



WestminsterResearch

<http://www.wmin.ac.uk/westminsterresearch>

Security as a service in smart clouds: opportunities and concerns

Vladimir Getov

School of Electronics and Computer Science, University of Westminster

This is an electronic version of a paper presented at IEEE COMPSAC 2012, 16th - 20th July 2012, Izmir, Turkey, and is reprinted here with permission.

The WestminsterResearch online digital archive at the University of Westminster aims to make the research output of the University available to a wider audience. Copyright and Moral Rights remain with the authors and/or copyright owners. Users are permitted to download and/or print one copy for non-commercial private study or research. Further distribution and any use of material from within this archive for profit-making enterprises or for commercial gain is strictly forbidden.

Whilst further distribution of specific materials from within this archive is forbidden, you may freely distribute the URL of the University of Westminster Eprints (<http://www.wmin.ac.uk/westminsterresearch>).

In case of abuse or copyright appearing without permission e-mail watts@wmin.ac.uk.

Security as a Service in Smart Clouds – Opportunities and Concerns

Vladimir Getov

School of Electronics and Computer Science
University of Westminster
115 New Cavendish Street, London, W1W 6UW, U.K.
V.S.Getov@westminster.ac.uk

Abstract—The cloud computing paradigm emerged shortly after the introduction of the “invisible” grid concepts but it has taken only a few years for cloud computing to gain enormous momentum within industry and academia alike. However, providing adequate security support by those complex distributed systems is of primary importance for the wide adoption of cloud computing by the end users. This paper gives an overview of the main cloud security issues and challenges. Existing and proposed solutions are also presented with particular attention to the security as a service approach. Some of the available directions for future work are also discussed.

Keywords—cloud computing; trust and security; security as a service; smart clouds

I. INTRODUCTION

In recent years the concepts and implementations of modern distributed computing infrastructures have been developing in the area of cloud computing. Subsequently, a new discipline rapidly emerged based on cloud computing’s layered architecture with accepted succinct definitions for the software-as-a-service (SaaS), platform-as-a-service (PaaS) and infrastructure-as-a-service (IaaS) layers [17]. However, in order to tackle further the security challenges at global scale we need qualitatively new high-end computing infrastructures with smart properties. The cloud computing paradigm is suitable for this purpose but it needs significant improvements to implement the best trade-off between the autonomy necessary to scale up, and the trust of the users. On the one hand, users must have the guarantee that their requirements, in terms of management of their applications and data, will always be satisfied. On the other hand, autonomic management of applications and data would give the infrastructure greater flexibility and would enable better security, performance, robustness, energy consumption, etc.

While the basic cloud middleware tools and technologies such as virtualization, distributed storage, distributed execution, data and computation migration, service definition and implementation have reached industry-standard engineering level [14], cloud computing is not yet widely adopted. This may be for several reasons, but probably the most important one is related to security [5]. For example, many companies are subject to regulations concerning the way the data they own should be dealt with. This is

particularly important when the cloud adoption implies the disclosure of private, sensitive, and/or valuable data to the cloud service provider (CSP). In such cases a company will wonder: Where will my data be stored? In which countries will the infrastructure be located? What are the security regulations in those countries? Is the data going to be stored in a single physical place or distributed across different sites? Are intermediate results of the computation secured/encrypted? More generally, do the data storage and computation comply with the requirements expressed by the user?

Issues of the same nature also arise in more complex scenarios. In a number of significant cases, the application of the cloud service is inherently distributed. This typically happens in manufacturing and business domains, where different organizations interoperate and may to some extent share services or data. In these cases, the current tendency in addressing security and trust is to achieve security-by-obscurity, i.e. either the user trusts the CSP or the user owns the cloud itself. In other words, attempts to gain the users’ trust are aimed mainly at avoiding the issue by privatizing or embedding the data and services (e.g. private clouds), rather than solving it.

A promising way to address the current security concerns and problems is based on making the quality of service (QoS) a first-class concept, which is managed as a contract, i.e. known by partners and independently verifiable. This can be achieved by designing and building smart clouds [4] via a full exploitation of the autonomic computing paradigm as a distributed set of mechanisms for monitoring and control, equipped with a distributed, diffused, and dynamic management overlay. Up to now, autonomic computing has been industrially developed only at a syntactic level as a method to express application-specific concepts. An important next phase is to develop a fully-fledged methodology, middleware services, and development tools able to support the QoS-scalable design of services in smart clouds.

