A NEW PERSPECTIVE ON PROVIDING CLOUD COMPUTING SECURITY
A Position Paper

Ashraf Matrawy*, Clifford Liem, Michael Wiener, Yuan Xiang Gu and Andrew Wajs
Cloakware Advanced Research Center (CARC), Irdeto, Ottawa, Canada

Keywords: IPR, Cloud, Isolation, Multi-tenancy, Hypervisor, Mutual auditing, Extended trust, Perimeter defence.

Abstract: Security issues in cloud computing are ranked high on the list of reasons why many organizations delay considering the computing paradigm for their future plans. There have been numerous discussions about the security problem in cloud computing. In this paper, we take a different stand and discuss the impact that cloud computing will have on offering security. The paper points out the unique aspects of cloud computing architecture and operation and cloud-specific security issues. These cloud-specific issues could have an impact on a number of areas of offering information security. We discuss its impact on Intellectual Property Rights (IPR) protection, organizations extending protection beyond the traditional notion of perimeter defence, and on attack surfaces. We point out how providing security in the new computing paradigm may be affected and that there might be some benefits that the cloud brings to the information security scene. In summary, in this paper, we attempt to initiate discussions around these important issues.

1 INTRODUCTION

Cloud computing is *not* a new technology but rather a new paradigm for delivering computing power to different types of computing power consumers (Armbrust et al., 2010; Mell and Grance, 2009). While it provides significant economic and technical benefits, there are many security concerns surrounding it. This new paradigm encompasses a multitude of computing and communication technologies. Therefore, it is reasonable to expect cloud computing to inherit most of the security problems of these underlying technologies. Two cloud computing characteristic properties contribute to the creation of security problems that are distinct to this paradigm (1) *IT outsourcing* where an organization's data and software are hosted beyond its control and (2) *the sharing of physical resources* by multiple customers (competitors?) which is referred to as multi-tenancy. There are other factors that may create new problems but these two appear to be the most prominent at the moment.

Many of the cloud computing security issues that have been reported recently (e.g. (ENISA, 2009; CSA, 2009)) are not specific to cloud computing. Actually some reported cloud security problems date back to the time-sharing era in the 1960s (Chen et al., 2010). This paper focuses on the new aspects of cloud computing that are creating security problems that are *specific* to cloud environments. In particular, *we focus on how cloud computing will affect the offering of information security.*

Contribution. We intend this work to be a starting point for discussions about providing security in the cloud rather than a paper listing security problems in the cloud. This paper raises important questions about crucial issues such as changing the notion of perimeter defence, how the cloud impacts Intellectual Property Rights (IPR) protection, and the changing attack surfaces. We also point out how the cloud may be beneficial to security. Note that this paper addresses the technical viewpoint only. Legal, financial, and other aspects of the problem are outside the paper's scope.

Outline. This paper is structured as follows: Section 2 briefly explains the special features of the cloud environment that pertain to the discussion on the new security problems in this paradigm. Section 3 briefly

*Ashraf Matrawy is an associate professor at Carleton University. He did this work while working at CARC during his sabbatical. After June 2011 he could be reached at *ashraf_matrawy@carleton.ca*.

Matrawy A., Liem C., Wiener M., Xiang Gu Y. and Wajs A..
A NEW PERSPECTIVE ON PROVIDING CLOUD COMPUTING SECURITY - A Position Paper.
DOI: 10.5220/0003449106500655
In *Proceedings of the 1st International Conference on Cloud Computing and Services Science* (CLOSER-2011), pages 650-655
ISBN: 978-989-8425-52-2
Copyright © 2011 SCITEPRESS (Science and Technology Publications, Lda.)

discusses the different efforts to summarize the cloud security problems. Section 4 categorizes the new security problems that are special to the cloud. Section 5 discusses attack surfaces and Section 6 discusses the changing notion of perimeter defence. Section 7 tackles IPR protection in the cloud while Section 8 focuses on security life cycle and the cloud. Section 9 concludes the paper.

2 THE CLOUD ENVIRONMENT

A comprehensive overview of cloud computing can be found in the work by the UC at Berkeley (Armbrust et al., 2010). Examples of other overviews are the presentation by NIST at ACM CCSW 2009 (Mell and Grance, 2009) and the white paper by Sun Microsystems (Sun, 2009). In this section, a summary of the special features discussed in these references is presented.

