

Evidence-Based Security and Privacy Assurance in Cloud Ecosystems

Saul Formoso and Massimo Felici^(✉)

Security and Manageability Lab, Hewlett Packard Labs,
Bristol BS34 8QZ, UK
{saul.formoso,massimo.felici}@hpe.com

Abstract. This paper is concerned with the problem of security and privacy assurance in cloud ecosystems. Different controls are deployed in order to guarantee security and privacy across cloud ecosystems. It is yet challenging to assess their effectiveness in operation. Therefore, it is necessary to devise methodologies and technologies for providing assurance of whether security and privacy controls are effective and appropriate for specific cloud ecosystems. This paper discusses the rationale and requirements for evidence-based security and privacy assurance. It also discusses the underlining mechanisms shaping a software defined storage system for gathering evidence drawn from a cloud ecosystem. It explains such requirements and mechanisms in the context of a sample use case. In conclusion, it provides insights for evidence-based security and privacy assurance of cloud ecosystems.

1 Introduction

Cloud computing provides an alternative way of providing and using Information Technology (IT) that differs from traditional systems [1, 2]. It can be characterized in terms of its main features [3], i.e. on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service. Alternative cloud computing deployments (e.g. public cloud, private cloud, and hybrid cloud) can accommodate different customer needs [1]. Alongside the new opportunities offered by cloud computing, there are new challenges related to security and privacy of data stored in the cloud [4]. Cloud customers as well as cloud providers are concerned with such challenges. Security and privacy are therefore shared responsibilities among the parties involved in cloud supply chains – that is, all actors that participate in the provision and consumption of cloud services.

In order to mitigate emerging security and privacy threats, consequently reducing risks in the cloud, different security and privacy controls are deployed across cloud supply chains [2]. Such security and privacy controls diffuse organisational boundaries as emphasised by the security conservation principle: *“for a particular service migrated to the cloud, the full set of necessary Security Components and controls that must be implemented to secure the cloud Ecosystem is always the same; however, the division of responsibility for those Components and controls changes based upon the characteristics of the cloud, particularly the service deployment.”* [2].

This paper is related to the Cloud Accountability Project, which has defined an approach and related mechanisms for accountability in the cloud [5]. In particular, this paper looks into how different technologies are deployed in order to support security and privacy in a cloud ecosystem – “*A cloud computing business ecosystem (cloud ecosystem) is a business ecosystem of interacting organisations and individuals – the actors of the cloud ecosystem – providing and consuming cloud services*” [6]. Although the security and privacy controls are easy to identify, it is yet challenging to assess their effectiveness operationally. There is yet a lack of support for assessing how such technologies work as a whole and how they are appropriate for the specific cloud ecosystems they are deployed in. It is necessary to understand how security and privacy controls enable organisations to comply with high level policies (drawn from relevant regulatory regimes).

This paper is structured as follows. Section 2 highlights relevant security and privacy research and industry practices, which provide a background to the research and development work presented here. The problem is then to provide evidence-based assurance that security and privacy controls are appropriate and operationally effective for the cloud ecosystems they are deployed in. Section 3 characterises the problem of security and privacy assurance in cloud supply chains. Section 4 discusses mechanisms that are necessary in order to implement evidence-based assurance in practice. These mechanisms have been used for tailoring a software defined storage system for the cloud in order to gather evidence related to the security and privacy controls deployed in the ecosystem. The evidence collected supports assurance and eases auditing the cloud supply chain. This aims to map the security and privacy controls deployed in the cloud supply chain for monitoring and auditing purposes. Section 5 explains the resulting implementation for evidence-based assurance in the context of a use case demonstrator drawn from the Cloud Accountability Project. This shows how the implemented system is useful in order to support assurance. An example explains how two sample SLAs between different cloud actors can be monitored. To sum up, this paper provides practical insights for implementing and supporting evidence-based assurance for security and privacy controls in cloud ecosystems.

2 Security and Privacy Practices

Cloud environments, due to their different nature, have an additional set of security and privacy requirements compared to traditional systems [1, 2]. While traditional systems require a large emphasis on the infrastructure, cloud computing focuses mainly on the provision of the services, relegating the former to a secondary level. Cloud environments’ inherent ubiquity brings along security and privacy risks that need to be addressed (e.g. broad, network access, decreased visibility and control by customers, multi-tenancy, data residency, etc. [2]). These have originated practices, certification schemes (including guidelines), and technologies.

