



Title:	Deliverable D3.3 Requirement Analysis for Secure Cloud Services over IPv6	Document Version: 3.2
---------------	--	-------------------------------------

Project Number: 297239	Project Acronym: GEN6	Project Title: Governments ENabled with IPv6
----------------------------------	---------------------------------	--

Contractual Delivery Date: 30/04/2012	Actual Delivery Date: 04/02/2013	Deliverable Type* - Security**: R – PU
---	--	--

* Type: P - Prototype, R - Report, D - Demonstrator, O - Other
 ** Security Class: PU- Public, PP – Restricted to other programme participants (including the Commission), RE – Restricted to a group defined by the consortium (including the Commission), CO – Confidential, only for members of the consortium (including the Commission)

Responsible and Editor/Author: Gabriela Gheorghe	Organization: UL	Contributing WP: WP3
--	----------------------------	--------------------------------

Authors (organisations):

Gabriela Gheorghe (UL), Onur Bektaş (ULAKBIM), Emre Yüce (ULAKBIM), Antonio Skarmeta (UMU), Pedro Martinez (UMU), Kamil Seyhan (TURKSAT), Sami Yenice (TURKSAT), Alvaro Vives (Consulintel), Jordi Palet (Consulintel).

Abstract:

This deliverable presents an analysis on the security requirements of Cloud services in public infrastructures when they are IPv6-enabled, along with a description of the Luxembourg pilot during IPv6 migration.

Keywords:

Cloud, IPv6, security requirements, public sector, emergency incident management, e-gov infrastructures.

Revision History

The following table describes the main changes done in this document since its creation.

Revision	Date	Description	Author (Organization)
V0.1	11/04/2012	Document creation	Gabriela Gheorghe (UL)
V0.2	20/04/2012	Changed structure and edited contents	Gabriela Gheorghe (UL)
V0.3	20/04/2012	Integrated input from Onur Bektaş (ULAKBIM) and Emre Yüce (ULAKBIM)	Gabriela Gheorghe (UL)
V0.4	21/04/2012	Contribution on cloud security challenges	Antonio Skarmeta and Pedro Martinez (UMU)
V0.5, V0.6	23/04/2012	Edited and added more content	Gabriela Gheorghe (UL)
V1.0	24/04/2012	Edited and added content, also from TURKSAT	Gabriela Gheorghe (UL)
V1.1	25/04/2012	Proof reading and internal review	Emre Yüce (ULAKBIM), Murat Soysal (ULAKBIM), Onur Bektaş (ULAKBIM)
V1.2	26/04/2012	Added comments from Latif Ladid, and minor changes according to Onur Bektaş's and Kamil Seyhan's feedback.	Gabriela Gheorghe (UL)
V1.3	26/04/2012	Pre-Final version.	Gabriela Gheorghe (UL)
V1.4	14/05/2012	Pre-Final version, comments of Fraunhofer addressed.	Gabriela Gheorghe (UL)
Final –M6	24/05/2012	Final version M6	Gabriela Gheorghe (UL)
V2.0	27/05/2012	Final review.	Jordi Palet (Consulintel)
V3.0	19/09/2012	Updated version to include clear description of the Luxembourg pilot.	Gabriela Gheorghe (UL)
v3.1	25/09/2012	Review	Uwe Kaiser (Fraunhofer)
Final M12	04/02/2013	Final version M12.	Gabriela Gheorghe (UL)

Disclaimer

The GEN6 project (number 261584) is co-funded by the European Commission under the ICT Policy Support Programme (PSP) as part of the Competitiveness and Innovation framework Programme (CIP). This document contains material that is the copyright of certain GEN6 partners and the EC, and that may be shared, reproduced or copied “as is”, following the Creative Commons “Attribution-NonCommercial-NoDerivs 3.0 Unported (CC BY-NC-NC 3.0) licence. Consequently, you’re free to share (copy, distribute, transmit) this work, but you need to respect the attribution (respecting the project and authors names, organizations, logos and including the project web site URL “<http://www.gen6.eu>”), for non-commercial use only, and without any alteration, transformation or build upon this work.

The information herein does not necessarily express the opinion of the EC. The EC is not responsible for any use that might be made of data appearing herein. The GEN6 partners do not warrant that the information contained herein is capable of use, or that use of the information is free from risk, and so do not accept liability for loss or damage suffered by any person using this information.

Executive Summary

The administration infrastructure of the government network in Luxembourg needs safe tools to achieve the transition to IPv6. UL's aim is to pilot the transition to IPv6 of an innovative Cloud-based tool that can be useful in testing IPv6 behaviour of services within the government administration network. This approach complements TUBITAK ULAKBIM's IPv6 tool that monitors the IPv6 availability of HTTP, DNS and SMTP services, but only based on domain names. Luxembourg's pilot uses an internal (or private) Cloud deployment to achieve internal interoperability and fault management monitoring, which can help survey the IPv6 readiness of any node internal to the network. In this endeavour, UL is collaborating with two governmental institutions: Centre de Communications du Gouvernement (CCG Luxembourg) that works with the network of all ministries and government agencies in the country, and the Ministry of Commerce of Luxembourg.

The main contribution of the pilot is the transition from an IPv4 Cloud to an IPv6/Dual-Stack Cloud deployment. This deployment is of an in-house Cloud that constitutes a completely contained environment, with heavy use of virtualization. The full transition to IPv6 is poorly documented for most open-source Cloud distributions; hence GEN6 can have a valuable impact in the open source Cloud community. Since the purpose of the Cloud environment in discussion is of monitoring IPv6 interoperability, security, and resilience, a secondary contribution of this pilot can be to develop a test suite for IPv6 connectivity and operation in the internal government cloud network. UL plans to validate the efficiency of this test suite and make the results available to governments willing to take on this approach.

Further, this deliverable aims to analyse the requirements for secure Cloud-based infrastructures when transitioning to IPv6. In order to do so, it follows some common lines from several requirements engineering methodologies. First, it is important to understand generic security notions in Cloud-based systems: the typical stakeholders, their conflicting interests, the different regulations that affect operation, security capabilities in Cloud systems, the risks and benefits that Cloud computing brings to the security of an enterprise. The different regulations that affect cross-institutional IT systems are of particular interest, since they influence the security goals and expectations of the typical stakeholders, and hence the security controls that can be ultimately used.

Moreover, the deliverable discusses requirements from two other case studies of GEN6 partners: ULAKBIM's academic network, and TURKSAT's e-Gov system. Analysing requirements on security for existing IPv6-enabled networks and IPv6-enabled e-Gov systems, gives a better picture over the set of guidelines to follow for a safe migration that can apply for a Cloud environment as well.

297239	GEN6	D3.3: Requirement analysis for Secure Cloud Services over IPv6
--------	------	--

Table of Contents

1.	Introduction.....	10
1.1	Context and Motivation	11
1.2	Contents Overview	12
2.	Terminology and Approach	14
	<i>Figure 1 Security Requirements Taxonomy by Open Security Architecture</i>	<i>14</i>
3.	State of the Art in Securing Cloud Infrastructures	17
3.1	Stakeholders and Regulatory Challenges	17
3.2	Security in Cloud Computing	19
3.2.1	Cloud Computing Benefits for Security	19
3.2.2	Cloud Computing Security Risks	20
3.3	Trust and Reputation in Cloud Computing	22
4.	Luxembourg's CC Pilot	24
4.1	UL's CC Environment	24
	<i>Figure 2 High-level architecture of the CC environment.....</i>	<i>25</i>
	<i>Figure 3 Pilot configuration</i>	<i>26</i>
4.2	Transition of CC Environment to IPv6	27
	<i>Figure 4 OpenStack Cloud components and IP flows [OpenStack12]</i>	<i>28</i>
4.3	Testing IPv6 Operation and Security	29
4.3.1	Motivation	29
4.3.2	Test Suite for IPv6 Readiness and Security	30
4.4	Usage and Dissemination	30
4.5	Related EU projects	31
4.5.1	SECRICOM	31
4.5.2	U2010	32
5.	Case Studies Analysing Security in IPv6 Infrastructures	33
5.1	ULAKBIM's National Academic Research and Education Network	33
5.1.1	ULAKNET Security Features	33
5.1.1.1	ULAKNET Blackhole Attack Detection System	33
5.1.1.2	OLTA Incident Tracking System	34
5.1.1.3	ULAKNET NetFlow Statistics	35

5.1.2	Worm Propagation in IPv6 Networks.....	35
5.1.3	Security Requirements of the Turkish IPv6 Traffic Exchange Point	36
5.2	TURKSAT's e-Government Infrastructure.....	36
5.2.1	Security Concerns for the e-Gov Gateway	37
5.2.2	IPv6 and IPv4 for Gateway Security	38
6.	<i>Requirements for Secure Cloud Services over IPv6</i>	39
6.1	Provisioning IPv6 Equipment	39
6.2	Potential Security Issues with IPv6 Transition Technologies	40
6.3	Known IPv6 Security Issues and Mitigations	43
6.4	Management Aspects for secure Cloud Services	45
6.4.1	Cloud Time Services	45
6.4.2	Identity Management	46
6.4.3	Access Control Management	46
6.4.4	System and Network Auditing.....	47
6.4.5	Security Monitoring.....	47
6.4.6	Incident Management	48
6.4.7	Security Testing and Vulnerability Remediation	48
6.4.8	System and Network Controls.....	49
6.4.9	Configuration Management	49
7.	<i>Conclusions.....</i>	50
8.	<i>References.....</i>	51

Figure Index

Figure 1 Security Requirements Taxonomy by Open Security Architecture	14
Figure 2 OpenStack Cloud components and IP flows [OpenStack12]	28
Figure 3 High-level architecture of the CC environment	25
Figure 4 Pilot configuration	26

Table Index

No table of figures entries found.

1. INTRODUCTION

Integrating IPv6 in governmental Cloud services allows these services to benefit from the inherent addressing advantages of IPv6. The efficiency of communication flows and application scalability are certainly going to improve, however the security production deployment is not yet mature enough to warrant a smooth transition to IPv6 with the security assurance expected from public sector Cloud services.

Cloud computing imposes several security challenges that range from data confidentiality and privacy (e.g., should the infrastructure provider be allowed to inspect communication flows in case of potential security incidents?) to compliance with contracts and agreements with ambiguous terms and conditions (e.g., all communication must be encrypted – with no mentioning of how and where to encrypt it). Broadly, most of these security challenges stem from the fact that perimeter security concepts no longer work naturally in Cloud services. Yet, it is left to the various service providers to secure data and operations in their application domain; what is worse is that usually actors have their own understanding of the meaning of “secure” and how to achieve “secure”. To take matters further, the lack of mature security solutions of an IPv6-enabled infrastructure is likely to add to those of the Cloud services that are deployed on that infrastructure. Clearly, the lack of solid IPv6 firewall deployment makes possible some IPv6 attacks such as Denial of Service, packet drop and redirection; these attacks can affect the proper functioning of the Cloud application that is using an IPv6 network.