The elastic nature of the cloud is a distinct characteristic. The ability to scale horizontally is more important than efficiency on a single Virtual Machine (VM). This is especially important in cases where "surge computing" (Sun, 2009) is needed. Applications could be programmed to request more virtual resources from the elastic infrastructure. This could be done through APIs specified by the cloud provider (Sun, 2009). This means that creating and managing VMs is crucial in cloud computing. The security of this process is one of the important new aspects.

In terms of cost and exposure, clouds could be *private*, shared within a certain *community*, or *public*. This order reflects an increasing level of exposure and a decreasing cost. Clouds could also be a *hybrid* of the models above. In organizations requiring a high degree of isolation (e.g., government), the cost to have a private cloud might be justifiable. A community cloud might be shared among organizations with similar requirements. *We note that most of the discussion in this paper applies to shared clouds (public or community) and not the case where an organization runs its own private cloud.*

Clouds provide services at different levels usually defined as SaaS (Software as a Service), PaaS (Platform as a Service), and IaaS (Infrastructure as a Service). The boundaries between these services may not be clear and their definition may change depending on the provider and the customer of the service.

3 TRADITIONAL SECURITY PROBLEMS

This massive concentration of resources and data in the cloud presents a more attractive target to attackers. The vulnerabilities that exist in the underlying technologies naturally migrate to the cloud. The work by the Cloud Security Alliance (CSA, 2009) and ENISA (ENISA, 2009) includes long lists of these security problems.

Vulnerabilities. According to the ENISA report (ENISA, 2009) there are numerous vulnerabilities that could affect the cloud environment. The report lists 31 cloud vulnerabilities and 22 additional vulnerabilities that are not specific to the cloud. Some of these vulnerabilities are not technical. The most noted technical ones include hypervisor vulnerabilities, user and application isolation, internal probing and encryption-related vulnerabilities.

Benefits. There are arguments that cloud-based defences can be more robust, scalable and cost-effective and that these defences will be a market differentiator between cloud providers. Additionally they could provide (ENISA, 2009): (1) standardized interfaces for managed security (2) rapid scaling of security resources to activities such as filtering and encryption. Another activity is auditing and evidence gathering through images of VMs although this aspect of cloud security is proving to be very challenging (3) timely and effective updates (4) resource centralization (cheaper, easier, etc.).

However, some of the above mentioned benefits could be the source of problems as well such as the case of centralization where it becomes a single point of failure.

Another benefit would be making some costly security features more affordable where the cloud provider makes up-front investment in a certain tool that will make the cost per-user per-transaction much more affordable for smaller users. An example is forensic investigation tools (Molnar and Schechter, 2010).

4 CLOUD-SPECIFIC SECURITY PROBLEMS

4.1 IT Outsourcing Problems

In the cloud, an organization will rely on applications and systems that are not under its IT staff's control and that may not fully comply with its own policies.

Moreover, the organization's data will be hosted outside. This has the potential to create security problems in the following areas:

- **Process.** The systems used could be applying risky approaches that could be unacceptable in certain industries. Examples are the level of required testing, frequency of backups, etc.
- **Data.** The loss of control of data creates obvious security and privacy challenges.
- **Compliance/Legal Issues.** Certain industries require more rigorous IT processes than others. This could create new challenges in compliance and perhaps complex legal issues.
- **Extended Trust.** Users will have to trust providers. They actually have to extend this trust to multiple (chain of) providers. The issue of extended (chain of) trust (Chen et al., 2010) arises from the fact that users will have to use an application that is offered by a SaaS provider, who in turn runs it on a PaaS provider running on a VM provided as IaaS. The end user does not have much of a choice but to trust the security in this chain.

4.2 Multi-tenancy

Another main issue is the sharing of physical resources where multi tenants (customers who could even be competitors) are sharing the same physical machine or storage system. The following are some of the security issues that arise from multi-tenancy:

- **Data Deletion.** Due to the shared data storage resources, safe deletion problems are a concern.
- **VM Attacks.**. Compromising one VM could lead to the compromise of other VMs on the same physical machine. An example of this is the work on VM-to-host escape presented by Kortchinsky (Kortchinsky, 2009) in Blackhat 2009.
- **Side/Covert Channel Attacks.** Another interesting attack arising from mutli-tenancy is published by Ristenpart et al. (Ristenpart et al., 2009). Side and covert channels could occur between two VMs existing on the same physical server. As demonstrated (Ristenpart et al., 2009) on EC2, it is feasible to co-locate an attacking VM on a specific server that hosts a target VM. Attacks could be mounted by measuring the cache load which could be used to create side channels to verify VM co-location on the same server, estimate traffic rates, launch keystroke timing attacks, and in general trace co-tenants' activities. Chen et al. (Chen et al., 2010) argue that *tenant activity* (not only data) is something that requires protection as it could reveal confidential information about the users.

