Security SLAs for Cloud Services: Hadoop Case Study

Massimo Ficco and Massimiliano Rak

Abstract Cloud paradigm is currently one of the most remunerative segments of Information Technology. It has gained the interest of a very large number of corporates and organizations. However, despite the promising features, security is the major concern for businesses that want to shift their services to the cloud. On the other hand, business critical systems must be certified against a set of security controls to be compliant to security standards, as well as to mitigate potential security incidents. Therefore, cloud service providers must employ adequate security measures that conform to security controls expected by the information systems they host; moreover, they should be able to grant the correct application of such controls to their customers. Security service level agreements (SLAs) are a way to face such issues, through the definition of contracts among cloud service providers and customers that clearly state the security grants applied to the offered cloud services. This chapter illustrates a case study that describes how it is possible to implement such security SLAs on a concrete cloud service, which offers Apache Hadoop services over public cloud providers. The chapter outlines how to write and assess security SLAs on such services.

Keywords Cloud security • Service level agreement • Security controls

1 Introduction

Cloud computing is nowadays a largely adopted technology for providing any kind of services. Its success is due to the on-demand self-service, which enables user to acquire cloud service and resources according to a pay-by-use business model. In general, cloud service providers (CSPs) offer guarantees in terms of service availability and performance during a time period of hours and days. The provisioning contracts regulate the cost that customers have to pay for provided services and

M. Ficco (✉) • M. Rak
Department of Industrial and Information Engineering, Università degli Studi della Campania "Luigi Vanvitelli", Via Roma 29, 81031, Aversa, CE, Italy
e-mail: massimo.ficco@unina2.it; massimiliano.rak@unina2.it

K. Corsi et al. (eds.), *Reshaping Accounting and Management Control Systems*,
Lecture Notes in Information Systems and Organisation 20,
DOI 10.1007/978-3-319-49538-5_7

resources. On the other hand, due to their openness to the Internet, cloud services are prone to cyber attacks, which aim at violating security and privacy of the targeted enterprise systems. Several works proposed in the literature present models and mechanisms for monitoring and assuring service privacy and security guarantees in the cloud computing context [1, 2]. In particular, several works explore SLAs for security and analyze security metrics in new paradigms like cloud computing [3, 4]. By incorporating security parameters in the SLA could improve the quality of the service being offered. This objective has profound implications in the security solution to be implemented and delivered. Moreover, it the last years, many security standards and requirement frameworks have been developed in order to address risks to enterprise systems and critical data. On the other hand, most of these efforts are essentially exercises in reporting on compliance and defining security program resources to face evolving attacks that must be addressed.

The *security controls* are guidelines to identify and prioritize security actions, which are effective against cyber threats, with a strong emphasis on "what works," i.e., tools, processes, architectures, and services that have been used and demonstrated real-world effectiveness. However, the available standards leave the process of security controls selection to the organizations. Moreover, the type of security controls to be applied depend on the asset to be protected and are identified on the basis of a risk analysis, which provides a set of significant risks and data to assist in the treatment of these risks.

In this chapter, we propose a method for security controls selection for a cloud-based service. In particular, we consider an Apache Hadoop service as case study. Apache Hadoop is an open-source software framework for distributed storage and distributed processing of very large datasets on cloud. We preset model to manage the SLA life cycle, which can be used to cover the semantic gap among CSC security requirements and security controls offered by CSPs, as well as adopted to compare the services offered by different CSPs. Moreover, we perform an asset evaluation to determine the most critical security controls to be implemented to protect the provided cloud service.

The rest of the chapter is organized as follows: Sect. 2 introduces the system model, as well as the definition of the problem we are focusing on. Section 3 presents the related work in the field of security controls applications. Section 4 introduces the adopted security SLA model, whereas Sect. 5 describes the risk assessment model to be used by the cloud customers. Section 6 illustrates the proposed approach on the Hadoop case study. Section 7 presents a short summary and future work.