For a government infrastructure, security assurance is essential for its users because of its national consequences. It follows that on this background, it is essential to study the security implications of using the IPv6 technology in tandem with Cloud services. In this vein, the purpose of this deliverable is to elicit requirements for the security of Cloud services over IPv6, and to describe Luxembourg’s Cloud-based pilot on the way to a safe IPv6 migration.

1.1 Context and Motivation

The adoption of Cloud Computing by governments in the world has started since 2010¹. The biggest advantage of moving services to the Cloud is the massive lowering of costs to develop and maintain software, as well as energy and administration costs. Because users of the Cloud no longer need to invest in their in-house infrastructure, they can delegate the administration costs to specialized third-parties, and the savings can be substantial. In the US, with support from NIST, examples of public agencies that have started the process already include the state of Utah (US), with estimated ²savings of the order of billions of dollars, and the U.S. Department of Transportation's Federal Aviation Administration started in 2011 a Cloud initiative ³estimated to save 100 million dollars per year. In the UK, the government Cloud movement⁴ aims to save 50 million pounds. Conversely, in continental Europe, moving to the Cloud has started in the private sector with large companies such as CERN, European Space Agency, SAP, and a number of European banks. As per a recent report⁵ of the European Commission, organisations moving to the Cloud are likely to cut down between 20 and 50% of IT spending.

It has been claimed that European data protection laws are the main hurdle in the face of adopting a computing paradigm essentially based on delegating the running of the software lifecycle. This claim alludes to a regulatory and legislative framework that applies solely to the financial sector and that in some countries imposes certain restrictions on the distribution and storage of financial data abroad. Nevertheless, outsourcing of administration (even fund administration) is already a well-spread practice for today's organisations. Moreover, national or cross-border regulations on data protection do not affect the possibility to outsource IT functions. It can hence be inferred that by using the Cloud model, governments do not lose their autonomy. Without particular concerns to financial data, Luxembourg's pilot in GEN6 aims to offer governments a Cloud-based framework to be more in control of their IT infrastructure rather than the data that they operate on. Being in control of the infrastructure covers important aspects such as an internal safe migration to IPv6 of the network as well as the applications, and the internal assessment of existing fault and security management mechanisms. The Luxembourg pilot can be used as an extensible tool to address these issues.

¹ U.S. General Services Administration, <http://www.gsa.gov/portal/content/208417>

² <http://www.nascio.org/awards/nominations/2010/2010UT6-Nascio%20Utah%20Cloud%202010.pdf>

³ Federal Aviation Administration [initiative](#) on a Software-as-a-Service Cloud

⁴ <http://gigaom.com/2012/02/01/britain-unleashes-gov-uk-its-google-for-government/>

⁵ <http://euobserver.com/news/117695>

1.2 Contents Overview

This deliverable follows the terminology used by Open Security Architecture (OSA) [OSA]. OSA is an organisation that provides an enterprise architecture framework based on open source principles and methodologies. One of the major goals of OSA is to put together the requirements of various standards in IT security, governance and related legislation. The OSA framework is based on the patterns of the security controls in NIST Special Publication 800-53 entitled “Recommended Security Controls for Federal Information Systems and Organisations”. Chapter 2 of this deliverable explains the OSA terminology and the approach that we follow in eliciting requirements for secure Cloud services over IPv6.

Based on the OSA terminology, Chapters 3 and 4 examine the security needs that the state of the art work has already described in Cloud computing, and that can be seen in the context of the transition to IPv6. In order to understand these security needs, one has to consider the different stakeholders (section 3.1), the current regulatory constraints, and how some of these constraints are addressed with current Cloud security tools (section 3.2). Using the Cloud computing paradigm comes with some benefits for IT security that are described in section 3.2.1, while the Cloud computing security risks are described in Section 3.2.2. Section 3.3 gives an overview of the existing efforts in trust and reputation for Cloud systems.

Chapter 4 describes Luxembourg’s Cloud Computing environment: what it is for and what part of it supports or can migrate to IPv6. The general migration of a CC environment based on open-source software is so far undocumented in open-source Cloud distributions, and actually performing it would greatly help both the government to use more Clouds over IPv6, and Cloud communities to adopt IPv6. An important aspect that section 4.3 mentions is how the pilot results can be used and disseminated within the framework of GEN6 and afterwards. Chapter 4 also enumerates some recent projects where UL was involved, that related to IPv4 and IPv6 technologies, and how their approaches are different or similar.

Concrete requirements on security related to other existing IPv6-enabled systems are given in Chapter 5. These requirements stem from ULAKBIM’s National Academic Research and Education Network, and TURKSAT’s e-Government infrastructure. First, ULAKBIM’s national network connects universities and research institutions in Turkey, and along with this network, ULAKBIM has also set up an IPv6 traffic exchange point to help Turkish ISP’s in the IPv6 transition. Second, TURKSAT’s e-Government infrastructure coordinates cross-institutional efforts in a secure public system, and uses an e-Gov gateway that helps integrate new services. In all, these two individual systems come with their own security concerns. Such concerns stem from the use of IPv6 transition technologies, as well as the need for integration of services across several domains. Examining these concerns will help us understand what are the risks and goals in transitioning to IPv6 and at the same time maintaining security assurance.

With the above examples in mind, Chapter 6 describes a set of higher-level requirements for securing Cloud services over IPv6. These requirements exist in the community and literature and concern vulnerabilities and mitigations for IPv6 technology, as well as the management of various essential aspects in a public infrastructure, such as time services, identity services, access control, network auditing services, incident handling and remediation services. Further UL efforts should investigate how these security requirements blend with IPv6 security technologies, for the Luxembourg pilot.

2. TERMINOLOGY AND APPROACH

Security requirements are those needs that a system must satisfy in order to achieve its security goals. They are sometimes termed “quality” requirements, and have been separated by OSA into several categories⁶:

- **Secure functional requirements** specify what shall not happen in the target system; they are usually derived from misuse cases.
- **Functional security requirements**, that are security services to be offered by a system.
- **Non-functional security requirements**, that relate to architecture requirements, e.g. “robustness” or “scalability”.
- **Secure development requirements**, that describe required development activities so that in the end, the target application would not present (many) vulnerabilities. Examples can be: “data classification”, “coding style”, “test methodologies”, etc.




Figure 1 Security Requirements Taxonomy by Open Security Architecture⁷

In this deliverable, OSA’s first two types of requirements are concentrated together: secure functional requirements, and functional security requirements. For ease, hereafter these two

⁶ OSA requirements http://www.opensecurityarchitecture.org/cms/definitions/it_security_requirements

⁷ Taxonomy to be found at http://www.opensecurityarchitecture.org/cms/definitions/it_security_requirements

requirements will simply be referred as “security requirements”.

Figure 1 gives OSA’s view over how to raise security requirements, and how they are later incorporated into the system to be secured. OSA’s approach resembles the approach of more established methodologies for eliciting security requirements such as Security Quality Requirements Engineering (SQUARE) [Mead05], and Secure Tropos⁸. The SQUARE process of eliciting security requirements has been developed by Carnegie Mellon University in 2005, and it consists of nine basic steps:

1. Agree on definitions (done by the stakeholders and the requirements team).
2. Identify security goals, business drivers, policies and procedures (done by the stakeholders and the requirements team).
3. Develop artifacts (i.e., scenarios, misuse cases, models, forms) to support security requirements definition (done by the requirements engineer).
4. Perform risk assessment (done by risk engineer, stakeholders, requirements engineer).
5. Select elicitation techniques, like goals, definitions, expertise of stakeholders, organizational style, etc. (done by the requirements engineer).
6. Elicit security requirements, by means of interviews, surveys, model-based analysis, checklists, etc. (done by the stakeholders and the requirements team).
7. Categorize requirements (done by requirements engineers, and other specialists).
8. Prioritize requirements (done by the stakeholders and the requirements team).
9. Requirement inspection, by means of peer reviews or special methods (done by an inspection team).

Secure Tropos is another relevant methodology, aimed at identifying the security needs of the stakeholders, the security entities that guarantee the satisfaction of those needs, and assigning capabilities to the system to help satisfy the identified requirements. There are four main steps in Secure Tropos:

1. Early requirements analysis phase, where high-level requirements are identified, along with the secure goals and entities that can satisfy constraints.
2. Late requirements analysis phase, where more refined security constraints are imposed on a “system to be”.

⁸ Secure Tropos, <http://securetropos.org/>

3. Architectural design phase, analysing any possible security constraint and secure entity that can be introduced by new actors in the system.
4. Detailed design phase, that designs the entities identified in the previous phase, with the help of an UML-based tool.

In line with these methodologies, we need to elicit security requirements, starting with identifying stakeholders, actors and goals and the concrete system to be transitioned (Chapter 4). Related requirements from GEN6 partner scenarios (ULAKBIM, TURKSAT) are given in Chapter 5). With the constraints of the particular cases in mind, as well as with hints from some existing literature reports, a list of generic security requirements that are relevant to the project were compiled in Chapter 6. What makes these requirements realistic is their tight connection with the legislative context of government-managed infrastructures. This context is described in Sections 3.1 and 3.2.

3. STATE OF THE ART IN SECURING CLOUD INFRASTRUCTURES

This section presents state of the art issues that are to be considered when developing Cloud security infrastructures. First, the stakeholders and the main challenges for public infrastructures are defined. Second, the state of the art in security solutions for these systems is overviewed, and then the trust and reputation issues are discussed.

3.1 Stakeholders and Regulatory Challenges

In public infrastructures that use Cloud computing, the analysis will focus on the following two categories of actors:

- Governments and governmental institutions that, in their cooperation, use a common set of Cloud applications or services. In terms of security, they are primarily interested in how classified or confidential data is being handled across domains. Other important needs are the anonymity of communication flows, and the compliance of their own systems with existing laws and regulations that relate to computer or telecommunication security (e.g., European directives on data privacy).
- Cloud infrastructure and service providers can rent systems or services to other entities (e.g., governments). These providers need to prove their systems are safe and preserve customer privacy according to laws and regulations. These providers need to prove:
 - Compliance checking tools, used to measure security parameters of their services or systems.
 - Reporting and accounting tools.

