

Chapter 14

KEY TERMS FOR SERVICE LEVEL AGREEMENTS TO SUPPORT CLOUD FORENSICS

Keyun Ruan, Joshua James, Joe Carthy and Tahar Kechadi

Abstract As cloud adoption grows, the importance of preparing for forensic investigations in cloud environments also grows. A recent survey of digital forensic professionals identified that missing terms and conditions regarding forensic activities in service level agreements between cloud providers and cloud consumers is a significant challenge for cloud forensics. This paper addresses the challenge by specifying standard terms for service level agreements that support cloud forensics.

Keywords: Cloud forensics, service level agreements

1. Introduction

Cloud computing has the potential to become one of the most transformative technologies in the history of computing. Gartner [4] estimates that, by 2015, 20% of non-IT Global 500 companies will be cloud providers. However, the rapid growth and adoption of cloud computing as a “non-standard system” [3] is raising significant challenges with regard to digital forensics. Cloud organizations, including cloud providers and cloud consumers, must establish forensic capabilities. Otherwise, they will face tremendous difficulties in investigating incidents such as policy violations and criminal intrusions.

Ruan, *et al.* [10] have introduced a three-dimensional model that explores the organizational, technical and legal aspects of cloud forensics. They also analyzed some of the major challenges and opportunities regarding cloud forensics. This knowledge was used to develop a survey of cloud forensics and critical criteria for cloud forensic capabilities, which was submitted to digital forensic experts and practitioners from around

the world. The survey received a strong response from the forensic community. Of the 156 survey responses, 75% of the respondents agreed or strongly agreed that “missing terms and conditions in service level agreements (SLAs) regarding investigations” is a challenge for cloud forensics [10]. Additionally, 77.61% of the respondents agreed or strongly agreed that the “tools provided, techniques supported and access granted regarding forensic investigations should be included in SLAs.” This paper attempts to address the needs by suggesting key terms that should be included in SLAs between cloud providers and cloud consumers regarding the organizational, technical and legal aspects of cloud forensics.

2. Cloud Computing

The NIST definition of cloud computing specifies three deployment models: software as a service (SaaS), platform as a service (PaaS) and infrastructure as a service (IaaS); and four delivery models: private (internal) cloud, community cloud, public cloud and hybrid cloud [5].

The segregation of duties between cloud providers and cloud consumers regarding forensic activities depends on the deployment and delivery models. For example, in the SaaS model, the cloud provider has administrative control over the application layer and total control over the middleware, operating system and hardware layers; while the cloud consumer has limited administrative control over the application layer and no control over the middleware, operating system and hardware layers. In community, public and hybrid clouds, multiple tenants share the middleware, operating system and hardware layers.

In the PaaS model, the cloud provider has administrative control over the application and middleware layers, and total control over the operating system and hardware layers; while the cloud consumer has limited programmability for the application and middleware layers, and no control over the operating system and hardware layers. In community, public and hybrid clouds, multiple tenants share the operating system and hardware layers.

In the IaaS model, the cloud provider has no control over the application, middleware and operating system layers, administrative control over the hypervisor (virtualization) and hardware layers; while the cloud consumer has total control over the application, middleware and operating system layers, and no control over the hypervisor and hardware layers. In community, public and hybrid clouds, multiple tenants share the hypervisor and hardware layers.

3. Forensic Data Access

Before discussing the organizational, technical and legal terms in SLAs regarding cloud forensics, it is important to review the specific cloud offering(s) that the cloud consumer may sign up for with the cloud provider. Also, it is important to understand the governance over the data and operations that are inherent to the offerings.

3.1 Encryption Keys

Currently, major cloud providers such as Amazon encourage cloud consumers to encrypt everything, or at least all their sensitive data [1]. As far as forensic investigations are concerned, it is currently impossible to analyze and examine data that is strongly encrypted.

In the survey by Ruan, *et al.* [9], 83.34% of the respondents agreed that “a procedure and a set of toolkits in the cloud organization to obtain keys for encrypted data in the cloud” is important or very important. The cloud provider and the cloud consumer must agree on the circumstances under which an investigation team may access the encryption keys, and stipulate the policies and procedures pertaining to key management, key access and collaboration with law enforcement during an investigation.