2 Problem Definition

Cloud computing paradigm involves many use cases (see [5] for an overview), each of them implying different types of security issues and different ways of involving security and SLAs. Existing standards [6–8] offer a clear classification of the main

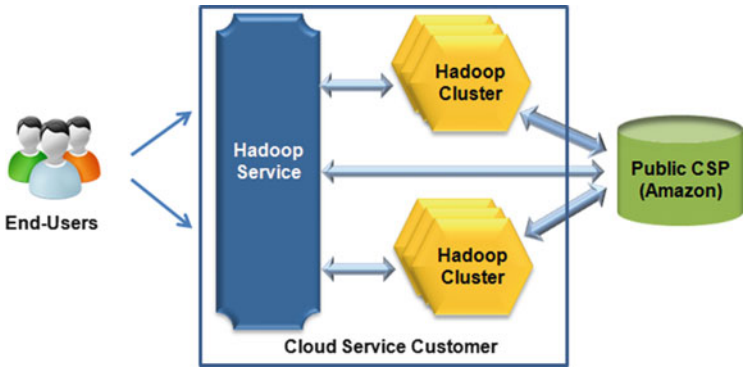

Fig. 1 The system model

concepts associated with cloud computing and of the roles that parties may assume in cloud scenarios.

In this chapter, we assume the common scenario, in which a cloud customer (CSC) wants to know the security grants offered by a public CSP, such as Amazon, in order to decide whether to acquire cloud resources, which will be used to provide a service (in our case study a Hadoop service) to its end users. Figure 1 depicts the scenario we are focusing on.

Therefore, this study focuses on the typical security issues related to the services offered by the CSP to the CSC: *How can the CSC rely on CSP services? How reliable is the offering?* In order to well outline the issues, few considerations are useful: the CSC is not a big cloud service provider, whose reliability is (ideally) granted by its dimension and relevance in the market. The CSC is a cloud service reseller that focuses on a specialized market with well-identified needs, differently from big CSPs, like Amazon, which has no interest in offering services *customized* for a specific audience. The CSC, on the other hand, has the need to evaluate the security risks associated with the usage of the cloud service, especially in case of management of critical data. In such a context, he needs detailed information about the security offered by the CSP, which often is not granted by big CSPs. Thus, we focus on the adoption of security SLAs as a way: (1) to allow CSC to be able to make a concrete risk assessment of adoption of cloud services and (2) to enable CSCs to add value to their services in a well-defined market niche. In order to obtain such a result, we propose that the CSP offers a security SLA able to represent, in a transparent manner, the security grants offered by the cloud provided to its CSCs. Moreover, we propose simple risk model that enables the CSC to compare the security SLA offered by the CSP in order to evaluate the cloud service that best fits his security needs.

3 Related Work

The main problem in adopting security controls is the lack of a clear representation in the cloud computing context, which makes it difficult to connect organizational certification efforts to the services offered by CSPs. In this direction [9] proposes a compliance vocabulary, which creates a set of security SLA terms that are derived from security controls in governance documents, including the NISTSP800-53 [2], the Common Criteria Part 2 [3], the DISA Secure Application Security Technical Implementation Guide (STIG), and the Cloud Security Alliance Cloud Control Matrix (CCM) [10]. Existing services would rely on the compliance vocabulary to represent the controls it must satisfy and embed the corresponding terms in its SLA. In [11], authors propose a methodology to evaluate the information security controls. They rank the controls quantitatively in accordance with given criteria. Peláez [12] describes how to measure the effectiveness of security controls. In particular, a qualitative risk assessment method is adopted. It assigns a huge amount of metrics to each security control in order to measure its quality.

4 Security SLAs and Security Controls

The main goal of security SLAs is to represent the security level offered by the cloud service in a clear way, in order to cover the gap between the CSC, which focuses on his own requirements, and the CSP, which focuses on the security mechanisms he is able to implement [2].

In order to characterize the security in a service, Lindskog [13] defines four dimensions, including type of protection service (e.g., confidentiality), protection level (e.g., number of assets that must be encrypted), adaptiveness (i.e., the ability of a service to change protection levels at run-time), and protection level specification (i.e., the security policy). Bernsmed et al. [14] develop a framework that supports the security SLA management in federated clouds. In this work, we adopt the security model proposed in the SPECS project [15]. Such an SLA model is founded on an SLA life cycle, based on all the up-to-date standards, which includes five main phases: negotiation, implementation, monitoring, and remediation.