In terms of challenges, security frameworks for Cloud services and infrastructures need to deal with several aspects that complicate security enforcement and evaluation:

- The existence of different legal and regulatory constraints that govern the way data and systems should function. These constraints often have specific country-based differences, and this contrasts with the cross-country nature of Cloud computing.
- As a consequence of the previous point, when data and systems are operated across national borders and across enterprises, it can be a challenge to decide the laws and regulations apply. For instance, there can be requirements on data protection (e.g., preserving client anonymity), that may clash with usage regulations that impose accountability. To make matters more complicated, it is sometimes difficult for legal authorities to decide on responsibility when security incidents occur, especially because the same data can be owned by several parties at the same time.
- In the Cloud there are different types of data being handled: there is control data (e.g., configuration parameters, monitoring data and logs) and there is application data (i.e.

the data that the Cloud application works with, directly). From a security perspective, both should be treated in such a way as to respect the confidentiality of the issuing party. However, this data should be handled differently depending on its life-cycle and volatility.

- Outsourcing is a well-known practice in Cloud computing, and it can significantly blur the lines of security responsibilities and assurance. Outsourcing security services to a specialized third-party exposes application managers to the clash between data disclosure and data security.

The bulk of security requirements for the Cloud domain stem from IT security standards and regulations. The target of these regulations is to protect against misuse of private or confidential data or of information systems that process such data. As briefly overviewed in [Ghe11], in Europe there are several security and privacy regulations for the IT sector:

- Directives 95/46/EC and 97/66/EC aim to protect personal data without the subject's consent, either in healthcare or in the telecommunications sector.
- Directive 2002/58/EC concerns the way personal data is processed by services in the electronic communications sector so that users' data receives the same protection level irrespective of the technologies use. This directive refers to authorization to access data, protecting data against loss and tampering, and the enforcement of a security policy to process sensitive data.
- Directive 2006/24/EC, also known as the "Data Retention Directive", stipulates a maximum retention period of between 6 and 24 months for users' telecommunication data. The communications data that needs to be retained is not the content of the communication but rather the duration, and identification information of the parties. This regulation has been subject to controversy.
- In the healthcare domain, Health Level Seven (HL7) defined a series of standards for the electronic exchange of medical data, of which there is emphasis on secure bidirectional communication flows that also ensure non-repudiation.
- In Italy, the Italian Data Protection Code (IDPC) has been adopted since 2004, and concentrates on electronic data protection in Italy. One of its main features is the principle of data minimization (Section 3), in that public data about individuals should be used as much as possible, while personal data, only when absolutely required.

Up to now, the implementation of such practices and regulations is left to the application providers, and to our knowledge there is no methodology of how to make the application compliant with these constraints. Nevertheless, the task of the security auditors is to ensure that some key security elements have been met in the overall design of the audited system.

3.2 Security in Cloud Computing

Security is a requirement imposed by all network architectures and frameworks. In the particular case of Cloud Computing, more complex security requirements are imposed and it becomes more sensible due to the multi-tenancy, the difficulties to ensure trust, the difficulties to apply data encryption, and the legislation compliance [Subashini11].

Strongly related to the issues concerning legislation and data distribution is the concern of data protection and other potential security holes arising from the fact that the resources are shared between multiple tenants and the location of the resources being potentially unknown. In particular, sensitive data or protected applications are critical for outsourcing issues. In some use cases, the information that a certain industry is using the infrastructure at all is enough information for industrial espionage.

Whilst essential security aspects are addressed by most tools, additional issues apply through the specifics of cloud systems, in particular related to the replication and distribution of data in potentially worldwide resource infrastructures [Wang09]. However, the data protection mechanisms should addresses legislative issues with respect to data location.

In addition, many usages of cloud systems and the variety of cloud types imply many different security models and requirements by the user. Classical authentication models may be insufficient to distinguish between the aggregators/vendors and the actual user. In particular, that is the case of infrastructure-as-a-service (IaaS) cloud systems, where the computational image may host services that are made accessible to other users.

In other cases of aggregation and resale of cloud systems, the mix of security mechanisms may not only lead to problems of compatibility, but may also lead to the user not to trust in the model due to lack of insight [Shubert10].

3.2.1 Cloud Computing Benefits for Security

The Cloud Computing technology permits to aggregate security implementations. All kinds of security measures are cheaper when implemented on a larger scale. Therefore the same amount of investment in security buys better protection [Catteddu09]. This includes all kinds of

defensive measures such as filtering, patch management, hardening of virtual machine instances and hypervisors, etc. Other benefits of scale include: multiple locations, edge networks (content delivered or processed closer to its destination), timeliness of response, to incidents, threat management.

Although this concentration of resources undoubtedly has some disadvantages for security, it has the obvious advantage of cheaper physical perimeterisation and physical access control (per unit resource) and the easier and cheaper application of many security-related processes.

That said, security is a priority concern for many Cloud Computing customers. They will make buying choices on the basis of the reputation for confidentiality, integrity, and resilience of a provider as well as the security services it offers. This is a strong driver and market differentiator for Cloud Computing providers to improve security practices.

Cloud Computing providers should offer standardised, open interfaces to manage security services. This creates a more open and readily available market for secure services and security oriented services. Also, the ability of the cloud provider to dynamically reallocate resources for filtering, traffic shaping, authentication, encryption, etc., to defensive measures (e.g., against Distributed Denial of Service or DDoS attacks) has obvious advantages for resilience, which is a very important aspect impacting overall security.

For the auditing and evidence gathering, current Cloud Computing approaches, which are heavily based on virtualization, are able to provide dedicated, pay-per-use forensic images of virtual machines which are accessible without taking infrastructure off-line, leading to less down-time for forensic analysis. It also provides more cost-effective storage for logging information, allowing more comprehensive logging without compromising performance.

Default virtual machine images and software modules used by customers can be pre-hardened and updated with the latest patches and security settings according to fine-tuned processes. IaaS Cloud APIs also allow snapshots of virtual infrastructure to be taken regularly and compared with a baseline. Updates can be rolled out much more rapidly across homogeneous platforms than in traditional client-based systems that rely on the patching model.

3.2.2 Cloud Computing Security Risks

We have discussed the benefits of Cloud Computing regarding security, but it also has risks that should be considered. Thus, it has “unique attributes that require risk assessment in areas such as data integrity, recovery, and privacy, and an evaluation of legal issues in areas such as e-discovery, regulatory compliance, and auditing” [Gartner08]. Here, Gartner considers some cloud specific security risks.

First, Cloud Computing technologies have data protection issues for both cloud customers and

providers. In some cases, it may be difficult for the cloud customer (in its role as data controller) to effectively check the data handling practices of the cloud provider and thus to be sure that the data is handled in a right and secure way. Also, a similar issue is related to the insecure or incomplete data deletion. A request to delete a cloud resource is made, as with most operating systems, may not result in the true wiping of the data. Adequate or timely data deletion may also be impossible, either because extra copies of data are stored but are not available, or because the disk to be destroyed also stores data from other clients. In the case of multiple tenancies and the reuse of hardware resources, this represents a higher risk to the customer than with dedicated hardware.

To overcome this risk, Gartner recommends users to get as much information as they can about the people who manage their data, since sensitive data is processed outside the enterprise making outsourced services a ground likely to permit bypassing the "physical, logical and personnel controls" that are usually in place over in-house programs. Since user's data is released to the cloud leaving the protection sphere of the data owner, it is important to consider and increase the authentication and administrative controls first, but afterwards, a monitoring of user and user application activities is needed to detect and react as soon as possible in face of possible attacks.

In cloud infrastructures, the client necessarily cedes control to the Cloud Computing provider on a number of issues which may affect security. The client abstracts from infrastructure management and complexity but loses its governance. At the same time, SLAs may not offer a commitment to provide such services on the part of the cloud provider, thus leaving a gap in security defences. In relation to [Gartner08], it takes this issue into account but still makes cloud customers the ultimate responsible for the security and integrity of their own data, even when it is held by an outsourced service provider.

Multi-tenancy and shared resources are defining characteristics of cloud computing. But it implies to be risked to isolation failure. This is the failure of mechanisms separating storage, memory, routing, and even reputation between different tenants (so-called guest-hopping attacks). However it should be considered that attacks on resource isolation mechanisms (e.g., against hypervisors) are still less numerous and much more difficult for an attacker to put in practice compared to attacks on traditional operating systems.

From the management point of view, customer management interfaces of a Cloud provider are accessible through the Internet and mediate access to larger sets of resources (than traditional hosting providers) and therefore pose an increased security risk, especially when combined with remote access and web browser vulnerabilities. To this extent, in [Gartner08], it is stated that investigating inappropriate or illegal activity may be impossible in cloud computing, since logging and data for multiple customers may be co-located and may also be spread across an ever-changing set of hosts and data centres.

That said, the monitoring is the main mechanism used to protect security but, at the same time as users activity is monitored, their privacy needs to be preserved. Very recently, the W3C has published two first drafts for standards that enable users to set up their preferences about on-line tracking. They deal with “Tracking preference expression (DNT)” [W3c11a] and “Tracking compliance and Scope specification” [W3c11b]. DNT defines mechanisms for users to express cross-site tracking preferences and for sites to indicate whether they honour these preferences while tracking “Compliance and Scope Specification defines the meaning of a “Do Not Track” preference and sets out practices for websites to comply with this preference. Also, in [W3c11a] we find the balanced relation that should exist between users’ privacy and context-aware, personalized web services. It declares:

“We know there are many types of users. Some eagerly welcome the benefits of personalized web services, while others value their privacy above all else. Do Not Track puts users in control, so they can choose the trade-offs that are right for them”.

Finally, there is the risk that the cloud customers may be locked to a specific Cloud Computing provider because today not all providers offer standard/open tools, procedures, data formats, or services interfaces that could guarantee data, application, and service portability. This can make it difficult for the customer to migrate from one provider to another or migrate data and services back to an in-house IT environment. This introduces a dependency on a particular provider for service provision, especially if data portability is not enabled.

3.3 Trust and Reputation in Cloud Computing

Without the assistance of an accurate trust and reputation management, users are “blind” when selecting the most reliable cloud-service to access. Such management performed in an appropriate way might help to find out which cloud-services are actually benevolent and which might have a malicious or fraudulent behaviour, isolating this way the latter.

