Journal of Indian Research, Volume 1 No. 1, 2013, 91-97

Computer Science

ESTABLISHING TRUST IN CLOUD COMPUTING

Shweta Agrawal*

ABSTRACT

Cloud computing delivers IT capabilities as services-on-demand. This scalable and plastic model provides advantages like faster time-to-market, no capex and pay-per-use business model. While there are several such benefits, there are challenges in adopting public clouds because of dependency on infrastructure that is not completely controlled internally and rather shared with outsiders. Several enterprises, especially large ones that have already invested in their own infrastructure over the years are looking at setting up private clouds within their organizational boundaries to reap the benefits of cloud computing technologies leveraging such investments. This paper describes the different options available, highlighting the key advantages and challenges posed by each and the approach enterprises should be taking in adopting cloud computing with minimal risk. In this article, we discuss the need for asking critical questions about the security implications of cloud computing. Answers to our questions are not readily apparent, even though viewing computing as a utility, similar to that of providing water or electricity on a for-fee basis, dates back to at least the 1960s.

Keywords:

DOD	-	Department of Defence
DISA RACE	-	Defense Information Systems Agency's Rapid Access Computing Environment.
VoIP	-	Voice over Internet Protocol
DHS	-	Department of Homeland Security
CSA	-	Cloud Security Alliance

1. INTRODUCTION

In the aptly titled article, "Cloud Assurance Still Missing," Allan Carey wrote, "The security problems that organizations face related to cloud computing are the same as those related to virtualization—but even more so." [1] He goes on to say, "Information assurance practitioners already have most of what is needed to make an informed set of decisions about cloud computing." [2] We would argue that the security problems go well beyond the use of virtualization in distributed systems. In this article, we discuss the need for asking critical questions about the security implications of cloud computing. Answers to our questions are not readily apparent, even though viewing computing as a utility, similar to that of providing water or electricity on a for-fee basis, dates back to at least the 1960s. [3]

A recent technical report published by the University of California, Berkeley, states that there is no commonly agreed upon definition of cloud computing. [5] Instead, a definition is emerging as the various organizations that are developing cloud service evolve their offerings. In addition, there are many shades of cloud computing, each of which can be mapped into a multidimensional space with the dimensions being characteristics, service models, and deployment models. [6]

*Shweta Agrawal is Assistant Professor at Northern Indian Engineering College in New Delhi.
Email: Shweta.airen@gmail.com

Cloud computing is a metaphor for giving Internet users a growing collection of computer system resources and associated software architectures to provide application services. [7] The applications include processing and application integration, storage, and communications services. Cloud services are typically available on demand and are charged on a usage basis. Often, what the user sees is an application instead of a particular computer.

The services are commonly described as:

- **PaaS (Platform as a Service)**—the cloud provides hardware resources, typically virtual machines, which can be loaded with the users, operating system and software;
- **IaaS (Infrastructure as a Service)**—the cloud provides an infrastructure including (virtual) platforms, networking, etc. on which applications can be placed;
- **SaaS (Software as a Service)**—the cloud provides software applications.

Amazon's Elastic Compute Cloud (EC2) is an example of these services. [8] Google also provides enterprise-level integrated application services such as email, appointment calendars, text processing and spreadsheets. [9]

The claimed advantages for an enterprise are that it does not require an investment in computer resources, infrastructure, administration, etc.: the purveyor of the cloud provides these resources. The user or enterprise only pays for the resources "consumed." In the Department of Defense (DoD), we have seen the introduction of infrastructure services on demand provided by the Defense Information Systems Agency's Rapid Access Computing Environment (DISA RACE). [10] Where available, the cost of developing and maintaining specialized applications can be shared among the users of that application. In theory, there is an advantage in having large-scale resources shared among a large class of users.

However, this has yet to be borne out. [11] There are, of course, applications that require a large number of resources. Google Search is one such example. It appears that Google, Amazon, and others are attempting to leverage their ability to construct such a system into other environments.

We can argue that it is not a matter of whether cloud computing will become ubiquitous but rather what we can do to improve our ability to provide cloud computing users with assurance that the cloud services and infrastructure provide appropriate security functionality. Cloud computing providers should supply their customers with an appropriate level of security transparency to alleviate customers' reservations about the security and privacy afforded by the cloud. [12] How much transparency is enough? How do we provide for transparency of cloud resources (i.e. determining the cloud in which customer data resides)? Is there a tipping point at which additional levels of transparency would only serve to help malefactor's compromise services and data centers?