In order to cover the semantic gap among CSCs and CSPs, the SPECS SLA model adopts the concept of security controls. Security controls can be physical, technical, or administrative [16]. Each category of controls can be further classified by using either preventive or detective approaches. Preventive controls attempt to avoid the occurrence of unwanted events. They inhibit the use of unauthorized computing resources. Detective controls attempt to identify unwanted events after they have occurred. Examples of detective controls include audit trails, intrusion detection methods, and checksums. Other types of controls are usually described as deterrent, corrective, and recovery, which do not belong to either preventive or detective categories. Deterrent controls are used to discourage malicious users from

intentionally violating information security policies or procedures. These are usually constraints that make it difficult to perform unauthorized activities or influence a potential intruder to not violate security. Corrective controls remedy the circumstances that allowed the unauthorized activity. They could result in changes to existing physical, technical, and administrative controls. Recovery controls restore lost computing resources or capabilities caused by a security violation. Deterrent, corrective, and recovery controls are considered to be special cases within the major categories of physical, technical, and administrative controls. For example, deterrence is a form of prevention because it induces dissuasive effect to the intruder. Corrective controls can be assimilated to technical controls, when antivirus removes a malware, or with administrative controls, when backup procedures enable restoring critical data. Finally, recovery controls can be considered as administrative controls when they implement disaster recovery and contingency plans.

The SPECS SLA model reports, for each service covered by SLA, the security controls that the CSP offers on top of it, as represented in Fig. 2. The model assumes that security is expressed in terms of (1) cloud resources, i.e., the description of the resources obtained by the cloud customer, (2) security capabilities, i.e., set of security controls granted on the cloud resources, (3) security metrics, which are the measurable (and externally verifiable) part of the security offered on the cloud service, and (4) service level objectives (SLOs), expressed as thresholds on security metrics, which represent the concrete grants offered to CSCs. Such model is built in order to be perfectly compatible with the WS-Agreement standard, and SPECS offers a WSAG extension to represent the model in a machine readable format.

According to the above model, CSP can build up a security SLA associated with its own service. In particular, the SLA implementation requires:

- A description of the cloud service
- The identification of the implemented security controls
- The identification of the security metrics that can be granted
- The formalization of the security SLA

The inclusion of the security controls in the SLA favors a comparison of offered service and shifts some certification burden to the CSP-based contractual SLA terms. Finally, on the basis of a risk assessment, the CSC can choose which CSP best meets their strict compliance requirements.

5 Security Risk Assessment

As presented in [17], a security model has to be considered three interconnected dimensions: *asset* is anything that has value to the organization; *threat* can inflict damage to assets of an organization; and *security control* is a management, operational, or technical mechanism, which allows defining assets against threats. It is clear that the main property of an asset is its importance for the organization.

