Project title: Enforceable Security in the Cloud to Uphold Data Ownership
Project acronym: ESCUDO-CLOUD
Funding scheme: H2020-ICT-2014
Topic: ICT-07-2014
Project duration: January 2015 – December 2017

D4.1
First Report on Security Metrics and Assessment

Editors: Neeraj Suri (TUD)
Stefan Winter (TUD)
Ahmed Taha (TUD)
Reviewers: Stefano Paraboschi (UNIBG)
Giovanni Livraga (UNIMI)

Abstract

D4.1 develops techniques to conduct qualitative and quantitative assessment of security for the information flow across users and Cloud Service Providers (CSP). The report first documents examples of typical Cloud operational chains to outline commonly encountered failures and security breaches in Cloud environments. It then documents the state of the art in trust metrics and outlines their linkage to the ESCUDO-CLOUD use cases. In practice, Service Level Agreements (SLA) are used to represent the negotiated terms of service delivery across the users and CSPs based on these metrics from standards bodies. Hence, utilizing these metrics, the deliverable proposes techniques to specify, reason and aggregate SLA-based trust specifications into trust metrics, and their usage to compare the trust levels across CSPs.

Type	Identifier	Dissemination	Date
Deliverable	D4.1	Public	2015.12.31

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 644579. This work was supported in part by the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract No 150087. The opinions expressed and arguments employed herein do not necessarily reflect the official views of the European Commission or the Swiss Government.

ESCUDO-CLOUD Consortium

1.	Università degli Studi di Milano	UNIMI	Italy
2.	British Telecom	BT	United Kingdom
3.	EMC Corporation	EMC	Ireland
4.	IBM Research GmbH	IBM	Switzerland
5.	SAP SE	SAP	Germany
6.	Technische Universität Darmstadt	TUD	Germany
7.	Università degli Studi di Bergamo	UNIBG	Italy
8.	Wellness Telecom	WT	Spain

Versions

Version	Date	Description
0.1	2015.11.23	Initial Release
0.2	2015.12.14	Second Release
1.0	2015.12.31	Final Release

List of Contributors

This document contains contributions from different ESCUDO-CLOUD partners. Contributors for the chapters of this deliverable are presented in the following table.

Chapter	Author(s)
Executive Summary	Neeraj Suri (TUD)
Chapter 1: Introduction	Neeraj Suri (TUD), Pierangela Samarati (UNIMI)
Chapter 2: Cloud Chain and Security Breaches: State of the Practice	Silvio La Porta / Andrew Byrne / David Bowden (EMC), Mercedes Castano Torres / Ignacio Garcia Vega (WT)
Chapter 3: SLA Based Metrics and Assessment - Approaches	Neeraj Suri / Ahmed Taha (TUD), Pierangela Samarati (UNIMI)
Chapter 4: Conclusions and Next Steps	Neeraj Suri (TUD), Pierangela Samarati (UNIMI)
Appendix A: Full List of C-SIG Trust Metrics	C-SIG WG, NIST

Contents

List of Figures

List of Tables

Executive Summary

Deliverable D4.1 comprises the *"first report on security metrics and assessment"*. Representing the M1-12 activities of T4.1, D4.1 outlines the security metrics applicable to assessing security on the information flow across the users and the Cloud Service Providers (CSP). Using actual Cloud scenarios as examples, and also relating to the ESCUDO-CLOUD use cases, it develops techniques to specify, reason and compare across the trust levels offered by CSPs. The scope of D4.1 is on fundamental concepts as applicable to the single CSP model. D4.2 will subsequently extend the approaches to the multi-CSP environment. The development of D4.1 (and also the subsequent D4.2) will develop on the methodology outlined below.

Security parameters for Cloud services are typically specified in the form of security Service Level Agreements (termed as security SLA's or simply secSLA's). These are typically a mixture of qualitative and quantitative attributes making it hard to quantitatively assess the Cloud Service Provider's actual provisioning of security fulfilling the user's needs.

Although the state of the art predominantly focuses on the methodologies to build and represent Cloud secSLAs, there is still a conspicuous lack of techniques that quantitatively evaluate Cloud secSLAs to provide security assurance. This lack of assurance along with the existence of multiple CSPs offering similar security services, often results in Cloud users being unable to trust and assess the security of the CSPs provided services they are paying for. Hence the user is hard pressed to find the most suitable CSP that fulfills their security requirements. Therefore, it is essential to develop techniques that can assess the security claims from the CSPs to select the provider that can fully satisfy the users security requirements.

To achieve this goal, D4.1 addresses trust metrics and especially the SLA-based solutions for enabling users to express and reason about the trust associated with different providers for matching their requirements and also to compare the relative trust offerings across the CSPs. This is done by developing two evaluation techniques, namely Quantitative Policy Trees(QPT) and Quantitative Hierarchical Process(QHP), for conducting the quantitative assessment and analysis of the secSLA based security level provided by CSPs with respect to a set of Cloud Customer security requirements. These proposed techniques improve the specification of security requirements by introducing a flexible and simple methodology that allows Customers to identify and represent their specific security needs. Following the D4.1 developed guidance on the standalone and collective use of QPT and QHP, these techniques are validated using two use case scenarios and a prototype, leveraging actual real-world CSP SLA data derived from the publicly available Cloud Security Alliance's Security, Trust and Assurance Registry.

1. Introduction

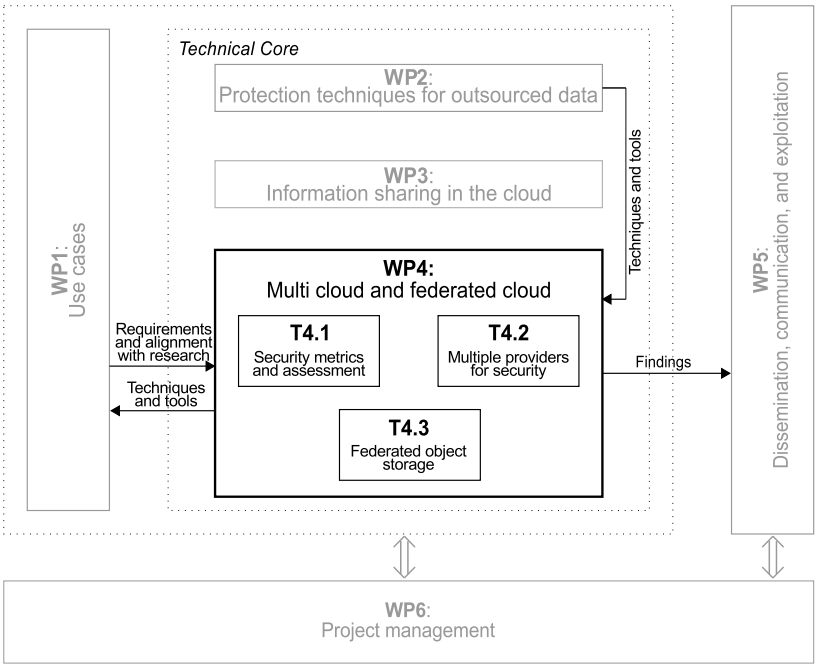

Figure 1.1: WP4 Tasks - T4.1 is the basis for D4.1

While the use of the Cloud opens a multitude of resources and services to the user, one also needs to assess the trust facets of the Cloud, i.e., characterizing the security, privacy and dependability levels available to the user. The diversity of Cloud providers, and also the diverse dimensions of trust[1] (security, privacy and dependability) offerings by each of them, complicates the users' task to select the provider that can best fulfil their security requirements. In the Cloud area, security parameters and mechanisms are typically specified in the form of Service/Security Level Agreements (termed SLA and SecLA respectively) to model security among providers and users. As the primary interfaces between the CSP and the user are these SLAs (which are primarily textual legal contracts of a mixed qualitative and quantitative nature), the need exists to be able to: a) quantitatively specify and assess user/CSP trust requirements based on SLAs and b) compare the trust services offered by the different CSPs. The objective of D4.1 is to investigate trust metrics and especially the techniques to express and reason about them, as well as to perform SLA-based comparative security assessments across the CSPs.

[1]The term "trust" is being used in literature with a variety of meanings [JIB07]. It has multiple definitions and usage spanning security, privacy and dependability. Fundamentally, trust refers to the quantifiable degree of faith a user can put (or a provider can assure) for the delivery of the service meeting the desired specifications. Security and privacy are often a mixture of quantitative and qualitative attributes while dependability (also covering reliability and availability) is mostly quantitative.

For the coverage of D4.1, two aspects are worth highlighting, namely:

1. Trust is an end-to-end attribute. Accounting for this, D4.1 considers failures and security breaches across the entire operational lifecycle of the Cloud chain covering the user, the network, cloud interfaces, and the implementation layers (e.g., the software tiers and back-end protocols).

2. While there exists considerable variety in the type and nature of attributes comprising an SLA, the specification of a Cloud SLA explicitly needs to conform to the guidelines laid by NIST, ENISA, ISO/IEC and similar standards bodies. Accounting for this, T4.1 will track SLA standards to ensure that the metrics developed in ESCUDO-CLOUD will comply with ongoing SLA standard guidelines, for better acceptance by the community.

1.1 Scope and Structure

D4.1 is structured into the following chapters:

- Chapter 2: This chapter documents a typical Cloud's operational chain outlining commonly encountered failures and security breaches in Cloud environments. The coverage extends to: a) documenting the "industry-oriented" State of Practice metrics/techniques and b) documenting the requisite standards that industry has to conform to for SLA based security metrics. In the final part of the chapter, it is documented the state of the art in trust metrics spanning the dimensions of security, privacy and dependability. As a first step this task compiles the SLA metrics proposed by standards bodies such as CIS, NIST, ENISA, ISO/IEC and especially the C-SIG WG on SLA's that considers inputs from existing EC SLA and Cloud Security/Privacy projects. The linkage to the ESCUDO-CLOUD use cases is elaborated here.

- Chapter 3: This chapter utilizes the metrics from Chapter 2, and proposes the techniques to specify, reason and aggregate SLA based trust specifications into trust metrics. The processes from this chapter are used to develop approaches to compare the trust levels across CSPs.

- Chapter 4: This chapter presents the conclusions and next steps for D4.1.

2. Cloud Chain and Security Breaches: State of the Practice

This section outlines a typical Cloud "system model" which will be utilized for the D4.1 discussions. As trust is an end-to-end attribute, the Cloud chain will present the user, the network, the requisite Cloud interfaces and the backend Cloud infrastructure layers. Given the high amount of technology dependent implementation details at each stage of the Cloud chain, an initial decision was made to have D4.1 abstract the Cloud chain to develop trust metrics and assessment processes at the level of "services" between the user and the CSP. This allows us flexibility to be agnostic to changing technology implementations and also allows flexibility of adding new trust metrics as relevant and adapting to the changing services.

Following the presentation of the Cloud chain, this section develops the background for D4.1 by documenting (a) the industry viewpoint on typically encountered failures and security breaches that are considered important in practice, (b) the emerging threat areas, (c) typical mitigation schemes, (d) commonly used assessment resources for evaluating CSPs, (e) introduce the SLA concept with the related state of the art, and (f) the relevant standards that threat metrics have to conform to. This last aspect is highly important as only SLAs conforming to standards get accepted and used by CSPs. It is also important to mention that this chapter is primarily an illustrative section highlighting typical threats and metrics, and is not intended for completeness. Also as there is no single "typical" Cloud scenario, D4.1 presents two illustrative example scenarios with corresponding Cloud chains and the relevant security/dependability breaches. These scenarios are developed using UC4 as the example reference Cloud chain although the discussion also applies to the other ESCUDO-CLOUD use cases.

2.1 The Cloud Chain

There exist three stakeholders in the ESCUDO-CLOUD scenario, namely: the end user, the broker and the CSP. The communication between these actors is always the same: the user (data owner) will be able to manage his/her files hosted in the Cloud through the ESCUDO-CLOUD middleware. This communication is shown in the figures below, but there are different aspects depending on the type of middleware (with or without agent) and the CSP.

The data owner will have access to the data files stored (and encrypted) in the Cloud through a browser or an agent. The single Cloud model illustrating how the user can access the ESCUDO-CLOUD middleware (hosted on the client side) is shown in Figure 2.1. Hence, the user's data will be encrypted and transmitted through the Internet and remain encrypted in the CSP.

When a user wants to access their data through a web browser, he/she will provide their username and password to the web portal. If the authentication is successful, the documents will be displayed in cleartext. The decryption of the files will be performed transparently by the

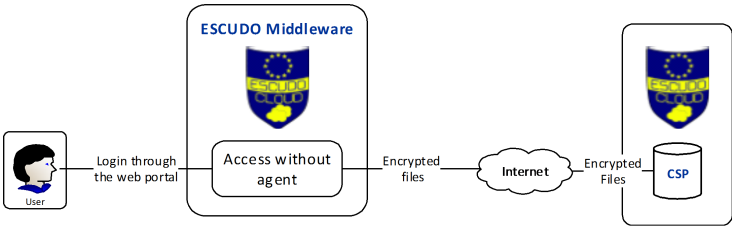

Figure 2.1: Single Cloud with web browser

ESCUDO-CLOUD middleware on the CSP side. In this way, ESCUDO-CLOUD will be running activities on the CSP (indirectly).

The user can also access his/her data through an agent such as an application on their device or as a synchronised local folder in their computer. Consequently the documents will be transmitted across the Internet (from the CSP to the agent), so the decryption must be performed on the user/client side. Furthermore, a synchronization between the CSP and the agent will be required prior to this interaction. In the previous case, the communication channel was established between the user and a single Cloud provider, but what is the implication on this model if the CSP employs a Multi-Cloud model? The scenario for accessing the data through a web browser or an agent in a Multi-Cloud model is illustrated in Figure 2.2. The front end user experience remains the same. The core difference here is the inclusion of several additional CSPs in the backend of the service that provide additional hosting for encrypted files. This model poses additional risks to the data as the CSP extends its service into third party CSPs. Through ESCUDO-CLOUD, the data remains protected as it is encrypted throughout the network of CSPs.

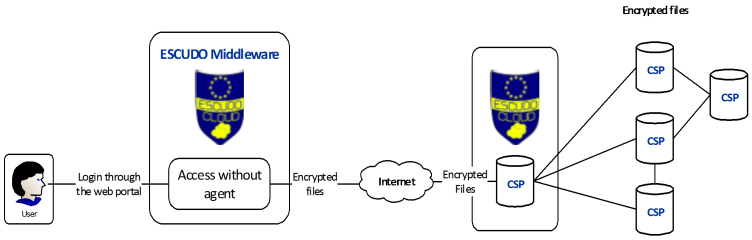

Figure 2.2: Multi-Cloud with web browser

The process to access the data is similar to the single Cloud model. When a user accesses his/her data files stored in the CSP, it does not matter if it is an elastic Cloud provider because the decryption will be performed in the ESCUDO-CLOUD middleware on the client side. Third party Cloud providers will store the files encrypted by ESCUDO-CLOUD, so that it is not possible for the information to be leaked or intercepted in plaintext.

2.2 Typical Failures and Security Breaches

There are numerous security issues for Cloud computing as it encompasses many technologies including networks, databases, operating systems, virtualization, resource scheduling, transaction management, load balancing, concurrency control and memory management. Therefore, security issues for many of these systems and technologies are applicable to Cloud computing.

The growth of Cloud computing has introduced new security challenges to the IT landscape. A lack of international consensus on best practices, standards and terminology, along with an industry still in its infancy, is central to the problem. Enormous amounts of data are stored, processed and accessed globally over the Cloud without clearly defined boundaries, security policies or access policies in place.

The lack of transparency in the Cloud makes it more difficult to properly assess the risk to user data. What processes and procedures are in place with the CSP? Does the CSP sub-contract any of the process? How does that affect the security of the user data?

As the user cedes control of their data to the CSP, there is an increased risk to the user's data due to breaches in the principal areas of Confidentiality, Integrity and Availability (CIA):

- The CSP may not conform to proper regulation & standards.
- The CSP may transfer, process or store data in a way that might compromise its security.
- The CSP may not provide adequate tools, such as auditing and logging, to enable the user to monitor the service and their data.
- Data may cross borders and move into regions with different legal jurisdictions, which have different requirements and rights of access.
- On termination of contract with CSP, assurances are needed that the CSP has taken the appropriate steps to delete, or put beyond use, the user's data.

A list of the threats (failures and security breaches) for Cloud customers categorized according to the CIA security model are presented below. The intent is not to provide a complete set of all possible threats (which is impossible given the diversity of Cloud models) but to provide a representative idea of common occurrences.