As this is a topic of increasing interest, there are research projects putting effort into complying with the aforementioned requirements. The main point is that in order to support security and privacy, it is necessary to adopt diverse mechanisms, from technical tools to process-oriented approaches, including certification and auditing. A brief

overview of security and privacy practices, in particular, frameworks, certification schemes, technologies, and research projects, is provided hereunder:

- **Frameworks:** different frameworks capture industry best practices that guide stakeholders in order to enhance security and privacy in operations. For example, at the architectural level, it is possible to identify different security controls that organisations need to implement [2]. From a management viewpoint, it is possible to identify critical processes (e.g. security risk assessment and privacy management) that address the mitigation of security and privacy threats [7].
- **Certification Schemes:** ENISA released the Cloud Certification Schemes Metaframework (CCSM) that classifies the different types of security certifications for cloud providers [8]. This metaframework is used to compare and compile a list (CCSL, Cloud Certification Schemes List) of different cloud certification schemes and map detailed security requirements to security objectives existing in them. The overall objective is to make the cloud transparent for cloud customers, in particular, in the way cloud providers meet specific security objectives. Relevant examples are the CSA STAR Certification and the ISO/IEC 27017 (cloud security) and ISO/IEC 27018 (cloud privacy) standards.
- **Technologies:** among the various security and privacy technologies, Security Information and Event Management (SIEM) technologies have a critical role in monitoring operational security and supporting organisations in decision-making. These can be deployed to monitor computational resources in a cloud ecosystem, generating evidence that can be used to prove that security and privacy controls are complied with. Gartner reviews the most widely adopted SIEM technologies (e.g. Hewlett Packard Enterprise’s ArcSight, IBM Security’s QRadar, McAfee’s Enterprise Security Manager, etc.) in industry [9].
- **Research:** recent and ongoing research projects (with a particular attention to European projects) have been concerned with providing the conceptual and technological foundations underpinning the European Cloud Computing Strategy [10]. For example, research has focused on accountability (A4Cloud¹), data sharing agreement (CocoCloud²), security assurance (MUSA³), certification, and many other aspects of security and privacy. This highlights an increasing interest in providing assurance in the cloud.

3 Assurance in Cloud Ecosystems

This section points out the role of evidence in supporting assurance. In particular, it recalls the concept of accountability which highlights the responsibilities of an organisation in order to be accountable [5]. This is central to the concept of accountability [6]: *“Accountability for an organisation consists of accepting responsibility for*

¹ <http://www.a4cloud.eu/>.

² <http://www.coco-cloud.eu/>.

³ <http://www.tut.fi/musa-project/>.

data with which it is entrusted in a cloud environment, for its use of the data from the time it is collected until when the data is destroyed (including onward transfer to and from third parties). It involves the commitment to norms, explaining and demonstrating compliance to stakeholders and remedying any failure to act properly". Underpinning the concept of accountability is the provision of an account, which involves the gathering of evidence supporting organisational practices. This section then discusses the problem of assurance in a sample cloud supply chain. This discussion helps clarifying the requirements for supporting security and privacy assurance in cloud ecosystems.

3.1 Evidence-Based Accountability

The Cloud Accountability Project points out the need for evidence-based accountability in order to support the assessment of whether adopted security and privacy solutions (e.g. technologies, processes, etc.) are suitable for the specific cloud ecosystems, and hence provide assurance [5]. Cloud ecosystems involve various actors with different responsibilities. Emergent relationships among cloud actors give rise to the need for chains of evidence – *"A process and record that shows who obtained the evidence; where and when the evidence was obtained; who secured the evidence; and who had control and possession of the evidence"* [6] – and evidence in terms of organisational practices. On the one hand, it is necessary to validate gathered evidence and trace its source. On the other hand, evidence (is transformed and) propagates across system and organisational boundaries. From a technical viewpoint, evidence is considered among the three fundamental capabilities of an accountable system [11]:

- **Validation:** "It allows users, operators and third parties to verify a posteriori if the system has performed a data processing task as expected"
- **Attribution:** "In case of a deviation from the expected behaviour (fault), it reveals which component is responsible"
- **Evidence:** "It produces evidence that can be used to convince a third party that a fault has or has not occurred".