Trust and reputation management is a body of research that has been recently widely investigated in the literature [Kamvar03, Gupta08, Zhou03, Sun05]. Successful examples are found in the eBay on-line auction system, where buyers and sellers can rate each other after each transaction, and the EigenTrust system [Kamvar03] that allows peers in a P2P environment to decrease the number of downloads of inauthentic files based on a unique global trust value assigned to each peer. However, all the mentioned works are used in P2P environments for computing the reputation of peers for file sharing purposes. Typically, in these systems there is no contract used to regulate relationships among peers; the system is “flat” and the reputation is not controlled and differentiated on the basis of the type and requirements specified for the service provision.

The trust problem is further exacerbated in Cloud Computing systems, where service providers

and customers are oblivious to each other. As such, defining which provider is trusted to produce or receive information is not a trivial task. The works found in [Fiege04, Bacon10] strive to address this issue either modifying the internal architecture of an existing middle-ware, by organizing providers, customers, and brokers in small overlays, or using Information Flow Control (IFC) techniques and Role-Based Access Control (RBAC) to monitor the data dissemination within a large scale (cross-border) system. However, these solutions restrict the information flow just to authorized parties, but they do not define a trust mechanism to infer the level of trustiness of cloud publishers, customers, and clients.

Meanwhile, a first approach made in [Vu05] suggests applying trust and reputation management to efficiently filter dishonest ratings when evaluating the Quality-of-Service (QoS) of providers. It proposes to use such certified QoS as the criterion for selecting one or another provider.

In [Malik09] a method to address a certainly important although commonly neglected issue regarding trust and reputation management can be found: bootstrapping the system (i.e., dealing with newcomers with lack of historical information). It presents different mechanisms in order to bootstrap the reputation of newcomers in a service-oriented environment in a fair and accurate fashion.

In turn, [Hang09] presents a trust model built upon Bayesian networks aimed to help users to capture the dynamism from not only non-functional QoS properties but also service composition in service-oriented environments.

Finally, a survey of trust and reputation based web service selection is presented in [Wang07]. It proposes a typology to classify those selection mechanisms from three different aspects, namely: centralized vs. decentralized, persons/agents vs. resources and global vs. personalized.

4. LUXEMBOURG'S CC PILOT

4.1 UL's CC Environment

The architecture of UL's CC application includes a governmental subsystem running services on a particular platform, a computing cloud, an application monitor, and a testing manager. These elements are described as follows:

- The governmental administrator has a service or platform whose properties, e.g., IPv6-readiness, resilience, security vulnerabilities, and compliance will have to be tested. This administrator would specify some details about its setup and the property to be measured. In the end the government will receive the result of the tests that were run and could take an informed action based on these results.
- Another architectural element is the computing Cloud, consisting of a number of hosts that will run the given tests in a controlled and contained environment – for this reason, we term them 'sandboxes' here. The computing cloud also contains one or more hosts for storing: configuration settings that can be enacted on the sandboxes, software patches downloaded for the Internet and that can be tested in the controlled environment, readily-made system images on which the government service or platform will be duplicated and tested.
- A monitoring application (or application monitor) will collect all local data produced by the sandboxes, data that can be relevant for the analysis. This data can cover system logs, low-level system activity, network traffic among sandboxes, etc.
- A testing manager will make a decision on whether the government service supports or not the properties that the tests were aimed at. The way this decision is made can also be influenced by information provided by the government administrator, e.g., patterns to search for in the system logs that are considered dangerous. This conclusion is then forwarded to the human user.

These high-level elements are depicted in Figure 2 below.




Figure 2 High-level architecture of the CC environment

Figure 3 gives the configuration of the deployed system. The main components are:

- A Cloud gateway server that also performs user management, scheduling, monitoring, and provides the web front end,
- An image server to store the images that, once customized, will be deployed in the virtual machines attached to the backend LAN,
- A data server, or data store, that keeps user data,
- An iSCSI store that keeps cloud objects, virtual machine images, virtual disks for sandboxes/VMs, any common data of VMs,
- Compute hosts on which are one or more virtual machines running. These virtual machines will host the system images with user-run services. Moreover, these virtual machines can be organized in virtual networks, as shown by the blue dotted buses in Figure 3.




Figure 3 Pilot configuration

The gateway, image service, data store and iSCSI are all connected directly to the backend LAN, which supports both IPv4 and IPv6; yet as of now, the network is using IPv4 addressing. Similarly, IPv4 addressing is also used by the virtual networks. If an Internet end-user wants to access a service in one of the VM1-VM48, he or she will contact the gateway and the gateway will do NAT to allow the access.

Therefore UL plans to reconfigure this Cloud environment to allow IPv6 among the VM nodes. The advantages of IPv6 here would be that the virtual system could scale and face user demand without impacting the gateway, that the user authentication and authorization would be sensibly simplified, and allow transparent user-to-VM interactions.

4.2 Transition of CC Environment to IPv6

Open source distributions of Cloud environments are seldom properly documented when it comes to IPv6 configuration. Two of the most popular distributions - OpenStack ⁹ and OpenNebula ¹⁰ – are claimed to support IPv6 and yet they provide no help in how to configure an IPv4 cloud into actually running IPv6.

To illustrate this point, Figure 4 shows the common components for an OpenStack deployment:

- Standard component categories in the distribution: API servers to be accessed by the Cloud users, network managers, schedulers, object stores, etc. – all connected by a message queue system that it tailored to the distribution,
- Cloud users and internet end-users, both accessing the Cloud environment from the outside, by going through an administrative network and a public network respectively,
- Virtual images of systems deployed internally in the Cloud environment.

Note here that end-users use services provided by the Cloud's virtual guest images. Cloud users are administrators, and they use OpenStack tools and APIs to manage virtual guest images.

While the IP traffic between Cloud users is HTTP based and can be IPv6 depending on the ISP provider, there are three main issues in terms of the platform's current IPv6 support:

- The traffic between internet end-users and virtual guests running on the Cloud is now IPv4 only and it should be upgraded to IPv6,
- The traffic between virtual guests that are running in the Cloud is IPv4, yet even if it is possible to migrate it to IPv6, this process is yet undocumented,
- OpenStack Internal management resources (e.g., the user authorisation system) are not documented in terms of how well they provide support for IPv6.

UL plans to address these issues and in the end deliver a methodology for IPv6 transition that governments can use to migrate in-house CC services, without losing their existing autonomy. Section 4.4 discusses how the methodology and guidelines can be used and disseminated during and after the project lifespan.

⁹ OpenStack Cloud Software, <http://www.openstack.org/>

¹⁰ OpenNebula.org, <http://opennebula.org/>




Figure 4 OpenStack Cloud components and IP flows [OpenStack12]

4.3 Testing IPv6 Operation and Security

4.3.1 Motivation

Governmental IT services are aimed at serving citizens and have to provide a high level of assurance from several points of view. These services are long lived and have to be resilient, must be secure, and must be able to prove that they comply with some standards of quality (e.g., ISO27001, COBIT, and ITIL). Resilience is needed in the face of changing software. For instance, upgrading to a new software version should not normally influence a public service running on the same host. The security of a government service is needed when the service deals with confidential information, or performs confidential or safety-critical processes. Complying with quality standards is equally important when the software belongs to an entity that itself controls other bodies' compliance with regulations.

However, government services have to keep up with the continuous evolution of IT technologies and tools. These services have to adapt to varying conditions: transition to the much needed IPv6, integration of patches for an application or tool that is in use, integration of new applications on existing platforms in use, varying security assurance criteria. These conditions have to reflect in the government infrastructure and yet cause minimal disruptions, since disruptions can impact public safety or privacy. Disruptions are very likely for governmental services when migrating to IPv6, and can be caused both at the network and at the application level.

In the face of unwanted disruption, there is the need to test the government system before making any of the above mentioned changes. For large industrial settings, testing is usually done in-house and manually by highly qualified administrators who need not only know details of the existing system, but also of the potential risks when migrating to new technologies or tools. Governments seldom have this kind of in-house expertise. This problem – of checking resilience, IPv6 readiness, security, etc. – is in fact even more difficult because of unpredictable interactions or dependencies among sets of applications or technologies. These interactions complicate manual testing, and make disruptions more likely.

The cloud is the ideal environment where checking these properties can be done automatically and in a contained environment, hence without any business impacting consequences. Testing for an application's IPv6 readiness, resilience, integration level with other applications, security or even standards compliance, can be now delegated to isolated testbeds, that can be controlled by external experts and that can also be massively virtualised. Computing power to perform a lot of tests, as opposed to performing a few tests manually, can be delegated to a cloud built up on demand. Moreover, with the cloud it is now easier to deploy certain software configurations or images, and subject them to varying load parameters. In other words, using the cloud paradigm for the problem of governments to evolve their software services has the advantages of scale, configurability, and ability to outsource computation. We envisage the following scenarios that are relevant for a government's system: testing of resilience/integration of new services, security testing, and an audit of the government services.

4.3.2 Test Suite for IPv6 Readiness and Security

Having a test suite for IPv6 readiness and/or security would be of an invaluable help to any system administrators. Assuming the administrator of a government Cloud manages public services that are currently deployed on a system with a certain configuration (e.g., a Windows XP machine, and currently using IPv4), it is likely that he or she meets with situations such as:

- The network is upgraded to IPv6;
- A new service is installed in the system;
- The network is upgraded to IPv6 and a new service is installed in the system;
- The OS or an existing application other than the service are patched/updated;
- The OS or an existing application is patched and the network is upgraded to IPv6.

Checking by hand if the Cloud network is unaffected in these cases would employ a lot of effort, and automating (parts of) this task would be a great improvement for both system administrators, and managers who need to know of the Cloud's performance.

The above problem hasn't been addressed before. Therefore, a possible approach would be that GEN6 bundles together a set of appropriate tools, test suites and guidelines that, in a private IPv6 enabled Cloud context, can test IPv6 readiness and security of individual services that are likely to be deployed in a CC environment.

A variation on the same idea can be that of testing of security vulnerabilities of known IPv6-enabled applications or application updates.

4.4 Usage and Dissemination

Luxembourg's Cloud Computing framework is designed by using open source tools (i.e., Cloud distributions – OpenStack/OpenNebula) and hence will contribute with IPv6 guidelines and expertise to the open source cloud community at large, unlike existing proprietary solutions. It

is important for GEN6 to make its results usable for governments, and this will be achieved in two ways: first, in that all guidelines and toolkits for the transition of Cloud environments to IPv6 -- that will result from this pilot will be available to interested EU member states as soon as possible by dissemination activities and publication of deliverables; and second, in that the Cloud Computing framework described above will be available to other governments to use or replicate internally. Therefore in no way will governments lose existing autonomy when replicating and using this framework internally.