- **Confidentiality**: Aims to limit access to the information to authorized users only and deny access to unauthorized ones. Some security breaches associated with confidentiality are described below:
  - Insider user threats:
    - Malicious Cloud provider employees and service users potentially have access to confidential information which they can freely view, download and/or distribute.
    - Similarly, malicious third party users (supporting either the Cloud provider or customer organizations) may have access to confidential information.
  - External attacker threats:
    - Remote software attack of Cloud infrastructure and applications.
    - Remote hardware attack against the Cloud.

        - Remote software and hardware attack against Cloud user organizations' endpoint software and hardware.
        - Social engineering of Cloud provider users, and Cloud customer users.
    - Data leakage:
        - Failure of security access rights across multiple domains.
        - Failure of electronic and physical transport systems for Cloud data and backups.
        - Vulnerabilities in the key management infrastructure, which could allow an attacker to circumvent the encryption, regardless of the strength of the cryptography used.

- **Integrity**: Aims to verify the trustworthiness of the information. Some security breaches associated with the integrity are described below:
    - Data segregation:
        - Incorrectly defined security perimeters and incorrect configuration of VMs and hypervisors can allow an attacker to exploit weaknesses to gain access to data and manipulate it. Note, this is also a concern for confidentiality.
    - User access:
        - Attackers can exploit poor identity/access management procedures to eavesdrop on transmissions between locations or on data at rest. This can enable undetected manipulation of data through Man in the Middle (MitM) attacks.
    - Data quality:
        - Introduction of faulty application or infrastructure components presents another vulnerability that can be exploited by an attacker.

- **Availability**: Aims to ensure that data and services (resources) are available when they are required. Some security breaches associated with the availability are described below:
    - Change management:
        - Customer penetration testing impacting other Cloud customers.
        - Infrastructure changes on Cloud provider, customer and third party systems impacting Cloud customers can degrade or disrupt services on the CSP.
        - Without proper patch management processes (in particular for testing), services are at risk of inadvertently being degraded or disrupted due to a patch.
    - Denial of Service (DoS) attack:
        - Network bandwidth distributed DoS.
        - Network DNS DoS.
        - Application and data DoS.
    - Physical disruption:
        - Disruption of Cloud provider IT services through physical access.
        - Disruption of third party WAN providers services.
    - Exploiting weak recovery procedures:
        - Invocation of inadequate disaster recovery or business continuity processes.

Many of the security concerns in the Cloud are as a result of software vulnerabilities present in Cloud technologies. In fact, these kinds of vulnerability cause the greatest concern in SaaS[1]. Currently, many users are afraid to move their business workloads and data into a SaaS model due to their security concerns, such as the loss of visibility and control over their data, workloads and processes. In particular, the loss of control over the location of their data and who can access it is a key concern. Some software breaches defined by NIST in CWE (Common Weakness Enumeration) are listed below:

- Buffer Errors
- Code
- Code Injection
- Command Injection
- Configuration
- Cross-Site Request Forgery (CSRF)
- Cross-Site Scripting (XSS)
- Cryptographic Issues
- Data Handling
- Format String Vulnerability
- Improper Access Control
- Indicator of Poor Code Quality
- Information Leak / Disclosure
- Information Management Errors
- Injection
- Input Validation
- Insufficient Information
- Insufficient verification of Data Authenticity
- Link Following
- Location
- Numeric Errors
- OS Command Injections
- Path Equivalence

[1] Software as a Service

- Path Traversal
- Permissions, Privileges and Access Control
- Race Conditions
- Resource Management Errors
- Security Features
- Source Code
- SQL Injection
- Time and State

Table 2.1 shows a sample list of countermeasures used to address failures and security breaches.

Table 2.1: Techniques to address failures and security breaches

Security threats	**Countermeasures**
Spoofing identity	Authentication Protect secrets Don't store secrets
Tampering with data	Authorization Hashes Message authentication codes Digital signatures Tamper-resistant protocols
Repudiation	Digital signatures Time-stamps Audit trails
Information disclosure	Authorization Privacy-enhanced protocols Encryption (securing data at rest and in transit) Protect secrets Don't store secrets
Denial of Service	Authentication Authorization Filtering Throttling Quality of service (QoS)
Elevation of privilege	Run with least privilege
Access to sensitive or critical data	Security improvements for VMs Virtual network separation Trusted Storage Trusted platform module access techniques
Sudden failure	Automated response Automated notification

2.2.1 Data Confidentiality

In addition to traditional data confidentiality challenges, Cloud computing environments expose the user's data to additional threats that need to be identified and mitigated against. In a traditional data center setting that is controlled and monitored by the internal IT staff, the owner of the data and workloads is considered the same at the tenant. There are no third parties entities in this scenario. However, human error, misconfiguration or even the malicious actions of a disgruntled employee can lead to a compromise of the user's environment.

Despite these threats, there is security in the knowledge that all processes, policies and physical infrastructure is under the control of the organisations internal IT department. When an organisation or user moves their data and workloads to a CSP environment, there are additional concerns; they may have to share their environment with other users and organisations; the CSP may not adequately vet a user's integrity or competence; their data can now freely move across different geographical regions, potentially exposing the data to various forms regulation and legislation. Of primary concern to the confidentiality of the user's data is the new attack vectors posed by:

- The CSP and its staff.
- Other users within the CSPs environment.
- Third parties that the CSP sub-contracts services to.
- Third parties that have legal jurisdiction over the CSP or its environment.

Insider User Threats

In the Verizon 2014 Data Breach Investigation report, the 50 organizations surveyed cited 11,698 incidents of insider and privilege misuses [Tea15], 88% of which were privilege abuse. This shows that insider threats are not uncommon events. Insider threats are not restricted to CSPs; but there is one important difference, the user is not in direct control of the employment, vetting and supervision process [HS10] at the CSP. Thus the user must rely on the CSP to have robust processes in place that conform to industry standards [ISO13], for the recruitment and supervision of their staff, which is an important set of metrics to consider.

Physical security of a data center is one aspect of insider threats that is often overlooked. The simple expedient of plugging a USB device into the back of a server can circumvent multiple layers of security appliances and firewalls. Devices such as servers, consoles, network switches and storage arrays, needs to be physically secured to prevent unauthorized access. A lack of physical security can leave resources vulnerable to attackers, such as employees, contractors, industry partners or intruders. This can result in data theft, data tampering, further compromised security, etc. Examples of vulnerable physical resources include:

- Unencrypted hard drives
- Open network ports
- Server access consoles

CSPs often host multiple tenants on the same environment, and this can leave users vulnerable to attacks from inside that environment from other tenants; e.g., poorly hardened servers or VMs that can be hijacked or malicious third party users using stolen credentials. Some of the factors that should be considered when assessing the threat from other users include:

- How well does the CSP vet their tenant's authenticity or competency?
- How does the CSP isolate VMs and Servers from each other on the LAN?
- How does the CSP prevent VM escapes [Fer07]?
- How does the CSP prevent unauthorized access to shared resources, such as databases or data stores?

Tools such as *hping*, *nmap* & *wget* can be used to gain knowledge of the network infrastructure and what systems are connected on it. Furthermore, they can be used to scan for any open ports that can then be exploited in an attack. Ristenpart et. al. showed how these tools could be used to map a virtual network [RTSS09]. This information was used to co-locate a "malicious" VM on the same physical host as a target VM. The ability to do this opens the target system up to side channel attacks and DoS attacks.

Linux distributions such as *Kali Linux, ArchAssault, BackBox, etc.* can also be used to scan the network to identify live hosts and open ports on the network. This is useful in identifying potential targets for a subsequent attack. They even have functionality to scan for vulnerabilities on systems. *Ettercap* is another suite of tools that poison the cache of target clients in order to launch MitM attacks, allowing attackers to sniff network traffic [OV03]

External Attacker Threats

CSPs may provide a larger attack surface compared to a user's internal computing environment, as both client and management interfaces typically pass over the Internet. Therefore more opportunities exist to compromise data in transit. Some aspects of data security include:

- Does the CSP provide certificate management and how is this secured and audited?
- Are all management interfaces provided over secure channels?
- Does the CSP require the user to use certificates to provide mutual authentication?
- How robust are the CSPs security algorithms, and do they conform to industry standards?

To make an assessment of a CSPs vulnerability to external threats, the user would need access to the CSPs architecture and attack countermeasures. CSPs may be reluctant to publish these details as they could give valuable insight to malicious third parties, but they may be willing to discuss their architecture under a Non-Disclosure Agreement (NDA), or provide some other method, such as a trusted third party assessment. An alternative to assess the CSPs external vulnerability is penetration testing. Typically CSPs do not allow users to perform unauthorised penetration testing of their Cloud service environment, as these may be construed as actual attacks. However they may provide penetration testing as a service to the user, carried out either by themselves or a trusted third party.

It is important to ensure that all correspondence between the user and CSP, for contractual or operational purposes, can be verified to come from the other party. Malicious third parties may impersonate the user or CSP to gain confidential account information using targeted phishing attacks. Social engineering/phishing attacks need to be guarded against by providing effective and regular training, good access management and/or provision of risk management, and mechanisms to report suspicious correspondence. According to a report by Trend Micro, spear phishing attacks are the most common vector in targeted attacks [T$^+$12].

Data Leakage

CSPs may sub-contract some of their services to other third party vendors, such as SAN storage or data replication. It is important that any CSP sub-contractors are clearly identified as part of the assessment process, and the relevant trust metrics applied to them. A data breach at Target in 2013 that exposed 70 million user's payment card details cost the company $162 million [Fon15] [Cha14]. The breach was traced back to a third party vendor with access to Target's systems.

Data within the boundaries of a CSP must be protected throughout the lifecycle of the data. The most common mechanism employed to protect data at rest or in transit is encryption. With encryption, only users in possession of the keys and that are authorised to access the data may view or modify the data. Therefore, effective, secure key management in combination with strong access control rules and policies are critical to the protection of data.

Even when the CSP decommissions equipment or the user terminates their contract with the CSP, the user's data may still be at risk. Failure to ensure that a user's data is placed beyond use, may lead to the data being maliciously or accidentally accessed by, the CSP, its staff, other tenants on the CSPs environment, or CSP third party sub-contractors providing operational services or decommissioning equipment.

The CSP should document and be able to prove how and when the user's data was put beyond use; either by wiping the storage device or destroying encrypted data keys. When putting data beyond use, it is important that not only the original, but also any copies that have been taken for backup or disaster recovery purposes are included. Such data may be stored at a remote location, and processes must be in place to ensure all copies of the data are put beyond use.

Depending on the jurisdiction(s) of the CSP or its environment, it may be required to make the user's data or metadata available to one or more third party legal entities. Canada considered not allowing government data to be transmitted to data centers within the USA, due to concerns over the powers of the Patriot Act [Cat08].

2.2.2 Data Integrity

When the user's data is stored or transferred within their own data center, they can take measures to ensure their data has not been tampered with. When the data passes to the CSPs environment, it is potentially exposed to additional attack vectors. Data in transit can be hijacked and altered, either whilst traversing the Internet or within the CSPs internal network. Attacker could also modify data stored on the CSPs environment that have access to the same databases or data stores.

Similar to the mechanisms used for the protection of data confidentiality, it is important the strong access control rules and polices are in place. While encryption can reduce the risk associated with confidentiality and integrity, digital signatures can be used in conjunction with encryption to provide assurance that the data has not been manipulated. It is important for the users or organisations evaluating services provided by a CSP to assess the mechanisms the CSP has put in place to isolate data and ensure that it has not been altered.

Data Segregation

There are two main attack vectors to consider where data may be tampered with; in transit across a public network (such as the Internet) or within the internal CSP network, and data at rest in a database or data store.

MitM attacks are not specific to CSPs, however they may allow other malicious users to redirect network traffic from inside the CSPs environment, thus allowing them to tamper with the user's data from the inside. ARP Cache Poisoning attacks can be launched in virtualized environments just as they are in physical ones. Tools like *Cain* can be used to poison the ARP cache of target clients. Another type of attack is Session Hijacking [GSC06], which allows a malicious user to takeover and continue a session with their own data. These attacks allow data to be secretly modified, without the user or CSP being aware of the change.

If another tenant can access the same storage device of the user on the CSPs environment, they may be able to tamper with the data at rest. iSCSI storage uses the CHAP protocol to provide security authentication, but CHAP is known to be vulnerable to MITM attacks with a brute force attack on the password. CHAP authentication is also vulnerable to message reflection attacks. Intercepting and impersonating the users credentials for the database can similarly breach database security.

The methods that the CSP utilizes to segregate data in transit and data at rest, can be used to assess the level of data segregation provided by the CSP.

User Access

Within an organisation's internal IT environment there is usually a high degree of trust that other users accessing the infrastructure are who they say they are and are carrying out activities without malicious intent. Internal IT security systems, staff security training, and the prospect of disciplinary action (in the event of a breach in company policy) help to ensure that the infrastructure, OSs and applications are verified and trusted. With the new paradigm of organisations and users moving data and workloads into a CSPs environment, authentication and authorization becomes more of a challenge. Users are now dealing with an external entity, and the CSP needs to verify that any user requesting access to resources is who they say they are and can be trusted.

The challenge is made more complex if users manage their CSPs systems, applications and accounts over a public channel, as opposed to an established VPN connection. By snooping on the traffic between a target user account and the CSP, an attacker might be able to gain enough Personally Identifiable Information (PII) or credential details to impersonate the user on the CSP. Another approach that does not require monitoring network traffic for unsecured transactions is to identify weaknesses in the security system (whether that is the cryptography used, the Identity and Access Management (IAM) system used, hand shake protocols, etc.) and launch targeted attacks on them.

There are a number of instances where impersonation of a user's identity over the Internet has led to data security breaches. In 2014, a collection of approximately 500 private images of various celebrities were stolen and posted publicly on the Internet. The images were obtained exploiting a weakness in the Apple Cloud services suite iCloud. The accounts hacked were part of a very targeted attack on the account usernames, passwords & security questions. A separate attack (*iDict*) was launched on iCloud accounts on Jan 1st, 2015. This attack exploited a security flaw that allowed a brute force crack using a dictionary of common passwords. This was done by overriding the maximum number of login attempts that normally block a brute force attack after a set number of tries. Apple was quick to release a patch, which implemented 2 factor authentication in iOS 8 [Rob15].

The CSP needs to have robust authentication and authorization mechanisms to validate users and their requests when managing resources on the CSPs environment. The CSP also needs secure

processes to reset details of a user's account, and most critically their password. Ideally this would be a multi-stage process, which would include emailing the user. Even changing the accounts email address or phone number could allow a malicious third party to hijack the user's identity within the CSPs customer database. The documentation and robustness of these procedures are an important metric when assessing the CSP.

Data Quality

In basic terms the quality of the data relates to the user getting the same, unaltered data out of the CSPs environment in the state it was when it was originally uploaded (or last modified by the user). The integrity of the data may be compromised either by a malicious third party, or by faulty CSP equipment, typically during data transfer. The standard way to ensure fidelity of the data to its original content is to use a one-way cryptographic hash, such as SHA-2. If two hashes are the same the data is considered to be unaltered. Thus the user can store a short hash code locally in their internal infrastructure and use it to ensure the fidelity of any remote data transferred or stored on the CSPs infrastructure. There are several reasons why one-way cryptographic hash may not guarantee the fidelity of the data [MKL09].

- If the hashing algorithm is weak, a malicious third party may be able to modify the data, whilst replicating the same hash value.

- The integrity of the hash value may be compromised, e.g., by substituting the hash value itself. Therefore the provenance of the hash value must be assured.

The user should assess the hash algorithms supported by the CSP and ensure they conform to industry standards, as well as how and where those hashes are stored, accessed and managed.

2.2.3 Data Availability

Data availability may not always have been considered an aspect of data security, but ensuring that data is accessible to an authorized user in a timely manner, is as critical as ensuring that access is denied to unauthorised users and that it is delivered in its original state. In a Cloud computing environment the user entrusts their data to the CSP, and the CSP should provide services to ensure the data is not lost or corrupted in event of a hardware failure, attack or natural disaster. The key provisioning for the availability of data is to provide redundancy so that if a component of the infrastructure delivering the service goes down, secondary systems can come online. It is also of particular importance in the case of natural disaster to maintain secondary data centers at geographically dispersed locations.

To reduce the risk of data loss, CSPs can provide data backup and disaster recovery facilities as a service to customers. Whilst historically it has been prohibitively expensive to guard against zero data loss, a CSP should provide two significant metrics for any data availability service, Recovery Point Objective (RPO) and Recovery Time Objective (RTO). RPO refers to the amount of data that can be lost, and RTO the time required to restore the data after loss. With the advent of in-memory data caches, clustered data stores and virtualized storage, it is possible to implement data availability solutions with very low RPO and RTO metrics.