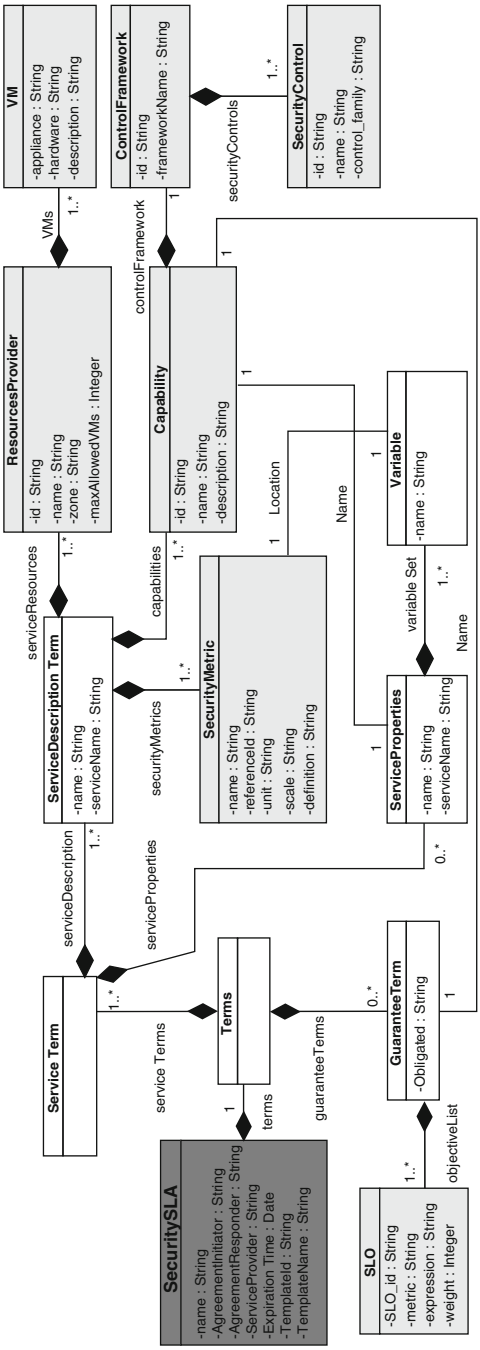

Fig. 2 SPECS SLA model

Therefore, in order to identify which CSP offers the service that best meets his requirements, an analysis of potential risks for the asset should be performed by the CSC. In particular, a risk evaluation matrix should be implemented. As Table 1 shows, the matrix represents the likelihood and consequences of each threat, which are used to compute the risk values.

Then, for each identified threat should be verified which kinds of security controls are offered by the considered CSPs. Such security controls represent mitigation means for the analyzed threat. Thus, on the basis of the level of risk the CSC accepts for the asset, it is necessary to select the CSP.

Definitely, providing security control compliance services can be economically advantageous for CSPs to attract CSCs with strict compliance requirements. Therefore, for each category of CSC, CSPs should choose the security controls to be implemented on the basis of CSC needs, considering also the costs (in terms of money), difficulty of implementation, and time consumption of maintenance that the CSC should waste to implement on its own the same security controls. This analysis would allow identifying the more appropriate security mechanisms to be implemented in comparison to their cost and the level of risk the CSC accepts. Additional security mechanisms can be contracted with the CSC in the security SLA.

6 Security Control Assessment

According to the proposed approach, in order to identify the security level to be offered to CSCs through an SLA, a CSP has to determine which, and how many, controls have to be implemented to protect the hosted service. In the context of this work, we assume that the only applied security controls are those implemented by a basic Hadoop cluster. We adopted NIST SP 800-53r4 guidelines to determine the implemented security controls. According to the NIST structure, the security controls are organized into 18 families, such as access control, security assessment

Table 1 Evaluation matrix with risks likelihood and consequences

	Insignificant	Minor	Moderate	Major	Catastrophic
Rare	Acceptable	Acceptable	Acceptable	**Acceptable**	**To be evaluated**
Unlikely	Acceptable	Acceptable	Acceptable	**To be evaluated**	**To be evaluated**
Possible	Acceptable	Acceptable	**To be evaluated**	**To be evaluated**	**To be evaluated**
Likely	Acceptable	**To be evaluated**	**To be evaluated**	**To be evaluated**	**To be evaluated**
Certain	**To be evaluated**	**To be evaluated**	**To be evaluated**	**To be evaluated**	**To be evaluated**

Table 2 Access controls for Apache Hadoop cluster

Name	Control	How	
Access Control Policy and Procedures	The organization develops and documents: (1) An access control policy that addresses purpose, scope, roles, responsibilities, management commitment, coordination among organizational entities, and compliance; and (2) Procedures to facilitate the implementation of the access control policy and associated access controls	The purpose of introducing a policy of access control is to increase the security with respect to external attacks, to ensure the functionality and integrity of our system The roles within the system are admin and users. The access control is assigned to the admin	Yes
Account Management	The organization: (a) Identifies and selects the types of information system accounts; assigns account managers for system accounts; establishes conditions for group and role membership (b) Assigns managers for information system accounts (c) Establishes conditions for group and role membership (d) Monitors the use of information system accounts	The types of accounts available in the system are admin who creates and manages accounts and users who are all users who use the system The account manager is an admin account Ubuntu All the users belong to the group user, which will have limited Hadoop permissions Monitor access to the cluster via the log files *auth.log* content in */var/log/* and the performed operations through the Hadoop log files	No
Access Enforcement	The system enforces approved authorizations for logical access to information and system resources in accordance with applicable access control policies	Only via SSH	Yes
Unsuccessful Login Attempts	The system enforces a limit of consecutive invalid login attempts by an user during a defined time period	Login takes place without the use of a password, with no limit on failed attempts	No

and authorization, personnel security, identification and authentication, system and communications protection, incident response, system and information integrity, etc. Each family contains a set of security controls related to the general security topic of the family. They can involve aspects of supervision, policy, oversight, manual processes, actions by individuals, and mechanisms. Tables 2, 3, and 4 list tree security control families and some related security controls applied to the considered case study. For each security control is described the name, a description of the control, how to apply it, and if it is already implemented by the Hadoop cluster. For example, as reported in Table 3, Hadoop does not support any security control to protect against Denial of Service (DoS) attacks [18–20], which could