Therefore, gathering evidence has a critical role in supporting assurance – *"Assurance is about providing confidence to stakeholders that the qualities of service and stewardship with which they are concerned are being managed and maintained appropriately"* [12]. This is also particularly important while dealing with emergent threats [13] due to a certain extent to the shift required while deploying new technological paradigms like cloud computing.

3.2 Example of Cloud Supply Chain

Figure 1 shows a sample supply chain involving different actors: a cloud customer and two cloud service providers. The emergent relationships among actors form cloud supply chains defined in terms of cloud roles [5]. From a data protection perspective [14],

cloud actors also have different roles and responsibilities (i.e. data subject, data controller, and data processor).

It is challenging to support operational compliance to policies and regulations. Security and privacy depend on the operational effectiveness and appropriateness of deployed controls and their dependencies. It is desirable to build and maintain dynamic assurance cases of security and privacy controls (providing security and privacy assurance of the cloud supply chain through continuous monitoring). The following points characterise some aspects of assurance in cloud supply chains (Fig. 1):

1. Different security and privacy controls are deployed across a cloud supply chain.
2. It is challenging to provide transparency and assurance to cloud customers.
3. It is necessary to provide technological solutions to support continuous assurance.
4. Operational evidence of security and privacy controls is required to provide assurance. This evidence can then be used for certification.




Fig. 1. Assurance in a cloud supply chain

Throughout the cloud supply chain, cloud actors share the overall responsibility of security and privacy. These objectives are achieved and supported by adopting and deploying different security and privacy technologies (as depicted in Fig. 1). Such technologies provide different support within and across cloud actors' domains. The problem then is how to provide assurance that the adopted technologies as a whole

support security and privacy objectives across the supply chain, that is, how to provide supporting evidence that the adopted security and privacy technologies are appropriate and effective for the specific cloud supply chain.

4 Implementing Assurance

Keeping in mind what has been introduced so far, this section discusses various aspects of implementing assurance in cloud supply chains, that is, emerging technical considerations to be addressed while implementing a system supporting assurance. System functionalities that support assurance for the whole cloud supply chain are discussed. Notice that specific technical points are not implementation steps to follow, but rather insights which aim to inform on how assurance can be implemented in a concrete scenario.

4.1 Evidence of Cloud Controls

Cloud service providers often work together (e.g. sub-contract services or relies on third-party resources constrained by specific service level agreements) in order to provide specific services to cloud customers. This may result in complex cloud supply chains involving several cloud service providers working jointly (as depicted in Fig. 1). In a cloud supply chain, security is therefore a shared responsibility among the actors involved. Cloud providers deploy different security and privacy controls in order to guarantee critical service features.

In order to support accountability, cloud providers need to gather evidence as proof that security and privacy controls are effective and suitable in addressing emerging threats. Cloud providers can then be entrusted with sensitive data. Table 1, for example, lists some controls drawn from the CSA Cloud Control Matrix v3.0.1 [15], in particular, controls from two different domains: *Data Security & Information Lifecycle Management*, and *Supply Chain Management, Transparency, and Accountability*. Similarly, the NIST Cloud Computing Security Reference Architecture identifies a list of controls (requirements) to mitigate security risks [2].

However, both NIST and CSA aim mitigating security risks from a high-level perspective, providing no guidelines on which operational aspects of controls should be supervised and which data should be stored in order to prove that deployed controls are effective and suitable in addressing emerging threats. Therefore, a specific set of controls and associated (type of) evidence should be defined for each specific cloud environment. However, independently of any cloud environment, it is possible to build a general framework that will ease the task of managing these controls and evidence.

It is necessary that each security and privacy control clearly defines which (type of) evidence it requires to be gathered in a cloud supply chain. Evidence should focus on operational aspects of deployed controls that need to be monitored. If such evidence is not produced, controls cannot be regarded as supporting security and privacy objectives (e.g. in terms of compliance with security and privacy policies). The proposed CloudTrust Protocol (CTP), for example, provides a basic mechanism for sharing evidence across cloud supply chains [16], hence supporting transparency in the cloud.