Change Management Risks

CSPs invest in training their operational staff, but there is always some level of risk due to human error. In 2014, a mis-configuration in a company web server exposed user data of MBIA Inc. including account numbers and balances. Administrative credentials were also exposed that allowed attackers to access data not yet made accessible via a simple web search [Ker14].

To reduce the risk of mis-configuration issues due to human error, the CSP should have an effective change management system defined and implemented to ensure the integrity and availability of their Cloud computing environment. The management of the infrastructure is typically far more complex for CSPs than a conventional data center, as the scale of the environment is significantly broader and takes in the management of multiple customers with differing requirements. This dynamic multi-tenant environment is a key challenge when developing an effective change management process. Some aspects to consider when evaluating a CSPs change management system include:

- Does the CSP have a Configuration Management Database (CMDB)?
- Does the CSP have documented procedures for specifying, sequencing and scheduling changes?
- Does the CSP use "gold standard" templates to standardize the configuration process?
- Does the CSP review proposed changes as part of the change management process?
- Does the CSPs change management system integrate with the sub-contractors change management system?
- Does the CSP document all changes so they can be analyzed in the case of a mis-configuration event?

Denial of Service Threats

In a user's traditional internal IT infrastructure, communications between clients and services are mainly behind the Internet firewall. In a CSP scenario, the user's clients typically access the hosted services via the Internet. This presents a whole new attack surface, which allows malicious third parties to launch DoS and Distributed Denial of Service (DDoS) attacks against the CSP.

End users running services on the CSP may not be the intended targets of these attacks, but the reduction in network bandwidth and processing power may adversely affect their VMs and servers running in the CSPs environment. In early 2015, GitHub servers were subjected to a targeted DoS attack that resulted in massive outages over a 5 day period. The source of the attack was a malicious piece of Javascript, injected into the traffic of sites that use a Baidu analytics service. The script loaded two GitHub pages on an endless loop [Goo15].

The ability to switch between the servers of different providers could help mitigate this type of attack in certain circumstances. In June 2014, Cloud security services provider Incapsula fought off a DDoS attack against an online gambling website that peaked at 100Gbps of malicious traffic. The attack used more than five DDoS attack vectors including SYN flood, Large SYN flood, NTP amplification, DNS flood, and DNS amplification [Gre14].

The CSP should be able (under an NDA) to specify the countermeasures they have in place against DoS and DDoS attacks. VPN tunnels are often used to implement communications channels between the user's internal network and CSPs data centers [TC08]. VPNs can be implemented

using a variety of technologies, but one of the more secure is IPsec using ESP in tunnel mode, which provides an encrypted channel between two mutually authenticated secure endpoints. Typically authentication is implemented using X.509 certificates, which provide strong authentication of the endpoints [Gut02].

Physical Disruption

A disaster recovery site set up as a hot site in sync with the main data center can provide rapid response times in order to re-establish services at the last known good state. This can provide some protection against local hardware failures or power outages, but if both data centers are subject to the same attack or natural disaster there is a very real threat to the system availability at both locations:

- A Gigaom report noted that in the US, the greatest numbers of data centers are located in states that also experienced the greatest number of disaster locations [Mol13].
- In 2012, Hurricane Sandy caused disruption to network infrastructures in the New York and New Jersey areas (mostly due to flooding and power outages).
- An outage at Amazon in 2012 was caused by a series of failures in the power infrastructure in a Virginia data center. The outage affected many AWS customers [Mil12].

2.3 Commonly Used CSP Assessment Resources

When assessing a CSP, it can be difficult for customers to accurately compare the services and features that are being offered, and to verify the implementation of those services and features. There is also a semantic mismatch between the services and features customers are looking for and how CSPs present their offerings. The following resources can be useful to customers looking for a standardised approach to qualify CSP offerings.

2.3.1 CloudAudit

This allows users of CSPs to perform automated audits, assertions, assessments and assurances via an open and extensible secure interface. CloudAudit addresses the issue of transparency, which is a major challenge when comparing a CSPs services[2].

2.3.2 Cloud Controls Matrix

This provides a baseline set of security controls aimed at helping companies assess the risk associated with a CSP. Guidance is offered in 16 domains including application security, IAM, encryption & key management [3].

[2] `http://cloudaudit.org/CloudAudit/Home.html`
[3] `https://cloudsecurityalliance.org/research/ccm/`

2.3.3 CloudTrust Protocol

This is a mechanism by which Cloud service consumers request and receive information about the elements of transparency as applied to CSPs. The purpose is to provide evidence-based confidence that everything that is claimed to be happening in the Cloud is happening as described [4].

2.3.4 Common CSP Standards

CSPs are required to conform to different industry standards depending on the jurisdiction they operate in or where the data centers are located. Some of the common standards include:

- A statement of the organization's commitment to privacy
- The type of information collected (name, address, credit card, phone numbers, etc.)
- Retaining and using e-mail correspondence
- Information gathered through cookies and Web server logs
- How information is shared with external partners
- Tools used to secure communications (e.g., encryption, digital signatures)
- Access control mechanisms
- Protection of PII stored by an organization
- Policy compliance & auditing

Three influential industry standards are HIPAA, HITECH Act (healthcare industry) and the Payment Card Industry Data Security Standard (PCI DSS). In the EU, the Data Protection Directive (95/46/EC) provides the regulations and standards required for data protection within EU and on its borders [Dir95].

In 2012 a major reform of the EU legal framework on the protection of personal data was proposed. The EU Data Protection Regulation framework aims to strengthen individual rights and find solutions to the problems of globalization and new technologies [Com12]. The new framework is still in development with it expected to come into force in 2018.

A major new requirement under the proposed rules is the "right to be forgotten". This means that if you no longer want your data to be processed and there is no legitimate reason for a company to retain it, the data must be deleted. Transparency is another key area reviewed by the new framework. The framework should allow users to control their personal data in a much easier way, so that they can view what personal information companies hold and so that they can easily transfer that data between providers.

2.4 Security Relevant Metrics: State of the Art and SLA Basis

The main objective of an SLA, in the context of Cloud computing, is to clearly define relationships and set expectations for service/security levels between the CSP and the Cloud service consumer (CSC). A traditional SLA is a rigid and custom contract with complicated legalese focused around

[4] `https://cloudsecurityalliance.org/research/ctp/`

operational metrics provided by the internet provider and using the providers' internal resources. A Cloud SLA is distinct from this traditional form mainly because Cloud customers leverage the Cloud as an extension of their internal IT infrastructure. They do not own the Cloud infrastructure, they do not maintain it, and they cannot control its provisioning or maintenance procedures.

The Cloud's shared responsibility model splits the responsibility between the Cloud provider and the Cloud customer. Depending on the service model (SaaS, PaaS, IaaS) the customer may responsible for the applications, OSs and even parts of the underlying infrastructure. Conversely, for the CSP the responsibility is primarily for the provisioning of the infrastructure. Again, for the different service models, the CSP may take on responsibility for the OSs and applications. Fundamentally, an SLA becomes a composition of operational attributes (e.g., throughput, latency, load characteristics), security attributes (e.g., encryption, data destruction, AAA) and finally, the economic attributes for the delivery of services and the remediation process in the event of a degradation (or complete loss) of services (e.g., "The CSC will accept Service X at Performance Level Y with Cost Z") across the Cloud providers and users.

Prior to considering new trust metrics, in particular techniques for assessing metrics to compare CSPs, the initial effort of D4.1 targeted compiling the State of the Art list of trust metrics offered by CSPs, standards bodies and from ongoing EC projects. There are four issues worth highlighting in order to meaningfully parse this list of metrics:

1. The list does not target completeness. It represents the state of the art that is currently proposed from the industry and standards bodies. As mentioned earlier, without clear adoption from the CSPs, a "novel" metric does not get accepted irrespective of its technical merits.

2. The listed metrics are notably diverse both for the technical elements they cover and for the level of abstraction they deal with.

3. The metrics are typically CSP-centric metrics that are difficult to parse and interpret for users lacking extensive security experience. There are very few consistent definitions that exist for the metrics.

4. A typical SLA includes only a small subset of the listed metrics. Three sub-issues arise:

   (a) Different CSPs offer different subsets thus making comparative assessments hard.

   (b) There is no single trust assessment that applies across CSCs. Each CSC's SLA with the CSP represents unique trust levels.

   (c) There are an insufficient number of guidelines available to ascertain which "set of metrics" is relevant for different use cases, e.g., for data at rest, data in transit, authentication etc.

Sources behind the security/privacy metrics compilation:

- NIST SP 800-55 v1 [5]
- Center for Internet Security (CIS) [6]
- NIST RATAX WG, ISO/IEC SC38, C-SIG WG PLA

[5] csrc.nist.gov/publications/nistpubs/800-55-Rev1/SP800-55-rev1.pdf
[6] benchmarks.cisecurity.org/downloads/metrics/

- EC FP7 A4Cloud [7]
- EC FP7 Cumulus [8]
- EC FP7 SPECS [9]

2.4.1 SoA for Metrics, and Associating Attributes Relevant to ESCUDO-CLOUD Use Cases

Document D1.1 analyzes the four use cases of ESCUDO-CLOUD and extracts their requirements. The four use cases cover the following areas a) facilitating user data security in the Open-Stack framework; b) managing secure data sharing in databases; c) secure federated data storage in the Cloud; d) secure user data storage in an elastic Cloud environment.

In order to relate the state of art in security/privacy metrics to the ESCUDO-CLOUD use cases, a table was created in D1.1 which lists the relevant state of the art metrics, and the standards or official bodies they are derived from. D1.2 elaborates on D1.1 and tries to synthesize the common elements of each use case to produce the common requirements. Each use case covers a real world problem and considers the security challenge to protect user's data in Cloud environment that are not addressed for the today technology.

The table in D1.1 maps the metrics to the ESCUDO-CLOUD use case requirements. It is worth noting that (a) not all trust metrics are relevant to the requirements of the ESCUDO-CLOUD use cases, and (b) that some metrics are prominent in relating to multiple use cases. In order to represent the requirements mappings between the metrics and the use cases, we utilize the format of REQ-UC♯-attribute to refer to the use case "Requirements-Use Case reference-UC attribute" in the table. Table 2.2 is an extract of the original created in D1.1. Only the metrics that are matched with ESCUDO-CLOUD use cases requirements are shown. For completeness, the full list of trust metrics appears in Appendix 2.

Table 2.2: ESCUDO-CLOUD UC Relevant Trust Metrics

ID	Trust Metric Name	Source	Use Case Ref
7	Configuration Changes Measure	NIST SP800-55-v1	REQ-UC2-AC-5 REQ-UC4-AC-5
9	User Accounts Measure	NIST SP800-55 v1	REQ-UC2-AC-1 REQ-UC2-AC-2 REQ-UC2-AC-6 REQ-UC4-AC-1 REQ-UC4-AC-3 REQ-UC1-IKM-2 REQ-UC1-TKM-2 REQ-UC3-KM-1 REQ-UC3-SO-1 REQ-UC3-KM-4
10	Incident Response Measure	NIST SP800-55-v1	REQ-UC4-SS-5

[7] www.a4cloud.eu
[8] www.cumulus-project.eu
[9] www.specs-project.eu

11	Maintenance Measure	NIST SP800-55-v1	REQ-UCI-SKM-2 REQ-UCI-SKM-3 REQ-UC2-AC-4 REQ-UC2-EQ-1 REQ-UC2-EQ-2 REQ-UC2-EQ-3 REQ-UC2-EQ-4 REQ-UC4-SS-1 REQ-UC4-SS-2 REQ-UC4-SS-3 REQ-UC3-KM-3 REQ-UC3-SO-2 REQ-UC3-SO-3 REQ-UC3-SO-4 REQ-UC3-SO-5
12	Media Sanitization Measure	NIST SP800-55-v1	REQ-UC1-IKM-4 REQ-UC1-TKM-4
16	Risk Assessment Vulnerability Measure	NIST SP800-55-v1	REQ-UC4-AC-6
18	System and Communication Protection Measurement	NIST SP800-55-v1	REQ-UC1-TKM-3 REQ-UC2-EQ-3
19	System and Information Integrity Measurement	NIST SP800-55-v1	REQ-UC2-AC-5 REQ-UC2-AC-6 REQ-UC2-KM-1 REQ-UC2-KM-2 REQ-UC2-KM-3 REQ-UC2-KM-4
24	Patch Policy Compliance	The Center for Internet Security	REQ-UC4-SS-4
26	Risk Assessment Coverage	The Center for Internet Security	REQ-UC2-KM-2 REQ-UC2-EQ-2
40	Percentage of Configuration Compliance	The Center for Internet Security	REQ-UC2-EQ-1 REQ-UC2-EQ-2 REQ-UC2-EQ-3 REQ-UC2-EQ-4 REQ-UC3-AC-3 REQ-UC3-AC-1
43	Percent of Critical Applications	The Center for Internet Security	REQ-UC4-AC-2
58	Tenant isolation level	EU FP7 Cumulus	REQ-UC3-AC-5
64	User authentication and identity assurance level	EU FP7 Cumulus	REQ-UC4-AC-1 REQ-UC4-AC-2 REQ-UC4-AC-3 REQ-UC4-AC-4 REQ-UC1-IKM-1

			REQ-UC1-TKM-1 REQ-UC3-KM-2 REQ-UC3-AC-2 REQ-UC3-AC-4
66	Password storage protection level	EU FP7 Cumulus	REQ-UC2-KM-3 REQ-UC3-KM-6
67	Cryptographic brute force resistance	EU FP7 Cumulus	REQ-UC2-AC-6 REQ-UC4-SS-5
68	Key access control level	EU FP7 Cumulus	REQ-UC3-AC-6
70	Data deletion quality level	EU FP7 Cumulus	REQ-UC4-SS-4 REQ-UC1-IKM-4 REQ-UC1-TKM-4
94	Level of confidentiality	EU FP7 A4Cloud	REQ-UC4-SS-4 REQ-UC4-DE-1 REQ-UC4-DE-2 REQ-UC4-DE-3 REQ-UC4-DE-4 REQ-UC3-DE-2
95	Key Exposure Level	EU FP7 A4Cloud	REQ-UC2-KM-1 REQ-UC2-KM-2 REQ-UC2-KM-3 REQ-UC2-KM-4 REQ-UC3-DE-1 REQ-UC3-DE-4 REQ-UC3-DE-5 REQ-UC3-AC-7 REQ-UC3-KM-5
123	Cryptographic strength	EU FP7 SPECS	REQ-UC2-KM-3 REQ-UC2-KM-4 REQ-UC1-IKM-3 REQ-UC1-TKM-3
131	FIPS compliance	EU FP7 SPECS	REQ-UC3-DE-3
136	E2EE Crypto Strength	EU FP7 SPECS	REQ-UC2-EQ-2 REQ-UC4-DE-4
140	Backup	EU FP7 SPECS	REQ-UC4-SS-5 REQ-UC1-SKM-1

This is an initial association of metrics to ESCUDO-CLOUD requirements arising from D1.1 and D1.2 which is in progress at the time of writing. It is intended that these requirements will be prioritised according to their relation to the ECUDO-CLOUD security dimensions. Hence, D4.1 will utilize the D1.1 prioritization to refine the metrics-requirements associations. A potential consequent approach is to synthesize higher level metrics as a composition of these initial metrics. The main objective of ESCUDO-CLOUD is to secure the end users data stored on the CSP using different technologies and processes. As shown in Tables 2.2 and 2.3, the encryption requirements are essential to assure the security of the end user data.

Table 2.3 shows the prioritized list of trust metrics identified for the ESCUDO-CLOUD use

cases. The number of use cases related to each metric is highlighted to clearly show the level of impact of each metric. The metrics with the highest number of use case requirements are built on the underlying security requirements of ESCUDO-CLOUD (encryption for confidentiality, access control mechanisms). Those metrics can be related to ESCUDO-CLOUD dimensions as well and, as expected, focus on security properties (confidentiality, integrity and availability).

All the use cases have some common functionality and rely on some shared infrastructure capabilities. For applications running on the Cloud, one of the main requirements outlined in the use cases is the acceptable level of availability of the service. As a direct correlation of this, the metric with higher number of references is Maintenance Measure (NIST SP800-55-v1).

