

What's New About Cloud Computing Security?

*Yanpei Chen
Vern Paxson
Randy H. Katz*



Electrical Engineering and Computer Sciences
University of California at Berkeley

Technical Report No. UCB/EECS-2010-5

<http://www.eecs.berkeley.edu/Pubs/TechRpts/2010/EECS-2010-5.html>

January 20, 2010

Copyright © 2010, by the author(s).
All rights reserved.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission.

What's New About Cloud Computing Security?

Yanpei Chen, Vern Paxson, Randy H. Katz

CS Division, EECS Dept. UC Berkeley

{ychen2, vern, randy}@eecs.berkeley.edu

ABSTRACT

While the economic case for cloud computing is compelling, the security challenges it poses are equally striking. In this work we strive to frame the full space of cloud-computing security issues, attempting to separate justified concerns from possible over-reactions. We examine contemporary and historical perspectives from industry, academia, government, and “black hats”. We argue that few cloud computing security issues are fundamentally new or fundamentally intractable; often what appears “new” is so only relative to “traditional” computing of the past several years. Looking back further to the time-sharing era, many of these problems already received attention. On the other hand, we argue that two facets are to some degree new and fundamental to cloud computing: the complexities of multi-party trust considerations, and the ensuing need for mutual auditability.

Categories and Subject Descriptors

C.2.0 [Computer-Communication Networks]: General—*Security and Protection*

General Terms

Design, Security, Reliability

1. INTRODUCTION

The economic case for cloud computing has gained widespread acceptance. Cloud computing providers can build large datacenters at low cost due to their expertise in organizing and provisioning computational resources. The economies of scale increase revenue for cloud providers and lower costs for cloud users. The resulting on-demand model of computing allows providers to achieve better resource utilization through statistical multiplexing, and enables users to avoid the costs of resource over-provisioning through dynamic scaling [12, 2].

At the same time, security has emerged as arguably the most significant barrier to faster and more widespread adoption of cloud computing. This view originates from perspectives as diverse as academia researchers [12], industry decision makers [35], and gov-

ernment organizations [29, 3]. For many business-critical computations, today's cloud computing appears inadvisable due to issues such as service availability, data confidentiality, reputation fate-sharing, and others.

To add to the confusion, some have criticized the term “cloud computing” as too broad [21]. Indeed, cloud computing does include established business models such as *Software as a Service*, and the underlying concept of on-demand computing utilities goes back as far as early time-sharing systems [17]. At the same time, the lack of consistent terminology for cloud computing has hampered discussions about cloud computing security. Thus, security criticisms of cloud computing have included a murky mix of ongoing and new issues.

This context frames the genesis of our paper. We recognize that security poses major issues for the widespread adoption of cloud computing. However, secure or not, cloud computing appears here to stay. Thus, our ambition is to get past terminology issues (Section 2) and attempt to sort out what are actually new security issues for cloud computing, versus broader and more general security challenges that inevitably arise in the Internet age. Our goal is to advance discussions of cloud computing security beyond confusion, and to some degree fear of the unknown, by providing a comprehensive high-level view of the problem space.

We ground the development of our viewpoint in a survey of contemporary literature on cloud computing security, coupled with a review of historical work on early time-sharing systems and virtual machine monitors. Contemporary discussions reveal security concerns that are indeed “new” relative to computing of the past decade (Section 3); however, looking back several decades, many contemporary challenges have quite similar historical counterparts (Section 4).

We build the case that few of the security problems arising in cloud computing are in fact new, even though satisfactory solutions for many still will require significant development. The combined contemporary and historical viewpoints allow us to identify a number of research topics that deserve more attention (Section 5). On the other hand, we argue that two facets are to some degree new and fundamental to cloud computing: the complexities of multi-party trust considerations, and the ensuing need for mutual auditability.