3.2 Logs and Forensic Artifacts

The survey by Ruan, *et al.* [9] noted that 74.24% of the respondents felt that “an agreement on access to and control over forensic data at all levels between cloud organizations” is important or very important.

Logs are valuable data sources in most forensic investigations. Accessing logs in the cloud is almost always a complicated scenario involving the provider and consumer. In the SaaS model, the provider has administrative control over the application logs, while the consumer has limited administrative access. In the IaaS model, the provider has no control over the application and system logs, while the consumer has total control. Also, the provider has administrative control over the hypervisor logs while the consumer has no control.

An SLA should address the right of a forensic team to access the logs at all levels from the provider and the consumer sides. Also, the agreement should address the segregation of duties that should occur when relevant logs from both sides are provided to external forensic entities and law enforcement. Language should be included in the SLA that describes the logging and retention policy, including retention duration, log storage location and the entities who are authorized to access the logs.

In a cloud environment, data is automatically mirrored to multiple locations to ensure data redundancy. Data redundancy is useful in forensic investigations because evidence that is lost or destroyed in one location may be found elsewhere. The agreement should stipulate the conditions under which the internal forensic team, external forensic entities and law enforcement can access mirrored artifacts during forensic investigations.

3.3 Data and Infrastructure Location

In the survey [9], 79.17% of the respondents agreed or strongly agreed that the loss of physical control of data is crucial to forensic investigations. The physical locations of data are needed to determine the jurisdiction of the investigation. The investigator must be aware of the physical locations of the data in order to not breach privacy or other privileges based on the jurisdiction where the data resides. Moreover, the physical location of the data should be kept confidential so as not to expose sensitive data to potential attacks. The agreement between the cloud provider and the cloud consumer must address the transparency of the physical locations of consumer data, virtual disk images and infrastructure. Example considerations include the conditions under which the cloud provider may provide information about the physical locations of consumer data and the procedures used by an internal forensic team, external forensic entities and law enforcement to obtain this information.

4. Organizational Dimension

The organizational dimension of an SLA covers staffing, certifications and the interactions between the cloud provider and the consumer forensic team, external forensic entities and law enforcement.

4.1 Forensic Staffing Structure

The forensic staffing structure depends heavily on the specific cloud offering. In an IaaS model, the forensic team is likely to comprise staff from the cloud consumer. In a PaaS model, forensics is a shared responsibility, so the forensic team would have personnel from the cloud provider and cloud consumer. In a SaaS model, the team would likely comprise staff from the cloud provider. It is important that the agreement between the cloud provider and cloud consumer address the segregation of duties when providing forensic staffing according to the specific cloud offering.

Some forensic activities in cloud organizations are carried out regularly, such as proactive forensic data collection and log monitoring; others are performed after the incident, such as the analysis of forensic

data. Forensic teams from the provider side and consumer side would perform the majority of these forensic activities. The SLA should, therefore, address the segregation of duties between the cloud provider and consumer during incidents, responses and activities.

4.2 Forensic Training and Certification

Practitioners who conduct forensic investigations should be appropriately trained and certified, and should be prepared to testify about the relevance and implications of their actions. Cloud providers and consumers should ensure that their internal forensic staff have forensic training at least annually to ensure up-to-date knowledge and skills. Likewise, when cloud consumers contract out for forensic services, they should ensure that the third-party entities are competent and knowledgeable about traditional and cloud forensic techniques and their implications.

4.3 Organizational Interactions

There is often a chain of dependencies in a cloud provider that can complicate a cloud forensic investigation [10]. It is important that an SLA mandates the transparency of the chain of dependencies regarding key operations and addresses the reaction plan in terms of communication, collaboration and access to forensic data within the chain of dependencies. Collaboration with law enforcement when multiple cloud providers are involved should also be considered. Additionally, it is necessary to specify the situations where external forensic assistance would be necessary and the entities who would be involved in these situations.

5. Technical Dimension

The technical dimension relevant to an SLA deals with the technologies that are used to prepare for and facilitate forensic investigations. These include data collection, forensic tools tailored to cloud environments, incident response policy and time synchronization.