Table 2.3: Prioritized List of UC Relevant Trust Metrics

ID	Trust Metric Name	Source	#UC Ref
11	Maintenance Measure	NIST SP800-55-v1	15
9	User Accounts Measure	NIST SP800-55-v1	10
64	User authentication and identity assurance level	EU FP7 Cumulus	9
95	Key Exposure Level	EU FP7 A4Cloud	9
19	System and Information Integrity Measure	NIST SP800-55-v1	6
94	Level of confidentiality	EU FP7 A4Cloud	6
123	Cryptographic strength	EU FP7 SPECS	4
70	Data deletion quality level	EU FP7 Cumulus	3
66	Password storage protection level	EU FP7 Cumulus	2
67	Cryptographic brute force resistance	EU FP7 Cumulus	2

2.5 Standards that Industry has to Conform to for Security Metrics

There exist multiple international standard development organizations (SDOs) in the domain of Cloud application and service deployments. Particularly with regards to security and privacy issues, the prominent organizations are listed below:

- NIST Cloud Standards (NIST); `www.nist.gov`
- Cloud Security Alliance (CSA); `cloudsecurityforum.org`
- Distributed Management Task Force (DMTF); `www.dmtf.org`
- Storage Networking Industry Association (SNIA); `www.snia.org`
- Open Grid Forum (OGF); `www.ogf.org`
- Open Cloud Consortium (OCC); `opencloudconsortium.org/`
- Organization for the Advancement of Structured Information Standards (OASIS); `www.oasis-open.org/`
- TM Forum (TMF); `www.tmforum.org/`
- International Telecommunication Union (ITU); `www.itu.int/`
- The European Telecommunications Standards Institute (ETSI); `www.etsi.org/`

- Object Management Group (OMG); www.omg.org/
- Association for Retail Technology Standards (ARTS); nrf.com/
- Institute of Electrical and Electronics Engineers (IEEE); www.ieee.org/
- Alliance for Telecommunications Industry Solutions (ATIS); www.atis.org/
- Internet Engineering Task Force (IETF); www.iso.org/
- International Standards Organization (ISO/IEC); www.iso.org/
- National Vulnerability Database (NVD) - Common Weakness Enumeration (CWE); nvd.nist.gov/cwe.cfm

It is difficult to establish appropriate security metrics in the Cloud networks, but it is possible to re-evaluate best practices and develop standards to ensure the deployment and adoption of secure Cloud. Some of these properties are described below:

- **Authentication and Identity Management**: These features permit distinguishing the types of users and allow/forbid the access depending on authorized or unauthorized users.
- **Access Control**: Authorization to allow/deny user access to the Cloud services.
- **Secure Interoperation**: Ensures communication between different components of the network.
- **Secure-Service Provisioning and Composition**: Use of virtualization technologies that separate application services from infrastructure.
- **Trust Management Framework**: Trust-based framework that facilitates the policy integration.
- **Information-centric security**: Ensure that the stored information is safe.
- **High-assurance remote server attestation**: Ensures that data is not being abused or leaked.
- **Privacy-enhanced business intelligence**: Method that encrypts all data stored in the Cloud.

References

http://www.cert.uy/wps/wcm/connect/975494804fdf89eaabbdab1805790cc9/Cloud_Computing_Vulnerability_Incidents.pdf?MOD=AJPERES[1]
http://csis.pace.edu/~marchese/SE765/Paper/security2.pdf
https://cloudsecurityalliance.org/guidance/csaguide.v3.0.pdf

3. SLA Based Metrics and Assessment - Approaches

As mentioned in Chapter 1, the SLA trust metrics need to be compatible with standards for acceptance in the community. Interestingly, Section 2.4.1 reveals no shortage of existing and proposed metrics - both for diversity of attributes and for the abstraction/implementation level of the metrics. Consequently, generating yet-another-metric is likely of limited value as it only adds to the plethora of metrics, and more importantly needs industry/standards buy-in to be successful, which is a tediously complex process.

However, we make a fundamental observation (to define D4.1's approach) that, as state of the art/practice, virtually all existing security attributes that are factored into composing an SLA are linear in nature, i.e., an SLA lists and considers trust attributes "*x, y, z*" as discrete attributes. Thus, the SLA is a simple representation of either the existence or absence of a trust attribute *x* or *y* or *z*. In reality, what is desired is a composition where the Cloud user can specify "combinations", i.e., some groupings of attributes that are mandatorily required (e.g., AND operator) and some groupings that are optional (e.g., OR operator). This simple observation allows us to utilize the wide variety of existing (and emerging) metrics from the standards bodies to compose complex logical aggregations to represent an extremely wide range of trust specification in SLAs. In addition, these compositions also allow us to consider matching groups of (actual or speculative) trust attributes to do compatibly fair comparisons across CSP capabilities. It is this ideology that D4.1 will develop, first for the classical Cloud models in D4.1, and subsequently expand for multi-Cloud models in D4.2.

The rest of the section organized as follows: Section 3.1 highlights the basic terminology related to Cloud SLAs. Section 3.2 outlines the SLA-based quantitative assessment methodology.

3.1 The SLA Elements

This section summarizes the basic Cloud SLA terminology, based (where applicable) on the latest version of the relevant ISO/IEC 19086 standard [ISO14]. A Cloud SLA is a documented agreement between the CSP and the CSC that identifies Cloud services and service level objectives (SLOs), which are the targets for service levels that the CSP agrees to meet. If an SLO defined in the Cloud SLA is not met, the CSC may request a remedy (e.g., financial compensation). If the SLOs cannot be (quantitatively) evaluated, then it is not possible for CSCs or CSPs to assess if the agreed SLA is being fulfilled. This is particularly critical in the case of SLAs, but it is also an open challenge on how to define useful (and quantifiable) security SLOs.
In general, an SLO is composed of one or more metrics (either quantitative or qualitative), where the SLO metrics are used to set the boundaries and margins of errors CSPs have to abide by (along with their limitations). Considering factors such as the advocated familiarity of practitioners with

security controls frameworks (e.g., ISO/IEC 27002 [ISO13], the Cloud Security Alliance's Cloud Control Matrix (CCM) [CCM14], and the National Institute of Standards and Technology NIST-SP 800-53 [NIS14]), the relevant workgroups (e.g., the EC's Cloud Select Industry Group on Service-Level Agreements C-SIG SLA in [SLA14]) have proposed an approach that iteratively refines individual controls into one of more measurable security SLOs. The elicited SLOs metrics can then be mapped into a conceptual model (such as the one proposed by the members of the NIST Public RATAX Working Group [NIS15]), in order to fully define them.
Based on our analysis of the state of practice, Cloud SLAs are typically modeled using the hierarchical structure shown in Figure 3.1. The root of the structure defines the main container for the SLA. The second and third levels represent the Control Category and Control Group respectively, and they are the main link to the security framework used by the CSP. The lowest level in the SLA structure represents the actual SLOs committed by the CSP, whose threshold values are specified in terms of security metrics.
For example, in Figure 3.1, let us suppose that a CSP implements the SLA Control "Entitlement

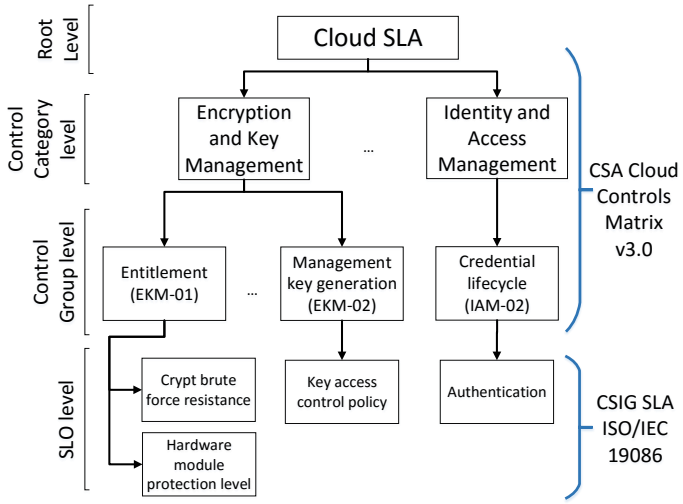

Figure 3.1: The Cloud SLA hierarchy

(i.e., EKM-01)" from the CSA CCM2. This control is actually contained within the group "Encryption and Key Management (i.e., EKM)". After selecting EKM-01, the same CSP then refers to the SLO list provided on the C-SIG SLA report [SLA14] (or any other relevant standard) and finds out that two different SLOs are associated with control EKM-01, i.e., "Cryptographic Brute Force Resistance" and "Hardware module protection level". Both SLOs are then refined by the CSP into of one of more security metrics, which are then specified as part of the SLA offered to the Cloud user. For example, a CSP can commit to a "Cryptographic Brute Force Resistance" measured through security levels such as $(level_1, \ldots, level_8)$, or through a metric called "FIPS compliance" defined as Boolean yes/no values. Therefore, the SLA could specify two SLOs: (Cryptographic Brute Force Resistance = $level_4$), and (FIPS compliance = yes). If any of these committed values is not fulfilled by the CSP, then the SLA is violated and the user might receive some compensation (this is the so-called SLA remediation process).
Using the presented approach, the security SLOs proposed by the CSP can be matched to the CSC's requirements before acquiring a Cloud service. Actually, these SLOs provide a common semantic that both CSCs and CSPs can ultimately use to automatically compare and negotiate

Cloud SLAs. As a note, the process presented in this section to elicit security SLOs (that will become part of the CSPs SLA) and is being also used by standardization bodies such as ISO/IEC and industrial working groups as C-SIG SLA.
In real-world Cloud scenarios, the process described in this section should take into account that most Cloud services have horizontal (e.g., Cloud supply chains) and vertical (e.g., different Cloud service model layers) dependencies. Thus, it does not suffice to understand how the Cloud service under one unique CSPs control may affect its own users, but one also needs to consider how the sub-services/CSPs contribute to the overall security level. Hence, there is a distinct need for the aggregation of security metrics guaranteed by individual Cloud services in order to get the values for a composite one. While practitioners have acknowledged the challenges associated with the composition of security metrics long before the "Cloud times", this topic is still mostly unexplored in Cloud systems. On this background, the following section presents two approaches to aggregate and evaluate Cloud security levels based on SLAs.

3.2 Quantitative Assessment of Cloud SLAs

The quantitative security-level assessment of CSPs based on SLAs (for their match to the CSC requirements) is the primary objective of the techniques to be developed in D4.1, namely the Quantitative Policy Trees (QPT) and the Quantitative Hierarchy Process (QHP). Using this assessment, the CSPs are ranked (as per their SLAs) for the best match to the CSC requirements. QPT utilizes a logical aggregation of security quantifiers, while QHP is based on multi-variable optimization techniques considering the various elements of a SLA as the optimization criteria. The coverage of these issues in D4.1 will follow the structure as:

- Detail the standalone operations of the QPT and QHP techniques from the SLA perspective, and subsequently discuss guidance for their usage discretely and collectively
- Apply SLA assessment and the ranking of CSPs in progressive stages (common to both QPT and QHP techniques)
- Demonstrate viability of SLA based comparing of CSP trust levels.

As an overview of the two techniques, the SLA assessment and the ranking of CSPs are performed in progressive stages (common to both QPT and QHP techniques), as shown in Figure 3.2.
In Stage (A), we express in a common way both the user's requirements and the CSPs committed SLOs using a standardised SLA template (e.g., based on ISO/IEC 19086 [ISO14]). In Stage (B), the user's requirements and CSPs SLA are quantitatively evaluated. This quantitative data is then used in Stage (C) as input to a ranking algorithm, in order to provide the final assessment result. We detail each of the two techniques (QPT and QHP) in the subsequent sections.

3.2.1 Quantitative Policy Trees

Luna et al. [LLS12] proposed the use of a tree-like data structure (i.e., the Quantitative Policy Tree), to model a CSPs security policy in order to numerically evaluate it with respect to a set of user's requirements. While the original QPT was designed to evaluate security control frameworks such as CSA CCM [CCM14], this section develops an extended QPT approach for the quantitative evaluation of Cloud SLAs.

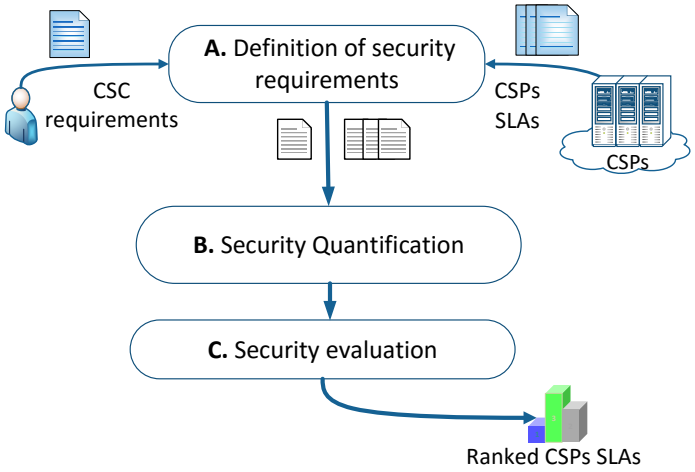

Figure 3.2: Stages comprising the quantitative SLA assessment.

Stage A. Definition of Security Requirements

The QPT is an AND/OR tree where the Cloud user's requirements are represented also as a security SLA (called *user SLA*). The Control Categories, and Controls are represented as intermediate nodes of the tree, while security metrics associated with SLOs are represented as *weighted* leaf nodes. Assigned weights are used to represent the relative importance of SLOs from the user's perspective (e.g., for some users the SLO metric "Encryption Key Size" might be more important than SLO metric "Backup Frequency"). The basic rules for setting weights on the *user SLAs* individual security SLOs are:

- Each user required security SLO will be associated with a quantitative weight ω_i ($0 \leq \omega_i \leq 1$).
- The sum of all the weights ω_i associated with a set of sibling security SLO metrics (i.e., those having the same parent Control) must be equal to 1.
- The user can choose specific elements of the SLA (cf., Section 3.1) to benchmark by assigning $\omega_i = 0$ to those not of interest.

To complete the customization of a *user SLA*, the user can also select the appropriate AND/OR relationships[1] between the different Control Categories, Controls and SLO metrics of the SLA. As inferred from their name, AND relationships will model *hard-requirements* where "Categories A, B and C are *all* required due to regulatory compliance", whereas OR relationships are more adequate to model *soft-requirements*, e.g., "Either A, B or C are needed to achieve my security goals". The overall QPT creation process is shown in Figure 3.3.

[1]Multi-level aggregations, correlations and complex logical operators are possible as per the needs of the security characterization. We have limited the presentation for ease of presentation of the concept to the basic case of binary tree with AND/OR operations. For complex logics, the aggregation rules of Table 3.1 need to be extended as needed for the desired logical composition.

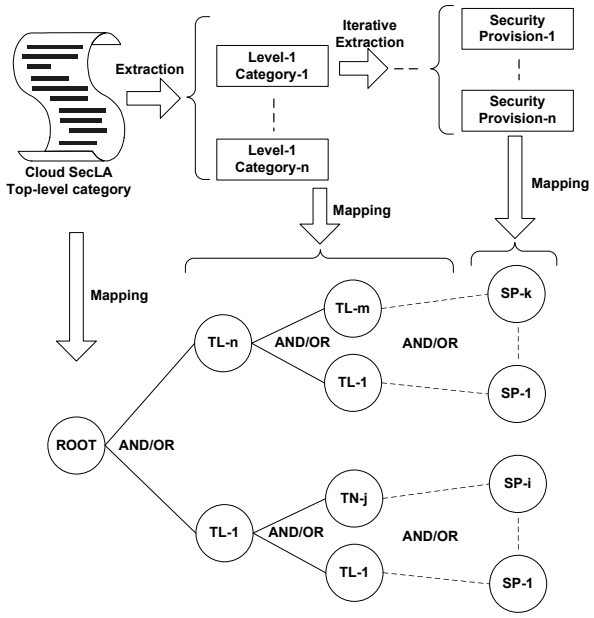

Figure 3.3: SLA-to-QPT: mapping a Cloud SLA into a QPT[LLS12].

Table 3.1: Aggregation rules for a QPT with n-sibling nodes [LLS12]

Parameter	Aggregation rule with $i = 1 \ldots n$	
	AND node	OR node
$Agg_{ParentL1}$	$\sum_{i=1}^{n}(LSL_i \times \omega_i)$	$min(LSL_i \times \omega_i)$
$Agg_{ParentL2}$	$\sum_{i=1}^{n} Agg_{ParentL1,i}$	$min(Agg_{ParentL1,i})$

Stage B. Security Quantification

In order to evaluate the *Cloud user SLA* (termed as the user QPT) with respect to the offered *CSP SLA* (CSP QPT), the QPT utilizes the notion of local security levels (LSL) [CPRT05] and two basic assumptions: *(i)* all the *i-leaf nodes* on the QPT have been already associated with a $LSL_i > 0$ and, *(ii)* there exists a maximum value LSL_{max} that is the same for *all* the leaf nodes of the QPT (i.e., $0 < LSL_i \leq LSLmax$).