The rest of this paper is structured as follows. Section II provides background overview information about the current status of security and interoperability in cloud computing. Section III presents four different categories of cloud security issues. Section IV describes the Security-as-a-Service (SecaaS) and finally, Section V concludes the paper.

II. BACKGROUND

The concept of a hybrid cloud is an attractive one for many organisations, allowing an organisation with an existing private cloud to partner with a public cloud provider. This can be a valuable resource as it allows companies to keep some of their operation in-house, but benefit from the scalability and on-demand nature of the public cloud. There are, however, a number of issues that organisations must consider before opting for a hybrid cloud set-up.

The single most pressing issue that must be addressed is that, by definition, the hybrid cloud is never 'yours' – part of it is owned or operated by a third party, which can lead to security concerns. With a true private cloud – hosted entirely on your own premises – the security concerns for an IT manager are no different to those associated with any other complex distributed system. However, when this private cloud is hosted by a third party, the security issues facing the customers become very complex and difficult to solve. Although this cloud is in theory, still private, the fact that it relies on external resources means that the users are no longer in control of their data. Security remains a major adoption concern, as many CSPs put the burden of cloud security on the customer, leading some to explore costly ideas like third party insurance.

It is a huge risk, as well as impractical, to insure the very expensive company data – potential losses from losing major trading or logistical applications are enormous. CSPs should offer greater assurance to reduce the idea that insurance is even needed. Another issue that organisations must consider is interoperability – internal and external systems must work together before security issues can be considered.

It could be said, therefore, that a true hybrid cloud is actually quite difficult to achieve, when interoperability and security issues are considered. One solution might be a regulatory framework that would allow cloud subscribers to undergo a risk assessment prior to data migration, helping to make service CSPs accountable and provide transparency and assurance.

Concerns with hybrid cloud are indicative of the anxiety that many companies feel when considering cloud computing as a viable business option. We need to see a global consensus on regulation and standards to increase trust and security in this technology and lower the risks that many organisations feel go hand-in-hand with entrusting key data or processing capabilities to third parties. Once this hurdle is removed then the true benefits of cloud computing can finally be realised.

III. SECURITY ISSUES IN CLOUD COMPUTING

Generally speaking, there are several categories of security attacks in a cloud. This section considers the most important ones.

A. Infrastructure Security

The analysis of infrastructure security can be split into three layers – network, host and application. In the case of a private cloud, there is no difference between the previous IT infrastructure and the cloud solution. The old security

measures can be kept in place and there are no risks of malicious insiders. Consequently, there are no new threats to the infrastructure security. Therefore, the following discussion is concerned only with the public and hybrid deployment models.

1) Network security

This is valid both for the physical network on which the cloud runs and the virtual network which is created by the virtual instances running in the cloud. There are three aspects which can be separated [7]:

a) *Securing data in transit (to and from the public cloud)*. This could be done by setting up a VPN at least for the virtual instances running in the cloud. The most common approach is to use HTTPS, thus implementing strictly in the browser as a thin client.

b) *Managing access control (authentication, authorization, auditing)*. The physical part of the network can be protected in the same way as before – using firewalls and IDS solutions. However, the management of the access control is different on the virtual part. Network zones disappear, so a virtual separation has to be envisioned. This is done using domains and rules are enforced by application (virtual) firewalls.

c) *Availability*. The network faces the same threats regarding its availability: routing protocol tampering (such as BGP prefix hijacking), DNS poisoning, DNS forgery or the infamous denial of service (DoS). In the context of the cloud, the attack can also originate internally.

2) Host security

Securing the hosts supporting the cloud requires the same operations as with protecting an enterprise IT infrastructure. The process of OS hardening should be strengthened and should only allow the processes that support the operation of the cloud. In this context the security of the hypervisor comes into play. The chosen virtualization tools should be the trusted ones, with as few successful attacks as possible in their background. Even these solutions have weaknesses which make them susceptible to attacks, but with time they are expected to mature enough to be considered secure.