Table 1. Examples of controls from the CSA Cloud Control Matrix v3.0.1

Control Domain	CCM V3.0 Control ID	Updated Control Specification
Data Security & Information Lifecycle Management: Classification	DSI-01	Data and objects containing data shall be assigned a classification by the data owner based on data type, value, sensitivity, and criticality to the organization
Data Security & Information Lifecycle Management: Handling/ Labeling/Security Policy	DSI-04	Policies and procedures shall be established for the labeling, handling, and security of data and objects which contain data. Mechanisms for label inheritance shall be implemented for objects that act as aggregate containers for data
Data Security & Information Lifecycle Management: Ownership/Stewardship	DSI-06	All data shall be designated with stewardship, with assigned responsibilities defined, documented, and communicated
Supply Chain Management, Transparency and Accountability: Supply Chain Metrics	STA-07	Policies and procedures shall be implemented to ensure the consistent review of service agreements (e.g., SLAs) between providers and customers (tenants) across the relevant supply chain (upstream/downstream). Reviews shall performed at least annually and identify non-conformance to established agreements. The reviews should result in actions to address service-level conflicts or inconsistencies resulting from disparate supplier relationships

4.2 Control-Evidence Data Model

As discussed in the previous section, it is necessary to associate controls to evidence about them. Such evidence can be gathered in a dedicated permanent storage platform (e.g. a software defined storage). There should be an entity in charge of creating, reading, updating, and deleting these controls and relating them to specific evidence. The same entity would also be in charge of the transactions to and from the storage platform. We will call this entity *Control Manager*. Our proposed implementation framework will hence include three main classes: (1) *Control Manager*, (2) *Control*, and (3) *Evidence Item*. Figure 2 provides a diagrammatic representation of the proposed data model.

Each *Control Manager* may handle several Controls, and each of these may have different types of Evidence Items associated with it. Note that the same type of Control may be configured differently in operation, hence, it may be necessary to store different types of evidence. The only way to associate Controls and Evidence Items should be through a Control Manager. While instantiating specific controls, changes will be immediately applied to the storage platform. A Control Manager may also include some metadata defined by the cloud actor using the Control Manager. This metadata, for example, may include information about specific instances of controls which it handles and their associated types of evidence. The Control Manager should be able to communicate directly with the storage platform via a dedicated Application Program Interface (API).

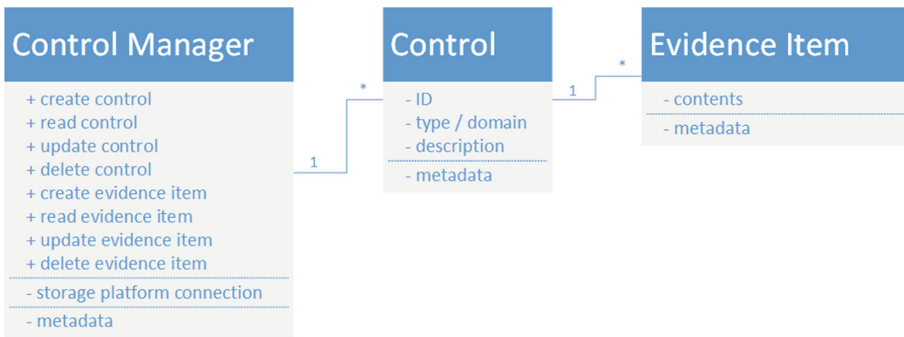


Fig. 2. Proposed Control-Evidence data model

A *Control* will be described by (at least) three fields, as listed by the CSA Cloud Control Matrix [15]: ID, control domain and description. Each Control should be supported by at least one Evidence Item. This evidence will support auditing of the Control (e.g. in terms of policy compliance). Each Control should keep track of its associated Evidence Items. It can also include user-defined metadata (e.g. what type of evidence it is associated with, timestamps like when was the last time this control was audited, etc.).

Finally, an *Evidence Item* is a collection of information that needs to be kept for a Control to support auditing. It can be regarded as a wrapper for the required information. Its contents are, a priori, not of interest for the Control Manager. On the contrary, they will be necessary for an auditor to grant that the deployed set of controls is suitable and effective in order to mitigate security and privacy threats. As with Controls, an Evidence Item may include user-defined metadata (e.g. type of evidence stored such as log file, configuration file, performance metrics, who generated it, etc.).