Table 3 System and communications protection for Apache Hadoop cluster

Name	Control	How	
Denial of Service Protection	The system protects against or limits the effects of Denial of Service, by employing some protection mechanism.	Using Secure Copy for the initial handshake will have problems of denial of service because an attacker could send a lot of files and then to consume system resources; No protection mechanism is provided.	No
Cryptographic Protection	The information system implements cryptography policies in accordance with applicable federal laws, directives, policies, regulations, and standards.	There is no encryption on the data stored on the distributed file system. The only encryption in the system is the encryption key that can be RSA or DSA.	No

Table 4 System and information integrity controls for Apache Hadoop cluster

Name	Control	How	
System and Information Integrity Policy and Procedures	The organization: (a) Defines system and information integrity policies; (b) Defines procedures to facilitate the implementation of the identified policies.	To recover data from a damaged Data Node, a client implements a checksum on the file HDFS, which compute a checksum for each Tile and stores it in a separate hidden file. When a client retrieves file, it verifies that the data received from each Data Node match the checksum. Otherwise, the client can choose to retrieve that block from another Data Node that has a replica of that block.	Yes
Malicious Code Protection	The organization: (a) Employs malicious code protection mechanisms at information system entry and exit points to detect and eradicate malicious code; (b) Addresses the receipt of false positives during malicious code detection, and the resulting potential impact on the availability of the system.	The system does not provide mechanisms to protect from malicious code if not one already present on the Linux below.	No

exhaust cloud resources used to run the CSC's services, whereas in Table 1 it is shown that the Hadoop framework provides mechanisms to manage the access control policies. Therefore, on the basis of hypothetical security requirements of the market niche to be covered, the CSP has to assess which controls have already been implemented by the hosted service, as well as identify which should be added to

satisfy CSC's security requirements. Then, for each identified control, it has to define the possible metrics to evaluate the control, as well as provide the assessment tool for supporting the CSC audit (monitoring). For example, in order to protect against the DoS attacks, CSP could deploy a prevention mechanism, such as mOSAIC-IDS [21], SNORT [22], and OSSEC [23], which are intrusion detection system to detect anomalous activities against the Hadoop cluster [24, 25]. Security metrics used to perform measurements of the correct delivery of the security capability during system operation could be "*false_positives*," "*true_positives*," "*detection_latency*," etc.

However, the process of security control assessment has to take into account the changes to the system and its operating environment, or the changes outside CSP direct control, which may introduce new security vulnerabilities, and may require a new assessment of some or all security controls. Moreover, new security controls could be added in order to cover possible new market niches.

7 Conclusions

Security is a key issue that inhibits many businesses and government organizations from moving to the cloud. For an organization to have cloud-based services with certification guarantees means increased service efficiency and reputation.

In this chapter, we propose an approach to perform the security assessment of the cloud services offered by CSPs. The results of this assessment are used in determining the overall security offered to the CSC, identifying residual vulnerabilities in the system, providing credible and meaningful inputs to the cloud security administrators, as well as enabling little CSPs to add value to their services in a well-defined market niche, by using security SLA able to describe the security offered on top of their services. On the basis of an accurate risk assessment of required cloud services, a CSC can compare the security SLAs offered by different CSPs in order to evaluate the cloud service that best fits his security requirements.

Acknowledgment This research is partially supported by the European Community's Seventh Framework Programme (FP7/2007–2013) under Grant Agreements no. 610795 (SPECS), as well as the POFESR Campania 2007/2013, Asse 2, 00 2.2, "Bando Sportello dell'Innovazione, Azione 3 e 4, Progetti di trasferimento Tecnologico Cooperativi e di Prima Industrializzazione per le Imprese Innovative ad Alto Potenziale" project ITINERE (ID 06-05-070138).

References

1. Ficco, M. (2013). Security event correlation approach for cloud computing. *Journal of High Performance Computing and Networking, 7*(3), 173–185.

