IN3 Working Paper Series
INTERNET INTERDISCIPLINARY INSTITUTE
ISSN 2013-8644

# Analysis of Security of Cloud Systems

## Comparative Analysis of Security features between Amazon Web Services and Windows Azure

**Ana M. Juan Ferrer** (ajuanf@uoc.edu)
*UOC*

Working Paper

Working Paper Series WP00-000

Research group: Security and Electronic commerce
Research group coordinator: Magdalena Payeras Capellà, Josep Lluís Ferrer Gomila (UIB)

Submitted in: June 2013
Accepted in: month yyyy
Published in: month yyyy

Universitat Oberta de Catalunya
IN3 INTERNET INTERDISCIPLINARY INSTITUTE
http://in3.uoc.edu


**Internet Interdisciplinary Institute (IN3)**
http://www.in3.uoc.edu
Edifici MediaTIC
c/ Roc Boronat, 117
08018 Barcelona
Espanya
Tel. 93 4505200

**Universitat Oberta de Catalunya (UOC)**
http://www.uoc.edu/
Av. Tibidabo, 39-43
08035 Barcelona
Espanya
Tel. 93 253 23 00

# Table of contents

# Analysis of Security of Cloud Systems

## Comparative Analysis of Security features between Amazon Web Services and Windows Azure

Ana M. Juan Ferrer (ajuanf@uoc.edu)
*UOC*

**Recommended citation:**

JUAN FERRER, Ana (2013). "Analysis of Security of Cloud Systems: Comparative Analysis of Security features between Amazon Web Services and Windows Azure [online working paper]. (Working Paper Series; WP00-000). IN3 Working Paper Series. IN3 (UOC). [Accessed: dd/mm/yy].<url>

**Abstract**

Analysis of current research on Cloud Security demonstrates that main source of security concerns related to cloud link directly to its inherent multi-tenant nature. However, multi-tenancy is the essential enabler to achieve commercial benefit for Cloud providers, through economies of scale, which, by means of exploiting this characteristic together with efficient management of resource utilization and overprovision techniques, can provide services at economic prices. Given the diverse range of services offered as "Cloud" today, in order to assess its security it is essential to concentrate at least on the specific cloud layer and, preferable on the specific service intended to be used, given that implementation of services can strongly differ among different providers. This paper concentrates in doing so: by providing a extensive literature review of Cloud security in general, and per Cloud layer, as well as, a comparative analysis of security features provided by two of the most representative Cloud providers: Amazon Web Services and Microsoft Azure.

**Keywords**

Cloud Computing Security, SaaS, PaaS, IaaS, Amazon Web Services, Azure

# 1. Introduction

Cloud computing first emerged as Infrastructure as a Service, having the Amazon Web Services as its de facto figureheads. The evolution of the term has brought us to a more generic approach becoming an alternative delivery and acquisition model in which anything, and everything, is offered as a service. Cloud computing is perceived now as a general purpose IT utility that is accessible by anyone, from anywhere at any time as-a-Service.

Cloud computing is the convergence of several trends in the past years, it joins a set of technologies and concepts that emerged over time: Software as a Service (SaaS), Grid computing, Virtualization, Utility computing and Hosting. More importantly, it represents a change in IT user's behaviour; users centre of attention is now on what the service offers rather that in how it is implemented or hosted, changing the focus from buying tools to enable a functionality, towards to the contracting of a third-party that delivers this functionality in an elastic, on-demand, and pay-per-use model. Of course, it is not new, Grid, SaaS and Utility models were already doing this for a long time, but it is a clearly a different approach than the classical on-premises infrastructure, license based software models.

Cloud Computing and its underlying "everything as a service" terminology refers to elastic Internet provision of X resources or capabilities. The US National Institute of Standards and Technology (NIST) (Mell & Grance, 2011)has provided the following definitions for the different elements on the Cloud Stack:

- Cloud Software as a Service (SaaS). The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through a thin client interface such as a web browser (e.g., web-based email). The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.
- Cloud Platform as a Service (PaaS). The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly application hosting environment configurations.
- Cloud Infrastructure as a Service (IaaS). The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, deployed applications, and possibly limited control of select networking components (e.g., host firewalls).



IN3 Working Paper Series is a monograph series promoted by the Internet Interdisciplinary Institute (IN3) of the UOC
IN3 Working Paper Series (2010) | ISSN 2013-8644 | http://in3wps.uoc.edu

The following have been identified as potential benefits of Cloud computing:

- Reduction of capital investment: Given its outsourcing nature, Cloud converts IT into an operational expense, paying per use. Also, from a customer perspective, for infrastructure provisioning, the risk of over provisioning or under-provisioning in the datacentre is reduced.
- Scalability on-demand: The elastic capacity provided by Cloud Computing avoids forecasting on compute capacity or compute demand; it can be swiftly and on-demand adapted to business needs, with no need of over- or under-provisioning.
- Lower Operating costs: Cloud providers achieve economies of scale in their shared infrastructure management due to greater resource sharing, greater levels of architectural standardization and operation as well as a better consolidation. These benefits are passed along to the customers of these services, which obtain significantly reduced prices in front of traditional offerings. In addition it has to be considered, the fact that the on-demand nature of Cloud offering results in a linearly priced business model, as the use increases cost scales directly, but without any increase on management complexity or any additional overhead.
- Metered usage and billing: Differently to many outsourcing models based on a fee or a flat rate, Cloud is a transparent pay-per-use pricing model. It is not a recurring bill; it is based on the real consumption of the service, allowing a fine granular IT costs assessment.

The wider adoption of the Cloud model, despite of all previously identified benefits, still faces various challenges with no sufficient mature solutions:

- Security: Data security is the principal concern in the adoption of Cloud services. Many users do only trust systems they have physical control over; systems with corporate firewalls, with known processes and audits; to outsource to any other model is not perceived as a secure model.
- Regulations: Some regulations require of tracking, logging and auditing of enterprise data that for the moment, are not offered by Cloud providers.
- Reliability: Nearly all public Cloud providers have suffered episodes of service level failure or unavailability. This severely concerns in enterprise environments, preferring not to outsource services where they lose control.
- SLA Limitations: Service Level Agreements provided by current public Cloud providers have very limited and not adequate SLAs for enterprise environments.
- Existing investments: Already made investments make the many companies are reluctant to abandon current systems and outsource to Cloud providers.

These facts are illustrated by a recent report from IDC published by the European Commission, '"Quantitative Estimates of the Demand of Cloud Computing in Europe and the likely barriers to uptake"[1] shown in Figure 1.

FIGURE 13 RELEVANCE OF CLOUD BARRIERS

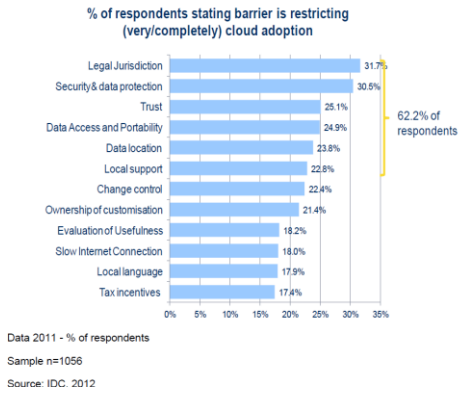


**Figure 1 Relevance of Cloud Barriers, IDC 2012**

[1] October 2012, "Quantitative Estimates of the Demand of Cloud Computing in Europe and the likely barriers to uptake",



IN3 Working Paper Series is a monograph series promoted by the Internet Interdisciplinary Institute (IN3) of the UOC
IN3 Working Paper Series (2010) | ISSN 2013-8644 | http://in3wps.uoc.edu

In particular, the dominant concerns on security issues are:

- Inherent loss of control over enterprise data;
- Confidentiality and integrity of data hosted in the Cloud;
- Inability to enforce security policies due to the de-perimeterisation of the enterprise boundary;
- Inability to perform audits and evaluation of the Cloud offerings;
- Absence of security standards and certification;
- Regulatory compliance with data privacy and mobility of data across multiple legal boundaries;
- Availability and resilience of the services and the associated need for infrastructure protection along with the security of the various technology layers of the Cloud model.

Cloud Security Alliance has recently published "The Notorious Nine" report (CSA,2013) that describes top nine cloud computing security threats for 2013. In it, the two first threats in the list refer to data confidentially, integrity and availability: data breaches and data loss.

# 2. State-of-the-art on Cloud Security

Research in Cloud computing security is a new area that is evolving rapidly. The more extensive use of Cloud computing technologies bring to the user's concerns on its data security and privacy issues, especially for public Cloud adoption. These concerns, in many cases are associated to intrinsic factors in the Cloud computing model, such as multi-tenancy.

As an introductory background, Cloud's essential characteristics as defined by NIST (Mell & Grance, 2011) are hereby presented:

- On-demand self-service. A consumer can unilaterally provision computing capabilities, such as server time and network storage, as needed automatically without requiring human interaction with each service provider.
- Broad network access. Capabilities are available over the network and accessed through standard mechanisms that promote use by heterogeneous thin or thick client platforms (e.g., mobile phones, tablets, laptops, and workstations).
- Resource pooling. The provider's computing resources are pooled to serve multiple consumers using a multi-tenant model, with different physical and virtual resources dynamically assigned and reassigned according to consumer demand. There is a sense of location independence in that the customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction (e.g., country, state, or datacenter). Examples of resources include storage, processing, memory, and network bandwidth.

http://ec.europa.eu/information_society/activities/cloudcomputing/docs/quantitative_estimates.pdf

- Rapid elasticity. Capabilities can be elastically provisioned and released, in some cases automatically, to scale rapidly outward and inward commensurate with demand. To the consumer, the capabilities available for provisioning often appear to be unlimited and can be appropriated in any quantity at any time.
- Measured service. Cloud systems automatically control and optimize resource use by leveraging a metering capability at some level of abstraction appropriate to the type of service (e.g., storage, processing, bandwidth, and active user accounts). Resource usage can be monitored, controlled, and reported, providing transparency for both the provider and consumer of the utilized service.

These essential characteristics of cloud computing pose new challenges IT Security, as well as, it presents specific threats and vulnerabilities to be addressed by Public Cloud providers, in order to make Cloud computing a secure and reliable technology.

In section 2.1 we discuss about the actors and roles in responsible on Cloud security, after it in section 2.2 an analysis of existing literature on general security concerns associated with Cloud computing is provided. To finish this overview, section 2.3. focuses on security issues in Public Clouds, going deeper in each layer of the stack.