In line with having a pilot whose use is open to other administrations, the adherence to EC's Interoperability Solutions for European Public Administrations (ISA¹¹) can be considered as an opportunity to emphasize GEN6 results even more. ISA defines action 2.1 [EU-ISA11] that can bring possible design extensions to the CC environment to spread its adoption by other governments. ISA action 2.1 "Towards a European Interoperability Architecture", describes requirements and steps towards a common architecture for interoperable public services at European level, and assesses the need for common infrastructure services as part of the final architecture. GEN6's CC pilot can become an internal infrastructure service, when considering that it can be replicated by governments into their own IT system, with the goal to test system operation and security. By running individualized test-suites and running new or existing internal controls, governments can thus have an internal platform to evaluate to what extent they can adopt new technologies and new services in the context of their current tools and technologies. Such test-suites and controls, as well as results from running them, can be shared with other governments as long as they use a communication channel that is secure (ISA Action 2.3, "Towards secure digital communication across networks", focuses on this aspect) and it serves the purpose of supporting and improving online communities (ISA Action 2.5, "Communication and Information Resource Centre for Administrations, Businesses and Citizens"), that span also across administrations who need to be aware of the power and limitations of their IT systems. Designing for this kind of interoperability is thus centered on two elements: the CC pilot should be an open model, that can be deployed and used easily, and also that inputs and outputs from using this system should be made available in a format that is understandable by other entities that also adhere to the ISA roadmap.

4.5 Related EU projects

4.5.1 SECRIKOM

SECRIKOM (Seamless Communication for Crisis Management) is an EU project that started in 2008, aiming to set up a secure communication platform for the management of crisis situations in Europe. The project solves the problem of interoperability of heterogeneous

¹¹ http://ec.europa.eu/isa/actions/02-interoperability-architecture/index_en.htm

devices with various capabilities (e.g., mobile, radio) by creating a pervasive and trusted communication infrastructure to bridge different types of existing network bearers. This brings a real advantage in emergency situations, so that humans can truly collaborate and save lives. The approach was to create a secure, wireless fault tolerant wireless communication infrastructure that is secure and distributed. The security on this infrastructure is based on custom chip-level encryption. SECRIком is related with GEN6 in considering communication security in a massively distributed system, yet SECRIком did not focus on IPv6 migration.

4.5.2 U2010

U2010 (Ubiquitous IP Centric Government & Enterprise Next Generation Networks, Vision 2010) is an EU project ended in 2009, which aimed to provide best-effort communication and the most effective access to information to anybody in case of situations of *force majeure* (accident, crisis, etc.). U2010 emphasised on the interconnection of existing devices and networks, the automatic redirection mechanisms, and the usage of new wireless ad-hoc network research. Since it investigated cross-technology and cross-network interoperability and availability, U2010 looked into how IPv6 can be used to enhance these objectives. Unlike GEN6, U2010 did not investigate the security of IPv6 communication.

5. CASE STUDIES ANALYSING SECURITY IN IPV6 INFRASTRUCTURES

This section shows the current state of the affairs in terms of security for the systems of two GEN6 partners, namely:

- ULAKBIM's National Academic Research and Education Network.
- TURKSAT's e-Government infrastructure and e-Government Gateway.

The security concerns of these systems stem from the need to use IPv6 (transition) technologies, as well as the need for integration of services across several domains. Some of these challenges have already been addressed, while others need further attention. Later iterations over these scenarios will prove useful for completeness and accuracy.

5.1 ULAKBIM's National Academic Research and Education Network

5.1.1 ULAKNET Security Features

ULAKNET, managed by Turkish Academic Network and Information Center (ULAKBIM), provides network connectivity to universities and research institutions. The security measures of ULAKNET are mainly held by ULAKNET Computer Security Incident Response Team (ULAK-CSIRT). The CERT/CSIRT units are mainly established in the well-organized constitutions and superior enterprises managing big networks. ULAK-CSIRT is responsible for preventing the potential security violations of external networks to ULAKNET. ULAK-CSIRT also aims ascertaining the attacks and the people in charge and in the same way, preventing the attacks from ULAKNET to the outside world and if there is an attack, ascertaining the people in charge of the attack and sharing the information with the administrators of this network. ULAK-CSIRT has been accredited by the European CSIRT Union Trusted Introducer (TI) in July 2007 and it is the first accredited team in Turkey. There are several security services operated by ULAK-CSIRT, as described below.

5.1.1.1 ULAKNET Blackhole Attack Detection System

Blackholes are used for monitoring unused Internet space to characterize security threats. Blackholes are capable of identifying configuration errors, routing problems, denial of service attacks (DoS), Internet worms and botnets. Blackholes, from a network security perspective, are placed in the network where traffic is forwarded and dropped. Once an attack has been detected, blackholing can be used to drop all attack traffic at the edge of an Internet Service Provider (ISP) network, based on either destination or source IP addresses.

ULAKNET Blackhole Attack Detection System includes a blackhole application based on ULAKNET unused IP blocks and the traffic forwarded to this system is received by a honeypot. ULAKNET Blackhole Attack Detection System captures the attack data and analyses the source of the attack. The results are published on a publicly available web site¹². The web site provides the following statistics on an hourly basis:

- Top ten attacker networks amongst the ULAKNET members.
- Top ten attacking IP addresses.
- Top ten attacking countries.
- Top ten attacked ports.
- Top ten attacking operating systems.

Current Blackhole Attack Detection System only supports IPv4 since it uses honeyd¹³ as the honeypot. In order to overcome this obstacle, Kovan has been developed. Kovan¹⁴ is an IPv6 honeypot which was developed under the “Design of National IPv6 Infrastructure and Transition to IPv6 Protocol”¹⁵ project. Kovan is a virtual honeypot framework that mimics a real network behavior to attract attackers. The new IPv6 capable blackhole honeypot system has been established using Kovan infrastructure but the captured data are not shared to public and only used for research purposes.

5.1.1.2 OLTA Incident Tracking System

The security incidents reported to the ULAK-CISRT is handled by an incident tracking system. OLTA incident tracking system is a custom web interface which was developed by using RTIR¹⁶ as a base platform. OLTA is not depending on the IP protocol version since it is a server site application that runs on apache server.

¹² <http://istatistik.ulakbim.gov.tr/balkupu/>

¹³ <http://www.honeyd.org>

¹⁴ <http://www.ipv6.net.tr/kovan/>

¹⁵ <http://www.ipv6.net.tr>

¹⁶ <http://bestpractical.com/rtir/>

5.1.1.3 ULAKNET NetFlow Statistics

NetFlow is a network protocol developed by Cisco Systems for collecting IP traffic information. Flow information is collected from ULAKNET core routers. ULAKNET NetFlow statistics are mostly used for network analysis. However, it is also used for providing security related information since NetFlow can provide IP header information. As the NetFlow provides IP header information, this service depends on IP protocol version. During the IPv6 transition of ULAKNET, network monitoring tools based on previous version of NetFlow (e.g. v5) has been upgraded to v9. The routers are also configured to export the NetFlow v9 as well.

5.1.2 Worm Propagation in IPv6 Networks

Worm propagation may be analysed in two phases:

- 1) Finding a vulnerable target.
- 2) Infecting the target. The infection phase is the part where a payload for the specific vulnerability is applied to the application. Hence, this part is not depending on the IP protocol version.

Searching for a vulnerable target may be achieved in different ways such as scanning the address space, using targets resources (i.e. Email addresses) or using removable media. A worm spreading over emails or removable media may use the same method in both IPv6 and IPv4 networks as expected. The main difference will be observed in worms searching for vulnerable hosts by scanning the network. As the address space is expanded in IPv6, scanning the whole network is expected to be harder than it is in IPv4 networks. At first look 128 bit address space is seemed to be large enough to be secure for scanning; however, there are some ways to reduce the space that will be scanned.

Some ways to reduce the scan space are listed below:

- Using ordered or easy to remember addresses e.g. ::1, ::2 or ::face, ::cafe.
- As remembering IPv6 addresses will be harder, DNS records will contain more information about the nodes.
- Some transition methods have their own IPv6 address structure such as 6to4 and Teredo. These methods will make guessing IPv6 addresses easier.
- Neighbourhood discovery messages, router advertisement/solicitation messages and other messages that assumed to be sent over secure local connections may be deployed to find hosts.
- Pre-defined multicast addresses may be used.

5.1.3 Security Requirements of the Turkish IPv6 Traffic Exchange Point

The aim of the Turkish IPv6 Traffic Exchange Point is to provide an experimental infrastructure for Turkish ISPs to test their network and services for IPv6 transition. The security policy of the service is a part of “IPv6 Traffic Exchange Point Policy”¹⁷ document which must be signed for participation. With respect to this policy, security requirements of the participants can be summarized as follows:

- IPv6 packet data information (i.e. Deep packet inspection) cannot be investigated except from diagnostics and statistic purposes.
- Participating ISPs can filter the traffic according to their security policy.
- The use of ICMP redirect, directed broadcast, spanning tree, IGP broadcast, CDP is forbidden.
- Participating ISPs must store the connection log information regarding the related laws.
- Participating ISPs must take the necessary actions to prevent worm propagation, e-mail spam, DDoS and other network security threats.

The traffic exchange point is still operational. During the operation of the exchange point no security incidents were reported by the participants.

5.2 TURKSAT's e-Government Infrastructure

In Turkey, TURKSAT is in charge of the e-Government infrastructure, with the purpose of developing and coordinating cross-institutional efforts to measure citizen satisfaction and to minimize “digital divide”. TURKSAT develops and operates the e-Government Gateway with 4 million registered users and 200 e-services. This Gateway is a cross-domain enablement product designed to speed up and secure Web services integration spanning identity and security domains.

Integrating new services to this e-Government Gateway is performed on a secure network that is not open to the Internet; with the help of IPsec VPN, connections to public institutions are made more secure. In addition, the data traffic to public institutions is made via the HTTPS protocol and while consuming the services from institutions extra software security measures are taken. Also all remote services are secured with software that senses potential attackers and all software units are regularly tested for security.

¹⁷ To be found at <http://www.ulakbim.gov.tr/ulaknet/basvuru/IPv6PaylasimProtokolu.pdf>

5.2.1 Security Concerns for the e-Gov Gateway

At the level of the gateway, security constraints are determined by the information technologies chief executive and they are applied by the network security group. In the scope of the e-Government project, the system is monitored 7/24 by the network security personnel in the monitoring group. Threshold values are set to the mechanism of monitoring the processes and logs of security devices; an alarm control system is in place when such thresholds are surpassed. The gateway is designed to accept or decline authentication requests by validating the requester's certificate, in accordance with TURKSAT's security policy.