Regarding the virtual aspects, security provisioning depends on the chosen type of service. At the IaaS layer, the customer is responsible for security measures. Hosts should be protected as if they were operating in an untrusted network using software firewalls, antiviruses OS hardening, etc. Furthermore, the virtual hosts should be patched and maintained in a way similar to the physical resources. However, these virtual hosts are incomparably more dynamic than the physical ones. It should be understood that the only feasible way to scale the infrastructure to hundreds and thousands of running virtual machines (VMs), is to employ the use of templates. If any security flaws slip in here, the whole cloud network using the weak template is vulnerable to a quickly expanding attack.

At the PaaS later, the CSPs usually offer a set of APIs for interaction with the host abstraction layer which developers cannot avoid – their security is in the hands of the provider. However, applications developed here should consider the

Internet-threat model. Aside from this, there is no control that the customer has over security. In the case of SaaS, the user totally relies on the CSP.

3) *Application security*

Until now, it has been a common practice to secure applications using perimeter security controls and access management. In the case of a public cloud, applications face higher risks than the applications running on a private cloud or on a traditional IT infrastructure. The new family of applications are based on the browser and rely on net-centric technologies. Developing these applications should take into consideration the common threats regarding web applications such as OWASP's Top 10 [10]. Of particular importance is the browser used to run the application. As all applications basically are run by the browser and not by the host's OS, it is of no surprise that the browser can be seen as an OS for cloud applications. Therefore, the security policy should cover browser maintenance as well.

A very powerful attack that can be employed at this level is the application DoS attack. Such attacks often originate from infected hosts as it is the case with network DoS. The DoS attack is a real threat in a cloud because of the real potential to access virtually unlimited resources. These attacks can range from repeatedly refreshing web pages to loading the applications with specific tasks or protocol-specific requests supported by the cloud service.

A unique type of attack is the economic denial of sustainability (EDoS) [7]. This type of attack is directly connected with a DoS attack, but its target is to inflate the cloud services budget. Because of the pay-per-use business model, a lengthy application DoS attack has financial implications as well.

B. *Data Security*

Data security tries to ensure data confidentiality or privacy and integrity of data and to allow audit operations to be carried on it. There are three relevant aspects for this.

1) *Data encryption (confidentiality)*

The infrastructure security should prevent unauthorized access to data. Following a layered security approach, the next protective mechanism should assume that the infrastructure security mechanisms have failed – and try to prevent or control data loss or damage.

Encryption might be the obvious next choice, but the current encryption techniques are only functional for data at rest and data in transit (to and from the CSP). It is not possible to perform operations on encrypted data – aside from simple sum and some products. Therefore, dynamic data such as data used at SaaS or PaaS layers and non-storage IaaS services remains unencrypted [7] in order to support operations like indexing, searching or mathematical calculations. There is current work in progress on the next generation technology called *homomorphic encryption*. This technique will allow operations on encrypted data – but is very computing-intensive [11] and opens up a number of new problems such as development, debugging, and validation of software working with encrypted data.

Considering these, the current conclusion is that sensitive data should not be moved to a public cloud, unless it's for

Storage-as-a-Service and is encrypted. Even in this case, recent surveys [12] show that there still is a major concern among security professionals regarding encryption in the cloud. Therefore the following should be considered when opting for an encryption system in the cloud:

- The CSP stores metadata which describes what clients store in their system. This metadata should also be protected through at least some type of access mechanisms, if not possible to encrypt. In the case of encrypting this metadata, there is also the problem of key management in the cloud which is generally several orders of magnitude larger than an enterprise key management system [7].
- The most secure cloud encryption environment (considering it runs in the cloud) is to separate the three components involved in the encryption process – key, algorithm and data [8]. In this way, even if one of them is compromised, the attacker still needs to compromise the other two. This offers time to secure the data – provided the first attack was observed.

2) *Data integrity and retrievability*

In the case of Storage-as-a-Service – e.g. storing encrypted data in a public cloud – aside from confidentiality, it is also important to implement an integrity mechanism that confirms that no one tampered with the data. This is usually a mechanism which is implemented together with the encryption – for example, digital signatures.