4.3 Roles in Providing Assurance

A cloud supply chain will need to meet certain controls to prove its accountability. These controls require evidence as proof of their fulfilment. As it was mentioned

previously, it is necessary that there exists some permanent storage platform in the cloud supply chain where this evidence will be stored. This responsibility will be assigned to one cloud provider. This storage platform should be accessible by the other providers in the cloud supply chain, as this is where they will store their Evidence Items. It should count with the required security measures to guarantee confidentiality, integrity, and availability (e.g. access control, encryption, backups, etc.). Figure 3 shows a sample cloud supply chain in terms of actors and their associated responsibilities in sharing and contributing to an evidence storage for controls.

In this example, different controls (numbered 1 to 5) are deployed to guarantee security and privacy of data. The evidence associated with them is stored in specific locations which are managed according to the responsibilities in the cloud supply chain. In this example, Cloud Provider A is in charge of managing a software defined storage platform, as well as it is responsible for providing the evidence for controls 1, 2, and 3. On the other hand, Cloud Provider B (subprovider) is only responsible for providing the evidence for controls 4 and 5, which are the ones that affect it. Once that all the evidence is produced, Cloud Provider A is able to reason over it and, if everything is correct, eventually demonstrate to the Cloud Customer that all the controls are implemented adequately, hence providing assurance.

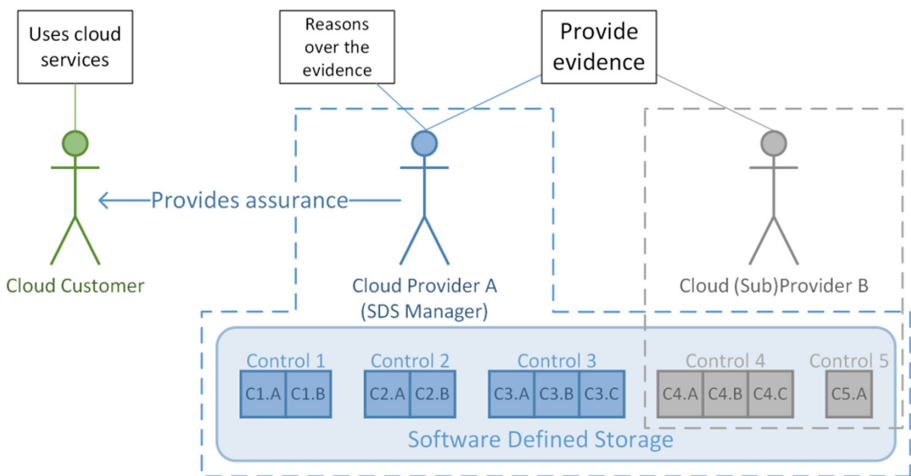


Fig. 3. Sample cloud supply chain

4.4 Evidence Access

As depicted in the previous section, specific Evidence Items are to be provided by specific cloud providers. The access to this evidence should be limited only to the providers who are responsible for them (and, when appropriate, to the auditors).

A Control may require several Evidence Items in order to be considered complied with. These Evidence Items could be supplied by different cloud providers. In this case, it would be desirable that each provider is only allowed access to its related Evidence

Items and no others, hence preventing them from being tampered with by unrelated providers. This scenario is shown in Fig. 4, where Control B requires evidence coming from two different sources. Evidence Items 4 and 5 should only be accessed by Cloud Provider A and Evidence Item 6 only by Cloud Provider B. In this case, an object-level access control is required.