## 2.1. Cloud Computing Security Stakeholders

Cloud computing differs from traditional IT security scenarios by its multi-tenant nature (Tianfield, 2012). Multi-tenancy refers to the ability for multiple customers (tenants) to share applications and/or infrastructure resources. This characteristic is the factor that allows Cloud providers to efficiently manage resource utilization among several users in a shared environment; and therefore, it is the enabler to achieve economies of scale and commercial benefit. However, it is the main source of concern for cloud users, if not sufficient protection mechanisms are in place to guarantee security and privacy for both data and applications.

In the secure provision and consumption of Cloud services, there is a difficult balance among confronted actors' interests: cloud service providers and cloud users, where none of them can provide an overall solution for the issue at a general level. Depending on the nature of the cloud service offering (IaaS, PaaS, SaaS) providers are not aware of the contents and security requirements for the applications, while users, at current state of development of Cloud, do not have sufficient vision of the security mechanisms and controls in place at providers' facilities neither for detecting security incidents or branches. This tension of forces is reflected, i.e., in the research findings (CSA, 2010), that provides recommendations for both providers and customers in identifying major security risks and also how-to proactively protect against these risks.

Another approach to this issue is provided by (Tianfield, 2012) that

TABLE I. RESPONSIBILITIES ON CLOUD SECURITY

| | Responsibility of | CP | SP | CC |
|---|---|---|---|---|
| IaaS | VM's security | | | R |
| | Secured VM images repository | R | | |
| | Securing VM boundaries | R | | |
| | Hypervisor security | SR | SR | |
| PaaS | SOA related security | SR | | SR |
| | API Security | | R | |
| SaaS | SaaS security | SR | SR | |
| | Web application security | R | | |

**Figure 2 Responsabilities on cloud security, source** (Tianfield, 2012)

divides the responsibilities associated to cloud security among the Cloud provider and the user, depending on the Cloud specific layer utilised. Although the provided list (shown in figure 2) is incomplete: i.e. IaaS does not consider security of interfaces, PaaS does not include neither development environments nor containers security. It is considered a very interesting approach, given that the different Cloud layers rely on a wide range of different technological approaches (from virtualisation to data management along with web applications and services) that makes it unattainable to address generic Security concerns and challenges in general for all layers with the necessary accurancy and specificity.

## 2.2. Cloud Computing Security challenges analysis

Numerous works provide address Cloud computing security from multiple perspectives:

(Jansen, 2011) from NIST, provides a detailed view on Cloud computing Key Security Issues, organized in the following categories: trust, architecture, identity management, software isolation, data protection and availability. In the context of Trust, it states that Cloud computing confers an unprecedented level of trust to the providers due to two main aspects: first due to the level of insider access, both for the provider, but also, for other users that share the infrastructure, and secondly, due to the lack of transparency of providers in their security practices, that do not offer users' means to assess and audit their services in execution. On Architecture, risks it analyses are limited to IaaS, considering hypervisor and virtual infrastructure main sources of vulnerabilities, such us: leak of sensitive data by accidental publication of VMs and lack of intrusion and detection systems in Virtual Networking infrastructure. With regards to Identity Management it identifies the issue of the need for cloud users' to manage two separate authentication and authorization schemas, the proprietary, and the one offered by the cloud provider, and proposes the use of Identity Federation as a potential solution. Specific issues related to multi-tenancy risks are exposed under software isolation category, such feasibility of injecting malicious code in VMs that share the same physical infrastructure. Again, multitenancy is identified as the main source of threats for Data protection, and it proposes data encryption and data sanitization as a means to protect sensitive information. Compliance is identified as a risk, due to nonexistence of means for users to track data location. With regards availability it presents examples of DDos Attacks and both permanent and temporal outages. It also presents the idea that attacks, not to customer applications, but to Cloud infrastructures themselves are expected to be harder and more common in the future, as cloud platforms gather more and more valuable user's information.

(Jensen & Schwenk, 2009) focuses on the technical analysis of security issues related to Cloud computing considering diverse technology foundations such as WS-Security and TLS, as well as, specific technologies to Cloud: XML Signature, Browser Security and integrity and binding issues, such as cloud malware injection attack (for IaaS, PaaS and SaaS) and metadata spoofing attack based on WSDL vulnerabilities exploitation. However modern PaaS, such as Beanstalk or Google App Engine, although based on services, do not make use of WSDL and UDDI SOA components, but they rely on REST web services, which may be exposed to security threats but not

the ones described in the paper. Other potential described attack are flooding attacks, which are described as an attacker sending so much workload to the provider, ending up with Denial of Service in the provider's hardware. It has to be noted, in this case, that although not completely unfeasible, many public cloud providers already consider this in their architectures, by establishing a maximum in the amount of services a user can request simultaneously ( i.e. AWS controls that a user cannot launch more than 20 VMs at the same time, Azure limits to 100 operations per user per day for non-identified users).

(Srinivasan & Sarukesi, 2012) present a taxonomy on cloud security challenges in which they are classified according to two main categories: architectural and technological aspects, including logical storage segregation & multi-tenancy security issues, identity management issues, Insider attacks, virtualization issues and cryptography and key management; and process & regulatory aspects; such as governance and compliance gaps, insecure APIs, cloud provider migration issues and SLA and trust management gaps. The paper presents in a generic manner potential issues related to these aspects. In addition, it includes a set of questions a user could use to evaluate providers' capacities in the area. The paper does not present solutions to address identified challenges but establishes security requirements and functionalities a cloud provider could offer.

Related to this approach, (Tianfield, 2012) shows a stranded view on Cloud security that poses requirements at every layer:

- Identity security: It proposed End-to-end identity management, third party authentication services, and federated identity in order to preserve integrity and confidentiality of data and applications while making access readily available to appropriate users.
- Information security: It proposed that security in this field becomes information-centric, data needs its own security that travels with it and protects it while in transit and in the Cloud by means of encryption techniques to protect data privacy and legal compliance.
- Infrastructure security: It includes not only securing the physical machine, but SANs (Storage Area Networks) and other hardware devices. Also, considering securing and monitoring the intangibles such as network, end points, traffic flowing among computers, software firewalls, etc…

A completely different approach is presented by (Grobauer, Walloschek, & Stocker, 2011) that focuses on understanding the specific vulnerabilities only applicable to Cloud computing, and inherent to its essential characteristics. To this end, it relates on demand self-service with unauthorized access to management interfaces, given to these APIs are open to be used to the public differently to traditional systems, where only systems administrators access to these interfaces. Ubiquitous network access is linked to internet protocol vulnerabilities, and susceptibility to man-in-the-middle attacks. Resource pooling and rapid elasticity are identified to Data recovery vulnerabilities, due to the fact that resources are reused over the time by different users and memory and storage outages among users are then potentially feasible. Measured service is connected to metering and billing evasions, by means of unauthorized metering and billing data manipulation.

(Subashini & Kavitha, 2011) elaborates on the different security issues and attacks at different layers of the cloud stack. However, a very detailed analysis is provided for

SaaS while for PaaS and IaaS identified risks and issues are analyzed with low level of detail. However, analysis of SaaS issues is very complete, therefore, used as basis for identifying concrete SaaS vulnerabilities later on in this document.

## 2.3. Security issues in Public Clouds: Analysis per Cloud Stack layer.

This section presents a more detailed analysis of the security issues and vulnerabilities in the different levels of the Cloud stack in public Clouds today.

### 2.3.1. SaaS

Software as a Service completely decouples application execution from the user's IT infrastructure. In this model, all application services are solely accessed by the user by a Web browser or thin client over the internet. While enterprise data is stored into the SaaS provider's infrastructure, which can be based on a PaaS or IaaS provider or in a traditional infrastructure provisioning model.

SaaS, although being the Cloud model in which user's information is more exposed to Cloud provider's threats, given the complete loss of control from the user, it is the lesser explored at research levels, accounting for only a few references addressing concretely this topic (Progress Software, 2008; Subashini & Kavitha, 2011). This can be motivated by the fact that SaaS applications are commonly delivered in the form of web applications for which security issues are a well-known and deeply analysed problem (OWASP Project, 2013). However, findings in previously mentioned research and development works demonstrate that specific security issues related to the SaaS Cloud model exists.

#### 2.3.1.1. Identity Management

(Subashini & Kavitha, 2011) identifies three methods for identity management and sign-on: SaaS provider's independent IdM stack, Credential synchronisation among SaaS provider and user's organisation and federated identity management. These mechanisms are exposed to many threats such as: insider threats, including password disclosures and fraud due to staff collusion as well as external online threats including zero-day attacks.

#### 2.3.1.2. Data Security

SaaS providers have to adopt additional security checks to ensure data security and prevent breaches due to security vulnerabilities in the application or thought malicious insiders that lead to unauthorized access to data through: cross-site scripting, access control weakness exploitation, OS and SQL injection flows, Cross-site request forgery, Cookie manipulation and hidden field manipulation (Subashini & Kavitha, 2011).

#### 2.3.1.3. Data Segregation and Data breaches

Data segregation is described as a mechanism to enable that although one organization's data is compromised, attackers cannot gain access to all data records from multiple organizations customer for the Cloud provider. The mechanism provided (Progress Software, 2008) to achieve this is the use of separated data base instances although adding an additional maintenance burden. The simplest mechanism to



IN3 Working Paper Series is a monograph series promoted by the Internet Interdisciplinary Institute (IN3) of the UOC
IN3 Working Paper Series (2010) | ISSN 2013-8644 | http://in3wps.uoc.edu

achieve it, it is data encryption.

An implicit motivation for data segregation is to avoid data breaches in a cloud provider, than in the case of not proper data segregation mechanisms in place could end up with massive attacks in all users' data.

#### 2.3.1.4. Network Security

In SaaS sensitive data is transported from the users organization to the SaaS provider, and therefore it is susceptible to men-in-the-middle attacks, IP spoofing, port scanning, packet sniffing. Providers, therefore need to put mechanisms in place in order to avoid these attacks, as they commonly do through i.e. TLS protocol (Subashini & Kavitha, 2011) .

