMs between servers, a feature intended to enhance reliability and/or to support maintenance. How dramatic might such moves be, and what other uses of such features are possible? Governments have come to appreciate that the physical points of presence of cyberspace provide loci for intervention, which limits the potential for havens (and may also drive policy that limits the efficiency of

---

[7] Stefan Savage was quoted as saying, "For providers, cloud infrastructure is a cyber-criminal's dream world, with plenty of ambiguity and anonymity behind which to hide. What could be more ideal for the cyber-criminal than paying for a huge amount of un¬traceable computing infrastructure with a stolen credit card?" See: SAVAGE (2009).

[8] See, for example, the discussion of secure introspection of VMs as a means of detecting malware: CHRISTODORESCU *et al*. (2009).

cloud services) (GOLDSMITH & WU, 2006). But those limits are as effective as the governments themselves, and investigations of cyber-attacks in China and Eastern Europe demonstrate that there are regions in which providers may operate under a blind or winking surveillance eye.

As the above examples suggest, users present the second, and arguably bigger, source of concern in contemplating the cloud as a platform for malice. Public clouds provide new places for malicious users to hide, and such users may undertake new and undesirable secondary uses of the data and processes originally generated by others. Indeed, perhaps the most striking illustration of possibility comes from recent research that makes clear that the cloud may be less cloudy than represented by advocates. First, there are possibilities for "cloud cartography" - for mapping the multi-tenant terrain, and then for manipulating the process for locating VMs (RISTENPART *et al*., 2009). Second, there are possibilities for monitoring what is going on in the cloud, after one has situated a VM, exploiting side channels (e.g., time-shared caches or keystroke activity) and covert channels (e.g., cache-load measurements where cooperative processes run on different VMs) to support reverse-engineering, infiltration, exfiltration, certain kinds of encryption cracking (GREENBERG, 2009), and other attacks (RISTENPART *et al*., 2009). More generally, technically skilled people are looking for ways to exploit whatever they find. As an analysis of hypervisor vulnerabilities observed, "VMware isn't an additional security layer - it's just another layer to find bugs in" (KORTCHINSKY, 2009).

Google's adoption of encryption for Gmail (the automatic https mode) in response to its Chinese attacks illustrates that defenses must both be available and used - the story of cybersecurity is one of known problems remaining untreated, and known solutions remaining unused. For this reason, optimism that reports of research demonstrating vulnerabilities, threats, and attacks will motivate the deployment of existing technology as well as development of new technology must be bounded.

Meanwhile, the cloud landscape is becoming more complex, which is likely to facilitate malice more quickly than defenses are likely to be mounted. Although the above discussion focused on issues presented by a given cloud, the security challenge is magnified by the prospect of a cloud ecosystem - different kinds of cloud with different kinds of interaction or intersection. At one level, there will be more efforts to facilitate interaction among applications within a given cloud. See, for example, GEAMBASU & LEVY (2009). At the consumer level, this can be seen in efforts to allow individuals to exchange information among different applications offered by a

single provider (such as Google's Gmail, Picasa, and YouTube). For enterprises, there is research into securing query processing for competitive users of cloud-based aggregation services, mitigating threats in the cloud environment relative to conventional Web portals (ZHOU *et al*., 2010). Even more challenging are the possibilities for interactions that bridge clouds, whether public cloud offerings, private clouds established by large organizations, community clouds that support specific groups of users, and/or hybrid clouds combining public and private aspects. Work has begun on standards to foster inter-cloud exchanges, and the debate about openness vs. proprietary technology has begun (OpenCloudManifesto.org, 2009).