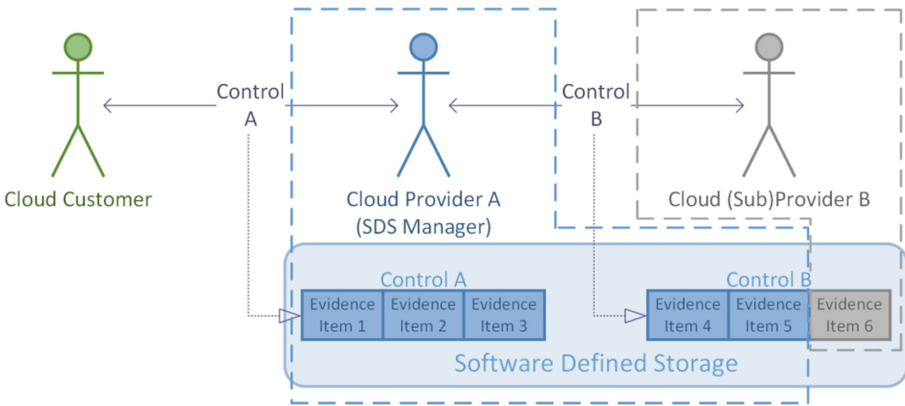


Fig. 4. Desirable requirements for access control

Three of the major software defined storage platforms – OpenStack Swift, Google Cloud Storage, and Amazon S3 – have a two-level hierarchy where the upper level serves as a container⁴ for the objects which contain the relevant information to be stored. One can think of a container as a folder where only files (objects) can be stored, not allowing nested folders. The finest granularity that some software defined storage platforms allow (for example, OpenStack Swift [17]) is per container. This means that a user who is granted access rights to a specific container (Control) may then access all its objects (Evidence Items) – depicted in Fig. 5. In order to support object-level access control, additional security mechanisms that allow finer access granularity are required.

Alongside access control, there are other security and privacy concerns that need to be addressed. As an example, integrity checks must be enforced in order to guarantee that the Evidence Items kept in the software defined storage platform have not been tampered with. Enabling monitoring of events at the object level could be useful in small scenarios. However, this may involve dealing with a remarkable amount of data in large scenarios, making it a hardly scalable solution.

In scenarios where two different cloud providers need to share the same Evidence Items, there is a risk of data aggregation. If this situation is likely to arise, additional mechanisms which filter the shared information to specific actors should be implemented – for example, transparency logs [18].

⁴ “Containers” in OpenStack Swift and Google Cloud Storage and “buckets” in Amazon S3.

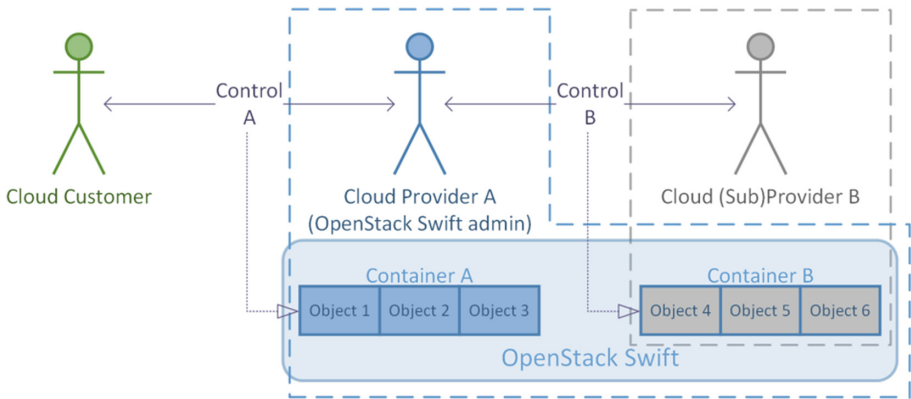


Fig. 5. Access control using OpenStack Swift

5 An Assurance Use Case

The main goal for the Cloud Accountability Project (A4Cloud) is to increase trust in cloud computing by devising methods and tools, through which cloud stakeholders can be made accountable for the privacy and confidentiality of information held in the cloud. Among other milestones, it has specified an accountability model for cloud supply chains ([5]) and several tools to support accountability have been implemented.

In order to prove the application of the accountability model and related tools, a demonstrator scenario has been developed [19]. In this section, we recall this scenario in order to explain a realistic situation where the system depicted in the previous sections can prove useful.

5.1 Wearable Service Use Case Explained

Wearable Co. is a manufacturer of wearable devices that collect well-being data from its wearers. It uses the SaaS⁵ provider Kardio-Mon to provide additional services to its customers. Kardio-Mon integrates Map-on-Web's services into their own. Kardio-Mon and Map-on-Web use the IaaS⁶ provider DataSpacer to run their services. This scenario is depicted in Fig. 6, where the interactions among the different actors have been numbered. For the sake of simplicity, only interactions between two actors have been considered.

These interactions are subject to be monitored (implementing controls), either continuously or occasionally. The evidence collected to support this process, supplied by the different cloud providers, will be stored in an OpenStack Swift server whose administrator will be Kardio-Mon. The reasons to use this platform are that the demonstrator scenario for the Cloud Accountability Project uses an OpenStack

⁵ SaaS: Software as a Service.

⁶ IaaS: Infrastructure as a Service.




Fig. 6. Wearable service use case: environment and storage platform with access permissions

deployment and also because it is open-source. In the event of having an external auditor to audit these controls, she will require access to read this evidence. Figure 6 shows also the access permissions⁷ for all actors involved in the demonstrator use case.