Once each leaf node in the QPT has been associated with the tuple $\{LSL_i, \omega_i\}$, it is possible to propagate these values to the rest of the tree using the aggregation rules shown in Table 3.1. Notice that QPT's AND/OR relationships allow modelling metrics/SLOs with direct dependencies, where low-level metrics (i.e., abstract metrics according to NIST [NIS15]) can be composed into more advanced/high-level ones (i.e., concrete metrics [NIS15]).

Stage C. Security Evaluation

Once the *Cloud user QPT* and the *CSP QPT* have been populated with the aggregated values (quantitatively computed from Table 3.1), it is possible to apply a ranking process to determine how different CSPs under-/over-provision a user's requirement. Luna et al. [LLS12] propose two different classes of benchmarks, namely $QuantB_{node}$ (cf. Definition 1) and $QualB_{node}$ (cf. Definition 2), both based on the quantitative security values already aggregated in the QPT.

Definition 1 *The quantitative benchmark $QuantB_{node}$ associated with a specific node of the QPT, is defined as follows:*

$$QuantB_{node} = \frac{Agg_{CSP,node} - Agg_{user,node}}{Agg_{max,node}}$$

Where:

- $Agg_{CSP,node}$ *is the aggregated security value for node in the* CSP QPT*, as computed with Table 3.1.*
- $Agg_{user,node}$ *is the aggregated security value for node in the* user QPT*, as computed with Table 3.1.*
- $Agg_{max,node}$ *is is the aggregated security value for node in either* user QPT *or* CSP QPT*, as computed with Table 3.1 and using the maximum Local Security Level (LSL_{max}).*

Definition 2 *The following expression defines $QualB_{node}$, the qualitative benchmark associated with a specific node of the QPT:*

$$QualB_{node} = \begin{cases} \lceil QuantB_{node} \times Ranks_{max} \rceil & if\ QuantB_{node} \geq 0 \\ \lfloor QuantB_{node} \times Ranks_{max} \rfloor & if\ QuantB_{node} < 0 \end{cases}$$

Where:

- $QuantB_{node}$ *is the quantitative benchmark in Definition 1.*
- $Ranks_{max}$ *is the total number of chosen qualitative labels minus one. For example, if the set of qualitative labels is {"Copper", "Silver", "Gold"} then* $Ranks_{max} = 2$

The result of the previous QPT metric is an integer number such that
$QualB_{node} \in \{-Ranks_{max}, \ldots, 0, \ldots, Ranks_{max}\}$.
In order to assign it a qualitative label from the set $Ranks = \{Label_1, \ldots, Label_n\}$ where $n = Ranks_{max} + 1$, we use the following mapping function:

$$f(QualB_{node} \mapsto Ranks) = \begin{cases} Label_1 & \text{if } QualB_{node} = 0 \\ Label_2 & \text{if } QualB_{node} = 1 \\ -Label_2 & \text{if } QualB_{node} = -1 \\ \vdots & \\ Label_n & \text{if } QualB_{node} = Rank_{max} \\ -Label_n & \text{if } QualB_{node} = -Rank_{max} \end{cases}$$

In the previous function, notice that a "negative" label such as $-Label_n$ literally represents the counterpart of the corresponding "positive" label $Label_n$. For example "A-" and "A+", "Silver" and "Silver-", and so forth.

3.2.2 Quantitative Hierarchy Process (QHP)

The quantitative security assessment of CSPs using control frameworks is the primary objective of the Quantitative Hierarchy Process (QHP), as originally introduced in [TTLS14]. By applying the QHP assessment technique, the CSPs can be ranked (as per their offered security controls) depending on how well they match the user's requirements. QHP allows Cloud users to (i) compare,

benchmark and rank the aggregated security level provided by two or more CSPs, (ii) provide a composite quantitative and qualitative security assessment technique based on the well-known Analytic Hierarchy Process (AHP) [Saa90] (depending on the user defined security requirements and priorities), (iii) allow users with different levels of security expertise to specify their security requirements at varied levels of granularity, and (iv) automate the overall assessment process.

Similar to the QPT, the SLA assessment and ranking of CSPs is a proposed extension of the original QHP [TTLS14], as developed in the progressive stages described below.

Stage A. Definition of Security Requirements

In this stage, the Cloud user creates his set of security requirements based on the same SLA template (structure) used by the CSPs to specify their security offers. The SLA template will have the structure presented in Section 3.1 (i.e., from Control Categories to individual security metrics associated with committed SLOs).
The user-defined requirements are distinctive elements of a Cloud SLA, where all the elements are weighted or evaluated in order to represent their relative importance from the user's perspective. For example, the (prospective) Cloud user might specify that some specific Control is "Very Important", or even request a specific key length value for an "Encryption Key" SLO metric. The output of this stage will be a set of user security requirements specified as a SLA.

Table 3.2: Used terms definitions

Term	Definition
k	security metric associated with the SLO.
CSP_i	Cloud provider i, such that $i \in \{1,\ldots,n\}$, where n is the total number of CSPs.
V_i	SLO value for based on metric k, and provided by CSP_i (CSP_i provides k with value V_i).
CSC	Cloud Service Consumer (for simplicity we use Cloud user in the case study calculations).
V_{CSC}	Consumer requested value for SLO metric k.
W	relative rank ratio.
CSP_1/CSP_2	Relative rank W of CSP_1 over CSP_2, regarding k. Or relative rank $1/W$ of CSP_2 over CSP_1, regarding k.
CSP_i/CSC	Relative rank of CSP_i over CSC, which specifies if CSP_i satisfies CSC requirements, with respect to k.

Stage B. Security Quantification

In order to evaluate the CSC requirements with respect to a CSP SLA, the so-called measurement model for different security SLO metrics needs to be defined. Over this stage, different comparison metrics for different types of requirements are defined, so they can be applied for the quantitative security assessment. The terms shown in Table 3.2 are used to present the QHP framework.

Definition 3 *The relationship across the CSPs with respect to a security SLO (k) is represented as a ratio:*

$$CSP_1/CSP_2 = \frac{V_1}{V_2}$$

Where CSP_1/CSP_2 indicates the relative rank of CSP_1 over CSP_2, regarding k as indicated in Table 3.2. The security SLOs metrics under evaluation can be Boolean (e.g., a yes/no representing the need of a security mechanism) or numbers (e.g., a cryptographic key length) such that:

- **Boolean:** In this case the CSPs yes/no SLO's metric values are defined as Boolean *true* and $false$ or 1 and 0, respectively. The relationship across the CSPs with respect to security SLO metric value (V) based on Definition 3 can be represented as:

$$CSP_1/CSP_2 = \begin{cases} 1 & \text{if } V_1 = 1 \\ 0 & \text{if } V_1 = 0 \end{cases}$$

- **Numerical:** Assume e.g., a cryptographic key length (in bits) defined as k and specified by $\{64, 128, 256, 512, 1024, 2048\}$, such that $64 < 128 < 256 < 512 < 1028 < 2048$, which is defined as $level_1, level_2, level_3, level_4, level_5, level_6$. The security levels are modelled as $\{1,2,3,4,5,6\}$ respectively, such that $1 < 2 < 3 < 4 < 5 < 6$. Thus, the relationship across the CSPs with respect to the security SLO value (V) based on Definition 3 can be represented as:

$$CSP_1/CSP_2 = \begin{cases} 1 & \text{if } V_1 \equiv V_2 \\ W & \text{if } V_1 > V_2 \\ \frac{1}{W} & \text{if } V_1 < V_2 \end{cases}$$

Stage C. Security Evaluation

Given the fact that a SLA might have a high number of individual security SLOs and that CSCs might specify their requirements with different levels of granularity, the challenge is not only how to quantify different metrics associated to these SLOs, but also how to aggregate them in a meaningful way. To solve these challenges, QHP's ranking mechanism is based on AHP [Saa90] for solving Multiple Criteria Decision Making (MCDM) [Zel82] problems.
The AHP-based methodology for CSP rankings consists of four main steps: (1) hierarchy structure (2) weights assignment (3) pairwise comparison and (4) SLOs aggregation to give the overall rank calculation. These steps are summarised next.

1) Hierarchy Structure

The SLAs are modelled as a hierarchical structure (cf., Figure 3.1), such that the top-most layer of the hierarchy structure defines the main goal and aims to find the overall rank (i.e., the root "SLA-level"). The lowest level is represented by the actual security metrics related to the committed SLO value.

2) Weights Assignment

CSC-defined weights are assigned to the different levels of the SLA hierarchy to take into account their relative importance. QHP considers two types of weights:

- *User assigned qualitative values.* Users assigns the desired weights to each SLO metric to indicate their priorities (High-Important (HI), Medium-Important (MI), Low-Important (LI)). These labels are transformed to quantitative values and assigned as normalized numbers to satisfy the AHP requirements.
- *Using AHP's standard method.* The user can assign numeric weights to each one of the SLA elements using values in some defined scale. For example, the AHP method proposes a scale from 1 to 9 to indicate the importance of one element over another.

3) Pairwise Comparison

In this phase, the relative ranking model defining the most important requirements and their quantitative metrics is specified. This ranking model is based on a pairwise comparison matrix of SLA elements provided by different CSPs as required by the users. Using a Comparison Matrix (CM) for each CSP, a one-to-one comparison of each CSP for a particular attribute is obtained, where CSP_1/CSP_2 indicates the relative rank of CSP_1 over CSP_2. This will result in a one to one comparison matrix of size n x n (if there are a total of n CSPs), such that:

$$CM = \begin{array}{c} \\ CSP_1 \\ CSP_2 \\ \vdots \\ CSP_n \end{array} \begin{array}{c} \begin{array}{cccc} CSP_1 & CSP_2 & \dots & CSP_n \end{array} \\ \begin{pmatrix} CSP_1/CSP_1 & CSP_1/CSP_2 & \dots & CSP_1/CSP_n \\ CSP_2/CSP_1 & CSP_2/CSP_2 & \dots & CSP_2/CSP_n \\ \vdots & \vdots & \ddots & \vdots \\ CSP_n/CSP_1 & CSP_n/CSP_2 & \dots & CSP_n/CSP_n \end{pmatrix} \end{array} \tag{3.1}$$

The relative ranking of all the CSPs for a particular SLO metric is given by the eigenvector of the comparison matrix. This eigenvector shows a numerical ranking of CSPs that indicates an order of preference among them as indicated by the ratios of the numerical values, which is called Priority Vector (PV).

4) SLOs Aggregation

In the final phase, the assessment of the overall security level (and consequently the final ranking of CSPs) is obtained using a bottom-up aggregation. To achieve that, the PV of each attribute is aggregated with their relative weights assigned in Step 2. This aggregation process is repeated for all the attributes in the hierarchy along with their relative weights.

$$PV_{aggregated} = \begin{pmatrix} PV_1 & \dots & PV_n \end{pmatrix} \begin{pmatrix} w_i \end{pmatrix} \tag{3.2}$$

where w_i is a Cloud user's assigned weight for criteria i.

3.2.3 QPT and QHP Comparison

Table 3.3 summarises the main features found in both the QPT and QHP methodologies presented in this section. The empirical validation presented in the following section, complements the features shown in Table 3.3 with a set of usage guidelines based on real-world use cases. In this

section, our focus is to introduce a set of criteria aiming to guide early QPT and QHP adopters in aspects related to the requirements of their specific application scenarios:

- As the QPT aggregation is based on AND/OR operations, and as the CSPs ranking (with respect to user requirements) is only executed at the root (highest) level, it clearly has the potential to outperform QHP's aggregation time. In QHP, the CSPs ranking is performed at each level of the SLA hierarchy structure, which means that by increasing the number of SLOs QPT shows better performance regarding aggregation time. This might be a useful feature in scenarios where low-latency is needed, e.g., Infrastructure-as-a-Service (IaaS) scheduling and automation. Section 3.3 will empirically demonstrate this assertion.

- QHP's ability to depict CSPs ranking at each level of the SLA hierarchy gives both CSPs and users the ability to determine which security SLOs are over/under provisioning the user's requirements. This is useful for CSPs to improve their provided SLAs match to the user's requirements.

- QHP's flexibility to represent user requirements at different levels of SLA hierarchy (i.e., from Control Category to individual SLO metrics), makes it more "user-friendly" and suitable for implementations where human-interaction is needed e.g., in a decision making dashboard. QPT can only evaluate security requirements specified at the SLO-level, thus it is better suited for scenarios where users can express security preferences at a very granular level (e.g., software agents negotiating SLAs).

- QPT and QHP can also be used complementarily. For example, QHP can be used by prospective users manually exploring different CSP offers through what-if scenarios. Once a SLA has been agreed upon, then applications can rely on QPT for dynamically negotiating new terms without user intervention.

- QHP relies on a mature set of techniques (i.e., multi-criteria decision analysis or MCDA), which eases its extensibility to add new features with few efforts. For example, we plan to use fuzzy MCDA techniques to add the notion of uncertainty to the security evaluation process.

The next section empirically demonstrates the features of both QPT and QHP based on two use case scenarios.

3.3 QPT and QHP Validation: Case Studies

This section has two main objectives: (a) an empirical validation of QPT and QHP, and (b) demonstrating the advantages and disadvantages of each approach.

The empirical validation is performed through two scenarios that use real world SLAs structured in compliance with the current draft version of the ISO/IEC 19086 standard [ISO14] and with data derived from the Cloud Security Alliance's STAR repository [STA11].

The associated metrics were extracted from the metrics catalog referenced in Section 2.4.1.

By following the refinement approach shown in Figure 3.1 and presented in Section 3.1, our validation approach created a dataset comprised of three Cloud SLAs [2] that were chosen to cover

[2] For confidentially reasons, the name of the CSPs have been anonymised.

Table 3.3: QPT and QHP — comparison of main features

Stage	Feature	QPT	QHP
Security Requirements	SLA granularity for expressing requirements	Weights and values only at SLO level	Weights and values at all levels
	Supported SLO values	Quantitative and Qualitative	
	Template for user requirements	SLA hierarchy	
	Model relationships among SLA elements	AND/OR among SLOs	None
Security Quantification	Base technique for aggregation	Ad-hoc	Multi-criteria decision technique
	Used SLA abstraction	AND/OR Tree	Matrix
Security Evaluation	Output	Ranked List, Overall Security Level	
	Format of resulting security level	Quantitative/ Qualitative	Quantitative

all possible conditions for each SLO (i.e., over/under provisioning or satisfying the Cloud user's requirements). Each SLA contained an overall of 139 SLOs (with both quantitative and qualitative metrics), and with real values corresponding to the CSP information found on the CSA STAR repository.

Figure 3.4 shows the process used to systematically perform the CSP comparison presented in the rest of this section. The overall process consists of four steps (common to both QPT and QHP techniques presented in Section 3.2) namely:

1. Step 1. Cloud user security requirements: in this step the (prospective) Cloud user defines his security requirements (SLO's thresholds and associated weights), and expresses them using a standardised SLA template (e.g., based on ISO/IEC 19086 [ISO14]).

2. Step 2. Quantitative Evaluation and Ranking: the user's security requirements (Step 1) are evaluated with respect to the CSPs SLAs. As shown in Section 3.3.1, QPT and QHP have different capabilities and the decision on which one to use will mainly depend on the Cloud user's degree of security expertise.

3. Step 3. CSP Selection: the output from Step 2 is a set of CSPs ranked with respect to the user's Security Requirements. In this step, any of the CSPs should be selected by the user, otherwise the whole process might be repeated with a refined set of security requirements (Step 4).

4. Step 4. Refine Requirements: this step is used in case the Cloud user decides to change his security requirements (e.g., with new weights assigned to selected SLOs) and repeat once again the whole comparison process, as shown in Section 3.3.2.

Our validation scenarios were designed taking into account real concerns from Cloud users (i.e., procuring Cloud services based on security requirements) and CSPs (i.e., maximising offered security levels). The used data set also consisted of different combinations of requirements and real SLA representing three different users (as shown in Table 3.4), and three different CSPs respectively.

It is important to notice that in compliance with the ISO/IEC 19086 standard [ISO14], the dataset used for our experiments only contained SLAs with elements (controls, SLOs) are independent but keep their compositional nature. Using terminology from this standard, the compositional nature of the SLA is based on top-level components (e.g., cryptography) comprising one of more measurable service commitments (e.g., cryptographic access control policy, key management, and data at rest). Both assumptions (lack of dependencies, compositional nature) are also consistent with the C-SIG SLA guidelines [SLA14]), and the NIST Cloud service metrics model [NIS15].