#### 2.3.1.5. Web Applications Security

Since SaaS applications often rely on a web application schema they are vulnerable to web application threats. For 2013, the Top Ten Open Web Application Security project(Kumar, 2013) has identified the following threats as the most critical for web apps: A1 – Injection (SQL, OS, ..), A2 – Broken Authentication and Session Management , A3 – Cross-Site Scripting (XSS) , A4 – Insecure Direct Object References , A5 – Security Misconfiguration, A6 – Sensitive Data Exposure, A7 – Missing Function Level Access Control , A8 – Cross-Site Request Forgery (CSRF) , A9 Using Known Vulnerable Components and A10 – Unvalidated Redirects and Forwards.

### 2.3.2. PaaS.

Nowadays, there is a truly diverse array of capabilities being offered as PaaS offerings. A PaaS cloud provides a container platform where users deploy and run their components. Diversities are present in supported programming tools (languages, frameworks, runtime environments and databases), in the various types of underlying infrastructure, and even on capabilities available for each PaaS.

However, most know platforms are Windows Azure and Google App Engine. These two platforms add to the execution environment a set of tools in order to facilitate development on top of the platforms. Google App Engine is a PaaS from Google. Developers can quickly build small applications locally (on developer machines) and deploy these to the cloud, on the same environments that power Google applications. It offers fast development and deployment and simple administration. The App Engine platform provides an execution environment where applications run on a virtualized technology foundation that scales automatically on demand. Google App Engine is often criticized for not providing transparency to the user to control infrastructure and how this infrastructure is used. Developers do not have direct control over resource allocation, because the underlying system and hardware resources are masked by the App Engine layer.

The Windows Azure Platform is a PaaS for applications with specific focus on the .NET framework- It is a part of Microsoft's Cloud computing strategy, along with their software as a service offering. The platform consists of various on-demand services hosted in Microsoft data centres and commoditized through three product brands. These are: Windows Azure: an operating system providing scalable compute and storage facilities, SQL Azure: a Cloud-based, scale-out version of SQL Server, and Windows Azure AppFabric: a collection of services supporting applications both in the Cloud and on premise.

Some existing PaaS platforms, such as Cloud foundry (Open Source initiative supported by VMware) and Stax.net automatize the application deployment to a set of template VMs, with complete and isolated platform stack. Vulnerabilities of these types of PaaS are the same than in IaaS environments.

(Rodero-Merino, Vaquero, Caron, & Muresan, 2011) focuses on the risks introduced by multitenancy in Platform-as-a-Service (PaaS) clouds where software components from different users are executed in a shared platform or container system. In order to do this it analyses the common container models a PaaS can employ to achieve isolation among user applications: Operating System (OS) level, Virtual Platform (VP) level (i.e. Java and .NET) and container level. Although common PaaS offering do not build container executions of directly on top of SO, although feasible, it presents issues to provide complete isolation. The common solution to these issues is the use of virtualization, that in this context leads to use of IaaS. At Virtual Platform level two very common platforms are analysed: Java, by far the most popular by the number of public PaaS offering it: and .Net, currently offered by Windows Azure. It has to be noted that also other stacks are common today as PaaS offering, such as Phyton or Ruby (CloudBees, Google App Engine, Amazon Beanstalk), however they are not considered in this analysis. Figure 3 summarizes the main open security issues at each level of a Java PaaS platform.

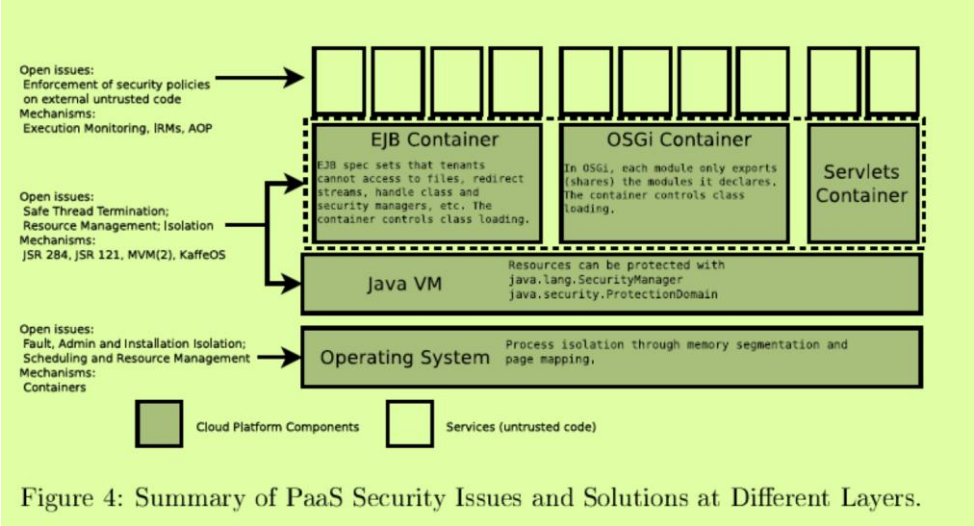


**Figure 3 Java PaaS security issues. Source (Rodero-Merino, Vaquero, Caron, & Muresan, 2011)**

Focussing in .NET, intra-process isolation using different application domains is recommended because they can be dynamically loaded and unloaded during the runtime of the application. For Accounting, similarly to Java, .Net does not provide any generic resource accounting functionality. Safe theat termination solution is based on a C#'s method that does not assure theat termination. With regards containers, no container system similar to Java's EJBs or OSGi exists in .NET.

It is important to note, that this work, although remarkable by the deeper analysis on PaaS, it only contains information about multitenancy in the execution environment of Cloud PaaS providers, and does not research on the data management and database systems they provide and security issues raining from a multi-tenant usage.

### 2.3.3. IaaS.

Infrastructure as a Service (IaaS) is by far the most analyzed Cloud layer with regards security.

In order to produce a systematic view on the question four different issues will be analyzed separately: network security, VM repositories security, Security of Cloud APIs, and general IaaS security demonstrated issues in public clouds. It has to be noted that security concerns on virtualization technologies, are intentionally not further elaborated but in the context of public cloud IaaS implications.

#### 2.3.3.1. Security of API and interfaces

Cloud APIs or Cloud control interfaces are the means that the Cloud providers offer to manage VM images in an IaaS environment. They provide the capacities to add VMs, to modify them, as well as to manage their lifecycle (start, stop and resume).

(Somorovsky, Heiderich, Bochum, Gruschka, & Iacono, 2011) analyses the security of these interfaces in a public Cloud environment, Amazon EC2, and a private Cloud management system, Eucalyptus. In it two different classes of attacks: XML Signature wrapping attacks and XSS attacks on browser front ends. It is important to notice that vulnerabilities reported in both aspects expose important security breaches of providers, given that the attacker get access to all virtual infrastructure of the user, and therefore its data. Focusing Amazon WS EC2 SOAP interface attacks, as the relevant for this work, it shows that knowledge of a single signed SOAP message was sufficient to compromise the security of the user's account. It exploited a vulnerability on the authentication of the SOAP request message that only checks that the XML Signature covers the timestamp header, not checking the overall structure of the SOAP message. Then, it is possible to add, remove or modify parts of the message, without the message validity being affected. The paper also demonstrate it is possible get knowledge on the EC2 SOAP interface implementation by sending manipulated messages with different construction errors, and by analyzing the different origins of responses, get insights about internal implementation.

#### 2.3.3.2. Security of VM repositories

Public VM repositories are a useful mechanism that both private cloud and public cloud providers can offer in order to simply to users the task of creating their own WM images from scratch. Regardless of the usefulness of the mechanism, research demonstrates it can be a source of security risks both for the publisher or the image, the consumer and even the provider in which an instance of this VM is executed. (Wei, Zhang, & Ammons, 2009) in their work identify that the publisher can release sensitive information inadvertently. From both the receiver and the provider that execute them, the result is that they get VMs, or host or execute VMs, that contains malicious or illegal content.

(Balduzzi, Zaddach, Balzarotti, & Loureiro, 2012) have performed an exhaustive analysis over a period of 5 months for all virtual images publicly available in the Europe, Asia, US East, and US West Amazon datacentres. In total 8.448 Linux images and 1202 Windows images were available, of which 5.303 images were analysed. The result of this analysis presents impressive figures:

- 98% of Windows AMIs and 58% contained software with critical vulnerabilities.
- Two Windows AMIs were infected with Trojan –Spy malware



IN3 Working Paper Series is a monograph series promoted by the Internet Interdisciplinary Institute (IN3) of the UOC
IN3 Working Paper Series (2010) | ISSN 2013-8644 | http://in3wps.uoc.edu

- 28% of the images contained leftover credentials
- Many of the VMs allowed recovering deleted files, even the official AMI image.

#### 2.3.3.3. Secure Networking of VMs

Once again, the main source of concern about networking in public clouds is multi-tenancy. VMs from different customers may reside in the same physical network through which data traffic generated by VMs is transported. In order to overcome this issue techniques as network virtualization, through VLAN or other logical network segmentation are applied, so it segregates and isolates traffic among different user groups or subnets(Komu et al., 2012). However, some authors claim that these techniques were designed for the context of an enterprise, and therefore not securely applicable in the context of a public cloud due to limitations in the scale i.e. firewall polices ability to support load or susceptibility to large scale DDoS attacks (Popa, Yu, & Ko, 2010).

#### 2.3.3.4. General IaaS security issues demonstrated

(Rocha & Correia, 2011) evidences the previously mentioned fact that Cloud computing requires an unprecedented level of trust among users and providers. It demonstrates how a malicious insider can easily obtain passwords, cryptographic keys, files and other data from a user through the execution of four attacks that suppose that the attacker has root access to the management VM of the servers that compose the cloud. It is important to remark, that these attacks are not performed on top of public Cloud, but they make use of Xen hypervisor, base technology for multiple public providers' VM format (such us Amazon's AMI). In the first case, the attacker gets a memory dump of the target VM and seeks for passwords in it. The second attack demonstrates the use of a memory snapshot to get an application server's password, stored in memory in the attacked VM. In the third case it uses LVM (Logical Volume Manager), to avoid memory snapshots, that results on a vulnerability of the volume disk. Finally, the last attack the user makes use of Trusted Platform Module (TPM) in order to avoid previously described attacks, and it demonstrate how a VM can be forced to be executed over a modified hypervisor, that would compromise confidentially of data in disk.