The activity on multiple fronts to promote the use of standards and interoperability among clouds points to the potential of an intercloud, a cloud of clouds as an internet is a network of networks. The intercloud today is a topic for speculation. Nelson sketches three scenarios: one with a few separate and unconnected platforms, one with proprietary platforms permitting data but not software interchange, and one that is maximally open and Internet-like, enabling data and software sharing (NELSON, 2009). Not only does an intercloud present technical interoperability challenges, it also raises questions about the interoperability of security policies across services (CREESE & HOPKINS, 2009). Regardless of how the future plays out in terms of structure and technology, it is clear that if it is hard to gauge risk in a given cloud, it is much harder in an interconnected cloud complex, [9] which would increase the potential number of interdependencies. The challenge will be even greater as that complex becomes more international, as is inevitably the case. [10] ENISA, for example, has recommended consideration by national governments and European Union entities of a "European Governmental cloud as a supra national virtual space" featuring interoperability and other standardization (CATTEDDU, 2011).

[9] The Government Accountability Office has suggested that the opportunity for attacks grows with interconnections. See: WILSHUSEN (2010).

[10] International coordination raises the spector of national policies limiting flow of data originating locally, especially data deemed privacy-sensitive. A balkanized, location-aware cloud is to some technologists not a true cloud, inasmuch as the most efficient use of the technology seems to imply ready movement of resources as demand and load evolve in real time.

■ The devil's in the data

The possibilities for the public cloud to be a platform for malice argue for more deliberate thinking about what we entrust to the (public) cloud and what we keep outside of it. Other things equal, [11] economics and the appeal of cloud functionality and dynamic scalability will make the choices steadily harder. They are also likely to change our judgments about what is secure enough - about how we gauge risk. This is evident when it comes to some aspects of personal use of the public cloud. For example, social media applications (e.g., Facebook, Twitter) are fundamentally about sharing information that might otherwise be kept private, and they involve personal decisions that the benefits of sharing outweigh at least some concerns about protecting some kinds of data.

The legal framework is both evolving and highly varied among nations, with varying attention to and protections for privacy [12] and the security of data generally. A risk-averse perspective might deem that whatever is in the public cloud - like whatever is e-mailed - is effectively public. Among the proponents for updating relevant U.S. laws are those (like the Digital Due Process coalition) who note that the laws originated when far less content was communicated or stored, let alone when the technology was less sophisticated. Many are hopeful that evolving technical and legal mechanisms will support higher expectations for data protection. In the meantime, contractual (procurement) mechanisms provide the frontlines for protection, and as discussed above, themselves may be targets for improvement. ENISA, for example, characterizes European perspectives in outlining how service-level agreements can be structured to promote greater security (CATTEDDU, 2011). The discussion in Europe focuses in part on issues arising from data-storage facilities that are outside of a given country or even the region, which is to be expected in the context of efficient, large-scale public cloud operation. The associated jurisdictional concerns provide impetus for efforts to harmonize law and policy across countries, if not globally - which, given the history of efforts to harmonize other instances of cybersecurity law and policy, may be easier said than done.

---

[11] Technology development should drive some progress on security, notwithstanding the challenges discussed above. Notable in the technical community are the attempts to use cryptography (See, for example: KAMARA *et al*. (2010). There is particular excitement about the prospects for homomorphic encryption, which would allow processes to act on encrypted data, but practical challenges to implementing this approach remain significant.

[12] Forrester Research developed an "Interactive Data Protection Heat Map" to illustrate this legal variation. See: http://www.forrester.com/cloudprivacyheatmap.

The strong appeal to the cloud (public or private) for organizations in part reflects the fact that storing data (or hosting applications) in a cloud can be cheaper than local alternatives. This is especially true for enterprises moving away from legacy applications with specialized data structures and associated databases, which require enterprises to address the inconsistencies (data "deconflicting"). Public cloud storage services (e.g., Amazon Simple Storage Service (S3)) can be an efficient substitute for customers building and operating private storage. And most simply, there can be security benefits where the cloud alternative results in less use of removable and therefore easy-to-steal media (e.g., CDs/disks, thumb-drives). But depending on a third party is inherently risky, and that is the point that needs explicit recognition:

> "Placing core business applications and data into the cloud doesn't really have a suitable backup plan unless you're maintaining local backups of all that data and can afford to bring the applications and data back online quickly during an outage - but what's the point of leveraging a cloud if you have to run all that gear locally anyway just in case? [...] If a third-party company falls down on the job and takes your data with them, your only failure was believing that you could safely farm out highly important data and applications and let them deal with it." (VENEZIA, 2010).