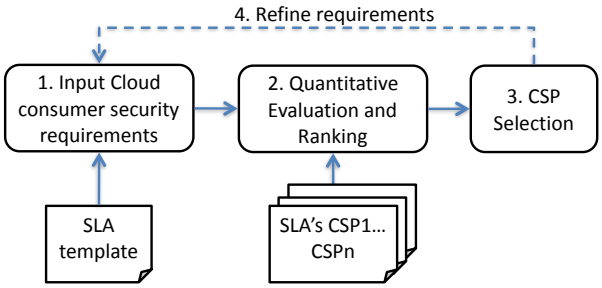

Figure 3.4: Selecting a CSP based on its SLA.

3.3.1 The User Perspective: Security Comparison of CSPs

This initial validation scenario demonstrates how a (prospective) Cloud user can apply the techniques presented in Section 3.2 to compare side-by-side three different CSPs based on their advertised SLAs, and with respect to a particular set of security requirements (also expressed as an SLA).

Table 3.4 presents a sample data set used for this scenario, where based on the information available in the CSA STAR repository [STA11], the values associated to 16 SLO metrics (out of 139) for the three selected CSPs are presented. In order to perform a comprehensive validation, the selected SLOs comprised both qualitative (e.g., yes/no) and quantitative (e.g., security levels from 1 to 4) metrics. The yes/no SLO's thresholds are modelled as Boolean $1/0$, whereas SLOs associated to security levels as $level_1, level_2, level_3, level_4$ are modelled as $\{1,2,3,4\}$. For example, the *CO3.3* SLO is defined using qualitative thresholds (None, Annually, Quarterly, Monthly) which are specified as $level_1, level_2, level_3, level_4$. Similarly, the *RI1.1* SLO is defined using qualitative (Internal, External) values.

Furthermore, Table 3.4 also shows three sets of Cloud user requirements used as baseline for comparing the selected CSPs. For validation purposes the user requirements are being expressed at different levels of granularity (as mentioned in Section 3.2):

Table 3.4: Case Study 1: Excerpt of CSPs SLAs and user requirements

Cloud SLA Element based on CSA STAR [STA11]			CSP_1	CSP_2	CSP_3	*CSC*		
Control Category	Control Group	SLO	Val_1	Val_2	Val_3	Case I	Case II	Case III
Compliance (CO)	audit planing (CO1)	CO1.1	yes	yes	yes	yes	High	High
		CO1.2	$level_3$	$level_2$	$level_3$	$level_3$		
	independent audits (CO2)	CO2.1	no	yes	yes	yes	Low	
		CO2.2	yes	yes	yes	yes		
		CO2.3	yes	yes	yes	yes		
		CO2.4	yes	yes	yes	yes		
	Third party audits (CO3)	CO3.1	yes	yes	yes	yes	Medium	
		CO3.2	yes	yes	yes	yes		
		CO3.3	Weekly	Annual	Monthly	Monthly		
Facility Security (FS)	Secure Area (FS1)	FS1.1	no	yes	yes	yes	High	Low
		FS1.2	yes	no	yes	yes		
	Asset Management (FS2)	FS2.1	yes	yes	yes	yes		
		FS2.2	$level_3$	$level_2$	$level_3$	$level_3$		
		FS2.3	yes	yes	yes	yes		
Risk Management (RI)	Risk assessments (RI1)	RI1.1	Internal	Internal	External	Internal	Internal	Medium
		RI1.2	yes	yes	yes	yes	no	

- In column "Case I", user requirements are expressed at a very granular level (i.e., per-SLO). This represents a security-expert user.
- Column "Case II" shows a set of requirements expressed at three different levels of granularity (corresponding to the hierarchy shown in Figure 3.1) namely SLO, Control Group, and Control Category. Notice that at the Control Group and the Control Category level, the user expresses his requirements depending on the relative importance[3] of the SLA element (e.g., High, Medium, or Low).
- Finally, in column "Case III" are shown user requirements only at the Control Group and Control Category levels. This might be the case of a user that is not a security expert.

In order to evaluate the CSPs SLAs with respect to the user requirements we proceed to apply the techniques presented in Section 3.2 (cf., Step 2 in Figure 3.4).

Case I: An Expert User

The quantitative evaluation of the Cloud security SLOs defined in Table 3.4 regarding user Case I is detailed in this section.

Using the QPT

For comparison purposes, all the QPT analyses shown in this section considered (i) a maximum of 4 Local Security Levels (i.e., LSL = 4), (ii) all leaf nodes on the QPT having the same weight, (iii) only AND relationships on the QPT, (iv) yes/no values specified as LSL_{max} and LSL_{min} respectively (i.e., $LSL_{max} = 4$ and $LSL_{min} = 0$), and finally (v) security levels specified using LSLs from 1 to 4.

For QPT we performed two sets of evaluations, first with the three CSPs to show individually which CSP outperforms the other two. Then, evaluating the three CSPs with respect to the consumer (CSC) requirements.

[3] This is the typical result of a risk assessment.

Table 3.5 shows the CSPs SLA aggregation using the rules specified before in Table 3.1. The information shown in Table 3.5 is useful to analyse how individual Control Categories contribute to the overall security level of the CSP. For example, if control *CO* is the prime requisite from a business perspective, then the absolute evaluation will advise to initially choose CSP_3 followed by CSP_2 over CSP_1. Notice that this conclusion cannot be drawn directly from the overall SLA level benchmarks, where CSP_1 outperforms CSP_2.

Table 3.5: Absolute quantitative benchmarks obtained for three different CSPs SLAs

SLA	CSP_1	CSP_2	CSP_3
Total	0.85	0.83	1
CO	0.86	0.89	1
FS	0.8	0.7	1
RI	1	1	1

A second set of benchmarks was applied to the dataset of the three Cloud SLA regarding the user SLA requirements. Definition 1 is used to show the quantitative benchmark $QuantB_{node}$ associated with each *node* of the QPT as shown in Table 3.6. For example, CSP_1 is under-provisioning *CO2*, *CO3* and *FS1*. While CSP_2 is not fullfiling the user requirements for *CO1*, *CO3*, *FS1* and *FS2*. Only CSP_3 fulfils the user requirements as shown in overall SLA rank. The aggregated SLAs values are normalized with respect to the user requirement (cf., Figure 3.5).

Table 3.6: Quantitative benchmarks obtained for three different CSPs SLAs based on user's Case I requirements

SLA	CSP_1	CSP_2	CSP_3
Total	-0.176	-0.2	0
CO1	0	-0.33	0
CO2	-0.33	0	0
CO3	-0.1	-0.22	0
CO	-0.167	-0.129	0
FS1	-1	-1	0
FS2	0	-0.2	0
FS	-0.25	-0.43	0
RI	0	0	0

Using the QHP

For this evaluation technique, the user specifies his requirements at the lowest level of the SLA (i.e., SLOs) and considers the same relative importance (i.e., weights) for all of these. Prior to the calculation of the relative ranking matrix using Equation 3.1, the following considerations take place:

1. QHP uses qualitative weights to indicate the user's relative priorities, and these weights are normalized as to comply with AHP requirements.

2. All SLOs specified by the user as Boolean no, are assigned a relative rank value 0.

3. All SLOs specified by the user as Boolean yes, are assigned a relative rank value 1.

4. *High-Important* and *Low-Important* indicate a weight 1 and 0 respectively.

5. *Medium-Important* can be considered any intermediate values between 1 and 0. In this analysis *Medium-Important* indicates a weight 0.5.

6. All CSPs security SLOs are normalized to the user requirements to eliminate masquerading[4]

For the Compliance Control Category, there are three security Control Groups which are further divided into a set of SLOs (as shown in Table 3.4). Definition 3 is used to create the attribute pairwise relation, as for example in the case of *CO1.2*:

$$CSP_1/CSP_2 = 3/2 \quad CSP_2/CSP_3 = 2/3$$
$$CSP_3/CSP_1 = 3/3 \quad CSC/CSP_2 = 3/2$$

Thus, the CM of *CO1.2* is calculated using Equation 3.1 (shown below for ease of explanation).

$$CM_{CO1.2} = \begin{array}{c} \\ CSP_1 \\ CSP_2 \\ CSP_3 \\ CSC \end{array} \begin{array}{c} \begin{array}{cccc} CSP_1 & CSP_2 & CSP_3 & CSC \end{array} \\ \begin{pmatrix} 1 & 3/2 & 3/3 & 3/3 \\ 2/3 & 1 & 2/3 & 2/3 \\ 3/3 & 3/2 & 1 & 3/3 \\ 3/3 & 3/2 & 3/3 & 1 \end{pmatrix} \end{array}$$

The relative ranking of the CSPs for *CO1.2* is given by the priority vector for $CM_{CO1.2}$ ($PV_{CO1.2}$). Similarly, we precompute $CM_{CO1.1}$ and $PV_{CO1.1}$. PV_{CO1} is then calculated by aggregating $PV_{CO1.1}$ and $PV_{CO1.2}$ with user normalized weights (w_{CO1}) using Equation 3.2. Where PV_{CO1} reflects which of the CSPs provide the *CO1* security SLO relative to other CSPs and to the user requirements as shown in Figure 3.9, such that:

$$PV_{CO1} = \begin{array}{c} \\ CSP_1 \\ CSP_2 \\ CSP_3 \\ CSC \end{array} \begin{array}{c} \begin{array}{cc} PV_{CO1.1} & PV_{CO1.2} \end{array} \\ \begin{pmatrix} 0.25 & 0.2727 \\ 0.25 & 0.1818 \\ 0.25 & 0.2727 \\ 0.25 & 0.2727 \end{pmatrix} \end{array} \begin{pmatrix} 0.5 \\ 0.5 \end{pmatrix}$$

Therefore, PV_{CO1} is:

$$PV_{CO1} = \begin{array}{c} \begin{array}{cccc} CSP_1 & CSP_2 & CSP_3 & CSC \end{array} \\ \begin{pmatrix} 0.2614 & 0.2159 & 0.2614 & 0.2614 \end{pmatrix} \end{array}$$

This implies that CSP_1 and CSP_3 equally satisfy *CSC*'s requirement. However, CSP_2 does not fulfill that requirement. The priority vector for *Independent audits* (PV_{CO_2}) is calculated similarly, such that *CO2.1, CO2.2, CO2.3* and *CO2.4* priority vectors are aggregated. Similarly, we compute PV_{CO3} where *CO3.1, CO3.2* and *CO3.3* are specified by the user as *yes*, *yes* and *Monthly* respectively.

The three *Compliance* priority vectors *CO1, CO2, CO3* are aggregated to have the overall compliance priority vector PV_{CO} as shown in Figure 3.7 such that:

$$PV_{CO} = \begin{array}{c} \begin{array}{cccc} CSP_1 & CSP_2 & CSP_3 & CSC \end{array} \\ \begin{pmatrix} 0.2299 & 0.2301 & 0.27 & 0.27 \end{pmatrix} \end{array}$$

[4]The masquerading effect happens when the overall aggregated security level value mostly depend on those security controls with a high-number of SLOs, thus affecting negatively groups with fewer although possibly more critical provisions. Other methodologies for the Cloud security assessment (such as REM [CPRT05]) suffer from this effect.

Both CSP_1 and CSP_2 under-provision *CO2* and CSP_2 under-provisions *CO1* and *CO3*. As a result, only CSP_3 satisfies *CSC*'s *CO* requirement. In a similar way the *Facility Security* and *Risk Management* priority vectors are considered (Figure 3.7).

$$PV_{FS} = \begin{matrix} CSP_1 & CSP_2 & CSP_3 & CSC \end{matrix} \begin{pmatrix} 0.2121 & 0.1970 & 0.29545 & 0.29545 \end{pmatrix}$$

$$PV_{RI} = \begin{matrix} CSP_1 & CSP_2 & CSP_3 & CSC \end{matrix} \begin{pmatrix} 0.25 & 0.25 & 0.25 & 0.25 \end{pmatrix}$$

Finally, the priority vectors of *Compliance*, *Facility Security* and *Risk Management* security are aggregated to obtain the total SLA priority vector:

$$PV_{total} = \begin{matrix} CSP_1 & CSP_2 & CSP_3 & CSC \end{matrix} \begin{pmatrix} 0.2307 & 0.2257 & 0.2718 & 0.2718 \end{pmatrix}$$

Consequently, only CSP_3 fulfills the user's requirements, as shown in Figure 3.6.
The proposed framework allows users to visualize the differences between various CSPs with respect to user requirements. Both CSP_1 and CSP_2 under-provisions *CO* and *FS* . As a result, CSP_3 is the best matching provider according to user's requirements.

QPT and QHP

Figure 3.5 shows the results of applying both QPT and QHP to the set of SLAs and also the user Case I requirements presented in Table 3.4. As shown in Figure 3.5 the resulting ranking of CSPs is consistent for both QPT and QHP: CSP_3 is the provider that better fulfils the user requirements, followed by CSP_1 and CSP_2 respectively. Where as shown in Figure 3.9, CSP_1 is not satisfying user requirements for *CO2.1*, *CO3.3* and *FS1.1* SLOs. Also CSP_2 is not satisfying user requirements for *CO1.2*, *CO3.3*, *FS1.2* and *FS2.2*. For users specifying the SLO-level requirements, this means that both techniques result on the same/consistent ranking.

It is worth noting that QPT can only evaluate requirements specified at the SLO-level, therefore it cannot be applied either to Case II or Case III requirements.

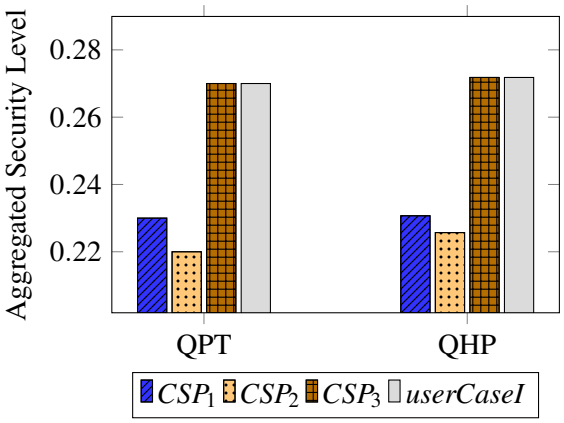

Figure 3.5: Comparing QPT and QHP for user Case I requirements.

The QHP evaluation technique allows users to evaluate CSPs security levels and perform comparisons at different levels of granularity. This can be observed in Figures 3.7 and 3.9. Figure 3.6

shows the overall security evaluation (i.e., at the top SLA-level) for each one of the three sets of user requirements. Figure 3.7 shows a different level of aggregation (i.e., Control Category) for the user Case I requirements. Figure 3.9 shows the CSPs ranking at the SLO-level. For example, during a procurement process Figure 3.6 can be used to provide preliminary guidance to select an initial set of CSPs, while a more detailed decision might be based on the more granular Figure 3.7 (or even by comparing at the SLO-level as in Figure 3.9).

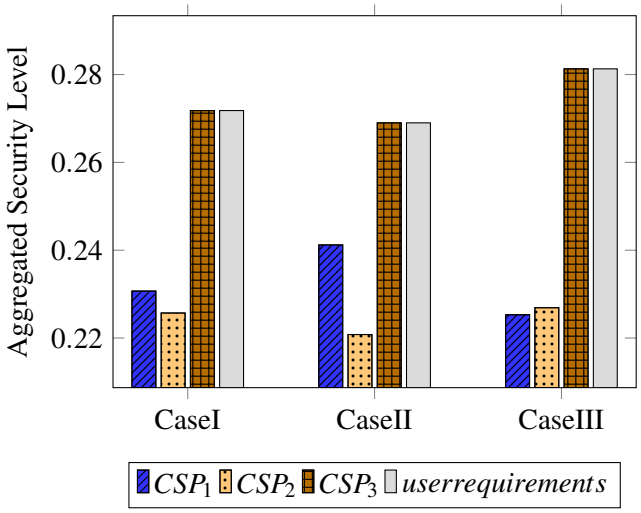

Figure 3.6: QHP-based evaluation showing the aggregated SLA level

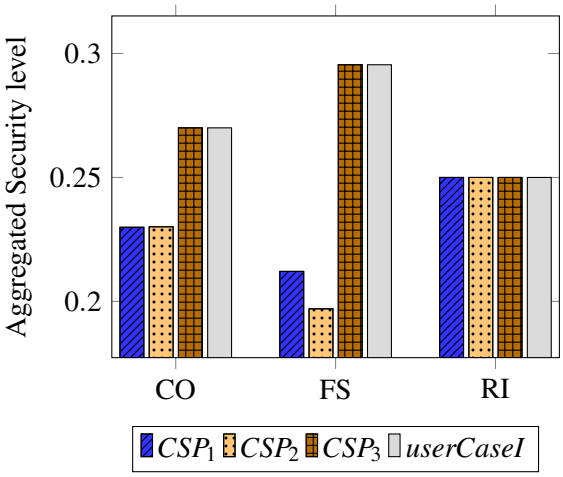

Figure 3.7: QHP-based aggregation at the Control Category-level (for user Case I requirements).

