Security concerns in commercial cloud computing

Jimenez, Jaime
Aalto University School of Science and Technology
jjimenez@cc.hut.fi

Abstract

Companies benefit from cloud computing services in terms of cost reduction and processes enhancement: enterprises are able to scale their systems easily and deploy new technologies more efficiently. Indeed, cloud computing allows businesses to become more flexible and to rapidly react to changes in the market, while focusing fully on their core businesses.

In this paper we focus on the new challenges that arise when companies partially or fully migrate their IT infrastructure into a third party cloud company. Some of these challenges are confidentiality and integrity, as well as availability of the data. Other challenges are addressed from a legal and organizational perspective.

In order for cloud systems to become more than a hype we state some solutions to paliate few of these problems. The focus is aimed at showing that losing physical control over data represents the main risk for users when migrating to the cloud. Thus, trust in cloud computing systems has to be built in order for them to be widely adopted.

1 Introduction

A 2009 survey [4] of more than 500 IT executives across 17 countries showed that "they trust existing internal systems over cloud-based systems due to fear about security threats and loss of control of data and systems", that "their current internal systems are too expensive" and that "they are increasing their investments in this technology (i.e. cloud computing)". Therefore, the main driver for the introduction of cloud computing in companies is cost-efficiency - both IT up-front costs and ongoing IT costs are lower. Security does not seem to be the driver for migrating resources to a third party cloud computing.

Although there are no major technological improvements with cloud computing, since this paradigm relies heavily on existing technologies, there are existing threats such as viruses, phising (masquerading as a trustworthy entity to steal sensitive information) or DDOS attacks that are magnified when centralization occurs [12]. Research has highlighted several dangers that might arise from cloud environment: some examples are the Distributed Denial of Service (DDOS) attack in May 2007 at the Parliament of Estonia, several of its banks and ministries [20]; the phising scam carried out in Salesforce [1] in 2007 [3]; the virus that in November 2008 affected thousands of military computers in the USA [20] or the Aurora attack on Google in 2010 [1].

According to Juniper Research, the number of Mobile Connected M2M and Embedded Devices will rise to almost 412 million globally by 2014 with several distinct markets, cloud computing technologies are believed to become the biggest rising industry in this ICT framework [18]. The current year (2010) is deemed as the year in which company-oriented cloud services will emerge [20], promising more security and data availability than current user grade systems provide. The boundaries between internal networks and external networks are transforming control of data poses new interesting challenges and its ultimate feasibility is not yet guaranteed.

2 Background

2.1 Definition of cloud computing

One of the several definitions of cloud computing [7] is the one made by the National Institute of Standards and Technology (NIST) [2]:

> Cloud computing is a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction. This cloud model promotes availability and is composed of five essential characteristics, three service models, and four deployment models.

The NIST also defines four different deployment scenarios. If the cloud infrastructure is operated by one organization, either managed by the organization itself or by a third party company then it is considered private cloud. If several organizations share the cloud infrastructure among a community for a specific purpose it is considered community cloud. In the case in which a company owns the cloud infrastructure and sells its services, we would be talking about public cloud. Last, a composition of two or more of the previous models linked together by standardized or proprietary technology is considered hybrid cloud.

For this paper we limit the scope to the third deployment scenario, in which cloud computing is a service offered by a

[1] Salesforce is the largest cloud based service provider, http://www.salesforce.com/

[2] http://www.nist.gov/

company over the Internet, usually in a pay-per-use manner although it can also be offered for free with other revenue streams. The services offered can be anything from web-based word processors or email clients, application development platforms like Google AppEngine or full virtualization like the offered by Amazon EC2 or Microsoft Azure. We chose the public cloud scenario since we are focusing on the possible dangers of cloud computing when the resources are managed by a third party and sold to other organizations or parties over the Internet.

2.2 Motivations

One of the main drivers, if not the most important, of the development of cloud computing services by the industry is that such development can help keeping costs down while gaining access to a wide range of software and IT expertise. The cost saving comes from the property of cloud computing to allow scalability and the deployment of services on demand without the problems and the costs that arise from provisioning a data center, not only hardware but also personnel, since the cloud computing provider is the one provisioning it [14]. The software and IT expertise comes from the cloud provider too, as the installation and use of the software is simplified and centralized, also enabling to share the data and store it in the cloud infrastructure [17].

On the other hand, one of the biggest deterrents for the introduction of cloud computing is that it is often perceived as unsafe and difficult to control; indeed, enterprises lose control over the infrastructure, which is now provided by the cloud, and switch to per-use service-based models. More importantly, they can easily lose control over their own data once it is in the cloud [21], but we will see more about this topic in the following sections.