Some of the most stringent security concerns when integrating a new institution's services and infrastructure are given below:

- *We are opening the Web service to you. How secure is the code on your side?*

TURKSAT's source code and the third-party software used passes safety tests on a regular basis. If during these tests vulnerabilities are discovered, they are corrected and tested before deploying the code to the live platform.

- *Is it possible to break into the secure network between us?*

TURKSAT is a Layer 1 operator in Turkey; the connections to the public institutions are made with dark fiber or using cable Internet infrastructures with MPLS VPN. At the same time, to raise the level of security IPsec VPN is being used. There have been no recorded attempts of penetration to this date.

- *How do you ensure physical security? Where are your servers in this environment?*

The data centre housing the e-Government Gateway servers is located in the TURKSAT campus. Entry is subject to retinal scanning, fingerprint reading, and face recognition. In addition, the TURKSAT campus is 40 km away from the city centre and it is protected by security teams 7/24.

5.2.2 IPv6 and IPv4 for Gateway Security

On the gateway, it is possible to authenticate and authorize requestors using any authentication mechanism appropriate for HTTP. Examples of such mechanisms include:

- Require SSL or TLS Transport Assertion.
- Require SSL or TLS Transport with Client Authentication Assertion.
- Require HTTP Basic Credentials Assertion.
- Require HTTP Digest Credentials Assertion.
- Require Windows Integrated Authentication Credentials Assertion.

The Gateway has a Security Token Service (STS) that can issue the following types of tokens: SAML Tokens (via the Create SAML Token Assertion), and Security Context Tokens (via the Create Security Context Token Assertion).

In terms of IPv6 Support, the Gateway field supports IPv6 literals for the Gateway host. The following formats are supported:

[2a01:1d58::7]

[2a01:1d58::7]:8443

The security policies mentioned above are for both IPv4 and IPv6. IPsec VPN will be used on both IPv4 and IPv6.

6. REQUIREMENTS FOR SECURE CLOUD SERVICES OVER IPV6

This section enumerates security requirements for IPv6 enabled Cloud services in the government sector. Several areas of investigation have been deemed as important:

- 1) When new equipment is purchased, provisioning support for IPv6-ready equipment should be checked thoroughly.
- 2) IPv4 to IPv6 Transition technologies and the security issues that they introduce.
- 3) Specific IPv6 security issues and how they are addressed by the community.
- 4) Management aspects from IPv4 that can impact system security when on IPv6.

The UL pilot will look into these requirements in order to come up with guidelines for governments transitioning their internal Clouds to IPv6.

6.1 Provisioning IPv6 Equipment

When searching for IPv6 support, it is essential for governments and other public organisations to have in place guidelines on what specific features to request of the equipment and of the tenders. A feature check is essential when equipment is bought, especially because security functionalities are very heterogeneous. Indirectly, this selection stimulates vendors to improve their products.

It follows that what is needed is a list of IPv6 security requirements to be satisfied. One of the most relevant helpers to this end is RIPE501 [RIPE501,RIPE501bis]. RIPE501 has been introduced by the Requirements for IPv6 in ICT Equipment (RIPE) working group, and is intended for IT managers. The document declares to provide a

“Best Common Practice (BCP) template that can be used by governments or large enterprises when developing tender documents. It can also serve as an aid to those people or organisations interested in tendering for government or enterprise contract”.

In line with USGv6 specifications, RIPE501 separates among several types of devices:

- 1) Host.
- 2) Switch, or “layer-2 switch”.
- 3) Router, or “layer-3 switch”.
- 4) Network security equipment.
- 5) Customer premise equipment.
- 6) Mobile device.
- 7) Load balancer.

For these devices, RIPE501 enumerates lists of security-related Requests for Comments (RFCs) to be supported in both mandatory and optional situations. In addition to this list, the document also has a section on the skill requirements of the system integrator.

RIPE501 is a mature document that has wide support from the community as well as from organisations such as Cisco. Its guidelines are deemed very useful when acquiring IPv6 resources, but with continuous technology and security evolution, governments should look for updated information for the above feature list. Therefore, the correct support of IPv6 security features, at the same level as for IPv4, should be required for network security devices, in the line of mentioned RIPE501 document.

In addition to the work of RIPE, governments should be strict about their resources and also require testing and conformance of certain equipments. This testing should be in line with:

- The "IPv6 Ready Logo Program" from IPv6 Forum, (<http://www.ipv6ready.org>).
- Similar initiatives from DoD in USA: <http://iitc.fhu.disa.mil/apl/ipv6.html>, or <http://www.nist.gov/itl/anttd/usgv6.cfm>.

GEN6 is represented in both the RIPE working group, and in the IPv6 Forum.

6.2 Potential Security Issues with IPv6 Transition Technologies

Adopting IPv6 has already started, while the existing IPv4 infrastructure is still in use. There are three main transition technologies currently employed (i.e. dual-stacks, tunnelling of IPv6 over IPv4, and translation), and security implications of the transition mechanisms are next examined.

Dual stack: As defined in [Min09], dual stack nodes manage two protocol stacks and allow the end system or router to use any of the two protocols. From a security standpoint, all security measures used in IPv4 networks should be implemented in the IPv6 networks. This can be difficult since the IPv6 topology may be different from the IPv4 topology, so applying the same rules as in IPv4 will not be possible. An example for this case is ICMP. Filtering all ICMP packets towards one host does not cause service interruptions in IPv4, but in an IPv6 network, *some* ICMP packets should be allowed in order to allow IPv6 communication at all.

Another issue of dual stack systems is that they could make users unhappy because of delays. Web browsers, email clients, instant message clients, etc. that are IPv6 enabled, can require several seconds of delay before falling back to IPv4. Such delays can slow the acceptance of IPv6, because IPv6 is disabled entirely on the end systems in order to improve the experience of the user. Clients could choose which protocol to use from the dual stack solution, yet usually IPv6 is used first by default. With different “types of IPv6” (native, encapsulated, and translated), in some specific scenarios users can suffer big delays because of using “bad IPv6” or because of having to switch from default-IPv6 to IPv4. A solution to the issue of this connection delay problem is Happy Eyeballs [RFC6555]. Happy Eyeballs specifies requirements and algorithms to reduce this delay, for instance by means of simultaneous connection attempts. There are currently some implementations for Happy Eyeballs on Google Chrome, Firefox 7, MacOS Lion X.

Tunnelling, also known as **Encapsulation**, is used to interconnect compatible nodes or domains over incompatible networks [Min09]. In terms of security, tunnelling has some major disadvantages, e.g., tunnel injection (an outsider can inject packets in the tunnel, which could lead to a reflection attack), packet sniffing or theft of service. In addition, automatic tunnels are a critical point that system administrators should care about. An IPv4 host may configure a tunnel interface and connect to the global IPv6 network without the system administrators knowing. Even worse, some Windows versions (e.g. Windows 7) configure their tunnel interfaces (e.g. 6to4, Teredo, ISATAP) automatically. In other words, the host machine makes a connection to the IPv6 network even without notifying the user.

[RFC6324] examines the security vulnerabilities of IPv6-in-IPv4 automatic tunnels, where there are inconsistencies between the IPv4 and the IPv6 routing states. In particular, this RFC focuses on the attacks when a routing loop is formed as a means of flooding (to end up in a Denial of Service), for those automatic tunnels that embed a packet’s exit IPv4 address with the destination IPv6 address. Some mitigation procedures are also described, and they range from checking caches of tunnel routers, to careful filtering of packets.

Therefore, for transition scenarios and mechanisms, the requirement is to implement measures to avoid loops attacks [RFC6324] or other tunneling attacks, as shown in RFC6169.

Translation: This transition technology has been proposed to be used where IPv6 support is not applicable for an IPv4 host or vice versa. Translation means the direct conversion of protocols, and can trigger transforming the headers and the payloads of the packets. Translator machine

between the IPv4 and IPv6 networks will lead to a single point of failure which will result in vulnerability to DoS/DDoS attacks. Moreover, end to end connectivity will be affected if a translation technique is applied. Translation is considered to allow access from an IPv6-only network to IPv4 Internet (NAT64/DNS64)[RFC6144-RFC6147]:

- RFC6144 describes a framework for IPv4/IPv6 translation, in the context of replacing NAT to NAT-PT (Network-Address-Translation – Protocol Translation).
- RFC6145 describes stateless IP-ICMP translation, that translates between IPv4 and IPv6 headers.
- RFC6146 describes stateful NAT64 translation (i.e., IPv6-only clients contacting IPv4 servers using unicast UDP, TCP, or ICMP).
- RFC6147 describes DNS extensions for NAT64.

Also, recently it has been standardized the Network Prefix Translation (NPT) that has several implications from the security point of view, because addresses are changed on transit.

Comparing these transition technologies from a security point of view, the dual stack method brings less vulnerability to the network with respect to tunnelling and translation methods. If the ISP has a native IPv6 connectivity, dual stack should be deployed throughout the network and necessary security procedures should be implemented.

6.3 Known IPv6 Security Issues and Mitigations

The new addressing scheme that comes with IPv6 has several security implications:

- First, since the number of addresses is much bigger than before, brute force or random scanning are more difficult. RFC5157 mentions this aspect, and observes that attackers might use other technologies to discover IPv6 addresses of interest in a target network. For instance, well known multicast addresses are defined so that services could be located, but this eases the work to find sensible services to attack (FF05::2 All routers, FF05::1:3 All DHCP Servers). One suggestion given by RFC5157 is that administrators should not use default or predictable numbering schemes for their hosts, which would make network scanning trivial.

- Second, the use of link-local addresses on an IPv6 interface allows for IP connectivity on a LAN segment without any external help. However, since each node could have several addresses and even random interface identifiers by means of stateless auto-configuration, it is difficult to control a host by its IP. RFC4941 describes an extension to IPv6 stateless auto-configuration that applies to interfaces whose identifier is derived from an IEEE identifier [RFC4941]. The point of this RFC is that having an interface identifier that changes over time makes eavesdropping more difficult, and enhances privacy.