2. DISTRACTED BY DEFINITIONS

The lack of a clear and widely accepted definition has posed a barrier to talking about cloud computing in general. Clearly “cloud computing” is an evolving term, defined more by usage than by

written documents. That said, overly broad use has led to criticism that cloud computing “include[s] everything that we already do” [21]. Similarly, splitting hairs on the precise definitions distracts us from the core technology issues. In this section, we briefly frame the definition we use for the remainder of our discussion.

An “early” (less than one year old!) effort at systematically framing cloud computing, “Above the Clouds: A Berkeley View of Cloud Computing,” defined cloud computing to include application software delivered as services over the Internet, and the hardware and systems software in the datacenters that facilitate these services [12]. Key characteristics of cloud computing include the illusion of infinite hardware resources, the elimination of up-front commitment, and the ability to pay for resources as needed.

This whitepaper spurred a flurry of follow-on cloud computing definitions and reports. For our purposes, the most notable of these is that published by the U.S. National Institute of Standards and Technology (NIST) [30]. NIST frames a broader definition, one that includes nearly all common terms used in cloud computing discussions and forms the basis for the NIST guide on cloud computing security [29]. It appears that other efforts may converge on a similar framing; most visibly, the European mirror effort to [29], a report from the European Network and Information Security Agency (ENISA), defines cloud computing in the same spirit as the NIST definition [3].

According to the NIST definition, key characteristics of cloud computing include on-demand self service, broad network access, resource pooling, rapid elasticity, and metered service similar to a utility. There are also three main service models—software as a service (*SaaS*), in which the cloud user controls only application configurations; platform as a service (*PaaS*), in which the cloud user also controls the hosting environments; and infrastructure as a service (*IaaS*), in which the cloud user controls everything except the datacenter infrastructure. Further, there are four main deployment models: public clouds, accessible to the general public or a large industry group; community clouds, serving several organizations; private clouds, limited to a single organization; and hybrid clouds, a mix of the others.

In keeping with this evolution, and because we believe the broad scope of the NIST definition enables us to encompass the full set of issues of interest, for the rest of this paper, we will talk about “cloud computing” in the spirit of the NIST definition.

3. CONTEMPORARY ASSESSMENT

In this section, we assess what appears new to cloud computing and what does not, so that we can identify the most challenging aspects of the cloud computing security threat model.

3.1 What is not new

With increased employment of cloud computing comes increasingly frequent cloud computing security incidents. Arguably many of the incidents described as “cloud security” in fact just reflect traditional web application and data-hosting problems. The underlying issues remain well-established challenges such as phishing [4], downtime [24], data loss [38], password weaknesses [31], and compromised hosts running botnets [20]. The Twitter phishing incident provides a typical example of a traditional web security issue now miscast as a cloud computing issue [4]. In contrast, we find the recent Amazon botnet incident noteworthy because it reflects one of the first known compromises of a major cloud provider [20],

highlighting that servers in cloud computing currently operate as (in)securely as servers in traditional enterprise datacenters.

In academia, cloud computing security has begun seeing the development of dedicated forums such as the ACM Cloud Computing Security Workshop, as well as dedicated tracks at major security conferences such as the ACM Conference on Computer and Communications Security (CCS). To date, most papers published on cloud security reflect continuations of established lines of security research, such as web security [40, 13], data outsourcing and assurance [14, 18], and virtual machines [41, 34]. The field primarily manifests as a blend of existing topics, rather than a set of papers with an exclusive focus on cloud security, though there are exceptions, such as [32], which we discuss below.

The “black hat” community has also discovered cloud computing exploits that reflect extensions of existing vulnerabilities, with a dedicated cloud security track emerging at *Black Hat USA 2009*. For example, username brute forcers and Debian OpenSSL exploit tools run in the cloud as they do in botnets [28]. Social engineering attacks remain effective—one exploit tries to convince Amazon Elastic Compute Cloud (EC2) users to run malicious virtual machine images simply by giving the image an official-sounding name such as “fedora_core” [28]. Virtual machine vulnerabilities also remain an issue [25], as does weak random number generation due to lack of sufficient entropy [37].