4.3 Auditing and Forensics

- **Mutual Auditing.** Cloud computing security problems are harder to investigate. Accountability, third-party auditing, and reputation sharing issues will continue to be a major obstacle to its deployment. There is a need to create mutual auditing techniques (Chen et al., 2010) which is unique to this paradigm.
- **New Forensics Capabilities.** Tools are needed to trace problems and help with legal investigations. Problems could arise when one (or more) system(s) is involved in a legal problem; all users of the same system might be disrupted when this system is under investigation.

5 ATTACK SURFACES

Cloud computing will result in the introduction of new attack surfaces and will change other traditional ones. An attempted taxonomy of attack surfaces is presented by Gruschka and Jensen (Gruschka and Jensen, 2010). The taxonomy defines six attack surfaces among three different entities: the user, the service, and the cloud. We map the taxonomy to the three entities [2] described in the Berkeley report (Armbrust et al., 2010). Surfaces 1 and 2 are what normally happens in a client-server environment between the SaaS user and the Software provided as a service. Surface 3 (from the SaaS provider to the cloud infrastructure) represents all the attacks that an application can mount against the infrastructure (hypervisor, VMs, network). Surface 4 (a very critical one) includes the attacks a cloud provider can perform against its direct users (SaaS providers) which could be very harmful given the control that the infrastructure provider has over the VMs hosting the applications. Surfaces 5 and 6 represent indirect relationships between the SaaS user and the cloud infrastructure provider since there always exists a service provider (SaaS) in between the two entities.

In non-cloud environments, attackers benefit from the existence of a large number of unprotected home machines that are running a full scale OS with many unneeded services enabled. Will the cloud change this by increasing the number of notebooks running a

[2]Where user is the SaaS user, service is the SaaS provider, and cloud is the cloud infrastructure provider.

lightweight OS with minimal cloud access tools (e.g. a browser)? Could this reduce the total number of available exploits and the number of machines running them? Will this reduce the success chances of new worms? Will it be more difficult for botnets? Will this force attackers to move to the heavily fortified clouds where the providers invested appropriately in security (hopefully)? It is still to be seen whether it is going to easier or harder for attackers in these new environments.

6 EXTENDING PROTECTION BEYOND TRADITIONAL PERIMETER DEFENCE

The notion of "Perimeter Defence" has been a part of information security that generated a significant amount of discussion. An organization seeking to protect its information infrastructure will create layers of defence starting with protecting its "perimeter" which traditionally denotes what physically separates its information infrastructure from the outside. Tools for protecting the perimeter include firewalls, filtering capabilities at routers, and the use of a DMZ (demilitarized zone) among others. The effectiveness of perimeters has been challenged as boundaries around an organization information infrastructure has become blurred. This in part is due to that fact that many employees connect from outside and with technologies such as HTTP that could allow attackers to pass through firewalls undetected.

In the cloud, an organization will have little control over how and where its data and applications are hosted. The notion of perimeter defence as known today will be challenged even more. In this section, we ask the question: *does moving an organization's data and applications into the cloud eliminate its capability to keep its perimeter defence at all?*

At this point, the answer seems to be that the notion of perimeter defence will change dramatically in the cloud. An organization could extend its security protection to *complement* what the Cloud Provider (CP) offers and not rely solely on the CP's provided security. Even when the CP offers all kinds of promised security, an organization that is serious enough about its security may want to provide additional protection to what the CP provides to extend its protection beyond its traditional perimeter. This could mean extending protection [3] beyond its perimeter to include various parts of the cloud that will dynamically change both in location and in nature (software, platform, and infrastructure).

In all cases, it is expected that the CP will offer some security tools and mitigation services. But how can an organization extend its protection differently depending on the type of cloud service provider?

- **The Case of IaaS.** When asking the CP (IaaS provider such as Amazon or GoGrid) for a virtual machine, the CP is expected to secure the network, the host OS, the hypervisor and isolate VM instances. The CP could also provide the clients with tools to use and customize. For example, Amazon offers Multi-Factor Authentication, customizable firewall, secure APIs, and Key management tools (Amazon, 2010). The organization could extend its protection beyond its perimeter to all the virtual machines that it has control over in the cloud. Leveraging the organization's security tools and expertise could be done in a number of ways including securing the guest OS, ensuring compliance with industry-specific requirements, creating multiple layers of security around the VM, etc.