5.1 Proactive Forensic Preparation

Proactive measures can simplify cloud forensic investigations. Measures include designing forensically-aware cloud applications and proactively collecting and retaining forensic data in the cloud. In the survey [9], 83.58% of the respondents agreed that “a procedure and a set of toolkits to proactively collect forensically-relevant data in the cloud” is important or very important.

An SLA should address the segregation of duties and procedures for proactive forensic preparation and data collection in order to facilitate investigations initiated by the cloud provider, cloud consumer or law enforcement. The agreement should cover the mechanisms used for logging, authentication and auditing of access to all the identified data sources, as well as the replication and verification of data sources. Proactive data collection should include not only the logging of connections external to the provided services, but also inter-cloud communications and management requests. Furthermore, lawful data collection should be made available at every layer of the cloud infrastructure up to the end user. The mechanisms must ensure that an investigator has timely access to the collected data after an event has occurred and also allow for the validation of the collected data.

5.2 Forensic Data Collection

Cloud forensic data collection is the process of identifying and acquiring data in the cloud using procedures that allow the results to be presented in a court of law. The data includes client-side artifacts that reside on client premises and provider-side artifacts that reside in the provider infrastructure. As in traditional digital forensics, the collection process should follow procedures that preserve the integrity of data without breaching laws and regulations in the jurisdiction(s) where data is collected. Also, the process should not compromise the confidentiality of other tenants who share the resources [10].

Due to access restrictions, it is possible that SaaS providers may not provide access to the IP logs of clients who have accessed content. Likewise, IaaS providers may not provide access to forensic data such as virtual machine and disk images. In the cloud, consumers have limited access to relevant log files and metadata at all levels. Also, cloud consumers have a limited ability to conduct real-time monitoring on their own networks and to audit the network operations of their cloud providers.

To support data acquisition by an investigator, the cloud provider, with the help of the consumer, should have previously identified all forensically-relevant data sources. Furthermore, the cloud provider and the cloud consumer, possibly with the support of a trusted third-party, should have the ability to independently take a snapshot of the current state of all consumer-related data when a malicious event is suspected.

The SLA should address the conditions under which cloud data is considered to be “deleted.” Data deletion in the cloud is difficult and may sometimes be impossible. However, techniques such as crypto shred-

ding [6] can make the recovery of the data impractical. For regulatory compliance, some data may have to be destroyed or rendered completely inaccessible. Stipulations should exist to ensure that the data management infrastructure of the cloud provider is able to fully meet all data destruction requirements; alternatively, acceptable options such as crypto shredding must be employed. It is also important to consider data retention policies when addressing data deletion issues.

5.3 Hybrid Forensic Tools

An SLA should specify the forensic tools that will be made available by the cloud provider to the cloud consumer and law enforcement. Rapid elasticity is one of the essential characteristics of cloud computing [5]. Cloud computing and storage resources can be provisioned on demand. As a result, cloud forensic tools must have dynamic scalability characteristics.

In most investigations, large-scale static and live forensic tools would be required for e-discovery, data acquisition, data recovery, evidence examination and evidence analysis. Scalable forensic data collection tools should be made available to consumers for use in investigations. Stipulations regarding data copying during investigations and the downtime and liability associated with the use of forensic tools should be clarified in an SLA.

Event correlation tools should also be made available to investigate incidents that affect multiple tenants. The correlation of log entries helps determine the scope of incidents and identify security flaws.

5.4 Incident Response and Recovery

The cloud provider and cloud consumer should work transparently to develop and agree on an incident response plan, and specify all the roles in the preparation, discovery, response, investigation, recovery and follow-up phases [7]. Forensic and security staff should work together to mitigate the impact of incidents, gather evidence and update security policies and procedures based on the lessons learned.

Virtualization is a key technology that is used to implement cloud services. Many security issues arise in the operation of virtualization technologies, especially in multi-tenant environments [8]. However, procedures and tools for conducting comprehensive investigations in virtualized environments have yet to be developed. Emphasis should be placed on the collection and analysis of artifacts related to incidents in virtualized environments.