3.2 What is new

For black hats, cloud computing offers a potentially more trustworthy alternative to botnets. While the recent brute-forcer presentation [28] claimed that using the cloud is presently more expensive than using botnets, another Black Hats presentation asserted that the botnet market likely suffers from the “lemon market” problem, where the lack of trust and the inability to verify the quality of goods leads to a minimal volume of goods being exchanged [22]. If this were the case, then attackers can find more reliable service in cloud computing at a premium price.¹ That said, botnets in the cloud are easier to shut down than traditional botnets.

Also, because cloud computing introduces a shared resource environment, unexpected side channels (passively observing information) and covert channels (actively sending data) can arise. One noteworthy paper [32] tackles precisely this problem. The exposed vulnerabilities include ways to place an attacker virtual machine (VM) on the same physical machine as a targeted VM, and then to construct a side channel between two VMs on the same physical machine, which enables the SSH keystroke timing attack outlined in [36]. This work also provides an example of research targeted exclusively at cloud computing.

Another new issue comes from reputation fate-sharing, which has mixed consequences. On the plus side, cloud users can potentially benefit from a concentration of security expertise at major cloud providers, ensuring that the entire ecosystem employs security best practices. On the other hand, a single subverter can disrupt *many* users. For example, spammers subverted EC2 and caused Spamhaus to blacklist a large fraction of EC2’s IP addresses, causing major service disruptions. There-

¹Note that the prices can be quite low. For example, we estimate that to reduce the brute force exploit in [36] to a single minute, rather than 1.3 PC-days, would require 200 extra-large EC2 instances, which at January 2010 pricing would total at about \$2 per exploit.

after, if someone wants to send email from EC2, they must fill out the form (<http://aws.amazon.com/contact-us/ec2-email-limit-request/>), provide a list of (static) EC2 addresses to authorize for sending, and document their use-case. Upon approval, Amazon forwards the EC2 addresses to Spamhaus for whitelisting [8].

A second noteworthy fate-shrining incident occurred during an FBI raid on Texas datacenters in April 2009, based on suspicions of the targeted datacenters facilitating cybercrimes. The agents seized equipment, and many businesses co-located in the same datacenters faced business disruptions or even complete business closures. One affected customer applied for a temporary restraining order, and was denied because the equipment concerned may have been used for criminal activities without the customer's knowledge [6].

3.3 Novelties in the cloud threat model

Putting together these discussions, we argue that the cloud computing threat model includes several novel elements.

First, data and software are not the only assets worth protecting. Activity patterns also need to be protected. Sharing of resources means that the activity of one cloud user might appear visible to other cloud users using the same resources, potentially leading to the construction of covert and side channels. Activity patterns may also themselves constitute confidential business information, if divulging them could lead to reverse-engineering of customer base, revenue size, and the like.

Business reputation also merit protection. When using shared resources to do business-critical computations, it becomes harder to attribute malicious or unethical activity. Even if there are ways to clearly identify the culprits and attribute blame, bad publicity still creates uncertainty that can tarnish a long-established reputation.

In addition, one must often accommodate a longer trust chain. For example, the application end-user could potentially use an application built by an *SaaS* provider, with the application running on a platform offered by a *PaaS* provider, which in turn runs on the infrastructure of an *IaaS* provider. While to our knowledge this extreme example cannot occur in practice today due to a lack of sufficient APIs, it illustrates that with any model of cloud computing, stakeholders' can find themselves with relationships considerably more complicated than simply a provider-user relationship.

Some participants could be subverters, who maintain the appearance of a regular cloud user or cloud provider, but in fact perpetrate cybercrime or other cyber attacks. Examples include cloud users who run brute forcers, botnets, or spam campaigns from the cloud; or cloud providers who scan cloud users' data and sell confidential information to the highest bidder.