(Ristenpart & Tromer, 2009) shows using Amazon EC2 how an attacker can map the internal cloud infrastructure to locate a specific target VM, and then by setting up numerous number of MVs, get one placed on the same physical host than the target VM. Once this is achieved the attacker could then mount cross-VM side channel attacks in order to get information from the target VM. This attack consists in two main steps, first placement, and then extraction of information. The paper demonstrates that the attacker gets a VM placed at the same host than the target VM by investing a few dollars in provisioning infrastructure with a probability of a 40%, and without having specific information about the attacked infrastructure or placement policies in place. In order to overcome this issue the authors recommend cloud providers to obfuscate their internal placement policies, and use blinding technics to minimise information leak.



IN3 Working Paper Series is a monograph series promoted by the Internet Interdisciplinary Institute (IN3) of the UOC
IN3 Working Paper Series (2010) | ISSN 2013-8644 | http://in3wps.uoc.edu

# 3. Comparative Analysis of Security features between Amazon Web Services and Windows Azure

This comparative analysis will focus on Security aspects related to three kinds of Services:

- Virtual infrastructure as a Service, considering both compute and network elements
- PaaS layer, for Java Web application deployment
- Storage Services (Structured Relational databases and Blob storage )
- Identity Management and Access Management Services

## 3.1. Amazon Web Services

Amazon Web Services(AWS) (Varia & Mathew, 2013) is a Cloud platform that offers a diversity of services that organizations can use to deploy applications. Services are classified as follows: Compute and networking Services, Database services, Storage and Content delivery services, Deployment and management services and Application Services. The overall vision of the platform is depicted in Figure 4 AWS Services Overview – March 2013 – Source (Varia & Mathew, 2013), at time of writing:

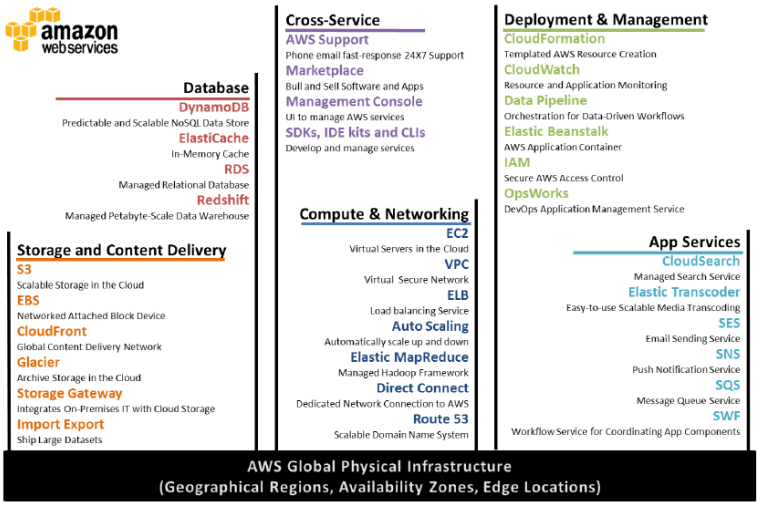


**Figure 4 AWS Services Overview – March 2013 – Source (Varia & Mathew, 2013)**

### 3.1.1. AWS Compute Services

AWS offers the following Compute services:

- Amazon Elastic MapReduce (Amazon EMR) enables the creation of a Hadoop Cluster to process large amounts of data making use of Amazon EC2 and Amazon Simple Storage Service (Amazon S3).
- Auto Scaling allows to automatically scale Amazon EC2 capacity up or down according to user defined metrics on Amazon Watch or based in a user's



IN3 Working Paper Series is a monograph series promoted by the Internet Interdisciplinary Institute (IN3) of the UOC
IN3 Working Paper Series (2010) | ISSN 2013-8644 | http://in3wps.uoc.edu

defined schedule.

- Elastic Load Balancing automatically distributes incoming application traffic across multiple Amazon EC2 instances.
- Amazon Elastic Compute Cloud (Amazon EC2), infrastructure as a service, for which we will analyze in deep security features.

**Amazon Elastic Compute Cloud (Amazon EC2)** offers virtual computing capacity. The virtual appliance format is Amazon Machine Image (AMI), an image template based on the Xen hypervisor with several options with regards to operating systems. Several preconfigured AMIs are available through the AWS Management Console. An instance type determines both the compute and memory capabilities of the VM, as well as, the hardware of the host computer used for its execution. Instance type determines its service execution cost per hour. At time of writing, available instance types are distributed in the following families:

**Table 1 AWS EC2 Instance Types.**

| *Family* | *Description* |
|---|---|
| **Cluster Compute** | Have a very large amount of CPU coupled with increased networking performance. Well-suited for High Performance Compute (HPC) applications and other demanding network-bound applications. |
| **Cluster GPU** | Provide general-purpose graphics processing units (GPUs), with proportionally high CPU and increased network performance for applications that benefit from highly parallelized processing. Well-suited for HPC applications as well as rendering and media processing applications. |
| **High CPU** | Have proportionally more CPU resources than memory (RAM). Well-suited for compute-intensive applications. |
| **High I/O** | Provide tens of thousands of low-latency, random I/O operations per second (IOPS) to an application. Well-suited for NoSQL databases, clustered databases, and OLTP (online transaction processing) systems. |
| **High Memory** | Have proportionally more memory resources. Well-suited for high-throughput applications, such as database and memory caching applications. |
| **High-Memory Cluster** | Have large amounts of memory coupled with high CPU and network performance. These instances are well suited for in-memory analytics, graph analysis, and scientific computing applications. |
| **High Storage** | Provide very high storage density and high sequential read and write performance per instance. Well-suited for data warehousing, Hadoop/MapReduce, and parallel file systems. |
| **Micro** | Provide a small amount of consistent CPU resources and enable you to burst CPU capacity when additional cycles are available. They're well-suited for lower throughput applications and websites that consume significant compute cycles periodically. |
| **Standard** | Have memory-to-CPU ratios suitable for most general-purpose applications. |

AWS EC2 uses a customised version of Xen open source hypervisor. Specifications on the customisations implemented by Amazon are not publicly available, however specific details are provided in (Amazon Web Services, 2013a). AWS EC2 uses both para-virtualisation (PV) and Hardware Assisted Virtualisation (HVM), also known as Full-virtualisation, specifically for supporting windows guest servers.

The isolation among VMs executing in the same physical host is managed by hypervisor in EC2. AWS firewall resides within the hypervisor, between the physical network interface and the VM instances virtual network interface, increasing Xen's default network security levels. In EC2 therefore, spoofing network traffic among two



IN3 Working Paper Series is a monograph series promoted by the Internet Interdisciplinary Institute (IN3) of the UOC
IN3 Working Paper Series (2010) | ISSN 2013-8644 | http://in3wps.uoc.edu

VMs that reside in the same physical host (Neighbour Discovery), is secured as if they were not. All traffic generated by a VM, goes through the AWS firewall, that controls inbound and outbound VM's network traffic. Similar approach is applied for separating physical RAM among VMs.

AWS firewall is configured by default in a deny-all mode, and customers have to explicitly open the ports necessary for in-bound traffic. Traffic to these ports can be restricted by protocol, service port and source IP address. Firewall management interface access is only authorised through customer's X.509 certificate and key. Access to raw disk devices is not allowed for customer VMs. AWS proprietary disk virtualisation layer, isolates virtual disks per VM, and automatically resets storage blocks used by the customer, so that, it is avoided unintentional exposure of customers' data.

Amazon EC2 provides guidelines for both securely use existing AMIs (Amazon Web Services, 2013b) and publishing VMs (Amazon Web Services, 2013c). In the first case, recommendations are to update AMI Tools at boot time, disable password-based logins for root, install public key credentials and validate cron jobs, among others. For publishing, recommendations go in the direction of removing all potential personal information from AMI's publisher.

AWS EC2 API for instances management are signed by AWS Secret Access Key (AWS account Secret Access Key or a user's created account in AWS IAM). API calls are protected through SSL to offer confidentiality.

Amazon takes the view that Security is a shared responsibility in a Cloud environment, where the service provider is in charge of providing physical security, as well as, mechanisms in the form of specific tools, features, guidelines and best-practices to be used by the customer in order to implement network and application level security applicable to their business. (Varia, 2010). As illustrative case, AWS EC2 provides root access to the customer, and even, administrative rights from the service provider can be revoked in the virtual infrastructure.

### 3.1.2. AWS Networking Services

AWS offers the following Networking services:

- AWS Direct Connect allows to establish a dedicated network connection from users' premises to AWS.
- Amazon Route 53 offers a Domain Name System (DNS) service.
- Amazon Virtual Private Cloud (Amazon VPC) allows the provision of a private, isolated section of the Amazon Web Services (AWS) Cloud in a user's defined virtual network.

**Amazon Virtual Private Cloud (Amazon VPC)** allows the user to create virtual networks that have private addresses in a selected range and that operate as distinct, isolated networks within the AWS Cloud(Amazon Web Services, 2013a). It allows the user to create multiple architectures such as:

- VPN with a single public subnet, where network ACLs and security groups can be used to control both inbound and outbound traffic to instances.
- VPN with private and public subnets: In this schema the private subnet cannot have access to internet, and instances in the private subnet can establish outbound connections through NAT.



IN3 Working Paper Series is a monograph series promoted by the Internet Interdisciplinary Institute (IN3) of the UOC
IN3 Working Paper Series (2010) | ISSN 2013-8644 | http://in3wps.uoc.edu

- VPN with private and public subnets and hardware VPN access: It adds IPsec connection between the VPN and on-premises user's network. It needs of a VPN appliance deployed in the user's in-house facilities.
- VPN with private subnet only and hardware VPN access: The instances running in AWS do not have direct connection to internet, however they are connected to the user's facility network via IPsec.