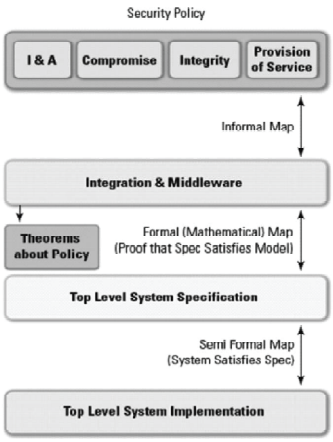

Fig : 1 Process for Integrating Security into the Cloud

In addition, as users and developers find new ways of applying cloud technologies, there will be new expectations about security and privacy. For instance, Twisted Pair Solutions of Seattle proposes to provide cloud computing resources for state and local agencies to link up disparate public safety radio systems (*e.g.,* police, fire, or ambulances)—a novel but difficult-to-predict usage of cloud computing, but also a usage that makes the cloud part of mission- and safety-critical systems. [13] The expectations for security, privacy, reliability, and quality of service and so on will be different in some respects for Voice over Internet Protocol (VoIP) radio systems than for the cloud's social networking aspects. This raises the question: how do we manage risk when we do not fully understand what we are trying to protect or guard against?

The fluid nature of cloud computing makes it a moving target, even when trying to determine the questions we should be asking to improve the security and privacy clouds afford. However, we can ask fundamental questions like: are the current architectures adequate for building trusted clouds? If not, what types of software system architectures do we need? Consider, for instance, the possibility that an organization might opt to fully outsource its computing infrastructure and data center to the cloud, retaining only thin clients within the organization. How do we make the thin client user terminals and the communications infrastructure secure?

2. DOD ENTERPRISE COMPUTING

What is our motivation for jumping feet first into asking hard questions about cloud computing? The growing importance of cloud computing makes it increasingly imperative that security, privacy, reliability, and safety communities grapple with the meaning of trust in the cloud and how the customer, provider, and society in general gain that trust. Consider the initiative of the DoD Enterprise Services & Integration Directorate to make the DoD Storefront Project a reality. The Storefront consists of a cloud-based set of core and specialized applications that users can discover through an application marketplace and which share an identity management framework. How will DoD provide security for the Storefront? It is more than a matter of having an identity management framework. The obvious security concerns include data integrity, data availability, and protection of personally identifiable information, data protection, data destruction, and communications security.

Moving beyond the Storefront concept, as the federal government migrates its data and applications to the cloud, issues regarding cross-domain resource sharing will arise within the cloud. For instance, how will DoD link its clouds to those of other agencies? Will a DoD user, authenticated to enter the DoD cloud sphere, be trusted to access services owned by the Department of Homeland Security (DHS)? Is there a need for a federal-wide cloud infrastructure and common set of security services? How will data be shared among the various different types of cloud?

3. INFORMATION ASSURANCE

At the Naval Postgraduate School, a major thrust of research on cloud computing is to investigate the security policies, models, and appropriate architectures to provide security for entities/users of cloud computing resources. Although cloud computing may appear to provide reasonably well understood operating system and application resources, cloud resources are distributed in space, time, and scale in ways that were never envisioned in the operating-system world. The current architectural approaches, especially those concerning security, may not scale to the much larger cloud computing approaches. In addition, the approaches for assuring operating system security functionality are not necessarily appropriate. It is unclear whether the current set of services is sufficiently secure and reliable for use in sensitive government environments. Current security claims are somewhat limited.

One of the fundamental problems with adopting cloud computing is providing not only security resources but also assurances that those resources are correctly implemented and maintained within the cloud. Several vendors have formed the Cloud Security Alliance (CSA). [14]

In the report titled *Security Guidance for Critical Areas of Focus in Cloud Computing V2.1,* CSA provides its take on some of the security issues related to cloud computing. [15] In the report, security properties are described as essentially the same set of properties that a user expects to see with a self-hosted system. These include the usual:

- Identification
- Authentication
- Privacy
- Integrity
- Provision of Service.

They view assurance as an audit of the function's implementation, that is, the cloud systems' administrators and implementers have used 'best practices'. Other than the notion that encryption is used to protect the data, there is little information that defines 'best practices.' There is, however, some form of key management included that provides potentially strong identification/authentication, as well as some form of data integrity/ recovery facility.

The security architecture proposed is essentially a layered operating system application. It consists of a network layer interposed between application programming interfaces (APIs) and the underlying operating system infrastructures. 'Trusted computing' is only mentioned at the hardware/operating system level. Additionally, the CSA paper enumerates several security issues that should be addressed by the cloud-style service provider, but does not provide any insight on security policies/models, interfaces or potential solutions.