Furthermore, competitive businesses can operate within the same cloud computing ecosystem: using the same cloud, or ending up in a provider-user relationship. This can lead to strong conflicts of interest, and creates additional motives to access the confidential information of a competitor.

These complications point to the need for *auditability* in cloud computing—already a requirement for health care, banking, and similar systems. What is new to cloud computing is *mutual* auditability. Because the system includes stakeholders with potentially conflicting interests, cloud users and providers both need re-

assurance that the other in a fashion that is both benign and correct (from a billing standpoint).

Mutual auditability can also significantly assist with incident response and recovery, since both the cloud provider and the cloud user could be either the source or the target of an attack. Auditability also enables the *attribution of blame* in search and seizure incidents, which can prove vital so that law enforcement agencies do not overreach in carrying out their duties.

Finally, a subtle difficulty with understanding cloud computing threats arises from potentially inaccurate mental models of cloud computing as an always-available service. This viewpoint—which arises from the general paradigm of drawing upon a commodity service with much the flavor of a utility—can create a false sense of security, leading to inadequate security good practices, such as regular data backups across multiple cloud providers. As such, we could find that while cloud computing fails at the same rate as other types of systems, the impact of those failures manifest more severely.

4. SOME DEJA VU

In this section we present three explorations of early computing systems that had characteristics similar to what we call cloud computing today. The profiles suggest that many contemporary cloud security issues will prove tractable, as their similar historical counterparts were indeed successfully tackled; but also that some assumptions from the past no longer wholly apply, and risk complicating our assessment of security issues due to out-of-date mental models. These historical approaches also offer us starting points to consider for current cloud security research, as we then develop in Section 5.

4.1 Multics

Multics introduced the “computing utility” concept as early as 1965 [17], in the same sense that cloud computing has taken off as providing today's computing utilities. Security considerations permeated all aspects of Multics design [33], and its security mechanisms influenced those of subsequent systems. Consequently, Multics was the first system to receive a Class B2 certification per the Orange Book [39].

A striking aspect of Multics was its security design principles [33], which deserve re-emphasis today. First, Multics used permission-based protection mechanisms, rather than exclusion-based. Every access to every object checked current authority. Second, Multics embodied a form of Kerckhoffs' principle, maintaining open design for its mechanisms, with only the protection keys secret. Third, the system always operated at least privilege. Finally, the design explicitly recognized the importance of human usability—especially relevant today with the proliferation of social engineering attacks.

Multics security design also framed the importance of preventing system administrators from becoming decision bottlenecks. Otherwise, users will bypass administrators by habit (in modern terminology, a form of “satisficing”) and compromise protection mechanisms [33]. Recall from Section 3.2 that the response to the Amazon EC2 spam blacklist incident involved imposing email limits that require administrator intervention to increase; this mechanism may become unscalable if EC2 users who wish to send email significantly increase.

Multics did not aim for security in an absolute sense, but allowed

users to build protected subsystems [33, 42]. Similarly, in cloud computing different users will have different security needs, so a good design would offer a choice of security levels and security mechanisms. Cloud providers have begun taking the first steps in this direction with offerings such as virtual private clouds, with dedicated resources and virtual private networks that “guaranteed” isolation [1]. The “spectrum of security” approach is worth advocating.

On a related note, key “Multicians” had a heavy influence on the Department of Defense Orange Book certification document [39, 9]. The Orange Book includes a treatment of covert channels quite similar to that of contemporary side channel work [32, 36]. In both cases, the risk assessment principles involve a quantification of the channel bit-rate, accompanied by an assessment of the bit-rate that constitutes a significant risk. The Orange Book sets this bit-rate as the level necessary to operate a computer terminal. In [36], the bit-rate corresponds to the workload reduction for a brute force password breaker. But even putting bit rate aside, in some settings the mere presence of a covert channel or side channel constitutes a significant risk, and more broadly both types of information leakage are fundamental concerns for cloud computing.