As the above quotation suggests, what data or applications are truly core to an organization (or an individual) needs to be thought through more explicitly than may have been the case with more centralized, local, and/or directly controlled infrastructure.

Given the proliferation of cloud types and applications, it is useful to differentiate the issues by kind of user - individual or enterprise / organization - and by kind of information - nonpublic and sheltered or at least semipublic and shared. See figure 1.

The traditional domain of cybersecurity (and privacy) is represented by the left column in the table - data that individuals and organizations seek to protect or shelter. As discussed in the earlier portions of this paper, the rise of the public cloud raises questions about how well the cloud can shelter data that its owners want sheltered. Because new kinds of applications and associated business models fundamentally involve sharing, even for organizations, there are also new questions about what should be shared and how that determination may evolve. For example, NASA, a U.S. government organization (agency), has a cloud pilot project called Nebula,

which shares scientific data after an initial review. [13] The government of Washington, DC, made a wide variety of data available to the public online, inviting the public to develop its own visualizations and applications using that data and facilitating certain visualizations via Google maps. These examples - and their architects - build on experience with open-source development of software, which has demonstrated benefits and business cases for sharing of technology insights and expertise. It seems that new kinds of data are being made public daily in the public cloud, supporting new uses and new ways of thinking about data, demonstrating benefits from relaxing some expectations for data sheltering.

Figure 1 - Data Status Taxonomy

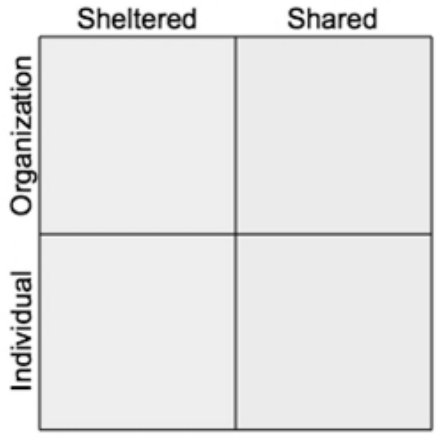

Meanwhile, the WikiLeaks saga may have two kinds of effect. First, since government data was at issue, a backlash that will roll back recent progress on sharing data among government organizations is likely; it is the most predictable risk-averse response. Second, what it should do is hone thinking about the kind of data that truly must be sheltered "at all costs." Spaulding has remarked on how difficult that can be for governments (and companies) accustomed to treating secrets as assets, even when it is counterproductive:

> "Moreover, a strategy based on keeping information from the prying eyes of your competitors often means not sharing information with those who could use that information to help you. An especially

[13] Currently a private cloud, the system is expected to become a hybrid cloud, with interconnections to other clouds. See: http://nebula.nasa.gov/about/.

> egregious example of this is when intelligence products based entirely on open sources are then stamped "classified." Limiting dissemination of information often means only your friends or potential collaborators don't have it while your enemies do." (SPAULDING, 2010).

Some data compromises seem intolerable and are likely to be so seen indefinitely. These would relate to truly core organizational data or perhaps certain health data for individuals (e.g., re infectious diseases). In theory, homomorphic encryption - a technology for supporting applications to act on data while it is encrypted - could provide an ideal compromise, enabling both sheltering and use of the cloud for applications (IBM Research, Homomorphic Encryption). But the concept, while proven recently in theory, remains short of practical implementation, and even with meaningful implementation other issues would remain. That ENISA (2009) points to the continuing need for research to support end-to-end data confidentiality in the cloud through encryption of search, processing, and tools for social applications is indicative of the limitations of encryption as a tool for sheltering data in the cloud. Under current conditions, the most critical data should remain out of the cloud.