5.5 Time Synchronization

Time synchronization in a network is important for security and reliability as well as regulatory compliance. The benefits of network time synchronization include prevention of operational failure, avoidance of data loss, improved security and mitigation of legal exposure. Time synchronization is important for auditing and transaction logging, a necessity for forensically-aware applications. Also, it helps ensure data reliability in the event of a breach.

In the survey [9], 62.5% of the respondents agreed or strongly agreed that “synchronization of timestamps” is a challenge in cloud forensics. It is, therefore, important that cloud providers implement secure network time synchronization throughout their infrastructures and provide audit information related to time synchronization for predetermined retention periods according to their agreements with cloud consumers. For time synchronization outside the control of a cloud provider, the provider should be responsible for auditing and logging the time differences over the predetermined retention periods.

6. Legal Dimension

The legal dimension of an SLA should address issues of jurisdiction, multi-tenant concerns, chain of custody, event notification, auditing and regulation compliance.

6.1 Jurisdiction

In the survey [9], 90.14% of the respondents agreed or strongly agreed that “jurisdiction” is important or very important. Also, 87.87% of the respondents agreed that “a procedure and a set of toolkits to retrieve forensic data involving confidential data under jurisdiction(s) and agreement(s) under which services are operating” are important or very important. The cloud consumer should be aware that it could be difficult, perhaps even impossible, to conduct an investigation when the data does not reside in jurisdictions with proper regulations, especially if the jurisdictions were not stipulated in the service agreement.

A cloud consumer should have the ability to choose the jurisdictions in which his/her data resides. The cloud provider should clarify the differences in laws regarding privacy, data protection and the legal conditions that could contribute to data loss or downtime, such as the ability of law enforcement to seize hardware during an investigation. The cloud provider should also accurately track the jurisdictions in which a cloud consumer’s data resides during a given retention period. The tracking of

jurisdictions enables the cloud consumer and law enforcement to more easily assess the legality of a claim.

6.2 Multi-Tenant Data Issues

Multiple tenancy is an inescapable feature of community, public and hybrid clouds. In a SaaS model, tenants share middleware, operating systems and hardware. In a PaaS model, tenants share operating systems and hardware. In an IaaS model, tenants share hypervisors and hardware. An SLA must address the ability of the cloud provider to accurately and comprehensively filter forensic data sources that contain data belonging multiple tenants and release only the data related to the specific tenant.

6.3 Data Ownership

Data ownership refers to the possession of data and the responsibility for the data. Ownership implies power as well as control. Control includes the ability to access, create, modify, package, derive benefit from, sell and remove data, as well as the right to assign these privileges to others. An SLA should address issues regarding data ownership in a cloud environment so that it is clear who owns the data that is being investigated.

6.4 Chain of Custody

In the survey [9], 77.6% of the respondents agreed that “a procedure and a set of toolkits to record and maintain the chain of custody in an investigation” is important or very important. The agreement should specify that the right to privacy exists so that only authorized parties can reliably access a consumer’s data, while others, including the cloud provider, cannot. Along with service-provider-independent authentication, auditing of access must be provided and agreed upon. External entities and their roles in the transmission and storage of consumer data should also be specified in an SLA.

In the cloud, much of the data that is of interest in an investigation is obtained via live forensic techniques because the system is critical and cannot be taken down or because the virtual system is not persistent. As a consequence, verifying live data using traditional *post mortem* methods such as hashing is impossible.

The investigator should have a clear understanding of the forensic techniques and tools used in an investigation and how they affect the system. The Association of Chief Police Officers (ACPO) guidelines specify that “[i]n circumstances where a person finds it necessary to

access original data held on a computer or on storage media, that person must be competent to do so and be able to give evidence explaining the relevance and the implications of their actions” [2]. Therefore, the terms of an SLA should stipulate that all data will be extracted and verified by competent, trained forensic practitioners using approved techniques and tools.

6.5 Event Notification

An SLA between a cloud provider and consumer should address the manner and the precise timeframe within which one party should notify the other party about critical incidents and law enforcement requests. All reaction plans should be made based on the notification protocol specified in the SLA and should include back-up solutions in the event that resources are seized.