The features available for enabling these architectures are the following:

- Security groups: Like in EC2, each Amazon VPC provides a complete firewall solution to manage security groups. The default groups allow inbound communications from other group members and outbound to any destination. Traffic can be restricted by protocol, service port and source/destination IP address.
- Network access control lists: Network access control list are stateless traffic filters that apply to all outbound or inbound traffic to the VCP. Similarly to security groups, Traffic can be restricted by protocol, service port and source/destination IP address.
- Dedicated instances: Within a VPN the user can launch instances that are not shared with any other existing workload, then avoiding completely all multi-tenancy issues.
- Virtual Private Gateway: Allows private connectivity with Amazon VPN and another network. Traffic in this gateway is isolated from network traffic in the rest of existing VPC, and connection among home network and VPC is secured by a shared secret key installed in the customer facilities in a VPC appliance.
- Internet Gateway: It enables direct connectivity to the rest of AWS and the internet. It can use Elastic IP or use a NAT instance deployed in the VPC. AWS provides reference NAT AMIs that can be extended to perform network logging, deep packet inspection, application layer filtering, etc…

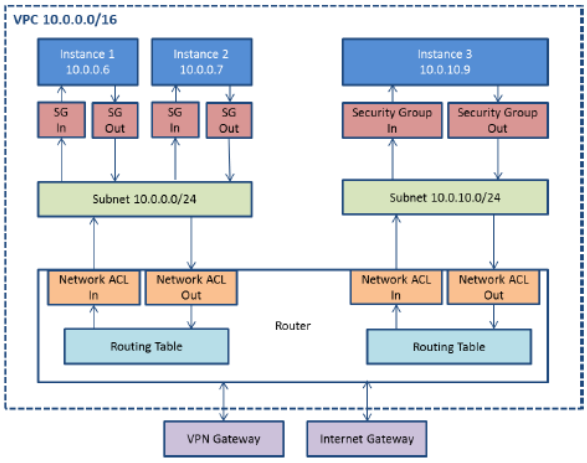


**Figure 5 Flexible Network Topologies in AWS VPC. Source**(Amazon Web Services, 2013a)

All these features are managed through the VPC API. Any call has to be signed AWS Secret Access Key (AWS account Secret Access Key or a user's created account in AWS IAM). API calls are protected through SSL to offer confidentiality.

### 3.1.3. AWS Storage Services

Amazon Simple Storage Service (S3) allows storing data as objects within buckets. An object can be any kind of file: a text file, a photo, a video, etc.

Data Access is restricted by default. Only the Data owner can access it. However, as often data has to be shared, there are three ways to control access to buckets and objects:

– Identity and Access Management (IAM) policies allow the user to manage



IN3 Working Paper Series is a monograph series promoted by the Internet Interdisciplinary Institute (IN3) of the UOC
IN3 Working Paper Series (2010) | ISSN 2013-8644 | http://in3wps.uoc.edu

access to Amazon S3 resources.

- Access Control Lists (ACLs) allow to give read or write access on buckets or objects to groups of users.
- Bucket policies can be used to add or deny permissions across some or all of the objects within a single bucket.

Data Upload/Download is performed over HTTPS.

Stored data is encrypted upon user's request. There are two ways to encrypt data: by using user's own keys and the Amazon S3 Encryption Client to encrypt data before upload, or, directly managing encryption through Amazon S3 SSE. It uses AES-256 and every protected object is encrypted with a unique encryption key.

A bucket can be configured to log access to the bucket and objects within it. The access log contains details about each access request including request type, the requested resource, the requestor's IP, and the time and date of the request. When logging is enabled for a bucket, log records are periodically aggregated into log files and delivered to the specified Amazon S3 bucket.

### 3.1.4. AWS Relational Database Services

Amazon Relational Database Service (Amazon RDS) allows setting up, operating, and scaling a relational database, current offerings include MySQL, Oracle or SQL Server databases.

Amazon RDS offers some basic security features that include DB security groups, permissions, SSL connections, as well as, more advanced capabilities such as: automated backups, DB snapshots, and multi-AZ deployments.

DB instances can also be deployed in an Amazon VPC in order to ensure network isolation. For the moment, this capability is only available for MySQL DBs.

DB Security Groups act like a firewall controlling network access to DB port of users' DB Instance. By default, Database Security Groups are configured in a "deny all" access mode, and in order to authorise network access, it has to be specifically authorised. Authorisation can be done in the basis of a network IP range or based on an Amazon EC2 Security Group.

Connections between applications and DB instances can be encrypted. However, again, this is only currently available for MySQL. In this case then, an SSL certificate is generated per each DB instance. Besides, DB instances can be configured so that they only accept encrypted connections.

For applications requiring data to be encrypted while stored in the DB, user has to manage it on its own.

## 3.2. Windows Azure

Windows Azure is Microsoft's cloud computing platform and infrastructure for building, deploying and managing applications and services through its global network of datacenters.

In May 2013 the available services in this platform are the following (Windows Azure, 2013a):

Information publicly available about technical insights of Security practices in Windows



IN3 Working Paper Series is a monograph series promoted by the Internet Interdisciplinary Institute (IN3) of the UOC
IN3 Working Paper Series (2010) | ISSN 2013-8644 | http://in3wps.uoc.edu

Azure in May 2013 is outdated and refers to 2010 available services and configurations (Kaufman & Venkatapathy, 2010; Meier & Enfield, 2010). For example, it is reported that the approach followed by Microsoft Azure, at the infrastructure management level, is to restrict user's access. In particular, remote debugging, remote Terminal Servers or remote access to VM file shares appears in the documentation as not permitted by default. However simple tests performed demonstrate that current configurations of available services do allow both root access and ssh to virtual infrastructure.

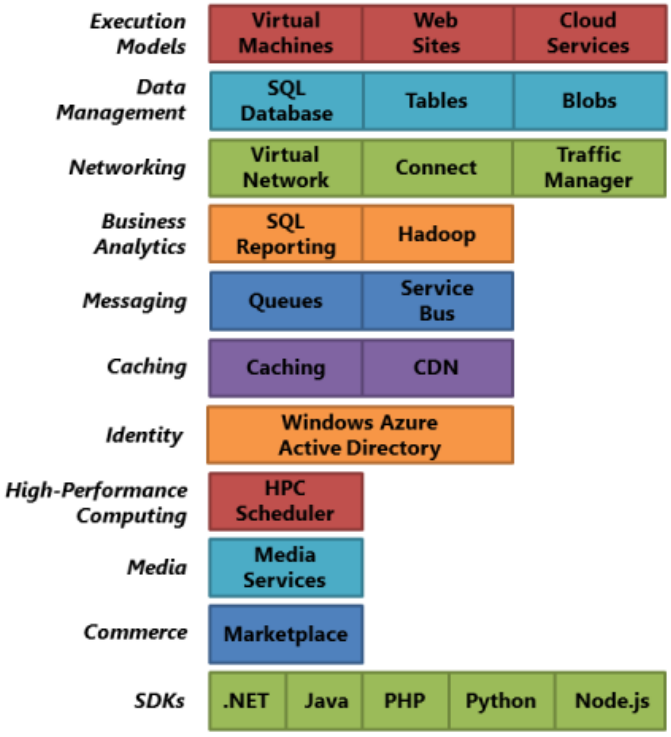


**Figure 6 Windows Azure Components. Source following (Windows Azure, 2013)**

## 3.2.1. Windows Azure Compute Services

**Virtual Machines** service offers virtual computing as a service (IaaS). Virtualization in Azure is based on the Hyper-V hypervisor, and supported operating systems are Windows, some Linux distributions (SUSE Linux Enterprise Server Red Hat Enterprise Linux versions 5.2-6.1 and CentOS 5.2-6.2), as well as, Unix Free BSD.

**Table 2 Windows Azure Virtual Machine Sizes. Source** (Windows Azure, 2013b)

| VM Size | CPU Cores | Memory | Bandwidth | # Data Disks |
|---|---|---|---|---|
| **Extra Small** | Shared | 768 MB | 5 (Mbps) | 1 |
| **Small** | 1 | 1.75 GB | 100 (Mbps) | 2 |
| **Medium** | 2 | 3.5 GB | 200 (Mbps) | 4 |
| **Large** | 4 | 7 GB | 400 (Mbps) | 8 |
| **Extra Large** | 8 | 14 GB | 800 (Mbps) | 16 |

Hyper-V is a hypervisor-based virtualization technology that was first introduced for x64 versions of Windows Server 2008. Isolation is supported in terms of logical units of isolation, called partitions. Host nodes run root or parent partitions enabled by

supported version of Windows Server Operating System (2008, 2008 R2, or 2012). The root partition is the single one that has direct access to the hardware devices, and creates child partitions by API calls.

Improvements to Windows Server 2012 Hyper-V (Microsoft, 2012) consider the following:

- Multitenant virtual machine isolation through private virtual LANs (PVLANs). PVLAN is a technique used to isolate VMs that share a VLAN. Isolated mode means that ports cannot exchange packets with each other at layer 2. Promiscuous ports can exchange packets with any other port on the same primary VLAN ID. Community ports on the same VLAN ID can exchange packets with each other at layer 2.
- Hyper-V Extensible Switch provides protection against a malicious virtual machine stealing IP addresses from other virtual machines through Address Resolution Protocol (ARP) spoofing.
- Protection against Dynamic Host Configuration Protocol (DHCP) snooping and DHCP guard, by configuring ports that can have DHCP servers connected to them.
- Isolation and metering though virtual port access control lists which enable to configure which MAC addresses can (and cannot) connect to a VM.
- Ability to trunk traditional VLANs to virtual machines: Hyper-V Extensible Switch trunk mode allows that traffic from multiple VLANs to be directed to a single network adapter in a virtual machine that could previously receive traffic from only one VLAN.
- Monitoring: Enable to monitor traffic from specific ports flowing through specific virtual machines on the switch.

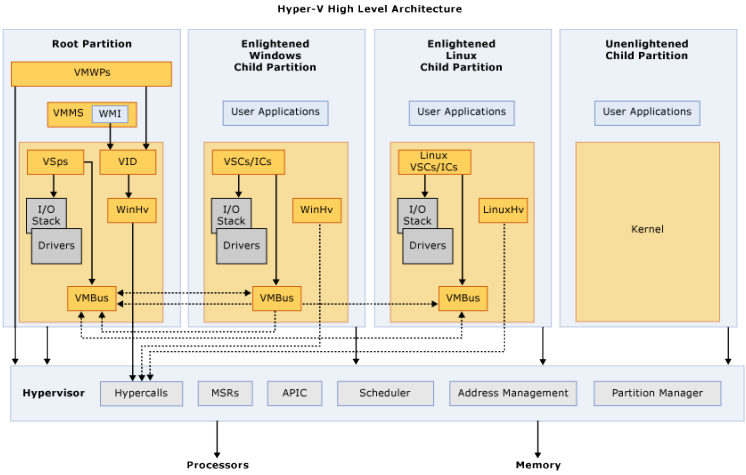


**Figure 7 Hyper-V high level architecture. Source** (MSDN Library, n.d.)

Given the out-dated information provided by Windows Azure with regards security controls in their virtual infrastructure management (Kaufman & Venkatapathy, 2010; Meier & Enfield, 2010), it is not clear nowadays if these interesting features are available, or not, in Virtual Machine services in Azure. What it can be assessed from direct usage is that Amazon Virtual Images offers Management Instrumentation through Windows Azure PowerShell. It is an interface scripting environment that allows controlling and automating deployment and management of workloads in Windows Azure. Authentication is over SSL for security and it can be used both user's own certificate or to generate a new one. With regards to Isolation of the Hypervisor, Root OS and Guest VMs information provided by (Kaufman & Venkatapathy, 2010) states that isolation is managed by the hypervisor and the root OS, without providing more details or evidences on how it is achieved.