One path forward may involve differentiating and acting upon different stages of the data lifecycle. [14] With the cloud, more data is in transit - up and down-loading or transfers (sometimes across jurisdictions) - and these moves have risks. Provenance of data may become more useful as a tool. Although cloud systems are not designed to store or use provenance or other kinds of meta-data (MUNISWAMY-REDDY *et al*., 2009), research has begun to address the challenge of distributed provenance (since data is distributed across nodes and applications) and the need to protect the integrity of provenance data, itself a potential target for malice (ZHOU *et al*., 2010). And there is new work on different architectural approaches to shape where and by whom information is held and accessed, holding out the promise of more user control (and less provider control) over data that is used in cloud-based applications. [15] Or, using decoy data or otherwise mixing less-valuable with high-value data could build on prior thinking about

---

[14] Encryption is commonly proposed (the Amazon Web Services Terms of Service recommend it, for example, as does CSA for data in transit, at rest, and on back-up media and ENISA as a vehicle for end-to-end data confidentiality). Of course, there is the well-known risk of seeing encryption as a panacea and overlooking the challenges of getting it right, avoiding compromises or end-runs.

[15] For example, the "Lockr" system uses encryption to separate content from other aspects of social networks, supporting choice by social media users over where to store data and whether to disclose their social networks. See: TOOTOONCHIAN, AMIN *et al*. (2009).

honeypots and about how adding noise may make finding the signal - the truly valuable data - harder. Yet another approach is to rethink storing large blocks of data and shift toward seeking out the data that is needed when it is needed. Such a model complements the rise of sensor-systems and sensor-nets, making use of more powerful processing systems in a cloud.

Meanwhile, the cloud adds to longstanding concerns about data durability. In particular, it generates new concerns about phantom deletions. Sometimes it is good to forget, and there is considerable uncertainty about whether or when deletions (of data or algorithms) actually happen. Research on self-destructing data may provide means for data owners to protect against retroactive disclosures and attacks. But 2009-2010 saw an interesting cycle of proposed approach to data self-destruction followed by successful attack and then revision of proposed approach, a familiar cycle of measure-countermeasure that underscores how difficult it is to secure data in a network-accessible system. [16]

■ Conclusions

Cloud computing seems to be advancing inexorably, with active support within the US government based on the economics, the benefits of aggregation, and the need to move beyond legacy systems; within organizations generally based on the economics; and for individuals, based on the appeal of applications such as social media and their interconnection. For both organizations and individuals, mobility - the ability to do transactions on, say, a smart-phone - is a big driver of the public cloud. As we do more and more using cloud technology, we should remember the 2007 cyber attacks on Estonia and the vulnerability that came with having so much online. A cloud-dependent society should be aware of the risks, including how the public cloud can be a platform for malice, rethink key decisions about data, and plan for contingencies. Policy can be expected to lag - it already has - and to be impelled, as it often is, by adverse experiences. One path forward, as suggested by ENISA, may be to begin to see at least some cloud infrastructure, notably that which supports e-government applications and services, as critical information infrastructure,

[16] This history, involving a system called Vanish, began http://vanish.cs.washington.edu/ and continues with "Unvanish" challenges http://z.cs.utexas.edu/users/osa/unvanish/.

subject to protection regimes that do or will exist for critical infrastructures (CATTEDDU, 2011).

If the cloud, specifically the public cloud, is a platform for malice, individuals may have the most to lose. From an individual's point of view, the cloud, if acknowledged at all, enables personal services - Web-based e-mail, social networking, and, increasingly, mobile services and various smart-phone applications. The distance that the public cloud interposes between a user and data and/or processes is hard for most people to understand. Individuals understand even less about the technology choices of entities with which they do business, to which they give their data. Hence they are unlikely to appreciate their full exposure to the public cloud and what that implies for personal or other sensitive information. The occasional system failure - which tends to get a lot of publicity if it involves a consumer system - is a helpful reminder not to trust the public cloud too readily and to be more intentional in the handling of the data one cares most about.