To provide an example of some of the potential issues, Google supports "Google Apps." [16] Google Apps applies the usual discretionary access controls to the resources it provides – files, calendars, address lists, *etc.* To make life easier, Google provides tools that integrate their identification and authentication systems into the enterprise providing single sign-on; the enterprise user need only log onto their home system. Once logged on, the enterprise user can automatically access the users' files and services on Google without an additional login. Although convenient, this functionality increases the security exposure to not only the weakness of the enterprise system, but also to the weakness of Google's infrastructure. If, for example, Google's infrastructure has a security flaw, then it may be possible for someone in one enterprise to access accounts from another enterprise.

On the other hand, security flaws in the enterprise system may lead to weaknesses in the access controls of the information managed by Google Apps. Additionally, connected applications may provide unintended connections among users, as was demonstrated with the introduction of Google Buzz. [17]

When each enterprise maintains its own infrastructure, a failure in one enterprise may cause failures across the cloud. Unless an enterprise uses a single cloud from a single vendor, integrating the various applications, infrastructures, and policies among many different clouds and cloud vendors will be a significant challenge. In fact, it will be a challenge to ensure that the different policies do not contradict and potentially permit access that should not be allowed at the system level.

Ultimately, the proof is in the pudding. Will the cloud vendors be willing to stand behind the security of their systems? In the case of Amazon's EC2 and Simple Storage Services (S3) services, Amazon suggests that their EC2 and S3 infrastructure not be used for systems that must satisfy the Payment Card Industry Security Standards [18], although it has published a paper on how Amazon Web Services can be used in a Health Insurance Portability and Accountability Act (HIPAA) compliant environment. [19]

In the HIPAA paper, Amazon essentially places almost all the requirements on the "user/enterprise"to encrypt all the data stored and to manage its keys. Amazon provides services to log safely into its systems and provide some data recovery and integrity.

In the realm of reliability, prior to the breakup, AT&T was required to build systems that had an up-time reliability of "five nines" (about 5.2 min/yr downtime). Part of the reason for this was to ensure services in case of national emergency. Current cloud based systems are advertised as providing "three nines" (almost 9 hrs/yr downtime). [20]

4. DETERMINING WHERE TRUST SHOULD BE PLACED

Clearly, there are many challenging security issues related to cloud computing. In our research, we are working on a formal, structured, possibly mathematical approach that will give users and cloud-developers deeper insight into what should be done, how it might be achieved, and where the trust should be placed. This research includes the investigation of implementation structures and assurance provisions for "security" in cloud-based systems. To do this, we will attempt to provide security architectures and models that satisfy the following: Awareness of the amorphous nature and scale of the cloud computing paradigm, inclusion of mathematical models of the security properties that can be used to help analyze those properties, provision of the underpinnings on which applications/enterprise/user level security policies/properties can be implemented. and provision of the foundations on which the implementation assurances can be ascertained.

Our hope is that the results of the research will provide a framework that can be at least partially applied to the current cloud architectures and may lead to new architectures with better defined, more assured security.

Over the past 30-plus years in the operating system security world, a lot of work has been done to provide highly assured components with trustworthy systems. Unfortunately, the commercial world has ignored a lot of this work. Recent efforts have focused on the use of separation kernels. For example, Green Hills has recently received a National Information Assurance Partnership (NIAP) certificate for its Integrity 178B Separation Kernel. [21] Separation kernels provide a minimal set of operating system services on which other trusted services and applications could be built. These may be thought of as slightly more functional than a Virtual Machine Monitor (VMM), although Green Hills and others are looking to implement high assurance VMMs using their technology.

Our approach to the problem involves separation of 'virtual' resources. This approach constructs an infrastructure that establishes (or reconstructs where appropriate) resources, identifies and authenticates users, and then controls access to the resources. Our focus is to provide a model and a security architecture that provides the infrastructure that will accomplish these goals.

4.1 An Example

For instance, consider PaaS. An enterprise might wish to run its own applications. These applications may only run on an intermittent basis and/or require a large number of resources. One way to achieve this is to use a cloud PaaS.

We use the term 'enterprise' to describe the organization requiring the platform and 'provider' for the organization providing the cloud platform resources. The PaaS provider would provide 'platforms,' either 'real' as part of a virtual environment (a means for downloading an operating system and for managing the platforms), or as a possible network interface(s) on the platform(s). The enterprise loads operating systems, applications, etc., onto the platform(s) and manages all the interfaces and resources provided. The example below assumes that multiple platforms will be used.

The security policy visible to the user includes:

- **Identification**—A set of platform names issued by the provider (unique to the enterprise)
- **Authentication**—A secure channel that can be used to load the operating system(s) onto the platforms—the provider is trusted to ensure that the only communication with the platforms is from or to the enterprise.