Some other IPv6 security points relate to:

- DHCPv6 can be used, yet its support is not homogeneous across providers like Microsoft and Unix/Linux. This can lead to security issues since different Cloud services are provided on various platforms.
- The Neighbour Discovery Protocol (NDP) is a protocol that helps the discovery of devices on an IPv6 network. RFC3971 shows several vulnerabilities of the NDP in IPv6 [RFC3971]. The original ND Protocol specification defines the use of IPsec to protect ND messages. However, this solution does not work in practice because of several bootstrapping problems with the Internet Key Exchange (IKE) protocol, and also because of the high number of security associations that need to be manually configured. SEcure Neighbour Discovery (SEND) [RFC3971], aims to mitigate the problems of the ND Protocol. Although nowadays SEND support on hosts is not yet available, some router vendors do support it. Relevant discussions and mitigations about these issues were discussed in the following RFC-s:

- RFC3756 specifies several trust models related to IPv6 neighbour discovery protocols and describes several threats, along with some requirements for securing this protocol.
 - RFC6104 describes that users or administrators can misconfigure or inject bogus IPv6 router advertisements. Also, if there are multiple nodes sending RAs with prefixes for stateless auto-configuration, this could result in a DoS attack. The RFC discusses some scenarios in which router advertisements can occur, and describes several possible solutions to this issue.
 - RFC6105 proposes a light-weight alternative, called RA-Guard, to SEND based on filtering at layer-two. It addresses the issue that SEND is not trivial to deploy.
- Therefore, in LANs, NDP problems should be taken into account and measures like RA GUARD or SEND should be applied. There are some attacks that could be triggered from outside the IPv6 network.
- Broadcast-Amplification attacks (also known as “Smurf attacks”) are DoS attacks where an ICMP echo is sent to the broadcast address of a prefix with the spoofed address of the victim. All hosts on the destination prefix in turn send an echo reply to the victim. In IPv6, there is no concept of broadcast so the destination address (DA) can at best be multicast. RFC2463 clearly states that an ICMP reply should not be generated for packets that have a multicast DA, a layer 2 multicast address, or a layer 2 broadcast address. If all stacks implement this RFC properly, this approach would provide a good level of protection against this type of attack. Therefore, ICMPv6 should be implemented in line with RFC2463 to avoid smurf attacks.
- RFC5095 observes that the IPv6 Routing Header Type 0 can be exploited for traffic amplification, that can, on a remote path, generate a DoS attack. This RFC suggests a update to the IPv6 specification that would deprecate the use of this type of headers. Hence, IPv6 devices should comply with RFC5095 regarding the Routing Header Type 0.
- IPSec is built-in in the extension header to enable end to end security. However its deployment is done differently from vendor to vendor such as Microsoft using Teredo for its IPv6 VPN solution called DirectAccess making it totally proprietary. Promoting IPSec usage should warrant better tools for the management of encryption keys for Cloud services.
- Security solutions should not be based on a complete IPcmpv6 filtering, because ICMPv6

is part of IPv6 and is used for vital mechanisms as the packet fragmentation (ICMPv6 Packet Too Big).

- Despite the flexible and powerful auto-configuration mechanisms available in IPv6, for some critical infrastructure based on servers, these should be more a security hole than an advantage. So the requirement is to use the most secure address configuration method for the cloud solution, starting with the more secure and less flexible manual/static address configuration to the DHCPv6 with static assignments and logs of address leases.

6.4 Management Aspects for Secure Cloud Services

In [Wink11], Winkler identifies several requirements for developing a security-aware Cloud that we find relevant in an IPv6 scenario. These security services transcend the network layer to the application layer, and are impacted by the IP technology below. They are as follows:

- 1) **Cloud-wide time services** should be in place in order to maintain a correct and reliable time, for the synchronisation of all the Cloud services and logs.
- 2) **Identity management** controls should be in place to protect the identity information of cloud tenants, users and personnel.
- 3) **Access control management** mechanisms should enforce access control of Cloud users over cloud data and systems.
- 4) **System and network auditing** mechanisms should be there to support the application's accountability and audit processes.
- 5) **Security monitoring** controls are essential to know at all times of the security status of the system (e.g., compliance to policies, security incidents).
- 6) **Incident management** mechanisms should be in place to rapidly and appropriately respond to events encountered in the security monitoring phase.
- 7) **Security testing and vulnerability remediation** controls should be performed onto the target application before its mass deployment.
- 8) **System and network controls** should be in place on the infrastructure systems, aiming at properly isolating various systems.

6.4.1 Cloud Time Services

Keeping a synchronised clock among computers in a distributed system is not a trivial matter, yet it is of extreme importance. If clocks drift between network devices, diagnosing failures based on logs reports based on different clocks becomes error-prone.

The Network Time Protocol (NTP) can definitely help to keep clocks synchronised, yet the links to national authoritative time servers should be secured and made via multiple paths. Cloud time systems should be based on several time source paths in order to be as resilient as possible. Administrators should know that attacks can aim at hindering synchronisation of Cloud computers at pre-established intervals. Since 2006, version 4 of the Simple Network Time Protocol has been known to support IPv6, therefore Cloud time services over IPv6 should implement the recommendations in RFC4330 [RFC4330] or newer.

6.4.2 Identity Management

Cloud services must preserve privacy when handling confidential data of clients, confidentiality, integrity and information availability when it comes to identity data. Cloud users must be authenticated at all times, and the identity management infrastructure should be scalable to afford managing new tenants, as well as to adapt when users are deprovisioned. A federated identity system can be useful to enhance identity portability and to ease the Cloud user experience. Such a federated identity system can also allow for easier interoperability with identity providers.

The authentication systems should be at all times compliant with identity management policies. However, it is an open issue how (long) identities of deprovisioned users should be stored, for later analysis; moreover, it is up to the application domain to decide if keeping “deprecated” identities is a privacy breach from the user’s perspective.

Off-the-shelf identity management tools should be compatible with IPv6, where applicable.

6.4.3 Access Control Management

Cloud services should provide a strong set of access control mechanisms, as well as to prove that they are compliant with existing security policies and regulations. In particular:

- Multiparty authentication as well as strong authorization is a must in order to use Cloud services and infrastructures. Role-based Access Control is a widely used practice.
- Least privilege should be implemented when permission assignment is done.
- Security policies should cover both normal operational states, but also emergency situations (e.g., break-glass policies that bypass normal access control decision flows).
- User data should always be encrypted in order to limit access of unwanted parties.

Of particular interest to UL, security compliance mechanisms should be in place to ensure that the system or the classified data are being handled in a way that adheres to laws, regulations, and policies. It is up to the application to define who is in charge of enforcing security

constraints in each case, and also who specifies security policies in the first place. Nevertheless, correct policy enforcement requires a set of tools that mediate security-sensitive events and step in when “something bad is about to happen”. Note here that enforcement of security policies helps *prevent* security incidents, if the security policies have been properly specified.

Access control management tools should be compatible with IPv6, where applicable.

6.4.4 System and Network Auditing

Log keeping is essential for the a-posteriori security of distributed systems. In the Cloud, different trust zones impose different levels of integrity for logs generated by third-parties. The generation and management of log (or audit) events should follow several important requirements, of which:

- Logs should be kept confidential, integer, and available for the auditing process.
- Logs should be collected and maintained centrally so that incident response has a minimal overhead.
- The lifetime of logs should be related to that of the security policies that requested them, yet it is good to keep them indefinitely for a thorough analysis.
- Logs should be appropriately “sanitized” when handed in to tenants or users for inspection.
- Audit logs should be generated in real-time.
- Log entries should contain all information about the event that has occurred, to allow for a minute a-posteriori analysis for that event (e.g., correct time, system location of event, user on behalf of which the event happened, etc.).

Note here that auditing helps *detect* incidents after they have already happened.

Where applicable, system and network auditing mechanisms should be compatible with IPv6.

6.4.5 Security Monitoring

Security monitoring should be performed on all Cloud services. Such monitoring allows the administrators to know that they are in control of their Cloud application, and also to react to events that can be security relevant. Monitoring compliance to security policies is also essential for timely reaction to software errors or security violations.

Monitoring can be performed on network traffic (e.g., IPv6 traffic inspection), on context data (i.e., the security status of the system, like emergency) or on security logs. Inspection of security/system logs helps in a-posteriori evaluation of the system, and for root cause analysis in the care of an incident. It should also be useful to have tools that allow users to build their own anomaly detection for PaaS / IaaS, and to trigger the right administrators depending on the type of incident that has occurred.

The security monitoring controls should be reliable under failure conditions, and correct with respect to laws, regulations, and security policies.

Where applicable, security monitors should be compatible with IPv6.

6.4.6 Incident Management

It is essential to know when and what kind of security events occur at runtime. This mandates the need for an incident management system that can reliably react to security incidents; such system should have processes for appropriate detection, identification, and response to unwanted security events (e.g., policy violations, or system misbehaviours), as per the client's specification or any security policies in place. The incident management system should be configured to detect the most damaging security incidents and take measures in order to isolate and remedy the situation that has occurred. The reaction to security events can be either automatic in some cases (e.g., blocking big messages or attachments), human-driven for instance in the cases of DDoS, or hybrid (i.e., partly automatic). It is however important to determine, for each security incident that has occurred:

- Whose is the jurisdiction, who should take the first measures to remedy the situation.
- How soon the reporting has happened, with respect to the time of the event.
- What technologies and tools can be used across enterprises for correcting the situation.

Where applicable, incident management mechanisms should be compatible with IPv6.

6.4.7 Security Testing and Vulnerability Remediation

Before mass deployment, the Cloud software must be thoroughly tested for security vulnerabilities. Special requirements at this level impose that:

- The environments for development, staging and testing should be kept separate. Yet in addition to [Wink11,page 99] we suggest that the testing and staging environments should be correlated with the final deployment environment, in terms of load,

resources, security context, etc. in order to better simulate the real application behaviour.

- A strategy for vulnerability remediation should be defined for security events of all types (less critical, to very critical).
- Patch management should be employed across the infrastructure, since in the Cloud resources are more easily allocated and re-allocated.

Where applicable, security testing and vulnerability remediation tools should be IPv6 compatible.

6.4.8 System and Network Controls

At the level of the infrastructure, special care should be paid to both virtualized and physical components so that they are properly isolated and configured for security purposes. Network separation should be done among different functional areas in the Cloud application, among management areas, and among security and network administration areas. Network controls and software firewalls should also be used to examine traffic among the above mentioned islands. Code vulnerability checking should be performed before deployment, and techniques such as black/whitelisting should be used to identify trusted entities to communicate with.

Where applicable, system and network controls should be IPv6 compatible.

6.4.9 Configuration Management

[Wink11] advises that it is critical to keep a list of all configurations, software packages, allocations and any other Cloud assets that can be relevant for security purposes, with their classifications in terms of sensitivity, criticality, and function.