In closing the Multics discussion, we note that a number of Multics security mechanisms, state-of-the-art at the time, remain prevalent today even though they do not work as well for modern computing environments. These mechanisms include access control lists (ACLs), machine-generated passwords, and weak encryption of the password file [33, 39]. Thus, while historical work can provide valuable insights into modern cloud security issues, naturally we must temper our assessment of those mechanisms with due consideration to how computing has changed over time.

4.2 Early VMMs

We find early work on virtual machine monitors (VMMs) noteworthy because different kinds of virtualization constitute a major facet of cloud computing. Here, we review the original argument of why VMMs are more secure than ordinary computing systems [26], and frame why the core assumptions of this argument no longer hold for today’s VMMs.

The argument has several parts. First, lower levels of multiprogramming (i.e. concurrent execution) lead to lower risks of security failures; in the extreme, a monoprogramming operating system (OS) has a much lower security risk than an OS running many concurrent programs. Thus, VMMs with low multiprogramming levels will prove more secure than OSs with high multiprogramming levels. Second, even if the level of multiprogramming is the same, VMMs are more secure because they are simpler and easier to debug. Third, for a guest OS that runs on a VMM that in turn runs on bare metal, security violation occurs only when both the guest OS and the VMM fail simultaneously. Thus, a VMM running k guest OSs with each OS running n programs fails much less easily than an OS running $k \times n$ programs. Fourth, the failure of each program is independent, and hence the failure probability is multiplicative. Thus, overall, any one program on a VMM running k guest OSs with each OS running n programs fails much less frequently than the same program on an OS with $k \times n$ programs. The multiplication effect amplifies the reduction in each failure probability.

The argument makes three crucial assumptions. First, VMMs are simple. Second, guest OSs have a lower multiprogramming level. Third, the VMM and guest OS have independent failures. Modern

VMMs undermine all three assumptions.

Modern VMMs are no longer “small” in an absolute sense. For example, Xen has approximately 150,000 lines of code [11]. While still considerably smaller than recent operating systems (e.g., \approx 12 million lines for Linux 2.6.32 [7]), this level is comparable to 176,250 lines of code for Linux 1.0 [5], which already constituted a fairly feature-rich general purpose OS.

Additionally, today a guest OS usually has the same level of multiprogramming as the native OS. Users treat guest OSs the same way they would treat a native OS, undermining the assumption that guest OSs have lower multiprogramming levels.

Further, some recent VMMs have the guest OS running on a VMM that in turn runs on a host OS [10]! In such a setup, clearly the VMM is as (in)secure as the host OS, and the host OS significantly enlarges the trusted code base.

Other researchers have raised similar concerns [23]. Thus, for cloud computing security, clearly we need to examine whether such assumptions hold for virtualizations at network or datacenter levels.

4.3 National CSS, Inc.

We finish our framing of historical perspectives with a case study examination of National CSS, Inc., a time-sharing company comparable to cloud providers today. The founders of the company envisioned moving upfront costs to variable costs, and the company succeeded due to the increased flexibility that their ready-to-use computing capability provided [15]. Cloud computing offers similar economic benefits today.

While the experiences of one company from the past clearly do not generalize to the experiences of others several decades later, we want to highlight two incidents on reputation fate-sharing which may prove illuminating for cloud computing today.

The first incident led to a negative outcome for National CSS. In 1979 an attacker stole a password directory from National CSS, compromising the security of all its corporate customers [27]. The company warned its 8,000 clients about a security problem, but did not provide additional details, which lead to a strong negative reaction. On the other hand, while their clients wanted more information, the company also “drew the wrath of many industry professionals for not covering up the incident.” Eventually, the FBI also became involved, creating even more negative publicity.

In contrast, another incident proved a major success. A hardware failure led to data loss for Bell Labs, a major National CSS customer. Contrary to standard procedure, there were no backups and the company deemed the data loss “irrecoverable”. National CSS conveyed the failure directly and honestly to Bell Labs. The message was that National CSS had screwed up, and would do all it could to help Bell Labs recover the data. After the initial shock, Bell Labs worked with National CSS, typing in data from stacks of printouts. The incidence response so impressed Bell Labs that they became a much bigger customer of the company [19].