- **Integrity**— The provider should guarantee that the resources are "empty" on first use and that none of the platform resources are modifiable by any party other than the enterprise. This includes any management functions; it is up to the enterprise to ensure that any network interfaces are appropriately protected.
- **Privacy**—The provider should guarantee that there is no third party access to the platform processor, memory, and/or disk files.
- **Provision of Service**—The provider should provide access to the resources on demand, per any service level agreements between the enterprise and the provider.

4.2 With at least two models of this kind of service

1. Resources are provided on an ad hoc, intermittent basis. In this version, there is no connection between consecutive uses of the resources. The enterprise uses the resources once. During subsequent uses, the enterprise assumes that all the previous data does not exist or has been erased by the provider. The only connection between the two usages is that the enterprise uses the "same identifiers" to access new instances of the resources. There is no guarantee that the same physical resources will be used for each run of the platform(s).
2. The enterprise 'turns off' the platform, but in subsequent use after turning it back on, finds the platform resources in the same state they were in after being turned off. As expected, the enterprise might pay more for this service. In this case, the provider must protect the information in the resources between runs from both modification and access by third parties. There is no guarantee that the same physical resources will be used in each run of the platform.

   (Note that in both cases, the provider provides access to platforms and associated data. The platforms are available to others when the enterprise is not using them. Any provider configuration data about the platforms must be protected from modification and, in the second case above, any enterprise information that will be reused must also be protected)

Informally, a portion of the model might then take the form of:

- **VPlatform**—The set of names of virtual platforms that will be provided to enterprises
- **VPlatformType**—Whether the VPlatform resources are persistent (type 2 above) or not
- **VPlatformResource**—The set of resources associated with a VPlatform
- **Enterprise**—The set of enterprises that use VPlatforms
- **Allocation**—An association of an Enterprise with a Platform, VPlatformType and VPlatformResources. The same Enterprise may have multiple VPlatforms, and VPlatformResources associated with it
- **PlatformCloud**—A sequence of sets of Allocations.

The security properties then become statements about the resources and platforms.

For example:

No pair of allocations shares any common VPlatforms or VPlatform Resources. As depicted in Figure 1, the security properties can be modeled on a collection of the statements above. Each of the statements should map back to some aspect of the system's user-visible security property. We could use our statements about the relationships of the entities (sets) we describe to prove additional properties of the system.

Following the security model's construction, a high-level execution model should be constructed and validated mathematically to determine that it satisfies our security model. Next, it is necessary to map our high-level model to varied cloud aspect implementations as documented by the vendors.

5. CONCLUSION

Cloud security is an ill-defined, little-understood area of distributed computing. However, we believe that progress can be made to provide a level of assurance that accommodates the resources needed to support government's information processing requirements.

6. References

1. IA Newsletter, vol. 13, no. 1, winter 2010, p. 34.
2. Ibid.
3. M. Campbell-Kelly. "The Rise, Fall, and Resurrection of Software as a Service: A Look at the Volatile History of Remote Computing and Online Software," Communications of the ACM, vol. 52, no. 5, pp. 28–30, May 2009.
4. B. Michael. "In Clouds Shall We Trust," IEEE Security & Privacy, vol. 7, no. 5, p. 3, September/October 2009.
5. M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. H. Katz, A. Konwinski, G. Lee, D. A. Patterson, A. Rabkin, I. Stoica, and M. Zaharia. "Above the Clouds: A Berkeley View of Cloud Computing, " EECS Department California, Berkeley. Technical University.
6. P. Mell and T. Grance, "The NIST Definition of Cloud Computing," Version 15, 7 October 2009, http://csrc.nist.gov/groups/SNS/cloud-computing/cloud-def-v15.doc.
7. http://en.wikipedia.org/wiki/Cloud_computing.
8. http://aws.amazon.com.
9. http://docs.google.com.
10. http://www.disa.mil/race.
11. H. G. Miller and J. Veiga. "Cloud Computing: Will Commodity Services Benefit Users Long Term? IEEE ITPro, vol. 11, no. 6, p. 67-69, November/December 2009.
12. http://www.opencloudmanifesto.org.
13. http://www.fcw.com/Articles/2009/04/16/Cloud-computing-moving-into-public-safety-realm.aspx.
14. http://www.cloudsecurityalliance.org.
15. http://www.cloudsecurityalliance.org/csaguide.pdf.
16. http://www.google.com/apps.
17. http://www.nytimes.com/2010/02/15/technology/internet/15google.html.
18. http://www.mckeay.net/2009/08/14/cannot-achieve-pci-compliance-with-amazon-ec2s3.
19. http://awsmedia.s3.amazonaws.com/AWS_HIPAA_Whitepaper_Final.pdf.
20. http://www.google.com/apps/intl/en/business/infrastructure_security.html.
21. http://www.niap-ccevs.org/cc-scheme/st/vid10119/maint200