2.3 Service Models

It is important to consider which is the right resource model for a cloud-based application. The model to use depends on the services and how effectively they can be provided by a cloud provider. In this sense it is considered that cloud computing has at least three service models [21, 8, 2] (See Fig. 1):

Platform (PaaS): Consist on the delivery of a platform or solution stack as a service, without the user having to manage nor control the infrastructure including network, servers, operating systems, or storage. Nevertheless, the user has control over the applications delivered with this service. Some implementations of this are Microsoft Azure [3] or Google's AppEngine [4].

Infrastructure (IaaS): Usually regarded as platform virtualization environment as a service, enabling the consumer to run arbitrary software, from operating systems to applications. It eases internal processes related with IT management, such as installing OS and configuring servers or other network devices. Well-known solutions are VMWare [5], or Amazon EC2 [6].

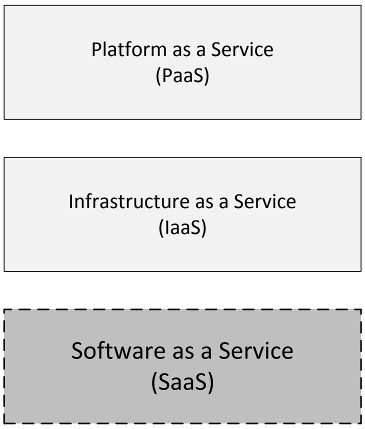

Figure 1: Service Models for cloud computing

Software (SaaS): Consists of computer software products specifically designed for the delivery of cloud services, the applications can be accessed via simple interfaces such as a web browser. The user does not manage any of the underlying infrastructure. Among the large ammount of common examples we can take Dropbox[7], Gmail[8] or even Wordpress [9]. It is this scenario in which we will focus since we believe that SaaS is the service that will be more appealing to end users with less technical background, thus being the one most likely to be widespread.

3 Current perimeter security

In order to compare the current security delivered in public cloud computing, we believe it is important to know how this security is delivered in current systems. Thus, it is out of the scope of this paper go deep into detail on network security matters, but to limit it to some of the measures that need to be explained in order to get a broader perspective and understand their implications in cloud computing (See Fig. 2).

One usual way to approach network security is to create different **access permissions** for different user types. Users should access content or services in the network or in an outside network depending on the responsibilities they have. This prevents from tampering with the inner systems and from abuse from malicious users inside the trusted network. Of course, this implies the necessity for an administrator role. The **administrator** is a set of experts with access privileges that configure and monitor the network, assign privileges and make sure that data and network are safe.

Another common method is the use of **firewalls**. Firewalls have three main design goals: to filter traffic that goes inside

[3] Windows Azure Platform, http://www.microsoft.com/windowsazure
[4] Google AppEngine, https://appengine.google.com
[5] VMware virtualization, http://www.vmware.gov
[6] Amazon elastic Compute Cloud, http://aws.amazon.com/ec2
[7] Dropbox for storage and file synchronization, www.dropbox.com
[8] http://mail.google.com
[9] http://wordpress.org

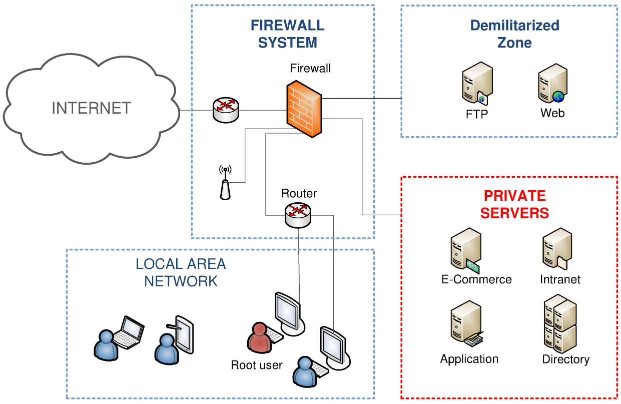

Figure 2: General components in perimeter network security

the network, to allow only authorized traffic and to be immune to penetration [16]. Firewalls create a barrier between the dangerous outside and the relatively secure inside of the network, they represent the first and most important obstacle for attackers and are the main component of any security perimeter [6].

There are several types of firewalls depending on what kind of control they are focusing on. There are firewalls focusing on services, i.e. they filter packets and monitor different applications; other firewalls monitor IP ranges and block or allow only certain trusted IP addresses; others allow only determined users that have the right set of privileges assigned by the local administrator. [6].