It could be relevant for a customer to verify that the CSP still has their data intact. If the CSP experienced a situation which led to client data loss – such as corrupted or deleted data – the integrity check alone could not reveal the problem. Furthermore, because of the pay-per-use model, it should be avoided to retrieve the data from the CSP just to perform this check.

There are some mechanisms called *data retrievability* that have been developed to allow the customer to confirm whether their data is still intact in the cloud [16]. From this, the mechanism was extended to allow stored data manipulation such as block deletion and insertion without the need to download a local copy [15].

3) *Data lineage and provenance*

In addition to protecting data, security also ensures that audit procedures can be operated. These procedures serve not only for forensic evidence in case of an attack, but can also help the customer to better understand how their data is being manipulated in the cloud – so that appropriate security measures are employed.

In such a dynamic environment as the cloud, it is desirable to keep track of the data lineage – e.g. where each particular piece of data was at any given moment of time. As the cloud infrastructure operates with virtual devices, this lineage tracking can be extended to virtual instances as well.

Because of the high dynamism in the cloud, it is very difficult to collect extensive information on data lineage such as the state of the systems dealing with each data piece. In fact, it is likely that the only information collected is limited to: IP addresses, country in which that specific host resides,

host name, host domain and time stamp. Nonetheless, this information alone is relevant and important in many cases.

In some applications it might also be relevant to consider data provenance. This information describes where and when data originated which is the first entry in the list of data lineage. If there is enough information describing where the information originated, the data can be considered as valid [7]. A relevant example is an application that operates with currencies in which case the exchange rate should originate from a trustworthy source.

4) *Data remanence*

Data remanence is not specific to cloud computing. Strong security policies cover the aspect of how data is eliminated from the physical media at the end of its life cycle. In the case of a private cloud, there are no new mechanisms which should be employed for this aspect.

In the case of a public cloud, it is important to note that the customer does not have access to the physical storage devices. Therefore, the client cannot control how these devices are disposed of or reused. The only protection in this case is a strict contract or SLA which stipulates how this process is to be implemented by the CSP.

There are currently no standards on how this operation should be performed. There are however guidelines from NIST which describe how data sanitization – e.g. removing the data from the media before reusing the media in an untrusted environment – should be performed [6]. Even if it is only intended for federal civilian departments and agencies, the document can be used as a starting point in formulating the contractual terms.

5) *Backup*

Even though this does not contribute to preventing unauthorized access to data, it is relevant to discuss the backup processes. One of the reasons to store data in the cloud is its capacity to survive in the case of an accident such as deletion or disaster such as fire or flooding.

As with data remanence, a private cloud should not change backup plans already in place in a system where redundancy was considered. In the case of a public cloud, the subject is again out of the customer's reach. Therefore, this aspect should also be included in the service contract.

C. *Identity, Authentication, Authorization and Auditing (IAAA) Management*

A rather important part of security provisioning is the management of *who* (authentication) can access the secured cloud service and *what can that person do* (authorization). Aside from this minimum, there are two other desired functions:

- having a more complex system which can rule out duplicated accounts and which could link accounts for different applications, belonging to the same person (identity management);
- for compliance and performance reasons, auditing the whole process.

All these four processes form the IAAA management of cloud services. When discussing IAAA, there are two orthogonal factors which should be considered in order to determine the appropriate solution: the size and capabilities

of the customer, giving the complexity of the IAAA mechanisms, and the cloud service utilised, giving the threats for the IAAA system.

1) *Customer*

The IAAA needs vary considerably based on the type of the customer. In the case of an individual consumer, their needs will usually be met by a simple “username and password” type of authorization and a Single Sign On (SSO) in services used.

In the case of an enterprise, the problem is more complex. First of all, there is the problem of the modified trust boundary. In the traditional model, the trust boundary encompasses the network, systems and applications hosted in the local trusted IT infrastructure. With cloud computing, the trust boundary extends to include the CSP's domain – and the customer does not have control over the IAAA mechanism provided in this domain.