Public cloud providers and their advocates would have people adopt the cliché of putting one's eggs in a basket (the cloud) and watching that basket. For the foreseeable future, it seems that we will continue to have trouble doing the necessary watching. Hence, the sister cliché about not putting all one's eggs in the same basket may be more apt. That is, absent better security mechanisms, being particularly careful about data or processes assigned to the public cloud is important. More attention to the public cloud as a platform for malice should motivate more research into better defenses, alternative architectures for data, meaningful economic comparisons of the costs and risks of traditional enterprise systems and cloud systems, and how to achieve control in the absence of the kind of control that is provided by direct ownership of infrastructure. Given governmental interests in both government uses of the cloud and the impact of everyone's uses of cloud infrastructure on the economy, governments should support relevant research. In the meantime, more awareness of risks associated with the public cloud should stimulate more careful choices about what people do with their data and the updating of legal frameworks for protecting data as information infrastructures evolve.

Might proper attention to security erode the apparent economic advantages of the public cloud? The history of cybersecurity is one of reluctance to pay the cost of security, whether that cost is in obvious dollar terms (e.g., more money for security features) or in utility (e.g., slower performance or loss of certain kinds of functionality); the market for insurance has lagged along with the market for security goods and services.

Although that history suggests a negative answer to the above question, this article has also suggested that it is possible to change the risk equation by changing choices about the use and valuation of data - changing how much of what is deemed to be at risk. It also illuminates a need for research into the economics of different scenarios, addressing alternative industry structures (notably the effects of concentration and different approaches to interconnection), the incidence of different kinds of costs (including for security expertise), and the valuation of different kinds of benefits.

References

ANDERSON, J. & STAJANO, F. (2009): "Not That Kind of Friend: Misleading Divergences Between Online Social Networks and Real-World Social Protocols" (Extended Abstract).
http://www.cl.cam.ac.uk/~fms27/papers/2009-AndersonSta-divergences.pdf
(accessed January 12, 2011).

ASAY, M. (2009): "Is cloud computing the Hotel California of tech?", *CNET News,* October 5. http://news.cnet.com/8301-13505_3-10367052-16.html
(accessed January 12, 2011).

CACHIN, C. (2009): "Trusting the Cloud", *ACM SIGACT News* (40:2), pp. 81-86, June. http://www.zurich.ibm.com/~cca/papers/trust-cloud.pdf
(accessed January 12, 2011).

CATTEDDU, D. (Ed.) (2011): *Security and Resilience in Governmental Clouds*, European Network and Information Security Agency (ENISA), January 17.
http://www.enisa.europa.eu/act/rm/emerging-and-future-risk/deliverables/security-and-resilience-in-governmental-clouds/ (accessed January 25, 2011).

CHRISTODORESCU, M., SAILER, R., SCHALES, D.L., SGANDURRA, D. & ZAMBONI, D. (2009): "Cloud Security Is Not (Just) Virtualization Security", *Proceedings of the 2009 ACM Workshop on Cloud Computing Security*,Chicago, IL.
http://portal.acm.org/citation.cfm?id=1655008.1655022&coll=ACM&dl=AMC&type=series&idx=SERIES320&part=series&WantType=Proceedings&title=CCS (accessed January 12, 2011).

CREESE, S. & HOPKINS, P. (2009): "Global Security Challenges of Cloud Computing – Extended Abstract", *Workshop on Cyber Security and Global Affairs*, August 309, St. Peter's College, Oxford, UK, Draft v0.7, July 27.
http://icc.ite.gmu.edu/sc/Global_Challenges_in_Cloud_Security_v07.pdf
(accessed January 12, 2011).

European Network and Information Security Agency (ENISA) (2009): *Cloud Computing: Benefits, risks and recommendations for information security*, Nov. 20.
http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment
(accessed January 12, 2011).

FOWLER, G. (2010): "Amazon Says WikiLeaks Violated Terms of Service", *The Wall Street Journal*, December 3.
http://online.wsj.com/article/SB10001424052748703377504575651321402763304.html
(accessed January 12, 2011).

FRIEDMAN, A.A. & WEST, D.M. (2010): "Privacy and Security in Cloud Computing", *Issues in Technology Innovation*, Number 3, October, Washington, DC: Center for Technology Innovation, Brookings Institution.

GEAMBASU, R., GRIBBLE, S. & LEVY, H. (2009): *CloudViews: Communal Data Sharing in Public Clouds*, USENIX, San Diego, June 15.
http://www.usenix.org/event/hotcloud09/tech/full_papers/geambasu.pdf
(accessed January 12, 2011).