- **The Case of PaaS.** If the required service from the cloud is a platform where an organization deploys its applications, the CP (PaaS provider) in this case should protect the network and the platform. The extended perimeter is this case protects the application within this provided platform. An example of such an environment is Google APP Engine. The platform offers a protected sandbox to the application with some restrictions on what could be accessed. However, it offers customers the option to integrate with a third-party Single Sign-On provider (Google, 2011). Another example is Microsoft Azure which offers identity management, authentication, and other mitigation services. However, for some threats, the application has to implement its own mitigation (Marshall et al., 2010).

  Suppose that the application provider does not believe that the CP provides sufficient security. This begins to raise the question of where more protection can possibly come from. The cloud offers the ability to scale at limitless proportions; however, all the application images are exactly the same. Once an attacker succeeds with one deployed instance, it is a simple extension to create an automatic exploit to attack the rest.

- **The Case of SaaS.** This is the situation where the service sought from the cloud is solely an application such as moving the enterprise email to a

[3] For example: to maintain data privacy and integrity, protect applications from tampering, and even to conceal its application activities from competitors.

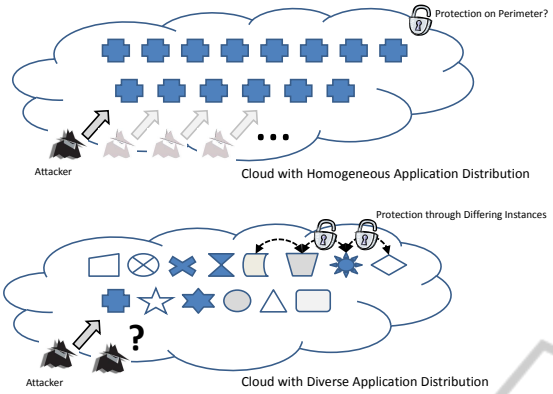

Figure 1: An Example of Protection by Diversity.

cloud-based email or managing human resources using a cloud-based service. In this case, the perimeter will extend to contain the data wherever it is hosted. This scenario gives the customer the lowest level of control. The importance of the data that resides in the cloud could range from incidental to competitor-critical to mission-critical. The possible attacks could range from automated searches to outright copying of data-base material. For data protection, homomorphic encryption (Gentry, 2009) has been suggested but there has been no practical solutions based on it yet. Related work on practical techniques could be in the area of white-box cryptography (Chow et al., 2002; Wyseur, 2009).

At this point, what are the possibilities for extra security beyond what the CP provides? Certainly, data encryption can play a part for mission critical information; however, is this overkill and potentially slow for other types of data? Can we consider a contrasting approach where data populations are encoded and varied, just enough to frustrate casual attacks and categorically defeat the automated attack. When diversity is applied to an application population in the cloud, it abates the threat that is accentuated by the ability to scale. Also in terms of forensics, diversity could help determine the location from which the source code had been lifted from the cloud. Refer to Fig. 1 for an illustration of diversity.

Any third party security provider (i.e., anyone other than the CP, regardless of what type of CP) in the cloud will have to deal with the problem of lack of control and/or access to the cloud. Furthermore, the possibility of rapid scaling and nearly limitless resources emphasizes this lack of control. Moreover, if a security problem arises, investigations will be limited to what the CP allows (Section 4.3). Given the nature of these three different environments discussed above, an IaaS environment provides its customers with more control over PaaS in terms of investigation capabilities. PaaS in turn provides more than SaaS.

7 PROTECTING INTELLECTUAL PROPERTY RIGHTS

It is expected that a great amount of digital content will eventually be placed in the cloud. There have been ongoing discussions on different ownership, jurisdiction, financial, and protection issues related to IPR in the cloud. It is expected the above issues will cause confusion and situations will arise where owners are not aware of where their content is hosted and what laws govern its access and protection. This section, however, focuses only on the technical issues related to IPR protection in the cloud. Reference to other issues could be found in (ENISA, 2009; CSA, 2009).

Content Exposure. The first question that could be asked about IPR is about exposure. Is the cloud increasing content exposure at the risk of increasing abuse? The answer depends on the type of content and the method of distribution. For a game that is traditionally distributed on a media (e.g., a CD) that the consumer buys and keeps, perhaps hosting this game in the cloud will reduce the end user access to the entire game. Reduced access to the entire game may allow cloud-based manufacturers an advantage to those that need to provide protection to a native application on a desktop. On the other hand, for contents that the owners traditionally offers from their own servers, moving to the cloud might increase its exposure. These servers will no longer be under the owners' control and this will increase chances of attacks on content where the servers are moved to the cloud.