Furthermore, as opposed to the traditional model where the infrastructure was trusted, in the case of cloud computing the threat model changes: IAAA mechanisms have to operate in an untrusted environment. This has deep implications in all four IAAA components [2]:

a) *Identity* – The number of user identities managed in a cloud might go from 50.000 in the case of an enterprise to millions as users no longer access a few local systems, but potentially many distributed systems.

b) *Authentication* – The process should be more complex lowering the risk of unauthorized access, implementing standard-based solutions and opting for multi-factor (usually two) authentication. Aside from username and password, in this case the authentication process should be validated by something unique for the user such as a smart card, or a token, or the person's location. The device used for authentication should provide strong confirmation of who the person using, for example, biometrics or somebody who the person knows – another person who can verify the identity of the user.

c) *Authorization* – Depending on the cloud service type (*aaS), the threat model must consider the shared characteristic of the environment – e.g. all users have access to the same application, platform or infrastructure. For example, data contained in a shared application should only be accessed by the desired users and the APIs used should each carry an authorization mechanism to block unauthorized access.

d) *Auditing* – Special care has to be taken in order to maintain the trustworthiness of the reports obtained in audit procedures. Furthermore, the customer does not have access to the auditing process that the CSP has in place – leaving some possible openings in the system, which should be covered by contractual terms and SLAs.

2) *Cloud services*

a) *SaaS* – The authorization system must be much finer – users must be allowed access to only a part of the

application, while most of the responsibilities are on CSP's side.

b) *PaaS* – In addition to SaaS, there are more possible roles with a whole new family of customers – the developers. Furthermore, as the cloud is usually regarded as an untrusted environment (particularly in the case of a public solution) and as developed applications could span multiple clouds, messages exchanged by APIs or new applications should be individually protected.

c) *IaaS* – The infrastructure IAAA mechanisms that are in place for a traditional (physical) solution no longer apply in the virtual world. The security mechanisms move to application-level security provisioning. Given the size of the users' base that a cloud solution has and the tendency to operate in multiple clouds, it is important to implement mechanisms which:

- avoid duplication of identity, attributes and credentials;
- allow automatic user management for customers;
- allow SSO without disclosing credentials.

In the end, both parties – the user and the CSP – have to understand that they play an important part in the IAAA mechanism. The CSP is expected to implement state-of-the-art, standard-compliant solutions and allow interoperability with other clouds including SSO; the user must secure access for its part of the stack and enforce policies. As this step of security provisioning in the cloud is rather complex and requires continuous management from the user, there is an emerging trend to move everything that has to do with IAAA mechanisms in the cloud as a specialised service called Identity as a Service or IDaaS.

D. *Nefarious Usage of Cloud Computing*

In a highly secure cloud environment, with unbreakable infrastructure and data security and with strong IAAA mechanisms, there is still the risk of somebody harnessing the power of a cloud system in order to perform attacks. Such an attacking cloud is called a *dark cloud* [7]. The attacking cloud could be a private one, but surprisingly it could also be part of a public one, as the cost-effectiveness is attractive to attackers as well.

A dark cloud could serve as the launching ground for very computing-intensive attacks. A cloud can also be used for creating dynamic attack sources, hosting malicious data or operating botnets. Furthermore, a cloud can also serve to launching DDoS attacks [3]. In the case of the computing-intensive attacks, there is virtually no protection that can be envisioned. Without continuously monitoring the actions that a customer performs which is almost always against privacy agreements, a CSP cannot tell what the computations are used for. The only solution is to create even more complex protection mechanisms such as longer encryption keys, stronger encryption algorithms and passwords.

Regarding the storage of malicious data or coordinating attacks from the cloud, there is very little that can be done without breaking privacy agreements and spying on the data

stored by the customer. In the case of a private cloud, there is literally no tool to prevent this from happening.

The case of DDoS attacks requires a closer look. If the target is not a cloud system, but a traditional IT infrastructure, then it can easily be denied its service. If the target is also residing in the cloud, then the consequences of a DDoS attack depend on the severity of the attack. A severe attack can also deplete the target's cloud resources and lead to a denied of service state, whereas a less severe one can lead to excessive billing – EDoS. This can also be the ground for a blackmailing case.

Defending against a DoS attack is not an easy task. There are two main approaches which could be envisioned:

1) *Protecting in the originating cloud (if the cloud is public)*. This could be accomplished by installing application IDS systems and monitoring the hosts of each client, but the level of resources required is so large, that this is not usually taking place.