Finally, in Figure 3.8 we compare both QPT and QHP from a performance viewpoint. For this experiment we measured the time consumed (in seconds) to evaluate an SLA comprised of an incremental number of SLOs (up to the 139 contained in our dataset) with respect to user requirements I. It can be observed that in the case of QPT the number of evaluated SLOs does not affects the performance, whereas for QHP the time required to evaluate an SLA increases exponentially depending on the number of SLOs (as explained in Section 3.3.1). In scenarios where performance is not important (e.g., decision-making dashboards), then QHP might be used because of the flexibility of showing the evaluations at varied levels of SLA. However, if SLA automation is required (e.g., a software agent deciding which Cloud storage provider to use), then

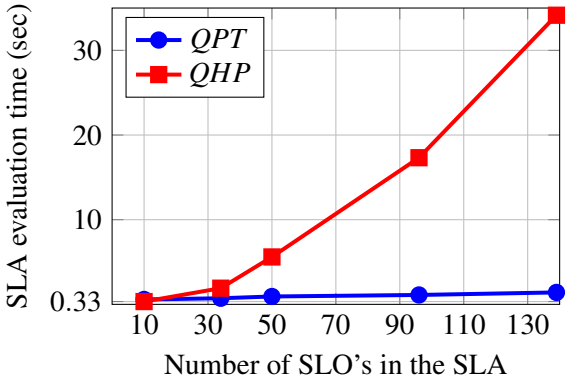

Figure 3.8: Performance comparison between QPT and QHP (evaluating user Case I and CSP_1).

QPT would provide the best results from the performance perspective.

Case II: A Semi-Expert User

As mentioned in Section 3.2.2, the QHP technique allows users to specify their security requirements at varied levels of granularity. This helps to remove the need for users to specify the value required for every single security SLO (which usually needs an extremely high level of expertise). Moreover, allowing users to specify their security requirements using qualitative labels, enables both basic and expert users to represent their needs according to their expertise and specific organisational context. This section shows a case study where security requirements are represented at different levels of granularity. In this case study we only considered qualitative weights to indicate the user's relative priorities (*High-Important*, *Low-Important* and *Medium-Important*) corresponding to the numeric values 1, 0 and 0.5 respectively.

We also assume a user denoting controls *Audit Planning*, *Independent Audits* and *Third Party Audits* as *High-Important*, *Low-Important* and *Medium-Important* respectively, *High-Important* for *Facility Security*, and specified low level requirements for *Risk Management* as shown in Table 3.4. Since *Audit Planning* is assigned *HI*, the respective weight is set to 1. On the other hand, *Third Party* is denoted *LI* by the user and the respective weight is set to 0. Therefore, PV_{CO1}, PV_{CO2} and PV_{CO3} are aggregated with user defined normalized weights (w_{CO}) using Equation 3.2 such that:

$$w_{CO} = \begin{matrix} CO1 & CO2 & CO3 \end{matrix} \begin{pmatrix} 0.67 & 0 & 0.33 \end{pmatrix}$$

Therefore, PV_{CO} is:

$$PV_{CO} = \begin{pmatrix} 0.2615 & 0.2154 & 0.2615 & 0.2615 \end{pmatrix}$$

This implies that CSP_2 does not fulfill *CSC* Compliance SLO and both CSP_1 and CSP_2 equally satisfy that requirement. For *FS*, the user specified *High-Important* which is assigned as 1 for all security SLOs.

$$PV_{FS} = \begin{matrix} CSP_1 & CSP_2 & CSP_3 & CSC \end{matrix} \begin{pmatrix} 0.2121 & 0.1970 & 0.29545 & 0.29545 \end{pmatrix}$$

Similarly, as Case I *Risk Management* is evaluated such that:

$$PV_{RI} = \begin{array}{c} \begin{array}{cccc} CSP_1 & CSP_2 & CSP_3 & CSC \end{array} \\ \begin{pmatrix} 0.25 & 0.25 & 0.25 & 0.25 \end{pmatrix} \end{array}$$

Subsequently, PV_{CO}, PV_{FS} and PV_{RI} are aggregated to obtain the total SLA priority vector:

$$PV_{total} = \begin{pmatrix} 0.2412 & 0.2208 & 0.2690 & 0.2690 \end{pmatrix}$$

Therefore, only CSP_3 satisfies the user needs while both CSP_1 and CSP_2 do not fulfill user requirements, as shown in Figure 3.6. That was expected, as CSP_1 is not providing *FS1.1* and CSP_2 is under-provisioning *CO1.2* and not providing FS1.2.

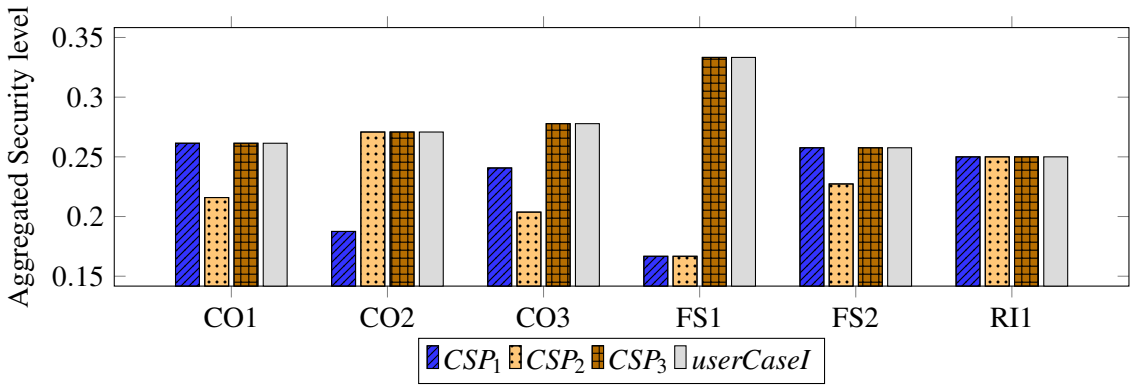

Figure 3.9: Using QHP to compare CSPs with respect to user Case I requirements at the SLO level.

Case III: A Non-Expert User

In this case study, the user represents his security requirements at a coarse-grained level (i.e., Control Category). For this purpose, the user associates the weights *High-Important* for *Compliance*, *Low-Important* for *Facility Security* and *Medium-Important* for *Risk Management* at the Control Category level. Similarly, as shown in previous cases, the priority vectors of *CO*, *FS*, and *RI* are aggregated with user normalized defined weights (w_{total}) using Equation 3.2:

$$W_{total} = \begin{pmatrix} 0.67 & 0 & 0.33 \end{pmatrix}$$

where:

$$PV_{CO} = \begin{pmatrix} 0.2254 & 0.2279 & 0.2734 & 0.2734 \end{pmatrix}$$

$$PV_{RI} = \begin{pmatrix} 0.2250 & 0.2250 & 0.2750 & 0.2750 \end{pmatrix}$$

Therefore, the total priority vector is:

$$PV_{total} = \begin{pmatrix} 0.2253 & 0.2269 & 0.2813 & 0.2813 \end{pmatrix}$$

As in the previous cases, only CSP_3 satisfies the user needs. However, as observed, the CSPs ranking was different than from previous cases. In this case, CSP_2 outperforms CSP_1. This result was expected as the user assigned weights only at the Category-level and Facility Security is assigned *Low-Important*, which affected the overall evaluation. Moreover, CSP_2 is under-provisioning *CO1.2* and CSP_1 is not providing *CO2.1*.

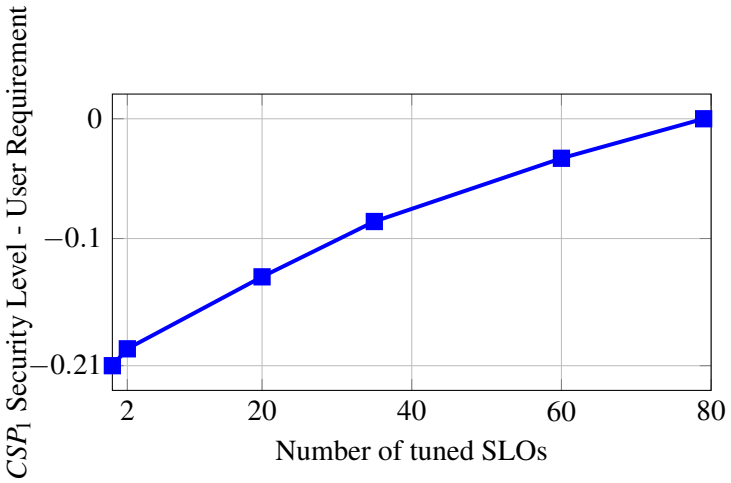

Figure 3.10: Sensitivity analysis: combined security effect of sets of SLOs.

3.3.2 The CSP Perspective: Maximising Offered Security Levels

The second validation scenario presented in this section applies the SLA evaluation techniques to solve problems faced by CSPs, i.e., (a) which specific security SLOs from the offered SLA should be improved in order to maximise the overall security level? And, (b) how to improve their service security level to meet the users' requirements? This might be the case of a well-established CSP deciding where to invest in order to achieve the highest possible security level, or a new CSP designing the SLA. To answer these questions, we performed two sensitivity analyses to ascertain the security benefits of improving one or more SLOs. The presented sensitivity analysis can be performed using QPT or QHP, however this section applies only QHP given the flexibility it offers for evaluating SLAs at different levels of granularity and its suitability for implementing what-if scenarios (cf., Section 3.2.3).
The experiments used the CSP_1 dataset described at the beginning of this section (139 SLOs based on CSA STAR), and applied the Case I requirements to setup the User's baseline for the security evaluation. From the existing 139 SLOs the CSP_1 is under-provisioning 80 of them. Figure 3.11 shows how the QHP technique can be used to analyse an existing SLA, and extract the individual SLOs that if enhanced would result on different improvements associated with the overall security level. In this case, the X-axis represents the improvement associated to the overall security level after enhancing any of the SLOs. It is shown as a percentage where 0% corresponds to the original SLA and 100% is the most effective SLO. For example, providing tenants with the security policies applicable to virtualised resources (RM2.2 in Figure 3.11), quantitatively increases CSP_1 security level better than improving the thresholds committed for any of the encryption-related SLOs IS18.4 or IS18.5. Also as observed in the figure, improving the SLO DG6.1 would result exactly in the same quantitative improvement as RM2.2's. In this case, the CSP might need to use additional criteria (e.g., economic cost associated with the proposed changes to the SLA) in order to take a decision related to the SLO to enhance.

The second sensitivity analysis considers the combined security effect of improving simultaneously two or more of the SLOs under-provisioned by CSP_1, based on the User requirements of the Case I. Results of the analysis are shown in Figure 3.10, where it can be observed how the security level of the CSP approaches faster to the User requirement (i.e., $Yaxis = 0$) if several of its offered SLOs are enhanced at the same time. Of course, if all 80 under-provisioned SLOs are

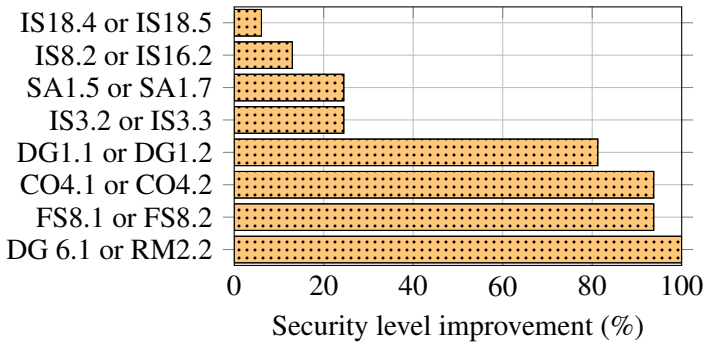

Figure 3.11: Sensitivity analysis: CSP_1 SLOs that maximise the overall security level.

improved then the security level of CSP_1 exactly matches the User requirement.

3.3.3 Summary

We have presented two security evaluation approaches (namely QPT and QHP), which build and extend upon the state of the art techniques, to quantitatively assess the security levels provided by Cloud SLAs. The proposed techniques were designed based on the specifics of SLAs from state of the art works and standardisation bodies. Furthermore, both QPT and QHP were empirically validated through case studies using real-world CSP data obtained from the Cloud Security Alliance. The validation experiments were useful not only to highlight the applicability of our techniques to real-world CSPs but also to highlight the advantages and limitations of these techniques, and to provide an objective comparison of both QPT and QHP in order to guide (prospective) adopters.

4. Conclusions and Next Steps

Covering M1-12 of T4.1, D4.1 has provided the first report on security metrics and assessment that (a) detailed two real-world Cloud scenarios (initial part of Chapter 2) outlining the information flow chains across the users and CSPs, (b) presented the state of the art trust metrics and related related the metrics to the ESCUDO-CLOUD Use Cases (Section 2.4.1), and (c) developed SLA-based techniques of QPT and QHP to specify, reason and compare trust levels across the CSPs (Chapter 3).

The use of QPT/QHP has been illustrated for three cases covering expert, semi-expert and non-expert users. Furthermore, the approach has been validated using publicly available SLA's. We re-emphasize that the intent is not to develop a standalone absolute security value, but to provide a common SLA basis and common set of SLO attributes to facilitate a comparative assessment of CSP trust levels.

Overall, D4.1 has developed the foundations for security assessment that can be utilized and validated in the UC's across WP2-4.

Also, D4.1 has built the concepts for the single CSP case. As ESCUDO-CLOUD develops the refinements on multi-CSP/federated Cloud requirements, D4.2 will extend the D4.1 assessment techniques to the multi-CSP models.