Thus, while cloud computing has complicated stakeholder relationships coupled with reputation fate-sharing, these incidents are suggestive with regard to the benefits of managing security risks by aligning business interests and building stakeholder partnerships.

5. NEW OPPORTUNITIES

Combining the contemporary and historical viewpoints, we arrive at the position that many cloud computing security problems are not in fact new, but often will still require new solutions in terms of specific mechanisms. Existing contemporary works already explore many pertinent topics; we highlight here several areas that deserve more attention.

First, cloud providers should offer a choice of security primitives with well-considered defaults. Cloud users know more about their applications, but cloud providers potentially know more about the relevant security issues due to a higher concentration of security expertise. The cloud user would ideally choose from a spectrum of security levels and security subsystem boundaries. We believe this flexibility could prove to be a major improvement if done well. One possible approach would be to formulate the security primitives around defending different stakeholders against different particular threat models. An additional feature might support “plug-and-play” services readily compliant with common standards such as those of HIPAA or Payment Card Industry.

Another important research area concerns determining apt granularities for isolation. Several are possible: isolate by virtual or physical machines, LANs, clouds, or datacenters. We at present lack a good understanding of the tradeoffs between security and performance for each of these options, but it would appear likely that cloud providers can fruitfully offer different granularities of isolation as a part of their spectrum of security.

Side channels and covert channels pose another fundamental threat, one which interplays with the granularities of isolation discussed above. While not a panacea (e.g., it takes very few bits to steal a password), a helpful analysis could include when appropriate a quantification of channel bit rates, coupled with an assessment of the bit rate required to do harm. The approaches in [32] and [36] provide good examples.

One important area that has yet to receive much attention is *mutual auditability*. The auditing capabilities of most existing systems focus on one-way auditability. In cloud computing, providers and users may need to demonstrate mutual trustworthiness, in a bilateral or multilateral fashion. As discussed above, such auditability can have major benefits with regard to fate-sharing, such as enabling cloud providers in search and seizure incidents to demonstrate to law enforcement that they have turned over all relevant evidence, and prove to users that they turned over *only* the necessary evidence and nothing more. Recent work notes that implementing thorough auditing is not a simple matter even for straightforward web services [16]. In cloud computing, it remains an open challenge to achieve thorough auditing without impairing performance. To complicate matters even further, the auditor fundamentally needs to be an independent third party, and a third-party auditor requires a setup quite different than today’s practice, in which cloud providers record and maintain all the audit logs. In short, mutual auditability needs significant work. On the plus side, achieving it robustly would constitute an important security feature.

More broadly, we see a need for research that seeks to understand the *ecosystem* of threats. Current work in the literature generally focuses only single aspects of the cloud security problem. As we begin to understand problems in isolation, we should also start to put together an understanding of how different issues and threats combine. For example, in web security we understand se-

curity problems at a high-level as an ecosystem involving the interplay between worms, bots, scams, spam, phishing, active content, browsers, usability, and other human factors. We argue that future work on cloud security needs to similarly bridge established topic boundaries.

Lastly, we would highlight that breaking real clouds makes them stronger. Such studies involve obvious ethical issues, but provide much more compelling results than breaking hypothetical clouds. For example, the EC2 information leak study in [32] triggered a highly visible security effort by Amazon Web Services, and serves as a model for similar future work in academia. Similarly, the Air Force Multics security enhancements [42] originated from a companion effort to find security exploits. Such coupled attack and defense approaches serve as a model for potential government cloud security projects today, and cloud providers should sponsor internal adversarial efforts to discover vulnerabilities before they become exposed in the wild. Needless to say, stakeholders also need to continue to track black-hat perspectives. Finally, research partnerships between different types of stakeholders will likely prove very beneficial to advancing the field.