Infringement Detection and Tracing. It is not clear if the cloud will bring any advantage to the area on IPR infringement detection and tracing. Due to the fact that content owners will have little control over the environment where content is hosted (or run in the case of software), owners will have limited abilities in detecting infringements. In other words, it might not be better than the case where the content was distributed on media. Even with the help of the CP, it may not be any easier since auditing and forensics capabilities for cloud environments are in a very early stage. Infringement detection might be done manually and using human expertise for some time before automated monitoring systems could be developed for the cloud.

8 SECURITY LIFECYCLE

In general, effective software security means more than an initial attack resistance. It also requires continual and active protection including: monitoring for attacks, breach-response, and mitigation. In the cloud, attention to the security lifecycle could mean facing some new challenges compared with running on a single computer. For example:

- Security lifecycle requirements may dictate the consideration of attacks that may be launched from the cloud environment itself. What kind of security lifetime strategy could prevent these types of attacks? For example, can updatable software take into account the shared nature of the underlying infrastructure?
- How is attack monitoring and detection of a compromise hindered by the limited control that applications owners are permitted by cloud providers? Will a CP be able to provide the necessary support for special security hooks for monitoring and detection?
- Is there any forethought for regular notification of the trustworthiness of applications running in the cloud? If so, who should do these validations of trust? The CP, the application itself, or both?
- Consider the detection of a breach. The remedial action might be limited since the problem may only be solvable with the CP's intervention. The application owner may only be able to stop the application or prevent access to certain users until the issue is resolved. This raises an interesting relationship issue between the application owner and the CP on renewability and restoring the integrity of the application.

9 CONCLUSIONS

In this paper, we attempted to raise questions on how cloud computing will affect the way security is offered. We discussed the cloud characteristics that create new challenges for security providers. Organizations should think of new ways to protect their assets and interests when they are deployed in the cloud. This protection may differ depending on the type of cloud environment (IasS, PaaS, or SaaS). Similarly, new ways of protecting IPR in the cloud are needed. While the cloud may bring some usability benefits where CPs (Cloud Providers) are expected to secure the cloud infrastructure, it also raises new questions where users and application providers have less control over how and where their data and programs are deployed. The future is uncertain at best.

REFERENCES

Amazon (2010). Amazon web services: Overview of security processes. *Amazon White Paper*.

Armbrust, M. et al. (2010). A view of cloud computing. *Communications of the ACM*.

Chen, Y. et al. (2010). What is new about cloud computing security? *Technical Rep. UCB/EECS-2010-5, Univ. of California at Berkeley*.

Chow, S. et al. (2002). White-box cryptography and an aes implementation. In *9th Annual Workshop on Selected Areas in Cryptography*.

CSA (2009). Security guidance for critical areas of focus in cloud computing v2.1. *Cloud Security Alliance*.

ENISA (2009). Cloud computing: Benefits, risks and recommendations for information security. *The European Network and Information Security Agency*.

Gentry, C. (2009). Fully homomorphic encryption using ideal lattices. In *Proc. of the 41st annual ACM symposium on Theory of computing*.

Google (2011). Saml single sign-on (sso) service for google apps. http://code.google.com/googleapps/domain/sso/saml_reference_implementation.html.

Gruschka, N. and Jensen, M. (2010). Attacksurfaces: A taxonomy for attacks on cloud services. In *IEEE 3rd International Conf. on Cloud Computing*.

Kortchinsky, K. (2009). Cloudburst, a vmware guest to host escape story. *A talk in BlackHat USA*.

Marshall, A. et al. (2010). Security best practices for developing windows azure applications. *Microsoft White Paper*.

Mell, P. and Grance, T. (2009). Effectively and securely using the cloud computing paradigm. In *ACM Cloud Computing Security Workshop*.

Molnar, D. and Schechter, S. (2010). Self hosting vs. cloud hosting: Accounting for the security impact of hosting in the cloud. In *Workshop on Economics of Information Security*.

Ristenpart, T. et al. (2009). Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. In *'Proc. of ACM CCS*.

Sun (2009). Introduction to cloud computing architecture. *Sun Microsystems White Paper*.

Wyseur, B. (2009). White-box cryptography. In *PhD thesis, Katholieke Universiteit Leuven, B. Preneel (promotor), 169+32 pages*.