In addition to this practice, a useful requirement should be to have mechanisms in place to automatically adapt the runtime security parameters of Cloud services to contexts that require either more security or better performance. A set of tools to centrally manage all security parameters across different subsystems has the advantage of maintaining security consistency across different domains, for one single application. This offers management flexibility in terms of security capabilities i.e. security mechanisms can be connected together differently depending on an intended trade-off between security and performance.

7. CONCLUSIONS

In this document, the migration of the Luxembourg Cloud-based environment has been presented, along with generic requirements and challenges about how to achieve secure Cloud-based services. The migration of Luxembourg's Cloud-based environment to IPv6 will mainly provide guidelines to governments about how to manage and transition their own private Clouds in a safe way, and can also have a positive impact in the open source Cloud community. This pilot is developed with the support of local government entities (CCG and the Ministry of Commerce), and using its results should not affect the autonomy of any existing CC deployment.

A number of security challenges over Cloud infrastructures have been overviewed when they shift to IPv6 technologies. Two more specific case studies (TUBITAK ULAKBIM, TURKSAT) are considered and several IPv6-related security requirements for secure Cloud infrastructures are suggested. Some of these issues are well known and documented in the community RFCs. IETF's RFCs are well-known, available and in some cases updated in line with new discoveries or solutions. The general requirements presented in this document are in line with several expert recommendations (e.g., IETF RFCs, Cloud Security Alliance, ENISA) and with a base-line legislative context in Europe. UL plans to analyse these requirements during the migration of the actual system presented above.

8. REFERENCES

[Bacon10]	J. Bacon, D. Evans, D. Evers, M. Migliavacca, P.R. Pietzuch, B. Shand, B. Enforcing End-to-end Application Security in the Cloud. In Proceeding of Middleware 2010, pp. 293—312, Springer, 2010.
[Catteddu09]	D. Catteddu, G. Hogben. Cloud Computing - Benefits, risks and recommendations for information security, European Network and Information Security Agency (ENISA), 2009.
[CSA10]	Cloud Security Alliance, "Top Threats to Cloud Computing", vol. 1, Mar. 2010; https://cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf .
[ENISA09]	European Network and Information Security Agency, "Cloud Computing; benefits, risks, and recommendations for information security", 2009, http://www.enisa.europa.eu/activities/risk-management/files/deliverables/cloud-computing-risk-assessment
[EU-ISA11]	European Commission - ISA Work Programme. (2011, November). <i>Phase 2 – Final Report: Common Vision for an EIA</i> . Retrieved October 2012, from http://ec.europa.eu/isa/documents/isa_2.1_eia-finalreport-commonvisionforaneia.pdf
[Fiege04]	L. Fiege, A. Zeidler, A. Buchmann, R. Kilian-Kehr, G. Muehl. Security aspects in publish/subscribe systems. In Proceeding of Third International Workshop on Distributed Event-based Systems (DEBS'04), 2004.
[Gartner08]	Gartner, Inc., The Security Risks of Cloud Computing, http://cloud.ctrls.in/files/assessing-the-security-risks.pdf , ID Number: G00157782. June 2008.
[Ghe11]	Gabriela Gheorghe, "Security policy enforcement in service-oriented middleware", PhD thesis, University of Trento, 2011, available at http://eprints-phd.biblio.unitn.it/673/1/PhD-Thesis-Gabriela.pdf
[Gupta08]	M. Gupta, P. Judge, M. Ammar. GossipTrust for Fast Reputation Aggregaion in Peer-to-Peer Networks. IEEE Transactions on Knowledge and adata Engineering (February 2008).
[Hang09]	C.-W. Hang and M. P. Singh, "Selecting trustworthy service in service-oriented environments," in The 12th AAMAS Workshop on Trust in Agent Societies, May 2009.
[Kamvar03]	P.S. Kamvar, M.T. Schlosser, H. Garcia-Molina. The EigenTrust Algorithm for Reputation Management in P2P Networks. In Proceeding of WWW 2003 Budapest, ACM, 2003.
[Malik09]	Z. Malik and A. Bouguettaya, "Reputation bootstrapping for trust establishment among web services," IEEE Internet Computing, vol. 13, pp. 40–47, 2009.
[Mead05]	N. Mead, T. Stehney, "Security Quality Requiements Engineering (SQUARE) Methodology", ACM Sigsoft Software Engineering Notes, Volume 30, Issue 4, July 2005
[Min09]	D. Minoli, J. Kouns, "Security in an IPv6 environment", Auerbach Publications, 2009
[OpenStack12]	<i>OpenStack Compute Administration Manual</i> . Retrieved in September 2012 from http://docs.openstack.org/trunk/openstack-compute/admin/content/index.html .
[OSA]	Open Security Architecture, http://www.opensecurityarchitecture.org/cms/en/about
[RFC3971]	J. Arkko, J. Kempf, B. Zill, P. Nikader, "Secure Neighbour Discovery (SEND)", March 2005, http://www.ietf.org/rfc/rfc3971.txt
[RFC3756]	P. Nikander, Ed., J. Kempf, E. Nordmark, "IPv6 Neighbor Discovery (ND) Trust Models and Threats", May 2004, http://www.ietf.org/rfc/rfc3756.txt

[RFC4330]	D. Mills, "simple Network Time Protocol (SNTP) version 4 for IPv4, IPv6, and OSI", January 2006, http://tools.ietf.org/html/rfc4330
[RFC4941]	T. Narten, R. Draves, S. Krishnan, "Privacy Extensions for Stateless Address Autoconfiguration in IPv6", September 2007, DRAFT STANDARD, http://tools.ietf.org/html/rfc4941
[RFC5095]	J. Abley, P. Savola, G. Neville-Neil, "Deprecation of Type 0 Routing Headers in IPv6", December 2007, PROPOSED STANDARD, http://tools.ietf.org/html/rfc5095
[RFC5157]	J. Chown, "IPv6 Implications for Network Scanning", March 2008, http://tools.ietf.org/html/rfc5157
[RFC6104]	T. Chown, S. Venaas, "Rogue IPv6 Router Advertisement Problem Statement ", Feb. 2011, http://www.ietf.org/rfc/rfc6104.txt
[RFC6105]	E. Levy-Abegnoli, G. Van de Velde, C. Popoviciu, J. Mohacsi, "IPv6 Router Advertisement Guard", Feb. 2011, http://www.ietf.org/rfc/rfc6105.txt
[RFC6555]	D. Wing, A. Yourtchenko, "Happy Eyeballs: Success with Dual-Stack Hosts", April 2012, PROPOSED STANDARD, http://tools.ietf.org/html/rfc6555
[RFC6324]	G. Nakibly, F. Templin, "Routing Loop Attack Using IPv6 Automatic Tunnels-Problem Statement and Proposed Mitigations", August 2011 http://tools.ietf.org/html/rfc6324
[RFC6144]	F. Baker, X. Li, C. Bao, K. Yin, "Framework for IPv4/IPv6 Translation", April 2011, http://tools.ietf.org/html/rfc6144
[RFC6145]	X. Li, C. Bao, F. Baker, "IP/ICMP Translation Algorithm", April 2011, http://tools.ietf.org/html/rfc6145
[RFC6146]	M. Bagnulo, P. Matthews, I. van Beijnum, "Stateful NAT64: Network Address and Protocol Translation from IPv6 Clients to IPv4 Servers", April 2011, http://tools.ietf.org/html/rfc6146
[RFC6147]	M. Bagnulo, A. Sullivan, P. Matthews, I. van Beijnum, "DNS64:DNS Extensions for Network Address Translation from IPv6 Clients to IPv4 Servers", April 2011, http://tools.ietf.org/html/rfc6147
[RFC5902]	D. Thaler, L. Zhang, G. Lebovitz, , "IAB Thoughts on IPv6 Network Address Translation", July 2010, INFORMATIONAL, http://tools.ietf.org/html/rfc5902
[RFC6296]	M. Wasserman, F. Baker, "IPv6-to-IPv6 Network Prefix Translation", June 2011, EXPERIMENTAL, http://tools.ietf.org/html/rfc6296
[RFC2463]	Conta, S. Deering, "Internet Control Message Protocol for the Internet Protocol version 6 Specification", December 1998, http://www.ietf.org/rfc/rfc2463.txt
[RIPE501],[RIPE501bis]	RIPE IPv6 working group, "Requirements For IPv6 in ICT Equipment", http://www.ripe.net/ripe/docs/ripe-501 ; RIPE-501(bis): http://go6.si/RIPE501bis/
[Rocha11]	F. Rocha, S. Abreu, M. Correia, "The Final Frontier: Confidentiality and Privacy in the Cloud", Computer Journal, Vol. 9, Issue 44, IEEE Computer Society, September 2011
[Schubert10]	L. Schubert, K. Jeffery, B. Neidecker-Lutz. The Future of Cloud Computing, Opportunities for European Cloud Computing Beyond, European Commission, Information Society and Media, 2010.

[Subashini11]	S. Subashini, V. Kavitha. A survey on security issues in service delivery models of cloudcomputing, In Journal of Network and Computer Applications, Volume 34, Issue 1, Pages 1-11, January 2011.
[Sun05]	L. Sun, L. Jiao, Y. Wang, S. Cheng, W. Wang. An adaptive group-based reputation system in peer-to-peer networks. In: Deng, X., Ye, Y. (eds.) WINE 2005. LNCS, vol. 3828, pp. 651–659. Springer, Heidelberg (2005).
[Vu05]	L.-H. Vu, M. Hauswirth, and K. Aberer, “Qos-based service selection and ranking with trust and reputation management,” in Proceedings of the Cooperative Information System Conference (CoopIS05), 2005, pp. 466–483.
[W3c11a]	W3C, Tracking Preference Expression (DNT) (http://www.w3.org/2011/11/dnt-pr.html.en), November 2011.
[W3c11b]	W3C, Tracking Compliance and Scope (http://www.w3.org/TR/2011/WD-tracking-compliance-20111114/), November 2011.
[Wang07]	Y. Wang and J. Vassileva, “Toward trust and reputation based web service selection: A survey,” International Transactions on Systems Science and Applications, vol. 3, no. 2, pp. 118–132, 2007.
[Wang09]	C. Wang, Q. Wang, K. Ren, W. Lou. Ensuring data storage security in Cloud Computing, In proceedings of the 17th International Workshop on Quality of Service, 13-15 July 2009.
[Wink11]	Vic (J.R.) Winkler, “Securing the Cloud: Cloud Computing Security Techniques and Tactics”, Syngress, 2011
[Zhou03]	R. Zhou, K. Hwang, C. Min. A Reputation System for Peer-to-Peer Networks, In Proceeding of NOSSDAV 2003, Monterey, California, USA, June 1-3 (2003).