Intrusion Detection Systems (IDS) also filter traffic, this time based on user behaviour. The IDS monitors users and adapts to their common use of resources, any suspicious activity is then logged and reported to the administrator. [6].

The main reason of a perimeter is to separate two regions, the outside (usually the Internet) and the inside, which would compress the local servers and workstations. A **Local Area Network (LAN)** is usually created within this perimeter and the firewall is usually located either between the LAN and the external network, at the network gateways that connect the internal network with the external or commonly at the end user machines. Packet-filtering firewall behaviour is modified by setting different policies that imply two possible actions, either blocking or forwarding the packets. This causes common known weaknesses such as application-specific attacks, or spoofing a legitimate user's identity. Statefull packet filters try to limit this threat by maintaining a record of all connections passing through, thus being able to determine whether a connection is part of an old one [6].

Within a perimeter it is common to place an **Application Level Gateway** between two firewalls, this configuration is commonly named **Demilitarized Zone (DMZ)**. The main purpose of a DMZ is to create a physical or logical subnetwork within the networks' network, that contains and exposes the organization's external services to the Internet. This adds an additional layer of security to the organization's Local Area Network (LAN); an external attacker only has access to equipment in the DMZ, rather than any other part of the network.

When a workstation connects to a server from the Internet it needs a secure channel to do so, **Secure Shell (SSH)** is the protocol used to secure this channel since it provides confidentiality and integrity via encryption. A **Virtual Private Network (VPN)** can be set in order to secure remote connections to a network or inter-networks.

Sometimes it is common to place **proxies** within the network as intermediaries. This is done not only for speeding up the network by means of caching, but also to monitor and control users' access in a similar way as firewalls do and to scan outbound content to prevent data leakage.

These measures are generally implemented at a software level, although some can also be implemented on hardware, as in the case of a firewall. In order to ensure data security certain **physical** measures have to be implemented as well. It is out of the scope of this paper to analyze physical vulnerabilities, but physical control of the servers that store data and physical precautions regarding whom can access those servers is part of security measures. Moreover, as it will be shown, the lack of physical control of data due to offshoring poses one of the major problems to cloud computing [11].

4 Cloud Perimeter: Is there any?

4.1 Cloud security

When considering perimeter protection in the public cloud, it is important to differentiate between two parties: the provider and the customer.

The provider's network security, which is ensured by a relatively similar perimeter to the one described in section 3. Depending on the cloud service provider, this perimeter security will be tighter or looser, but it is safe to say that cloud computing providers will generally follow stricter security policies. For instance user authentication is carried out within the firewall [15] instead of at an inner server.

Customer security is guaranted by the use of Secure Socket Layer (SSL), since we mention earlier that the connection to the cloud is carried out in a light application like a web browser. SSL is a cryptographic protocol standard that aims at providing security for communications carried out over the Internet. Confidentiality and integrity are ensured by the use of symmetric key cryptography and a keyed message authentication code for the message ensures authentication.

Nevertheless one factor that is extremely difficult to address in public cloud computing is that the user loses control over the data and over the processes that are carried out within its own organization, heavily relying on a third party company that will manage its own data [9]. A result of the centralized storage of information and monopolized management of security is the cost reduction but at the expense of making the whole structure more vulnerable to attacks, maximizing the danger of data loss or theft. Therefore, it is fundamental for public cloud computing to develop forms of gaining the trust and confidence of the potential users and addressing these vulnerabilities [19] (See Figure 3).

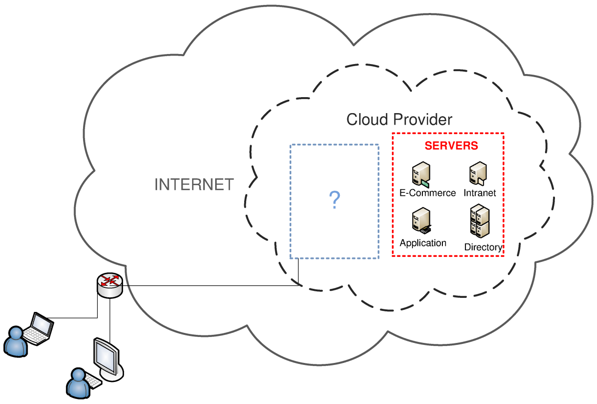

Figure 3: Cloud computing conceptual diagram

A risk assessment of different frameworks is mandatory in order to identify the threats of cloud computing. In this section we mention only the most probable and critical threats and obstacles, since analyzing all vulnerabilities of cloud computing is beyond the scope of this paper.