5.2 An Assurance Example: Implementing SLAs

Given the scenario presented in the previous section, let’s consider an example where service level agreements (SLA) among the different cloud providers are to be implemented and reviewed, as defined in control STA-07 from the CSA Cloud Control Matrix v3.0.1 (see Table 1). Each SLA will be considered as a separate Control.

For the sake of simplicity we will ignore one of the cloud providers (Map-On-Web) and we will focus on the SLAs between (1) Wearable Co and Kardio-Mon and (2) Kardio-Mon and DataSpacer. The case of having an external auditor in the system will also be considered.

As specified previously, Kardio-Mon will be the OpenStack Swift server administrator. Each Control – one per SLA – needs to be associated with a container in

⁷ Note that each control will be associated with a container in OpenStack Swift.

OpenStack Swift. These will be named STA-07-SLA1 and STA-07-SLA2. Kardio-Mon is responsible for creating them and for granting the expected access rights. Let’s consider that each Control requires only three Evidence Items to support its proper operation: (1) the SLA definition, (2) some performance metrics, and (3) some operation logs.

Let’s consider that Kardio-Mon is the cloud provider in charge of supplying the SLA definitions and the updated performance metrics. The logs are to be supplied by the cloud provider running the service. This means that Kardio-Mon is responsible for all the Evidence Items from Control STA-07-SLA1 and for the SLA definition and performance metrics for STA-07-SLA2. With respect to DataSpacer, it should only provide the logs for STA-07-SLA2. This scenario is depicted in Fig. 7.

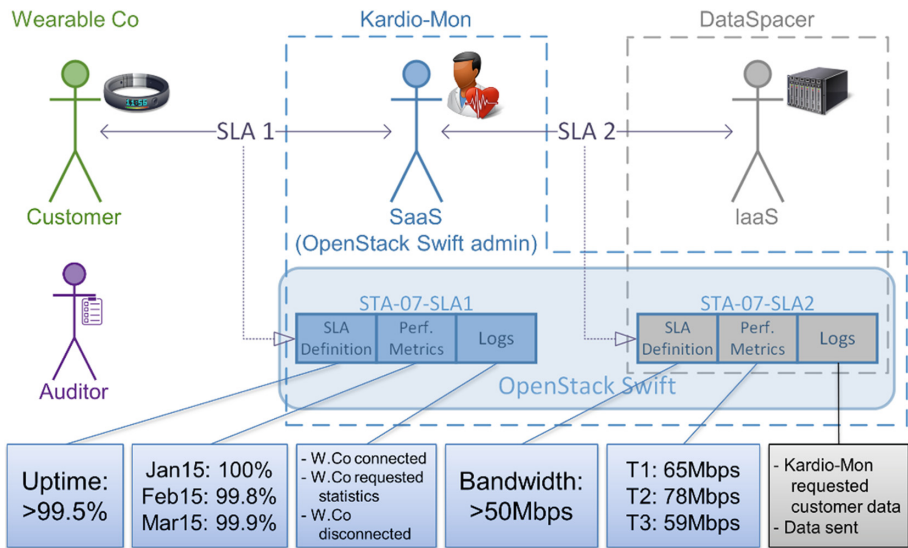


Fig. 7. Example of a cloud environment with SLAs in place

As pointed out in Sect. 4.4, OpenStack Swift does not support object-level access controls. This means that both Kardio-Mon and DataSpacer are able to access the Evidence Items from Control STA-07-SLA2. Consequently, it needs to be ensured that none of the cloud providers have modified – either intentionally or accidentally – the Evidence Items whose responsibility falls on the other provider.

As an auditor is to be expected to join the scenario, her access rights should be set in OpenStack Swift. She should be granted reading permissions to all the Controls. On a different note, the role of Wearable Co is limited to cloud customer, hence not being part of the cloud. Therefore, it should have no access rights whatsoever to the storage platform. All the access rights are collected in Fig. 8.

Actor	Wearable Co		Kardio-Mon		DataSpacer		Auditor	
Operation	Read	Write	Read	Write	Read	Write	Read	Write
SDS*	x	x	✓	✓	x	x	x	x
SLA1	x	x	✓	✓	x	x	✓	x
SLA2	x	x	✓	✓	✓	✓	✓	x

*SDS: Create/Delete containers and modify privileges

Fig. 8. Access rights for the different actors

Wearable Co is the one who, ultimately, is interested in receiving assurance that the data that it puts in the cloud will be adequately protected using privacy and security measures. This assurance may be provided by an external auditor or by an auditor within the cloud environment. In the latter case, one of the cloud providers should act as an auditor, providing comprehensive assurance about the cloud supply chain to the customer. Note that this system can be used as well for such internal auditing.