2. Cicotti, G., Coppolino, L., D'Antonio, S., & Romano, L. (2015). Runtime model checking for SLA compliance monitoring and QoS prediction. *Journal of Wireless Mobile Networks, Ubiquitous Computing, and Dependable Applications (JoWUA), 6*(2), 4–20.
3. Ficco, M., & Rak, M. (2012). Intrusion tolerance as a service: A SLA-based solution. *Proceedings of the 2nd International Conference on Cloud Computing and Services Science (CLOSER)* (pp. 375–384).
4. Ficco, M., & Rak, M. (2012). Intrusion tolerance in cloud applications: The mOSAIC approach. *Proceedings of the 6th International Conference on Complex, Intelligent, and Software Intensive Systems (CISIS 2012)* (pp. 170–176).
5. Cloud Computing Use Cases (2010, July). White Paper, ver. 4.
6. ISO/IEC 17789:2014, Information technology—Cloud computing—Reference architecture. Retrieved from http://www.iso.org/iso/catalogue_detail?csnumber=60545
7. NIST. (2013, April). Special Publication (SP) 800-53, Revision 4.
8. Common criteria (Part 2) for IT security evaluation V3.1. (2012, September). Retrieved from https://www.google.it/?gws_rd=ssl#q=Common+criteria+(Part+2)+for+IT+Security+Evaluation+V3.1
9. Hale, M. L., & Gamble, R. (2013). Building a compliance vocabulary to embed security controls in cloud SLAs. *Proceedings of the IEEE 9th World Congress on Services* (pp. 118–125).
10. CSA. (2012). *Cloud controls matrix*. Retrieved from https://cloudsecurityalliance.org/research/ccm/#_overview
11. Lv, J.-J., Zhou, Y.-S., & Wang, Y.-Z. (2011). A multi-criteria evaluation method of information security controls. *Proceedings of the 4th International Conference on Computational Sciences and Optimization* (pp. 190–194).
12. Peláez, M. H. (2010, April). *Measuring effectiveness in information security controls*. SANS Institute InfoSec Reading Room. Retrieved from https://www.google.it/?gws_rd=ssl#q=Measuring+Effectiveness+in+Information+Security+Controls
13. Lindskog, S. (2005). *Modeling and tuning security from a quality of service perspective*. Ph.D., Department of Computer Science and Engineering, Chalmers University of Technology Goteborg, Sweden.
14. Bernsmed, K., Jaatun, M., Meland, P., & Undheim, A. (2011). Security SLAs for federated cloud services. *Proceedings of the International Conference on Availability, Reliabilty and Security*.
15. SPECS project, Secure provisioning of cloud services based on SLA management. Retrieved from http://www.specs-project.eu/
16. Tipton, H. F. (2003). *Access control principles and objectives*. Retrieved from https://www.cccure.org/Documents/HISM/003-006.html
17. Breier, J., & Hudec, L. (2013, September). On selecting critical security controls. *Proceedings of the 8th International Conference on Availability, Reliability and Security (ARES)* (pp. 582–588).
18. Ficco, M., & Rak, M. (2016). Economic denial of sustainability mitigation in cloud computing. *Organizational Innovation and Change, 13*, 229–238.
19. Ficco, M., & Rak, M. (2015). Stealthy denial of service strategy in cloud computing. *IEEE Transactions on Cloud Computing, 3*(1), 80–94.
20. Ficco, M., & Rak, M. (2012). Intrusion tolerance of stealth DoS attacks to web services. In *Information Security and Privacy* (LNCS, Vol. 376, pp. 579–584).
21. Ficco, M., Venticinque, S., & Di Martino, B. (2012). mOSAIC-based intrusion detection framework for cloud computing. In R. Meersman, H. Panetto, T. Dillon, S. Rinderle-Ma, P. Dadam, X. Zhou, et al. (Eds.), *On the move to meaningful internet systems: OTM 2012* (LNCS, Vol. 7566, pp. 628–644).
22. OSSEC, an open source host-based intrusion detection system. Retrieved from http://www.ossec.net/
23. SNORT, an open source network-based intrusion detection system. Retrieved from https://www.snort.org/

24. Ficco, M., Rak, M., & Di Martino, B. (2012, November). An intrusion detection framework for supporting SLA assessment in cloud computing. *Proceedings of the 4th International Conference on Computational Aspects of Social Networks (CASoN 2012)* (pp. 244–249).
25. Esposito, C., & Ficco, M. (2016). Recent developments on security and reliability in large-scale data processing with MapReduce. *International Journal of Data Warehousing and Mining, 12*(1), 49–68.