### 3.2.2. Windows Azure Networking Services

Windows Azure Virtual Network provides the following capabilities (Windows Azure, 2013c):

- Creation and management of virtual private networks in Windows Azure with your user's defined address space to connect with Cloud Services (PaaS) and virtual machines. The address space follows the RFC1918 specification and public addresses are not allowed in the Virtual Network.
- Cross-site connectivity over IPsec VPN between the virtual network and on-premises network to enable hybrid Cloud and securely extend on-premises datacentre(Windows Azure, 2013d). This feature can be enabled by a VPN device or use Routing and Remote Access Service (RRAS) on Windows Server 2012. Windows Azure has validated a set of standard S2S VPN devices in partnership with device vendors, in order to ensure its compatibility[2].
  Windows Azure defined site-to-site VPNs can be either static, or dynamic:
  - Static routing VPNs are policy-based VPNs. Policy-based VPNs encrypt and route packets through an interface based on a customer-defined policy. The policy is usually defined as an access list.
  - Dynamic routing VPNs are route-based VPNs. Route-based VPNs depend on a tunnel interface specifically created for forwarding packets. Any packet arriving on the tunnel interface will be forwarded through the VPN connection.

Windows Azure recommendation is to use Dynamic routing VPNs when possible. There are different features available both for dynamic and static routing VPNs, among others:

| Property | Static routing VPN gateway | Dynamic routing VPN gateway |
|---|---|---|
| Site-to-site connectivity | Policy-based VPN configuration | Route-based VPN configuration |
| Computer-to-site connectivity | Not supported | Supported (coexists with S2S connectivity) |
| Authentication method | Pre-shared key | Pre-shared key for site-to-site connectivity, Certificates for computer-to-site connectivity |
| Maximum Number of Site-to-site connections | 1 | 1 |
| Maximum Number of Computer-to-site connections | Not supported | 250 |
| Key exchange | IKE v1 | IKE v2 |
| Encapsulation | ESP | ESP for site-to-site, SSTP for computer-to-site |
| Encryption Algorithms | 3DES, AES128, AES256 | 3DES, AES256 |

[2] Windows Azure Known Compatible VPN Devices, http://msdn.microsoft.com/enus/library/windowsazure/jj156075.aspx#bkmk_VPNDevice

| Property | Static routing VPN gateway | Dynamic routing VPN gateway |
|---|---|---|
| Hashing Algorithm | SHA1(SHA128), SHA2 (SHA 256) | SHA1(SHA128), SHA2 (SHA 256) (SHA 384) |

– Use users' DNS servers in the cloud:

Windows Azure provides two mechanisms for Name Resolution in VPN(Windows Azure, 2013e), to use the one provided by Windows Azure or to use users' own DNS server. The latter, is only recommended for cases where user only need name resolution of external DNS names.

An additional feature complementary to the VPN service is the Traffic Manager. Traffic Manager allows to load balance incoming traffic across multiple hosted Windows Azure services whether they're running in the same datacentre or across different datacentres. It can be used to ensure availability of applications, and to manage follow-the-sun policies, enabling by enabling the user of an application to be routed to its closest to him in terms of network latency. Details on specific security features of this service are not available at time of writing.

### 3.2.3. Windows Azure Storage Services

Windows Azure Blob storage is used to store unstructured binary and text data. This data can be accessed over HTTP or HTTPS. Based on user's preferences, data can be encrypted through the .Net Cryptographic Service Providers libraries. Through them, developers can implement encryption, hashing and key management for storage and transmitted data(Kaufman & Venkatapathy, 2010).

Azure Drive(Calder & Edwards, 2010) is a feature of Azure providing access to data contained in an NTFS-formatted virtual hard disk (VHD) persisted as a page blob in Azure Storage. A single Azure instance can mount a page blob for read/write access as an Azure Drive. However, multiple Azure instances can mount a snapshot of a page blob for read-only access as an Azure Drive. The Azure Storage blob lease facility is used to prevent more than one instance at a time mounting the page blob as an Azure Drive. When mounting a drive, the application has to specify credentials that allow it to access the Page Blob in the Windows Azure Blob service. Windows Azure Drive support two different authorization schemes, account and key, as well as Shared Access Signatures (SAS).

The Azure Storage Service supports to associate access permissions with a container through public access control. This allows public read access to the container and the blobs in it or public read access only to the blobs in the container and not to the container itself. The latter would, for example, prohibit the unauthenticated listing all the blobs in the container. The Azure Storage Service also supports of shared access signatures which can be used to provide a time-limited token allowing unauthenticated users a time-limited ability to access a container or the blobs in it. Shared access can be further managed through container-level access policy.

By default storage accounts are configured for Geo Redundant Storage (GRS), meaning that Blob data is replicated both within the primary location and also to a location hundreds of miles away (geo-replication).

In addition, durability for Windows Azure Storage is achieved through replication of data. The replication mechanism used is Distributed File System (DFS) where data is spread out a number of storage nodes. The DFS layer stores the data in what are

called “extents”. This is the unit of storage on disk and unit of replication, where each extent is replicated multiple times. The typical extent sizes range from approximately 100MB to 1GB in size. When storing a blob in a Blob Container, entities in a Table, or messages in a Queue, the persistent data uses one or more extents (Calder, Wang, & Ogus, 2011).

### 3.2.4. Windows Azure Relational Database Services

Windows Azure offers SQL Database based on Microsoft SQL Server.

It offers two types of access control: SQL Authentication and a server-side firewall that restricts access by IP address(Berry, 2011; Hoffman, 2011).

SQL Authentication: SQL Database only supports SQL Server authentication, this is, user’s accounts with strong passwords and configured with specific rights.

SQL Database firewall lets the user to allow or to prevent connections by specifying or IP addresses or ranges of IPs.

Along with access control SQL Database only allows secure connections via SQL Server’s protocol encryption through SSL protocol.

SQL Database supports Transparent Data Encryption (TDE).It performs real-time I/O encryption and decryption of data and log files. For encryption it uses a database encryption key (DEK), stored in the database boot record for availability during recovery. TDE protects data stored at DB, and enables software developers to encrypt data by using AES and 3DES encryption algorithms without changing existing applications.

## 3.3. Security Features Comparison

The following table provides a side-by-side comparison of security features available for each comparable service.

| | Amazon Web Services | Microsoft Azure |
|---|---|---|
| **Compute** | **EC2**<br>Hypervisor level Security(embedded firewall)<br>Default Deny-all mode for inbound traffic<br>Secret Access key for access<br>SSL Secured APIs | **Virtual Machines**<br>Outdated information. Windows Server 2012 Hyper-V characteristics (not confirmed they are applied in the provided service)<br>- Hypervisor level security through private virtual LANS and virtual port access control<br>- Hyper-V extensible switch provides protection to DHCP snooping and DHCP guard |



IN3 Working Paper Series is a monograph series promoted by the Internet Interdisciplinary Institute (IN3) of the UOC
IN3 Working Paper Series (2010) | ISSN 2013-8644 | http://in3wps.uoc.edu

| | Amazon Web Services | Microsoft Azure |
|---|---|---|
| **Network** | **Virtual Private Cloud**<br>**Direct Connect**<br>**Route S3**<br>VPN management (isolated from all other traffic)<br>IPSec Cross-site connectivity: static (policy based), dynamic (based on VPN Virtual appliance deployed) | **Virtual Network**<br>VPN management<br>IPSec Cross-site connectivity: static (policy based), dynamic (based on devices)<br>Usage of own User's DNS servers<br>Traffic manager |
| **Storage** | Access control<br>– Identity and Access management policies<br>– ACL<br>– Bucklet policies<br>SSL protected upload/Download<br>Data encryption on request | Access control<br>– Public access control<br>– Shared access signatures<br>Default data geo-replication |
| **Relational DB** | Access control: DB Security groups (by accessing IP, default deny –all)<br>SSL connection (mySQL only)<br>Isolation through VPN<br>No available data encryption (user has to manage it) | Access control: SQL Server Authentication and by accessing IP)<br>Transparent Data Encryption (AES and 3DES) |

# 4. Hands on AWS and Windows Azure Cloud

In order to make a minimal validation on the security features analyzed in previous section, here security testing performed in both AWS and Windows Azure in the context of IaaS and PaaS services is presented.

## 4.1. IaaS tests: AWS EC2 and Windows Azure Virtual Images

A similar virtual environment was deployed in both providers in order to execute IaaS tests:

| | Instance Type | Operating System | Virtual. | Av.Zone | Open Ports | Characteristics |
|---|---|---|---|---|---|---|
| **AWS** | Small | Installed from base image, Ubuntu 13.04 (64 bits) | Paravirtual | EU West | 8080,22, 8834 | 1.7 G memory, 1 EC2 Compute Unit |



IN3 Working Paper Series is a monograph series promoted by the Internet Interdisciplinary Institute (IN3) of the UOC
IN3 Working Paper Series (2010) | ISSN 2013-8644 | http://in3wps.uoc.edu

| | | | | | | |
|---|---|---|---|---|---|---|
| **Windows Azure** | Small | Ubuntu Server 13.04 (amd64 20130601) for Windows Azure. | Unknown | West Europe | 8080,22, 8834 | 1 core, 1,75G Memory |

First bunch of tests were done using nmap ("Nmap project," 2013) in order to determine open ports in the virtual server and services executing in each port. As expected, results obtained are identical for both providers, consisting in 4 ports where services are listening (22, 8080, 8005, and 8834), all corresponding to identifiable services (ssh, apache tomcat and other installed applications). It has to be noted that port 8005, is used internally by Tomcat.

Second bunch of tests were done using Nessus("Nessus Vulnerability Scanner," 2013). Nessus is a vulnerability scanner used to assess computers, computer systems, networks or applications for detecting security weaknesses. Two scans were performed: Internal network scan and External network scan.