6 Concluding Remarks

This paper has briefly discussed security and privacy assurance in cloud ecosystems and provided some guidelines on how it can be implemented throughout a cloud supply chain. The controls to be set should be associated with evidence that supports compliance with security and privacy policies. This evidence should be saved in a permanent storage platform accessible to the different cloud providers.

The discussion provides a rationale for the assurance problem in the cloud and highlights some preliminary requirements. In order to provide support for security and privacy assurance throughout the cloud supply chain, it is necessary:

1. to regard security and privacy solutions as deployed across the cloud supply chain rather than from a single organisation viewpoint,
2. to design and implement means for supporting assurance,
3. to understand emergent dependencies among security and privacy solutions deployed in cloud ecosystems,
4. to assess how security and privacy solutions comply with (or enable to comply with) organisational as well as regulatory policies,
5. to gather operational evidence that supports security and privacy assurance across the cloud supply chain.

A system that can help gather and classify assurance evidence and control which users can access it is also introduced. This system can ease auditing the cloud supply chain, eventually contributing to providing security and privacy assurance.

Future research and development activities will focus on continuous monitoring of the cloud supply chain in order to address the security and privacy risks as soon as they arise, hence avoiding jeopardizing its assurance.

Acknowledgements. This work has been partly funded by the European Commission's Seventh Framework Programme (FP7/2007-2013), grant agreement 317550, Cloud Accountability Project – <http://www.a4cloud.eu/> – (A4Cloud).

References

1. NIST Cloud Computing Reference Architecture, Special Publication 500-292
2. NIST Cloud Computing Security Reference Architecture, Special Publication 500-299
3. The NIST Definition of Cloud Computing, Special Publication 800-145
4. ENISA: Cloud Computing Benefits, risks and recommendations for information security
5. Felici, M., Pearson, S.: Accountability for data governance in the cloud. In: Felici, M., Fernández-Gago, C. (eds.) A4Cloud 2014. LNCS, vol. 8937, pp. 3–42. Springer, Heidelberg (2015)
6. Felici, M.: Cloud accountability: glossary of terms and definitions. In: Felici, M., Fernández-Gago, C. (eds.) A4Cloud 2014. LNCS, vol. 8937, pp. 291–306. Springer, Heidelberg (2015)
7. NYMITY Inc.: Privacy Management Accountability Framework (2014)
8. ENISA: Cloud Certification Schemes Metaframework, Version 1.0, November 2014
9. Gartner: Magic Quadrant for Security Information and Event Management, June 2014
10. European Commission: Unleashing the Potential of Cloud Computing in Europe (2012)
11. ENISA: Privacy, Accountability and Trust – Challenges and Opportunities. European Network and Information Security Agency (ENISA) (2011)
12. Baldwin, A., Pym, D., Shiu, S.: Enterprise Information Risk Management: Dealing with Cloud Computing. In: Pearson, S., Yee, G. (eds.) Privacy and Security for Cloud Computing. Computer Communications and Networks, pp. 257–291. Springer, London (2013)
13. CSA: The Notorious Nine Cloud Computing Top Threats in 2013. Top Threats Working Group, Cloud Security Alliance (2013)
14. Pearson, S.: Accountability in Cloud Service Provision Ecosystems. In: Bernsmed, K., Fischer-Hübner, S. (eds.) Secure IT Systems. LNCS, vol. 8788, pp. 3–24. Springer, Heidelberg (2014)
15. CSA: Cloud Control Matrix v3.0.1, October 2014
16. CSC: A Precipice for the CloudTrust Protocol (V2.0). Computer Sciences Corporation (CSC) (2010)
17. Arnold, J.: OpenStack Swift, O'Really (2015)
18. Pulls, T.: Preserving Privacy in Transparency Logging. (Doctoral dissertation). Karlstads Universitet (2015)
19. Wiktor Włodarczyk, T., Pais, R. (eds.): Deliverable D38.2 Framework of evidence (final), A4Cloud (2015)