6. FINAL THOUGHTS

Given the stakes, it strikes us as inevitable that security will become a significant cloud computing business differentiator. Furthermore, in addition to revisiting approaches for specific issues in securing shared computing, history teaches us that developing security architectures early in the process can pay off greatly as systems evolve and accrue more disparate functionality. On the other hand, the history of commercial Internet offerings repeatedly shows that time-to-market and undercutting prices can greatly sway customers even in the absence of sound security underpinnings. The situation may be somewhat different this time around, however, given that much of cloud computing targets customers who have extensive business reasons (and scars from the past) leading them to treat security as an elevated priority.

We close our discussion with what we find to be an interesting analogy. Companies such as National CSS began by offering affordable computation for businesses. Time-sharing eventually gave way to personal computers, which brought affordable computation to the general public. In a similar fashion, cloud computing currently offers affordable, large-scale computation for businesses. If the economic case prevails, then we may find that nothing—not even security concerns—will prevent cloud computing from becoming a consumer commodity. Just as the commodity PC and the Internet brought about the Information Revolution, and made information universally accessible, affordable, and useful, so too does cloud computing have the potential to bring about the Computation Revolution, in which large-scale computations become universally accessible, affordable, and useful. Let’s hope we can add to this outcome “and be reasonably safe”.

7. REFERENCES

- [1] Amazon virtual private cloud. <http://aws.amazon.com/vpc/>.
- [2] Amazon web services economics center. <http://aws.amazon.com/economics/>.
- [3] Cloud computing risk assessment. European Network and Information Security Agency. November 20, 2009.
- [4] Gone phishing. Twitter Blog. January 03, 2009.
- [5] Linux kernel. Wikipedia.