A. Full list of C-SIG Trust Metrics

Table A.1: Full list of Trust Metrics

ID	Trust Metric Name	Source	Use Case Ref
1	Security budget measure	NIST SP800-55-v1	
2	Vulnerability Measure	NIST SP800-55-v1	
3	Remote Access Control Measure	NIST SP800-55-v1	
4	Security Training Measure	NIST SP800-55-v1	
5	Audit Record Review Measure	NIST SP800-55-v1	
6	C & A Completion Measure	NIST SP800-55-v1	
7	Configuration Changes Measure	NIST SP800-55-v1	REQ-UC2-AC-5 REQ-UC4-AC-5
8	Contingency Plan Testing Measure	NIST SP800-55-v1	
9	User Accounts Measure	NIST SP800-55-v1	REQ-UC2-AC-1 REQ-UC2-AC-2 REQ-UC2-AC-6 REQ-UC4-AC-1 REQ-UC4-AC-3 REQ-UC1-IKM-2 REQ-UC1-TKM-2 REQ-UC3-KM-1 REQ-UC3-SO-1 REQ-UC3-KM-4
10	Incident Response Measure	NIST SP800-55-v1	REQ-UC4-SS-5
11	Maintenance Measure	NIST SP800-55-v1	REQ-UCI-SKM-2 REQ-UCI-SKM-3 REQ-UC2-AC-4 REQ-UC2-EQ-1 REQ-UC2-EQ-2 REQ-UC2-EQ-3 REQ-UC2-EQ-4 REQ-UC4-SS-1 REQ-UC4-SS-2 REQ-UC4-SS-3 REQ-UC3-KM-3 REQ-UC3-SO-2 REQ-UC3-SO-3 REQ-UC3-SO-4

			REQ-UC3-SO-5
12	Media Sanitization Measure	NIST SP800-55-v1	REQ-UC1-IKM-4 REQ-UC1-TKM-4
13	Physical Security Incidents Measure	NIST SP800-55-v1	
14	Planning Measure	NIST SP800-55-v1	
15	Personnel Security Screening Measure	NIST SP800-55-v1	
16	Risk Assessment Vulnerability Measure	NIST SP800-55-v1	REQ-UC4-AC-6
17	Service Acquisition Contract Measure	NIST SP800-55-v1	
18	System and Communication Protection Measurement	NIST SP800-55-v1	REQ-UC1-TKM-3 REQ-UC2-EQ-3
19	System and Information Integrity Measurement	NIST SP800-55-v1	REQ-UC2-AC-5 REQ-UC2-AC-6 REQ-UC2-KM-1 REQ-UC2-KM-2 REQ-UC2-KM-3 REQ-UC2-KM-4
20	Mean Time to Deploy Critical Patches	The Center for Internet Security	
21	Percent of Systems Without Known Severe Vulnerabilities	The Center for Internet Security	
22	Mean-Time to Mitigate Vulnerabilities	The Center for Internet Security	
23	Mean Cost to Mitigate Vulnerabilities	The Center for Internet Security	
24	Patch Policy Compliance	The Center for Internet Security	REQ-UC4-SS-4
25	Percent of Changes with Security Review	The Center for Internet Security	
26	Risk Assessment Coverage	The Center for Internet Security	REQ-UC2-KM-2 REQ-UC2-EQ-2
27	Security Testing Coverage	The Center for Internet Security	
28	Number of Incidents	The Center for Internet Security	
29	Configuration Management Coverage	The Center for Internet Security	
30	Current Anti-Malware Coverage	The Center for Internet Security	
31	Number of Applications	The Center for Internet Security	
32	Cost of Incidents	The Center for Internet Security	
33	Mean Time to Incident Discovery	The Center for Internet Security	

34	Mean Time Between Security Incidents	The Center for Internet Security	
35	Mean Time to Incidence Recovery	The Center for Internet Security	
36	Mean Cost of Incidence	The Center for Internet Security	
37	Mean Incident Recovery Cost	The Center for Internet Security	
38	Mean Time to Patch	The Center for Internet Security	
39	Mean Cost to Patch	The Center for Internet Security	
40	Percentage of Configuration Compliance	The Center for Internet Security	REQ-UC2-EQ-1 REQ-UC2-EQ-2 REQ-UC2-EQ-3 REQ-UC2-EQ-4 REQ-UC3-AC-3 REQ-UC3-AC-1
41	Mean Time to Complete Changes	The Center for Internet Security	
42	Percent of Changes with Security Exceptions	The Center for Internet Security	
43	Percent of Critical Applications	The Center for Internet Security	REQ-UC4-AC-2
44	Information Security Budget as a Percentage of IT Budget	The Center for Internet Security	
45	Information Security Budget Allocation	The Center for Internet Security	
46	Vulnerability Scan Coverage	The Center for Internet Security	
47	Number of Known Vulnerability Instances	The Center for Internet Security	
48	Patch Management Coverage	The Center for Internet Security	
49	Percentage of Incidents Detected by Internal Controls	The Center for Internet Security	
50	Mean Time from Discovery to Containment	The Center for Internet Security	
51	Percentage of uptime	EU FP7 Cumulus	
52	Percentage of processed requests	EU FP7 Cumulus	
53	Percentage of timely provisioning requests	EU FP7 Cumulus	
54	Service provider data access level	EU FP7 Cumulus	

55	Percentage of systems with time synchronization	EU FP7 Cumulus	
56	Maximum measured time difference	EU FP7 Cumulus	
57	Number of (successful) audits performed	EU FP7 Cumulus	
58	Tenant isolation level	EU FP7 Cumulus	REQ-UC3-AC-5
59	Data portability	EU FP7 Cumulus	
60	Mean time between incidents	EU FP7 Cumulus	
61	Percentage of timely incident reports	EU FP7 Cumulus	
62	Percentage of timely incident responses	EU FP7 Cumulus	
63	Percentage of timely incident resolutions	EU FP7 Cumulus	
64	User authentication and identity assurance level	EU FP7 Cumulus	REQ-UC4-AC-1 REQ-UC4-AC-2 REQ-UC4-AC-3 REQ-UC4-AC-4 REQ-UC1-IKM-1 REQ-UC1-TKM-1 REQ-UC3-KM-2 REQ-UC3-AC-2 REQ-UC3-AC-4
65	Mean time required to revoke a user	EU FP7 Cumulus	
66	Password storage protection level	EU FP7 Cumulus	REQ-UC2-KM-3 REQ-UC3-KM-6
67	Cryptographic brute force resistance	EU FP7 Cumulus	REQ-UC2-AC-6 REQ-UC4-SS-5
68	Key access control level	EU FP7 Cumulus	REQ-UC3-AC-6
69	Country level anchoring	EU FP7 Cumulus	
70	Data deletion quality level	EU FP7 Cumulus	REQ-UC4-SS-4 REQ-UC1-IKM-4 REQ-UC1-TKM-4
71	Percentage of timely effective deletions	EU FP7 Cumulus	
72	Percentage of tested storage retrievability	EU FP7 Cumulus	
73	Durability	EU FP7 Cumulus	
74	Vulnerability exposure level	EU FP7 Cumulus	
75	Percentage of timely vulnerability corrections	EU FP7 Cumulus	
76	Percentage of timely vulnerability reports	EU FP7 Cumulus	
77	Recovery point	EU FP7 Cumulus	
78	Recovery time	EU FP7 Cumulus	

79	Percentage of authorized personnel that received training	EU FP7 Cumulus	
80	Percentage of recovery success	EU FP7 Cumulus	
81	Configuration change reporting capability	EU FP7 Cumulus	
82	Percentage of timely configuration change notifications	EU FP7 Cumulus	
83	Percentage of compliant applications	EU FP7 Cumulus	
84	Authorized collection of PII	EU FP7 A4Cloud	
85	Privacy Program Budget	EU FP7 A4Cloud	
86	Privacy Program Updates	EU FP7 A4Cloud	
87	Periodicity of Privacy Impact Assessments for Information Systems	EU FP7 A4Cloud	
88	Number of privacy audits received	EU FP7 A4Cloud	
89	Successful audits received	EU FP7 A4Cloud	
90	Record of Data Collection, Creation, and Update	EU FP7 A4Cloud	
91	Data classification	EU FP7 A4Cloud	
92	Coverage of Privacy and Security Training	EU FP7 A4Cloud	
93	Account of Privacy and Security Training	EU FP7 A4Cloud	
94	Level of confidentiality	EU FP7 A4Cloud	REQ-UC4-SS-4 REQ-UC4-DE-1 REQ-UC4-DE-2 REQ-UC4-DE-3 REQ-UC4-DE-4 REQ-UC3-DE-2
95	Key Exposure Level	EU FP7 A4Cloud	REQ-UC2-KM-1 REQ-UC2-KM-2 REQ-UC2-KM-3 REQ-UC2-KM-4 REQ-UC3-DE-1 REQ-UC3-DE-4 REQ-UC3-DE-5 REQ-UC3-AC-7 REQ-UC3-KM-5
96	Data Isolation Testing Level	EU FP7 A4Cloud	
97	Type of Consent	EU FP7 A4Cloud	
98	Type of notice	EU FP7 A4Cloud	
99	Procedures for Data Subject Access Requests	EU FP7 A4Cloud	
100	Number of Data Subject Access Requests	EU FP7 A4Cloud	

101	Responded data subject access requests	EU FP7 A4Cloud	
102	Mean time for responding Data Subject Access Requests	EU FP7 A4Cloud	
103	Readability (Flesch Reading Ease Test)	EU FP7 A4Cloud	
104	Rank of Responsibility for Privacy	EU FP7 A4Cloud	
105	Certification of acceptance of responsibility	EU FP7 A4Cloud	
106	Frequency of certifications	EU FP7 A4Cloud	
107	Log Unalterability	EU FP7 A4Cloud	
108	Identity Assurance	EU FP7 A4Cloud	
109	Mean time to revoke users	EU FP7 A4Cloud	
110	Mean time to respond to complaints	EU FP7 A4Cloud	
111	Number of complaints	EU FP7 A4Cloud	
112	Reviewed complaints	EU FP7 A4Cloud	
113	Number of privacy incidents	EU FP7 A4Cloud	
114	Coverage of incident notifications	EU FP7 A4Cloud	
115	Type of incident notification	EU FP7 A4Cloud	
116	Privacy incidents caused by third parties	EU FP7 A4Cloud	
117	Number of Business Continuity Resilience (BCR) plans tested	EU FP7 A4Cloud	
118	Maximum tolerable period for disruption (MTPD)	EU FP7 A4Cloud	
119	Sanctions	EU FP7 A4Cloud	
120	Incidents with damages	EU FP7 A4Cloud	
121	Total expenses due to compensatory damages	EU FP7 A4Cloud	
122	Average expenses due to compensatory damages	EU FP7 A4Cloud	
123	Cryptographic strength	EU FP7 SPECS	REQ-UC2-KM-3 REQ-UC2-KM-4 REQ-UC1-IKM-3 REQ-UC1-TKM-3
124	Forward secrecy	EU FP7 SPECS	
125	HSTS (HTTP Strict Transport Security)	EU FP7 SPECS	
126	Secure cookies forced	EU FP7 SPECS	
127	Client certificates	EU FP7 SPECS	
128	Certificate status request (a.k.a. OCSP stapling)	EU FP7 SPECS	
129	Certificate pinning	EU FP7 SPECS	
130	DANE	EU FP7 SPECS	
131	FIPS compliance	EU FP7 SPECS	REQ-UC3-DE-3
132	Level of Diversity	EU FP7 SPECS	

133	TLS Cryptographic Strength	EU FP7 SPECS	
134	Vulnerability Report Max Age	EU FP7 SPECS	
135	Vulnerability List Max Age	EU FP7 SPECS	
136	E2EE Crypto Strength	EU FP7 SPECS	REQ-UC2-EQ-2 REQ-UC4-DE-4
137	dDoS Attack Report Max Age	EU FP7 SPECS	
138	Write-Serializability	EU FP7 SPECS	
139	Read-Freshness	EU FP7 SPECS	
140	Backup	EU FP7 SPECS	REQ-UC4-SS-5 REQ-UC1-SKM-1
141	Attack Detection Latency	EU FP7 SPECS	
142	Number of False Positives	EU FP7 SPECS	
143	Number of Detected Attacks	EU FP7 SPECS	
144	Number of Vulnerabilities (Family)	EU FP7 SPECS	
145	Number of Vulnerabilities (Gravity)	EU FP7 SPECS	
146	Number of Executed Vulnerability Tests	EU FP7 SPECS	
147	Number of Available Vulnerability Tests	EU FP7 SPECS	

Bibliography

[Cat08] F. Cate. Provincial Canadian geographic restrictions on personal data in the public sector, 2008. https://www.hunton.com/files/Publication/2a6f5831-07b6-4300-af8d-ae30386993c1/Presentation/PublicationAttachment/0480e5b9-9309-4049-9f25-4742cc9f6dce/cate_patriotact_whitepaper.pdf.

[CCM14] CSA CCM. Cloud Security Alliance Cloud Controls Matrix v3.0.1. Online:https://cloudsecurityalliance.org/research/ccm/, 2014.

[Cha14] M. Chawla. The biggest security hacks of 2014, 2014. http://articles.economictimes.indiatimes.com/2014-12-29/news/57494907_1_dropbox-cyber-security-professionals-icloud.

[Com12] European Commission. Commission proposes a comprehensive reform of the data protection rules, 2012. http://ec.europa.eu/justice/newsroom/data-protection/news/120125_en.htm.

[CPRT05] V. Casola, R. Preziosi, M. Rak, and L. Troiano. A Reference Model for Security Level Evaluation: Policy and Fuzzy Techniques. *Journal of Universal Computer Science (UCS)*, 11(1):150 – 174, 2005.

[Dir95] EU Directive. Directive 95/46/EC of the European Parliament and of the Council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data. *Official Journal the European Communities*, 281(23/11):0031–0050, 1995.

[Fer07] P. Ferrie. Attacks on more virtual machine emulators. *Symantec Technology Exchange*, page 55, 2007.

[Fon15] J. Fontana. Breach costs at $162 million, Target reports, 2015. http://www.zdnet.com/article/breach-costs-at-162-million-target-reports/.

[Goo15] D. Goodin. DDoS attacks that crippled GitHub linked to Great Firewall of China, 2015. http://arstechnica.com/security/2015/04/ddos-attacks-that-crippled-github-linked-to-great-firewall-of-china/.

[Gre14] A. Greenberg. Online gambling site hit by five-vector DDoS attack peaking at 100Gbps, 2014. http://www.scmagazine.com/online-gambling-site-hit-by-five-vector-ddos-attack-peaking-at-100gbps/article/355020/.

[GSC06] R. Gill, J. Smith, and A. Clark. Experiences in passively detecting session hijacking attacks in IEEE 802.11 networks. In *Proceedings of the 2006 Australasian workshops on Grid computing and e-research-Volume 54*, pages 221–230. Australian Computer Society, Inc., 2006.

[Gut02] P. Gutmann. PKI: it's not dead, just resting. *IEEE Computer*, 35(8):41–49, 2002.

[HS10] D. Hubbard and M. Sutton. Top threats to cloud computing v1. 0. *Cloud Security Alliance*, 2010.

[ISO13] ISO. ISO/IEC 27002:2013 Information technology – Security techniques – Code of practice for information security controls, 2013. http://www.iso.org/iso/catalogue_detail?csnumber=54533.

[ISO14] ISO. Information Technology - Cloud Computing - Service Level Agreement (SLA) Framework and Terminology. Technical Report ISO/IEC 19086, International Organization for Standardization, 2014.

[JIB07] A. Jøsang, R. Ismail, and C. Boyd. A survey of trust and reputation systems for online service provision. *Journal of Decision Support Systems (DSS)*, 43(2):618–644, 2007.

[Ker14] B. Kerbs. Huge data leak at largest U.S. bond insurer, 2014. http://krebsonsecurity.com/2014/10/huge-data-leak-at-largest-u-s-bond-insurer/.

[LLS12] J. Luna, R. Langenberg, and N. Suri. Benchmarking Cloud Security Level Agreements Using Quantitative Policy Trees. *Proc. of the ACM Cloud Computing Security Workshop (CCSW)*, pages 103–112, 2012.

[Mil12] R. Miller. Generator fan failure triggered AWS outage, 2012. http://www.datacenterknowledge.com/archives/2012/06/21/aws-outage/.

[MKL09] T. Mather, S. Kumaraswamy, and S. Latif. *Cloud security and privacy: an enterprise perspective on risks and compliance*. " O'Reilly Media, Inc.", 2009.

[Mol13] R. Molla. The states with the most data centers are also the most disaster-prone [maps], 2013. https://gigaom.com/2013/01/10/the-states-with-the-most-data-centers-are-also-the-most-disaster-prone-maps/.

[NIS14] NIST. Security and Privacy Controls for Federal Information Systems and Organizations. Technical Report NIST 800-53v4, National Institute of Standards and Technology, 2014.

[NIS15] NIST. Cloud Computing Service Metrics Description. Technical Report NIST 500-307, National Institute of Standards and Technology, 2015.

[OV03] A. Ornaghi and M. Valleri. Man in the middle attacks demos. *Blackhat [Online Document]*, 19, 2003.

[Rob15] T. Robinsons. Apple patches iCloud vulnerability exploited by iDict hacking tool, 2015. http://www.scmagazine.com/apple-patches-icloud-vulnerability-exploited-by-idict-hacking-tool/article/390922/.

[RTSS09] T. Ristenpart, E. Tromer, H. Shacham, and S. Stefan. Hey, you, get off of my cloud: exploring information leakage in third-party compute clouds. In *Proceedings of the 16th ACM conference on Computer and communications security*, pages 199–212. ACM, 2009.

[Saa90] T. Saaty. How to make a decision: the analytic hierarchy process. *European Journal of Operational Research*, pages 9–26, 1990.

[SLA14] C-SIG SLA. Cloud Service Level Agreement Standardisation Guidelines. Technical Report C-SIG SLA 2014, European Commission, 2014.

[STA11] CSA STAR. Cloud Security Alliance Security, Trust & Assurance Registry. Online: https://cloudsecurityalliance.org/star/, 2011.

[T+12] TrendLabs APT Research Team et al. Spear-phishing email: Most favored apt attack bait. *Last accessed September*, 2:2013, 2012.

[TC08] E. Torres and J. Clamen. Optimizing performance across the enterprise IP VPN, 2008. http://www.tatacommunications.com/vpn/cutTheComplexity/Hybrid_VPN_whitepaper.pdf.

[Tea15] Verizon Business Risk Team. Data breach investigations report, 2015. http://www.verizonenterprise.com/DBIR/2015/.

[TTLS14] A. Taha, R. Trapero, J. Luna, and N. Suri. AHP-Based Quantitative Approach for Assessing and Comparing Cloud Security. In *Proc. of the IEEE Conference on Trust, Security and Privacy in Computing and Communications (TrustCom)*, pages 284–291, 2014.

[Zel82] M. Zeleny. *Multiple Criteria Decision Making*. McGraw Hill, 1982.