2) *Protecting in the target's cloud*. The customer can opt for security services that include DDoS prevention. Solutions usually employ the approach of application IDS systems – a tried and tested approach with the traditional IT infrastructure. There is continuous work in improving the algorithms which protect against this type of attack as it is very likely in the cloud era.

IV. MOVING SECURITY TO THE CLOUD: SECURITY-AS-A-SERVICE

A relatively new approach to security is based on the Security-as-a-Service (SecaaS) concepts [13]. As discussed above, securing a cloud system is a complex and time consuming task. Some companies opted to outsource this process and the market experienced a necessity for such security services. The answer was the appearance of managed security services (MSSs). A MSS provider (MSSP) assigns security personnel to its clients to administer the security mechanisms (generally, related to cloud services), using a pay-per-use model. With this, the customer is in charge of the security policies and it is his responsibility to monitor the efficiency of the services provided by the MSSP.

As the client pool increased, MSSPs considered ways to centralize the service, in order to simplify and improve their results – hence the idea of security provisioning from the cloud. But MSS is not the only factor that leads to the development of the SECaaS. Companies that did not outsource the security provisioning and individual cloud users were still an attractive target and opportunity for cloud security services.

Regardless of the target customer, moving cloud security into the cloud itself does have a strong advantage – it can, as well, harness the power of the cloud. Furthermore, due to the proliferation of endpoints and their dynamicity, it is attractive to protect the endpoints from within the cloud. A second benefit for this approach is a unified view over the threats, which can lead to better response times in the case of a new type of attack. However, SecaaS cannot achieve complete cloud security provisioning – there are still security measures which must be taken locally by each customer or

their MSSP [9]. As observed in [7], there are four security mechanisms that could be provided from the cloud:

A. Email Filtering

Email was the first type of SaaS moved to the cloud. Protecting the email service from where it resides (and not from each endpoint) was regarded as normal – so this was also the first type of cloud SecaaS. This type of protection can employ multiple engines to scan for malware, spam and phishing threats carried in emails. Furthermore, eliminating these threats in the cloud reduces the bandwidth used with email, the load on the customer's email servers and helps client's existing anti-malware solutions if there are any.

B. Web Content Filtering

Having traffic routed through proxy servers of SecaaS providers allows for cleaning of web traffic while keeping the delay to a minimum. The capabilities for this type of service are quite advanced, from general URL filtering, HTTP header information screening, analysis of page content and embedded links to an advanced reputation system that is continuously being updated by users' traffic. Additionally, it could also involve monitoring of the outgoing traffic in order to block possible information leaking, such as IDs, credit card information or intellectual property.

C. Vulnerability Management

This represents the evolution of MSS. In an effort to alleviate the problems related to vulnerable VMs due to the shared environment, client vulnerabilities are managed from the cloud through the use of complex systems such as application firewalls (e.g. between VMs), virtual IDSs, cloud antivirus and VPNs over VMs. A recent study revealed that CloudAV (a cloud antivirus solution) obtained a 35% improvement over endpoint-residing antivirus solutions. The centralization of vulnerability management activities could enable interoperation of increasingly complex security measures.

D. Identity-as-a-Service (IDaaS)

The identity management can be moved into the cloud too with simple adjustments to existing enterprise architectures. Such a move promises to be beneficial for customers because there is a centralized point from where identification requests are managed. This ensures standard compliance, interoperability in federations (using SSO) and removes the complexity of IAAA mechanisms from customers' management.

The SecaaS technology is still evolving – with multiple approaches. Some providers offer solutions that require traffic to be directed through their systems, whereas others offer specialised SaaS solutions which can be installed over the services that are already in use from other CSPs.

Currently, SecaaS providers come from two backgrounds: they are either new entrants specializing in niche zones and proposing novel solutions, or well-established anti-malware companies which extend their services in the cloud world. With time, the trust relationship

between customers and CSPs on the one side and the SecaaS providers on the other side is expected to improve. This will lead to a proliferation of this type of services and a better penetration of these solutions.

V. CONCLUSIONS

Over the next decade, CSPs may expand the types of services delivered; new component technologies may create shifts in platform standardization, commoditization, and differentiation; and cloud adopters may gravitate to specific providers for certain workloads depending also on the prices. In each of these cases however, trusted security solutions and services are going to play primary role towards the wider adoptions in this fundamental paradigm shift in computing and data storage services.