- [6] Liquid Motors, Inc. v. Allyn Lynd and United States of America. U.S. District Court for the Northern District of Texas, Dallas Division. April 2009.
- [7] Summary of Linux 2.6.32. h-online.com.
- [8] Thread 37650: Email changes. Amazon Web Services Discussion Forums.
- [9] Trusted Computer System Evaluation Criteria (Orange Book). Department of Defense Standard. DoD 5200.28-STD. December 1985.
- [10] VMware Workstation. <http://www.vmware.com/products/workstation/>.
- [11] Xen hypervisor. <http://xen.org/products/xenhyp.html>.
- [12] Michael Armbrust et al. Above the Clouds: A Berkeley View of Cloud Computing. Technical report EECS-2009-28, UC Berkeley, <http://www.eecs.berkeley.edu/Pubs/TechRpts/2009/EECS-2009-28.html>, Feb 2009.
- [13] Robert Biddle, P. C. van Oorschot, Andrew S. Patrick, Jennifer Sobey, and Tara Whalen. Browser interfaces and extended validation ssl certificates: an empirical study. In *CCSW '09: Proceedings of the ACM workshop on Cloud computing security*.
- [14] Kevin D. Bowers, Ari Juels, and Alina Oprea. Hail: a high-availability and integrity layer for cloud storage. In *CCS '09: Proceedings of the 16th ACM conference on Computer and communications security*.
- [15] P. Ceruzzi. An Interview with Robert E. Weissman. Charles Babbage Institute. May 3, 2002.
- [16] Anton Chuvakin and Gunnar Peterson. Logging in the age of web services. *IEEE Security and Privacy*, 7(3):82–85, 2009.
- [17] Fernando J. Corbató and V. A. Vyssotsky. Introduction and overview of the multics system. *IEEE Ann. Hist. Comput.*, 14(2):12–13, 1992.
- [18] Chris Erway, Alptekin Küpçü, Charalampos Papamanthou, and Roberto Tamassia. Dynamic provable data possession. In *CCS '09: Proceedings of the 16th ACM conference on Computer and communications security*.
- [19] H. Feinleib. A Technical History of National CSS. Computer History Museum. April 2005.
- [20] M. C. Ferrer. Zeus in-the-cloud. CA Community Blog. December 9, 2009.
- [21] G. Fowler and B. Worthen. The internet industry is on a cloud—whatever that may mean. Wall Street Journal. March 26, 2009.
- [22] C. Herley. Economics and the underground economy. Black Hat USA 2009.
- [23] P. Karger. Securing virtual machine monitors—what is needed. Keynote address, ASIACSS 2009.
- [24] E. Knorr. Gmail follies and Google’s enterprise pitch. InfoWorld. September 8, 2009.
- [25] K. Kortchinsky. Cloudburst—a VMware guest to host escape story. Black Hat USA 2009.
- [26] Stuart E. Madnick and John J. Donovan. Application and analysis of the virtual machine approach to information system security and isolation. In *Proceedings of the workshop on virtual computer systems*. ACM, 1973.
- [27] V. McLellan. Case of the purloined password. New York Times. July 26, 1981. <http://www.nytimes.com/1981/07/26/business/case-of-the-purloined-password.html>.
- [28] H. Meer, N. Arvanitis, and M. Slaviero. Clobbering the cloud. Black Hat USA 2009.
- [29] P. Mell and T. Grance. Effectively and securely using the cloud computing paradigm. National Institute of Standards and Technology. October 7, 2009.
- [30] P. Mell and T. Grance. NIST definition of cloud computing. National Institute of Standards and Technology. October 7, 2009.
- [31] D. Raywood. The twitter hacking incident last week should be a call to better security awareness and not about cloud storage. SC Magazine. July 20, 2009.
- [32] Thomas Ristenpart, Eran Tromer, Hovav Shacham, and Stefan Savage. Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. In *CCS '09: Proceedings of the 16th ACM conference on Computer and communications security*.
- [33] Jerome H. Saltzer. Protection and the control of information sharing in multics. *Commun. ACM*, 17(7):388–402, 1974.
- [34] N. Santos, K. P. Gummadi, and R. Rodrigues. Towards trusted cloud computing. Hot Cloud 2009. http://www.usenix.org/event/hotcloud09/tech/full_papers/santos.pdf.
- [35] S. Shankland. HP’s Hurd dings cloud computing, IBM. CNET News. October 20, 2009.
- [36] Dawn Xiaodong Song, David Wagner, and Xuqing Tian. Timing analysis of keystrokes and timing attacks on SSH. In *SSYM'01: Proceedings of the 10th conference on USENIX Security Symposium*.
- [37] A. Stamos, A. Becherer, and N. Wilcox. Cloud computing security—raining on the trendy new parade. Black Hat USA 2009.
- [38] J. Stokes. T-Mobile and Microsoft/Danger data loss is bad for the cloud. Ars technica. October 2009.
- [39] T. van Vleck. How the Air Force cracked multics Security. multicians.org. May 21, 1993.
- [40] K. Vikram, Abhishek Prateek, and Benjamin Livshits. Ripley: automatically securing web 2.0 applications through replicated execution. In *CCS '09: Proceedings of the 16th ACM conference on Computer and communications security*.
- [41] Jinpeng Wei, Xiaolan Zhang, Glenn Ammons, Vasanth Bala, and Peng Ning. Managing security of virtual machine images in a cloud environment. In *CCSW '09: Proceedings of the ACM workshop on Cloud computing security*.
- [42] J. Whitmore, A. Bensoussan, P. Green, D. Hunt, A. Robziar, and J. Stern. Design for multics security enhancements. Technical report ESD-TR-74-176, Air Force Systems Command, <http://csrc.nist.gov/publications/history/whit74.pdf>, Dec 1973.

3.1 What is not new With increased employment of cloud computing comes increasingly frequent cloud computing security incidents. Arguably many of the incidents described as "cloud security" in fact just reflect traditional web application and data-hosting problems. The underlying issues remain well-established challenges such as phishing [4], downtime [24], data loss [38], password weaknesses [31], and compromised hosts running botnets [20]. The Twitter phishing incident provides a typical example of a traditional web