GOLDSMITH, J. & WU, T. (2006): *Who Controls the Internet? Illusions of a Borderless World*, New York: Oxford University Press (USA).

GREENBERG, A. (2009): "Why Cloud Computing Needs More Chaos", *Forbes*, July 30. http://www.forbes.com/2009/07/30/cloud-computing-security-technology-cio-network-cloud-computing.html (accessed January 12, 2011).

KAMARA, S. & LAUTER, K. (2010): *Cryptographic Cloud Storage*, Microsoft Research Cryptography Group, January.
http://research.microsoft.com/en-us/people/klauter/cryptostoragerlcps.pdf
(accessed January 12, 2011).

KORTCHINSKY, K. (2009): "Cloudburst: A VMware Guest to Host Escape Story", Presented at BlackHat USA 2009, Las Vegas.
http://www.blackhat.com/presentations/bh-usa-09/KORTCHINSKY/BHUSA09-Kortchinsky-Cloudburst-SLIDES.pdf (accessed January 12, 2011).

KRIGSMAN, M. (2008): "MediaMax/TheLinkup: When the Cloud Fails", August 27 [Web log post]. http://blogs.zdnet.com/projectfailures/?p=999 (accessed January 12, 2011).

McMILLAN, R. (2010): "China: Google Attack Part of Widespread Spying Effort", *Macworld*, January 13.
http://www.macworld.co.uk/digitallifestyle/news/index.cfm?newsid-28293
(accessed January 12, 2011).

MUNISWAMY-REDDY, K.-K. & SELTZER, M. (2009): "Provenance as First Class Cloud Data", *3rd ACM SIGOPS International Workshop on Large-Scale Distributed Systems and Middleware (LADIS '09)*, Big Sky, MT, October.
http://www.cs.cornell.edu/projects/ladis2009/papers/muniswamy-reddy-ladis2009.pdf
(accessed January 12, 2011).

NELSON, M.R. (2009): "Building an Open Cloud", *Science*, v.324, pp.1656-1657, June 26. http://www.sciencemag.org/cgi/content/short/324/5935/1656?rss=1
(accessed January 12, 2011.

OpenCloudManifesto.org (2009): *Open Cloud Manifesto*.
http://www.opencloudmanifesto.org/open%20cloud%20manifesto.pdf
(accessed January 12, 2011).

POULSEN, K. (2009): "Future of Cybersecurity: Hackers Have Grown Up", *Wired*, July 28. http://www.wired.com/dualperspectives/article/news/2009/07/dp_security_wired0728
(accessed January 12, 2011).

RISTENPART, T., TROMER, E., SHACHAM, H. & SAVAGE, S. (2009): "Hey, You, Get Off of My Cloud", *Proceedings of the ACM Conference on Computer and Communications Security*, Chicago, November 9-13, pp. 199-212. http://portal.acm.org/citation.cfm?id=1653662.1653687&coll=GUIDE&dl=&type=series&idx=SERIES320&part=series&WantType=Proceedings&title=CCS.
(accessed January 12, 2011).

SAVAGE, S. (2009) "Are Cloud Privacy and Security Possible?", *HotCloud09: Workshop on Hot Topics in Cloud ComputingI*, San Diego: USENIX, June 15). http://www.usenix.org/publications/login/2009-10/openpdfs/hotcloud09.pdf (accessed January 12, 2011).

SPAULDING, S. E. (2010,): "No More Secrets: Then What?", *The Huffington Post,* June 24. http://www.huffingtonpost.com/suzanne-e-spaulding/no-more-secrets-then-what_b_623997.html (accessed January 12, 2011).

TOOTOONCHIAN, A., SAROIU, S., GANJALI, Y. & WOLMAN, A. (2009): "Lockr: Better Privacy for Social Networks", *ACM CoNEXT (Conference on Emerging Networking Experiments and Technologies)*, Rome, December 3. http://conferences.sigcomm.org/co-next/2009/program.php (accessed January 12, 2011).