4.2 Organizational Framework

The main risks within this framework are: data lock-in, loss of governance and compliance problems.

1. Lock-in: there have been very few improvements in interoperability between different cloud service providers, there is no current standardization being actively carried out. Therefore, it is difficult to migrate from one provider to another, since providers have incentives to prevent portability of customers and data [13]. Also, moreover, cloud computing introduces the possibility of price speculation among providers, lack of reliability and even failure to provide the services in the case a provider goes out of business. **API standardization** could help to paliate the problem so that a SaaS can deploy services across various cloud computing providers [8].

2. Loss of Governance, specially in IaaS: when using cloud infrastructures, the client must cede control over security. Conflicts between the security measures taken by client and provider may arise, making it impossible to comply with different security standards.

3. Compliance Problems: different organizations comply with different, often incompatible Certifications; moreover, due to the very nature of the cloud itself, it may not be possible to audit or enable specific certification procedures.

4.3 Legal Framework

The main legal risk is lack of juridical transparency: the data is often stored in multiple jurisdictions and clients rarely are able to trace it. Some of the locations might be high-risk countries with different legal frameworks.

4.4 Technical Framework

Risks within the technical framework consist of: availability of service, performance unpredictability, data transfer bottlenecks, isolation failure, malicious user or administrator, compromised client interface or machine, data deletion and monitoring challenges.

1. Availability of a Service: in a cloud computing world, services need to have an adequate availability. Some cloud services have proven to be extremely reliable, take for example Google Search; it is likely that users will expect a similar availability from other similar services. Within this threat we have to include the vulnerability to Distributed Denial of Service (DDoS) attacks, that are likely to be more effective in some cloud providers, even though there is resource replication the attacks are aimed at just one larger target, and in case of success the consequences are disastrous. In addition to that, the cloud computing community has followed the motto of "no single source of failure" yet it remains unclear how the management of a cloud computing service by a single company does not enter in that definition. Single management introduces a single point of failure that can be addressed by using **hybrid cloud** systems or follow a **community cloud** approach.

2. Performance Unpredictability: while Virtual Machines can share CPUs and memory quite well [8] I/O sharing is more difficult since there are certain interferences between virtual machines. This problem in particular might be easier to solve, since wide introduction of **flash memory** might decrease I/O interference, but it is then a matter of time to see the outcome.

3. Data transfer bottlenecks: network congestion, misconnection and non-optimal use of network. Several of these threats arise from misconfiguration and OS vulnerabilities, since the resources are not isolated. Network breaks, however unlikely in this case, would have detrimental consequences affecting thousands of customers at the same time. Therefore it is important that cloud users and cloud providers understand the implications of **placement and traffic** at every level of the system if they want to minimize costs and attack risks.

4. Isolation Failure: in a cloud computing network, resources, storage and computing capacity are shared among clients, thus the so-called guest-hopping attacks or attacks directed to multiple users are more likely to succeed in a cloud environment, specially in public clouds.

5. Malicious Insider: a malicious user or administrator within the cloud provider is one of the most obvious and dangerous risk, especially if the different roles - access privileges and responsibilities - are not well defined nor enforced. Proper **role definition** is fundamental in cloud computing.

6. Compromised Interface: the client of a cloud Service usually connects to the service via an interface on his host machine, usually the web browser, this interface mediates in to access the remote resources. The client can access a larger set of resources than traditional hosting providers, thus the risk is increased.

7. Data deletion: if a client decides to change its provider, the physical resources are reallocated but in many cases the data might be available to the user or the provider for longer than expected by the customer. Full data deletion might not be possible in the cloud, since the hardware (hard drives) are shared by several different clients and full disk deletion might not be possible. In adition to that, cloud companies do have incentives not to delete that data.

8. Monitoring: as a Russian proverb says: "Trust, but verify". Customers demand more information and new security monitoring interfaces, that will enable them to verify the security checks in place [10]. A great challenge for the providers will be to provide the tools to re-empower the clients and allow them to monitor the security of their data.

5 Conclusion

In this paper, the notion of perimeter has been discussed and as it has been pointed out, this notion disappears when considering cloud computing. It could be argued that the perimeter disappears as soon as the hardware and the end devices are no longer under control of the organization. This loss of control does not necessarily imply a loss of security, provided that certain security measures are implemented.

Current methods in cloud computing do not efficiently tackle the main problems and solutions for some of them have been proposed. A standardization process should be arranged, for the sake of a common API to prevent user lock-in [5], the use of flash memories is recommended to decrease I/O interference and defining precise roles within the cloud organization is paradigm.