First test focused on discovering internal network vulnerabilities; in this scan the policy is tuned to scan large internal networks with many hosts, several exposed services, and embedded systems such as printers. CGI Checks are disabled and a standard set of ports is scanned for, not all 65,536. In this case, the system to probe was only the defined virtual server. The test was executed directly from the base image provided by each provider, without performing any update or applying any patch on top of the VM provided for the selected operating system, Ubuntu 10.4. Under these circumstances, any critical or high severity vulnerability was detected. Amazon Web Service Image presented five Medium severity vulnerabilities (CVE-2013-1959, CVE-2013-1979, CVE-2012-3544, CVE-2013-2067 and CVE-2013-2071), all of them related to the need of applying security patches to Ubuntu system.

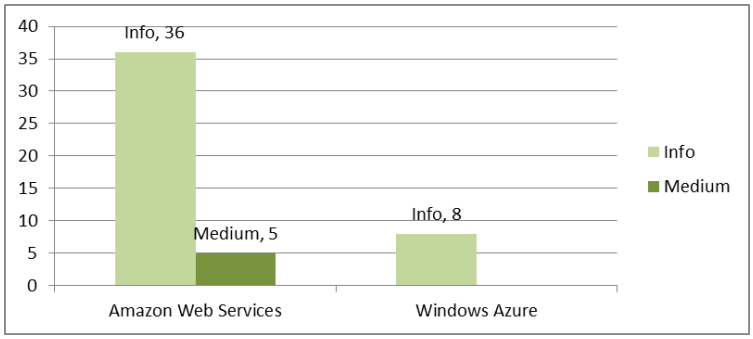


**Figure 8 Internal vulnerability network scan results**

The second test, analyzed the external network configuration for the virtual server. This policy is tuned to scan externally facing hosts, which typically present fewer services to the network. It scans well-known web application vulnerabilities (CGI Abuses: XSS plugin families) and all 65,536 ports (including port 0) on each analyzed virtual server. No additional vulnerability for both cases was discovered as part of these tests.

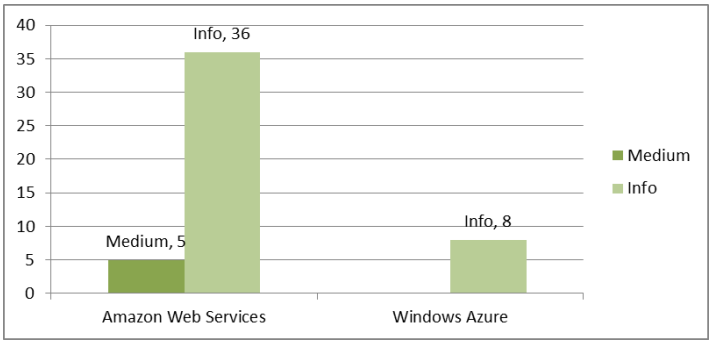


**Figure 9 External vulnerability network scan**

## 4.2. PaaS tests: Amazon Beanstalk and Windows Azure Cloud Services (Java)

PaaS tests were executed using a standard J2EE application developed using the Spring framework and using Tomcat 7 as application server. The analysis did simply consist in dynamic application scan performed through Nessus. Static analysis of code, such as the ones provided by Lapse+(OWASP, 2013) or Zep Attack proxy(OWASP, 2013.) OWASP projects, are considered much of interest for this project, but out of its scope given that the provision of secure code and applications is under the responsibility of the Cloud user in the PaaS layer, and therefore out of the control of the Cloud provider.

PaaS tests therefore, were intended to discover vulnerabilities on the software and virtual infrastructure automatically provisioned by the Cloud provider more than in the application itself.

An important remark to be done before exposing the tests results is that both Amazon Beanstalk and Windows Azure Cloud Services do provision separated virtual infrastructure for each user. Potential security issues identified for PaaS layer in section Section 2.3.2, related to the shared usage of platform or container by software components from different users, are not applicable in these providers, because multitenancy is managed through Virtualisation, and therefore, at IaaS level.

In Amazon Beanstalk application source code, packaging and deployment was done using the AWS Eclipse plug-in in a timely manner and without need of additional components installation. Based on the user's configuration and defined application environments, this plug-in automated the generation of the virtual infrastructure, together with Load balancers, scalability rules, security groups and container configuration. Tests for Amazon Beanstalk used a 64 Amazon AMI running Tomcat 7. Access to the virtual infrastructure was granted to the security group defined by the user. Tests over the web application were executed remotely, thought Nessus Web App tests over the automatically created domain. The Nessus Web App test spiders all discovered web content and then look for vulnerabilities present HTTP servers, Web mirroring among others.

Tests performed only revealed Informational level issues that are depicted in the table below.



IN3 Working Paper Series is a monograph series promoted by the Internet Interdisciplinary Institute (IN3) of the UOC
IN3 Working Paper Series (2010) | ISSN 2013-8644 | http://in3wps.uoc.edu

| Severity | Name | Severity | Name |
|---|---|---|---|
| Info | HTTP Server Type and Version | Info | Apache Tomcat Default Error Page Version Detection |
| Info | Traceroute Information | Info | HTTP Server Cookies Set |
| Info | Web mirroring | Info | CGI Generic Tests Timeout |
| Info | Web Server Directory Enumeration | Info | Backported Security Patch Detection (WWW) |
| Info | Nessus SYN scanner | Info | HTTP Methods Allowed (per directory) |
| Info | OS Identification | Info | Common Platform Enumeration (CPE) |
| Info | Host Fully Qualified Domain Name (FQDN) Resolution | Info | CGI Generic Injectable Parameter |
| Info | Nessus Scan Information | Info | Device Type |
| Info | Service Detection | Info | TCP/IP Timestamps Supported |
| Info | HyperText Transfer Protocol (HTTP) Information | Info | CGI Generic Tests Load Estimation (all tests) |

Tests performed with Windows Azure Cloud Service showed lack of maturity of the solution to provide expected features beyond “toy” applications that consist in a single jsp component, as illustrated in the available developer’s resources.

Windows Azure Cloud Service offers an eclipse plug-in for packaging and deployment of Java applications. This plug-in has a dependency with the .Net 4.0. Framework, that requires of at least 2G of free space to be installed in the user’s environment. The packaging of the application is performed in a separated project, requiring uploading both used jdk and container for deployment. Only Windows virtual hosts are available as candidates to perform the deployment. However, although with in these constrains, the virtual infrastructure generated fails repetitively on start up after inconvenient deployment time. The solution found to overcome these issues, was to manually set up a virtual server and install the application software dependencies, Tomcat 7, to execute the tests. The results gathered are not compared to Amazon obtained results, because in this manner, the service offered by Window Azure is just IaaS, and therefore the security of the software infrastructure provisioned cannot be studied, given that the user was in charge of this task.

Within this context tests over the Web Application using Nessus on *localhost* exposed the following Informative issues:

| Severity | Name | Severity | Name |
|---|---|---|---|
| Info | SSH Server Type and Version Information | Info | Nessus Scan Information |
| Info | SSH Protocol Versions Supported | Info | Apache Tomcat Default Error Page Version Detection |
| Info | Host Fully Qualified Domain Name (FQDN) Resolution | Info | Back ported Security Patch Detection (SSH) |
| Info | Authenticated Check: OS Name and Installed Package Enumeration | Info | Common Platform Enumeration (CPE) |

# 5. Conclusions

Analysis of current research on Cloud Security demonstrates that main source of security concerns related to cloud link directly to its inherent multi-tenant nature. However, multi-tenancy is the essential enabler to achieve commercial benefit for Cloud providers, through economies of scale, which, by means of exploiting this characteristic together with efficient management of resource utilization and overprovision techniques, can provide services at economic prices.

In addition, the hype around Cloud, and the fact that currently it is assimilated to a model in which any IT asset is accessible by anyone, from anywhere at any time as-a-Service, does not help to understand concretely risks and threats derivate from its use. The lack of definition that this fact generates, affects directly to its security assessment, and generates buzz around Cloud security. The literature analysis of Cloud security in general has shown that given the wide scope of technologies and models assimilated, it is difficult to assess concrete security risks and threats, beyond those, as multi-tenancy and need of transparency on security practices followed by providers. An important remark to be made is that often, literature does not provide concrete examples of providers in which theoretical vulnerabilities can be found, lacking of concreteness in their assessment.

In the secure provision and consumption of Cloud services, there is a difficult balance among confronted interests: cloud service providers and cloud users, where none of them can provide an overall solution for the issue at a general level. Depending on the nature of the cloud service offering (IaaS, PaaS, SaaS) providers are not aware of the contents and security requirements for the applications, while users, at current state of development of Cloud, do not have sufficient vision of the security mechanisms and controls in place at providers' facilities neither for detected security incidents or branches.

Given the diverse range of services offered as "Cloud" in order to assess its security it is essential to concentrate at least on the specific cloud layer and, preferable on the specific service intended to be used, given that implementation of services can strongly differ among different providers.

When moving to specific cloud layers, each cloud layer is a different state with regards security analysis. IaaS is by far the most explored area, where well-documented risks and vulnerabilities are exposed by literature. SaaS is just assimilated to Web application, often omitting the data aspect. PaaS security is an area hardly studied, maybe due to the diversity and range of offerings.

Tests performed, although not extensive in scope, show the need of the user to take the responsibility and follow provider's best practices and recommendations, with regards to maintenance of virtual infrastructure, applications and data in the Cloud.