VENEZIA, P. (2010): "McAfee's blunder, cloud computing's fatal flaw", *InfoWorld*, April 26. http://www.infoworld.com/t/software-service/mcafees-blunder-and-cloud-computings-fatal-flaw-742 (accessed January 12, 2011).

WAYNER, P. (2008): "Cloud versus cloud: A guided tour of Amazon, Google, AppNexus, and GoGrid", *InfoWorld*, July 21. http://www.infoworld.com/d/cloud-computing/cloud-versus-cloud-guided-tour-amazon-google-appnexus-and-gogrid-122 (accessed January 12, 2011).

WILSHUSEN, G.C. (2010): "Testimony Before the Committee on Oversight and Government Reform and its Subcommittee on Government Management, Organization, and Procurement, House of Representatives", GAO-1-855T, U.S. Government Accountability Office, July 1. http://www.gao.gov/new.items/d10513.pdf (accessed January 12, 2011). [relates to the May 2010 GAO report, *Information Security: Federal Guidance Needed to Address Control Issues with Implementing Cloud Architectures*. [http://www.gao.gov/new.items/d101513.pdf].

WINGFIELD, N. (2009): "Microsoft, T-mobile Stumble with Sidekick Glitch", *The Wall Street Journal*, October 11. http://online.wsj.com/article/SB10001424052748703790404574467431941990194.html (accessed January 12, 2011).

ZHOU, W., SHERR, M., MARCZAK, W.R., ZHANG, Z., TAO, T., BOON, T.L. & LEE, I. (2010): "Toward a Data-centric View of Cloud Security", Presented at *CloudDB 2010: Second International Workshop on Cloud Data Management*, October 30, Toronto. http://delivery.acm.org/10.1145/1880000/1871934/p25-zhou.pdf?key1=1871934&key2=0028684921&coll=DL&dl=ACM&CFID=5858714&CFTOKEN=74278797 (accessed January 12, 2011).

Web References:

Amazon Elastic Compute Cloud (Amazon EC2). http://aws.amazon.com/ec2/ (accessed January 12, 2011).

Amazon Simple Storage Service (Amazon S3). http://aws.amazon.com/s3 (accessed January 12, 2011).

Amazon Web Services Customer Agreement. http://aws.amazon.com/agreement/ (accessed January 12, 2011).

Cloud Security Alliance. http://www.cloudsecurityalliance.org/ (accessed January 12, 2011).

Digital Due Process. http://digitaldueprocess.org/index.cfm?objectid=37940370-2551-11DF-8E02000C296BA163 (accessed January 12, 2011).

District of Columbia Government, Office of the Chief Technology Officer. http://data.octo.dc.gov/ (accessed January 12, 2011).

Forrester Research, *Interactive Data Protection Heat Map*. http://forrester.com/cloudprivacyheatmap (accessed on January 12, 2011).

Google App Engine Terms of Service. http://code.google.com/appengine/terms.html (accessed January 12, 2011).

IBM Research, Homomorphic Encryption. http://domino.research.ibm.com/comm/research_projects.nsf/pages/security.homoenc.html (accessed January 12, 2011).

National Institute of Standards and Technology, Computer Security Division, Computer Resource Center. http://csrc.nist.gov/groups/SNS/cloud-computing/ (accessed January 12, 2011).

NEBULA Cloud Computing Platform. http://nebula.nasa.gov/about/ (accessed January 12, 2011).

Salesforce.com. http://www.salesforce.com/ (accessed January 12, 2011).

Twitter. http://search.twitter.com/ (accessed January 12, 2011).

Unvanish – Reconstructing Self-Destructing Data. http://z.cs.utexas.edu/users/osa/unvanish/ (accessed January 12, 2011).

U.S. General Services Administration, Apps.gov. https://apps.gov/cloud/advantage/main/start_page.do (accessed January 12, 2011).

Vanish – Self-Destructing Digital Data. http://vanish.cs.washington.edu/ (accessed January 12, 2011).

Windows Azure: Microsoft's Cloud Services Platform. http://www.microsoft.com/windowsazure/ (accessed January 12, 2011).