In this paper we argue that substantial improvements in cloud security can be achieved by adopting the autonomic computing approach for developing and building smart cloud systems.

REFERENCES

- [1] A. Chonka, Y. Xiang, W. Zhou, and A. Bonti, "Cloud Security Defence to Protect Cloud Computing Against HTTP-DoS and XML-DoS Attacks", *Journal of Network and Computer Applications*, vol. 34, 2011, pp. 1097-1107.
- [2] Cloud Security Alliance, "Domain 10: Guidance for Application Security", Version 2.1, July 2010, <https://cloudsecurityalliance.org/guidance/csaguide-dom10-v2.10.pdf>
- [3] Cloud Security Alliance, "Security Guidance for Critical Areas of Focus in Cloud Computing", Version 3.0, November 2011, <https://cloudsecurityalliance.org/guidance/csaguide.v3.0.pdf>
- [4] V. Getov and S. Srinivasan, "From Invisible Grids to Smart Cloud Computing", *Proc. EuroPar 2010 Workshops, LNCS*, vol. 6586, Springer, 2011, pp. 263-270.
- [5] M. Gregg, "10 Security Concerns for Cloud Computing", White Paper, Global Knowledge, 2010, http://viewer.media.bitpipe.com/1078177630_947/1268847180_5/WP_VI_10SecurityConcernsCloudComputing.pdf
- [6] R. Kissel, M. Scholl, S. Skolochenko, and X. Li, "Computer Security: Guidelines for Media Sanitization", NIST Special Publications 800-88, 2006, http://csrc.nist.gov/publications/nistpubs/800-88/NISTSP800-88_rev1.pdf
- [7] T. Mather, S. Kumaraswamy, and S. Latif, "Cloud Security and Privacy: An Enterprise Perspective on Risks and Compliance", O'Reilly, 2009.
- [8] R. Mogull, "Cloud encryption use cases", SearchCloudSecurity.com, 2011, <http://searchcloudsecurity.techtarget.com/tip/Cloud-encryption-use-cases>
- [9] J. Oberheide, E. Cooke, and F. Jahanian, "CloudAV: N-Version Antivirus in the Network Cloud", *Proc. 17th Security Symposium*, 2008, pp. 91-16; <http://dl.acm.org/citation.cfm?id=1496718>
- [10] OWASP, "Top Ten Most Critical Web Application Security Risks", 2010, <http://owasptop10.googlecode.com/files/OWASP%20Top%200%20-%202010.pdf>
- [11] B. Prince, "IBM Discovers Encryption Scheme That Could Improve Cloud Security, Spam Filtering", 2009, <http://www.eweek.com/c/a/Security/IBM-Uncovers-Encryption-Scheme-That-Could-Improve-Cloud-Security-Spam-Filtering-135413/>
- [12] M. Savage, "Cloud compliance, cloud encryption top enterprise security concerns", SearchCloudSecurity.com, 2011, <http://searchcloudsecurity.techtarget.com/news/2240031767/Cloud-compliance-cloud-encryption-top-enterprise-security-concerns>

- [13] SecaaS WG, Cloud Security Alliance, "Defined Categories of Service 2011", Version 1.0, Oct. 2011, https://cloudsecurityalliance.org/wp-content/uploads/2011/09/SecaaS_V1_0.pdf
- [14] S. Srinivasan and V. Getov, "Navigating the Cloud Computing Landscape – Technologies, Services, and Adopters", IEEE Computer, vol. 44(3), 2011, pp. 22-23.
- [15] C. Wang, Q. Wang, J. Li, K. Ren, and W. Lou, "Enabling public verifiability and data dynamics for storage security in cloud computing", Proc. 14th European Conference on Research in Computer Security, ACM Press, 2009.
- [16] C. Wang, Q. Wang, K. Ren, and W. Lou, "Ensuring Data Storage Security in Cloud Computing", Proc. 17th IWQoS, IEEE Xplore, 2009, pp. 1-9.
- [17] S.S. Yau and H.G. An, "Software Engineering Meets Services and Cloud Computing", IEEE Computer, vol. 44(10), October 2011, pp. 47-53.