IN3 Working Paper Series is a monograph series promoted by the Internet Interdisciplinary Institute (IN3) of the UOC
IN3 Working Paper Series (2010) | ISSN 2013-8644 | http://in3wps.uoc.edu

# Bibliographic references

Amazon Web Services. (2013a). Amazon Web Services : Overview of Security Processes, *2013*(March). Retrieved from http://awsmedia.s3.amazonaws.com/pdf/AWS_Security_Whitepaper.pdf

Amazon Web Services. (2013b). Sharing AMIs Safely. Retrieved May 5, 2013, from http://docs.aws.amazon.com/AWSEC2/latest/UserGuide/AESDG-chapter-sharingamis.html

Amazon Web Services. (2013c). Public AMI Publishing: Hardening and Clean-up Requirements. Retrieved May 1, 2013, from https://aws.amazon.com/articles/9001172542712674

Balduzzi, M., Zaddach, J., Balzarotti, D., & Loureiro, S. (2012). A security analysis of amazon's elastic compute cloud service. *Proceedings of the 27th Annual ACM Symposium on Applied Computing*. Retrieved from http://dl.acm.org/citation.cfm?id=2232005

Berry, W. W. (2011). Overview of Security in Windows Azure SQL Database. Retrieved from http://social.technet.microsoft.com/wiki/contents/articles/1573.overview-of-security-in-windows-azure-sql-database.aspx

Calder, B., & Edwards, A. (2010). Windows azure drive. *Windows Azure Platform*, (February), 1–13. Retrieved from http://download.microsoft.com/download/9/7/0/97097468-CD68-4FD6-967D-C6C99B39A498/Windows Azure Drive - February 2010v2.pdf

Calder, B., Wang, J., & Ogus, A. (2011). Windows Azure Storage: a highly available cloud storage service with strong consistency. *SOSP '11 Proceedings of the Twenty-Third ACM Symposium on Operating Systems Principles*, 143–157. Retrieved from http://dl.acm.org/citation.cfm?id=2043571

CSA. (2010). *Top threats to Cloud Computing V1.0* (pp. 1–14). Retrieved from https://cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf

Grobauer, B., Walloschek, T., & Stocker, E. (2011). Understanding cloud computing vulnerabilities. *Security & privacy, IEEE*, *9*(2), 50–57. Retrieved from http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=5487489

Hoffman, J. (2011). SQL Azure: Securing SQL Azure. Retrieved from http://technet.microsoft.com/en-us/magazine/hh351833.aspx

Jansen, W. a. (2011). Cloud Hooks: Security and Privacy Issues in Cloud Computing. *2011 44th Hawaii International Conference on System Sciences*, 1–10. doi:10.1109/HICSS.2011.103

Jensen, M., & Schwenk, J. (2009). On technical security issues in cloud computing. *Cloud Computing, 2009 IEEE International Conference on,* (2009), 109–116. doi:10.1109/CLOUD.2009.60

Kaufman, C., & Venkatapathy, R. (2010). Windows Azure™ Security Overview. *go. microsoft. com*. Retrieved from http://www.info.smartglance.com/smartglance/Marketing/WindowsAzureSecurityOverview.pdf

Komu, M., Sethi, M., Mallavarapu, R., Oirola, H., Khan, R., & Tarkoma, S. (2012). Secure Networking for Virtual Machines in the Cloud. *2012 IEEE International Conference on Cluster Computing Workshops*, 88–96. doi:10.1109/ClusterW.2012.29

Kumar, S. (2013). Owasp Top 10 - Most Critical Web Application Security Risks For 2013, 1–15.

Meier, J., & Enfield, P. (2010). Azure security notes. Retrieved from http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:Azure+Security+Notes#1

Mell, P., & Grance, T. (2011). The NIST Definition of Cloud Computing Recommendations of the National Institute of Standards and Technology.

Microsoft. (2012). *Whitepaper Windows Server 2012: Server Virtualization*. Retrieved from http://download.microsoft.com/download/5/D/B/5DB1C7BF-6286-4431-A244-438D4605DB1D/WS 2012 White Paper_Hyper-V.pdf

MSDN Library. (n.d.). Hyper-V Architecture. Retrieved May 5, 2013, from http://msdn.microsoft.com/en-US/library/cc768520(v=bts.10).aspx

Nessus Vulnerability Scanner. (2013). Retrieved from http://www.tenable.com/products/nessus

Nmap project. (2013). Retrieved from nmap.org

OWASP. (n.d.). OWASP Zep Attack Proxy project. Retrieved from https://www.owasp.org/index.php/OWASP_Zed_Attack_Proxy_Project

OWASP. (2013). OWASP Lapse Project. Retrieved from https://www.owasp.org/index.php/OWASP_LAPSE_Project

Popa, L., Yu, M., & Ko, S. (2010). CloudPolice: taking access control out of the network. *Proceedings of the 9th ACM SIGCOMM Workshop on Hot Topics in Networks*, (1), 1–6. Retrieved from http://dl.acm.org/citation.cfm?id=1868454

Project, O. (n.d.). The Open Web application Security project (OWASP). Retrieved from https://www.owasp.org/index.php/Main_Page

Ristenpart, T., & Tromer, E. (2009). Hey, you, get off of my cloud: exploring information leakage in third-

party compute clouds. *Proceedings of the 16th ACM conference on Computer and communications security*, 199–212. Retrieved from http://dl.acm.org/citation.cfm?id=1653687

Rocha, F., & Correia, M. (2011). Lucy in the sky without diamonds: Stealing confidential data in the cloud. *2011 IEEE/IFIP 41st International Conference on Dependable Systems and Networks Workshops (DSN-W)*, 129–134. doi:10.1109/DSNW.2011.5958798

Rodero-merino, L., Vaquero, L. M., Caron, E., & Muresan, A. (2011). Building Safe PaaS Clouds : a Survey on Security in Multitenant Software Platforms, (November).

Software, P. (2008). SaaS Security and privacy.

Somorovsky, J., Heiderich, M., Bochum, R., Gruschka, N., & Iacono, L. Lo. (2011). All Your Clouds are Belong to us – Security Analysis of Cloud Management Interfaces, 3–14.

Srinivasan, M., & Sarukesi, K. (2012). State-of-the-art cloud computing security taxonomies: a classification of security challenges in the present cloud computing environment. *Proceedings of the International Conference on Advances in Computing, Communications and Informatics*, 470–476. Retrieved from http://dl.acm.org/citation.cfm?id=2345474

Subashini, S., & Kavitha, V. (2011). A survey on security issues in service delivery models of cloud computing. *Journal of Network and Computer Applications*, *34*(1), 1–11. doi:10.1016/j.jnca.2010.07.006

Tianfield, H. (2012). Security issues in cloud computing. *2012 IEEE International Conference on Systems, Man, and Cybernetics (SMC)*, 1082–1089. doi:10.1109/ICSMC.2012.6377874

Varia, J. (2010). Architecting for the cloud: Best practices. *Amazon Web Services*, (January), 1–23. Retrieved from https://jineshvaria.s3.amazonaws.com/public/cloudbestpractices-jvaria.pdf

Varia, J., & Mathew, S. (2013). Overview of Amazon Web Services, (March).

Wei, J., Zhang, X., & Ammons, G. (2009). Managing security of virtual machine images in a cloud environment. Proceedings of the 2009 ACM workshop on Cloud computing security, 91–96. Retrieved from http://dl.acm.org/citation.cfm?id=1655021

Windows Azure. (2013a). Introducing Windows Azure. Retrieved May 5, 2013, from http://www.windowsazure.com/en-us/develop/net/fundamentals/intro-to-windows-azure/

Windows Azure. (2013b). Windows Azure Training Kit - Virtual Machines. Retrieved May 5, 2013, from https://github.com/WindowsAzure-TrainingKit/PRESENTATION-WindowsAzureVirtualMachines

Windows Azure. (2013c). Windows Azure Virtual Network Overview. Retrieved May 5, 2013, from http://msdn.microsoft.com/en-us/library/windowsazure/jj156007.aspx

Windows Azure. (2013d). About VPN Devices for Virtual Network. Retrieved May 5, 2013, from http://msdn.microsoft.com/en-us/library/windowsazure/jj156075.aspx

Windows Azure. (2013e). Windows Azure Name Resolution. Retrieved May 5, 2013, from http://msdn.microsoft.com/en-us/library/windowsazure/jj156088.aspx

***Summary***

Analysis of current research on Cloud Security demonstrates that main source of security concerns related to cloud link directly to its inherent multi-tenant nature. However, multi-tenancy is the essential enabler to achieve commercial benefit for Cloud providers, through economies of scale, which, by means of exploiting this characteristic together with efficient management of resource utilization and overprovision techniques, can provide services at economic prices. Given the diverse range of services offered as “Cloud” today, in order to assess its security it is essential to concentrate at least on the specific cloud layer and, preferable on the specific service intended to be used, given that implementation of services can strongly differ among different providers. This paper concentrates in doing so: by providing a extensive literature review of Cloud security in general, and per Cloud layer, as well as, a comparative analysis of security features provided by two of the most representative Cloud providers: Amazon Web Services and Microsoft Azure.

***Keywords***

Cloud Computing Security, SaaS, PaaS, IaaS, Amazon Web Services, Azure

***Ana Juan Ferrer***
*anajuanferrer@gmail.com*
*Information and Communication Technology Security Master*
*Universitat Oberta de Catalunya*

*Ana Juan Ferrer has a Bachelor Degree in Computer Science from Universitat Oberta de Catalunya. She works for Atos since 2006, in Atos Research and Innovation, currently she is Head of Lab of the Service Engineering & IT Platforms Lab, group that focus its research in Cloud Computing, Service Engineering and GreenIT. From 2012 she coordinates the OPTIMIS project, investigating platforms and architectures for scalable and trustful Cloud services platforms. She also participates in Helix Nebula, focusing in research on Cloud Service Management models. In the past, she has worked in NUBA; NEXOF-RA, definition of the NESSI Open Framework Roadmap; BEinGRID, industrial application of Cloud and Grid; and Crosswork resereach projects. Before Atos she has wide experience as consultant and software architect in the Internet environment and e-business. Ana also participates in the Atos Scientific Community, a network of some 90 members which aim is to foster innovation in the company.*

*Additional information about the author at: http://es.linkedin.com/in/anajuan*

***Ana Juan Ferrer***
*anajuanferrer@gmail.com*
*Information and Communication Technology Security Master*
*Universitat Oberta de Catalunya*

*Ana Juan Ferrer has a Bachelor Degree in Computer Science from Universitat Oberta de Catalunya. She works for Atos since 2006, in Atos Research and Innovation, currently she is Head of Lab of the Service Engineering & IT Platforms Lab, group that focus its research in Cloud Computing, Service Engineering and GreenIT. From 2012 she coordinates the OPTIMIS project, investigating platforms and architectures for scalable and trustful Cloud services platforms. She also participates in Helix Nebula, focusing in research on Cloud Service Management models. In the past, she has worked in NUBA; NEXOF-RA, definition of the NESSI Open Framework Roadmap; BEinGRID, industrial application of Cloud and Grid; and Crosswork resereach projects. Before Atos she has wide experience as consultant and software architect in the Internet environment and e-business. Ana also participates in the Atos Scientific Community, a network of some 90 members which aim is to foster innovation in the company.*

*Additional information about the author at: http://es.linkedin.com/in/anajuan*