How to solve some of the threats still remains unclear, this is the case of the lack of regulation and the lack of a common legal framework that can address this issues. In the case of loss of governance it is fundamental to ensure the protection of the client, some mechanism to regulate data control needs to be developed. Yet again, cross-organizational alliances and Open Standards are needed to ameliorate the wealth of different certifications that are bound to appear, due to individual approaches taking by the various players.

It is the task of cloud clients to to carefully consider the benefits that public cloud computing can bring to their organization, and decide wether the threats outweight the benefits. In order to avoid the trade-off between security and costs, alternatives like the community cloud are a possible solution.

More research and active testing is required since a major attack on one or several cloud computing providers where to occur, the damage for its customers would be fatal not only for the potential loss of data but because their activity would be impaired. Moreover, the image of the cloud provider would be severely damaged and the confidence on this type of services lost.

References

[1] Aurora attack on google. Cited: 09.10.2010, Available: http://www.wired.com/threatlevel/2010/01/operation-aurora.

[2] Benefits, risks and recommendations for information security. Cited: 09.10.2010, Available: http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment/at_download/fullReport.

[3] Salesforce.com security beached. Cited: 09.10.2010, Available: http://www.erpblogger.com/salesforce-hacked.htm.

[4] Survey on cloud computing. Cited: 15.10.2010, Available: http://www.circleid.com/posts/20090226_cloud_computing_hype_security.

[5] Survey on cloud computing. Cited: 10.10.2010, Available: http://www.opencloudmanifesto.org.

[6] *Network Security: Private Communication in a Public World, Second Edition*. In computer networking and distributed systems. Prentice Hall PTR, April 2002.

[7] Experts Define Cloud COmputing. In *SYS-CON Media, Inc., http://virtualization.sys-con.com/node/612375*, January 2009.

[8] 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. Technical Report UCB/EECS-2009-28, EECS Department, University of California, Berkeley, Feb 2009.

[9] R. Buyya, C. S. Yeo, and S. Venugopal. Market-Oriented Cloud Computing: Vision, Hype, and Reality for Delivering IT Services as Computing Utilities. Aug 2008.

[10] S. Charney. Establishing end to end trust. Cited: 09.10.2010, Available: http://download.microsoft.com/download/7/2/3/723a663c-652a-47ef-a2f5-91842417cab6/Establishing_End_to_End_Trust.pdf.

[11] R. Chow, P. Golle, M. Jakobsson, E. Shi, J. Staddon, R. Masuoka, and J. Molina. Controlling data in the cloud: outsourcing computation without outsourcing control. In *CCSW '09: Proceedings of the 2009 ACM workshop on Cloud computing security*, pages 85–90, New York, NY, USA, November 2009. ACM.

[12] D. E. Comer. *Security Guidance for Critical Areas of Focus in Cloud Computing*. Cloud Security Alliance, 2009.

[13] S. Creese, P. Hopkins, S. Pearson, and Y. Shen. Data protection-aware design for cloud services. In *CloudCom '09: Proceedings of the 1st International Conference on Cloud Computing*, pages 119–130, Berlin, Heidelberg, 2009. Springer-Verlag.

[14] I. Foster, Y. Zhao, I. Raicu, and S. Lu. Cloud computing and grid computing 360-degree compared. In *2008 Grid Computing Environments Workshop*. IEEE, 2008.

[15] A. Gopalakrishnan. Cloud Computing Identity Management. In *SETLabs Briefings*, volume 7, July 2009.

[16] M. Gouda and A. Liu. Structured firewall designâŸĘ. volume 51, pages 1106–1120, March 2007.

[17] C. Hoff. Cloud: Security doesnt matter. , January. Cited: 09.11.2010, Available: http://www.rationalsurvivability.com/blog/?p=1694.

[18] W. Holden. More than 130 million enterprise users in the mobile cloud by 2014. March 2010.

[19] Y. Khmelevsky and V. Voytenko. Cloud computing infrastructure prototype for university education and research. In *WCCCE '10: Proceedings of the 15th Western Canadian Conference on Computing Education*, pages 1–5, New York, NY, USA, 2010. ACM.

[20] S. Mullen. Cloud computing considerations and insight. Cited: 09.10.2010, Available: https://www.opengroup.org/conference-live/uploads/40/22077/.

[21] T. G. Peter Mell. Effectively and securely using the cloud computing paradigm, 2009. http://csrc.nist.gov/groups/SNS/cloud-computing/cloud-computing-v26.ppt.