6.6 Change of Cloud Provider

The cloud consumer has the right to change cloud providers. The SLA should clarify the responsibility of the cloud provider to retain consumer data for investigative purposes if and when the consumer chooses to migrate to another cloud provider.

6.7 Auditing

The cloud provider and consumer should agree on the auditing of the key terms in the SLA. Auditing should be performed by the client, by the provider and by third parties on a regular basis, and should include on-site inspections. Third parties should be approved by the cloud provider and consumer, along with the conditions under which an audit will occur. Also, the cloud provider and consumer should agree on the information that will be provided and who provides the information during an audit.

6.8 Regulatory Compliance

Regulatory compliance covers many areas such as the location of data, validity of timestamp information and privacy. The cloud consumer should determine the regulations that apply and specify the appropriate data transportation, storage, retention and management protocols. The cloud provider also has to adhere to regulations that are imposed on providers. The cloud provider must provide a satisfactory level of regulatory compliance and demonstrate its compliance in audits by a third party. The audit results should be made available to the cloud consumer, who should be able to clearly discern areas of non-compliance.

7. Conclusions

As the potential for criminal incidents involving cloud resources increases, it is imperative that cloud providers and consumers work together to create an environment that supports forensic investigations of the highest quality. This paper sets the stage for this process by specifying standard terms for SLAs that support cloud forensics. The terms cover the organizational, technical and legal dimensions of cloud forensics, and, as such, could help standardize and regulate the emerging area of cloud forensics. We hope that the concepts and terms described in this paper will be analyzed, refined and augmented by the various stakeholders to help create a strong foundation for cloud forensics.

References

- [1] Amazon, Amazon Web Services: Overview of Security Processes, Seattle, Washington (aws.amazon.com/articles/1697), 2008.
- [2] Association of Chief Police Officers, Good Practice Guide for Computer-Based Electronic Evidence, London, United Kingdom (www.7safe.com/electronic_evidence/ACPO_guidelines_computer_evidence.pdf), 2008.
- [3] N. Beebe, Digital forensic research: The good, the bad and the unaddressed, in *Advances in Digital Forensics V*, G. Peterson and S. Shenoi (Eds.), Springer, Heidelberg, Germany, pp. 17–36, 2009.
- [4] B. Gammage, D. Plummer, R. Valdes, K. McGee, K. Potter, S. Tan, D. Aron, R. Hunter, J. Heiser, B. Prentice, G. Alvarez, M. Basso, L. Fiering and K. Dulaney, Gartner's Top Predictions for IT Organizations and Users, 2011 and Beyond: IT's Growing Transparency, Document ID Number G00208367, Gartner, Stamford, Connecticut, 2010.
- [5] P. Mell and T. Grance, The NIST Definition of Cloud Computing, Recommendations of the National Institute of Standards and Technology, NIST Special Publication 800-145, National Institute of Standards and Technology, Gaithersburg, Maryland, 2011.
- [6] R. Mogull, Cloud Data Security: Archive and Delete (Rough Cut), Securosis, Phoenix, Arizona (securosis.com/blog/cloud-data-security-archive-and-delete-rough-cut), 2011.
- [7] T. Osborne, Building an Incident Response Program to Suit Your Business, InfoSec Reading Room, SANS Institute, Bethesda, Maryland (www.sans.org/reading_room/whitepapers/incident/building-incident-response-program-suit-business_627), 2001.

- [8] T. Ristenpart, E. Tromer, H. Shacham and S. Savage, Hey you, get off of my cloud: Exploring information leakage in third-party compute clouds, *Proceedings of the Sixteenth ACM Conference on Computer and Communications Security*, pp. 199–212, 2009.
- [9] K. Ruan, I. Baggili, J. Carthy and T. Kechadi, Survey on cloud forensics and critical criteria for cloud forensic capability: A preliminary analysis, *Proceedings of the Sixth Annual Conference on Digital Forensics, Security and Law*, 2011.
- [10] K. Ruan, J. Carthy, M. Kechadi and M. Crosbie, Cloud forensics, in *Advances in Digital Forensics VII*, G. Peterson and S. Sheno (Eds.), Springer, Heidelberg, Germany, pp. 35–46, 2011.