# Can Forensic Audit Help Address the Cloud Forensic Problem in Light of the Requirements of the Forthcoming European Union General Data Protection Regulation?

Bob Duncan*, Mark Whittington†
Business School
University of Aberdeen
Aberdeen, UK
Emails: *robert.duncan@abdn.ac.uk, †mark.whittington@abdn.ac.uk

***Abstract*—There is no doubt that the forthcoming European Union (EU) General Data Protection Regulation (GDPR), which comes into effect on 25th May 2018, will certainly concentrate many corporate minds. As for those who rely on cloud computing, there is likely to be even more consternation in the ranks, due to the issues surrounding dealing with the Cloud Forensic Problem. While it is the case that all computing systems are constantly under serious attack, this particular problem arises due to the fact that once an attacker gains a foothold in a cloud system and becomes an intruder, there is very little to prevent the intruder from gaining sufficient privileges to then completely delete all trace of their incursion, possibly deleting far more records than they need to in the process. Additionally, there is nothing to prevent them from then helping themselves to any amount of data covered by the GDPR, either by viewing it, modifying it, deleting it or ex-filtrating it from the victim system. This, then, will present a compliance nightmare to a great many cloud users, many of whom are poorly prepared to cope with this serious practical and financial challenge. In this paper, we consider how the use of robust forensic audit techniques from the accounting world might be applied to mitigate this serious challenge for such companies.**

***Keywords–Forensic audit; GDPR compliance; cloud forensic problem.***

## I. INTRODUCTION

Achieving information security with conventional distributed network computer systems presents a significant challenge, but this challenge increases exponentially when we introduce cloud computing to the mix, due to the multiplicity and complexity of hardware and software layers and the number of actors with differing agendas, involved in any cloud ecosystem. The principal reason for the difficulty of this challenge is the so called "Cloud Forensic Problem".

This arises once an attacker gains a foothold in a cloud system and becomes an intruder. Once this happens, there is little to prevent the intruder from helping themselves to any amount of data, either by viewing, modifying, deleting or exfiltrating it from the victim system. Worse still, there is nothing to prevent the intruder from gaining sufficient privileges to completely delete all trace of their attack.

The forthcoming EU General Data Protection Regulation (GDPR) [1] comes into effect on 25th May 2018, and a principal requirement is the protection of personally identifiable information held by any organisation, anywhere in the world, on pain of severe financial penalties. Given that the cloud forensic problem presents a potentially insurmountable compliance problem, a great many organisations are likely to be exposed to incalculable potential penalties for the inevitable string of cyber breaches that are likely to ensue.

We start in Section II, by considering the cloud forensic problem and the challenges it poses. We turn to the accounting world to see which techniques we could implement to help address these serious challenges in Section III, where we look at accounting, audit and forensic accounting to see how it works for the accounting world, and in Section IV, we consider how we might develop some of these well established techniques to help us address this significant cloud security problem. In Section VI, we look at how we might use the immutable database as the core of this approach. In Section VIII, we discuss our conclusion and future work.

## II. THE CLOUD FORENSIC PROBLEM

Cloud systems are extremely popular with companies due to the flexibility offered by cloud. Speed of start-up, ease of scalability to match the demand curve and the revenue nature of the costs involved all provide a strong incentive for companies to use cloud services. Cloud computing has been with us now for over 10 years, and while much of the early research concentrated on usability [2] [3] and performance [4]–[6], it was not long before thoughts of security [7]–[9] and privacy [10] [11] started to surface.

While the US National Institute for Standards and Technology (NIST) were one of the first organisations to propose standard definitions [12] [13], interest in security [14]–[17] and privacy [18]–[20] started to grow.

Thoughts also started turning to accountability [8] [21]–[23], given the evolving complexities of cloud ecosystems.

While there have been some really positive advances in both security and privacy during this time, there remains one fundamental weakness that has not been resolved, namely the "cloud forensic problem". All computer systems are subject to continuous and serious attack, and cloud systems are no exception. It would be realistic to state that no system is

 ISBN: 978-1-61208-607-1 

immune to attack, and this is particularly true for cloud systems.

The main focus of an attacker is to breach a system, which can involve a considerable amount of work on their part. The more diligent will first perform surveillance, compile many analyses of how the various company systems are structured and how they interact with each other. Often, they will also carry out huge amounts of work to understand the people of the organization, since they are usually the weak link in the chain [24]. This extensive intelligence gathering will usually cover every conceivable aspect of all company systems to ensure they discover everything they need to know about the company. This why it is so important to analyse system logs, in order to gain a better understanding of who is actually attacking their systems.

Other attackers, will be much less organised, simply trying to hack in to company systems, without a thought of the overview of the company concerned. They will merely look for known vulnerabilities and try to attack them. There are other attackers who will specifically attack the people of the company through social engineering and other similar approaches. The first objective of all attackers is the same — to penetrate the system in order to become an intruder.

The aim is not just to get in, and out, as quickly as possible, but to develop a long term foothold, secrete themselves into corporate servers and other systems which will allow them to return to help themselves whenever they want. The longer they remain in the system, the more they are likely to try to escalate privileges to give them access to more and more possible information. All too often, they are helped along the way by the companies themselves, often through an element of laziness on the part of system administrators [25].

If we look back five years ago, at previous cyber breach reports [26], there was a global average time of 6 months between breach and discovery. With more rigorous attention paid to reading and analysing their server logs, it is obvious they could have discovered intruders much more quickly. By 2016, the time between breach and discovery had dropped to a matter of weeks rather than months [27], however, this is still not good enough to keep on top of what is going on in corporate systems.

Companies often contribute to their own downfall by failing to update security patches to both operating systems and software systems, complexities from legacy applications applications and risks of outages being reasons or excuses for slow implementation [28]. All of these issues conspire to lead inexorably to the, as yet unresolved, cloud forensic problem — namely, that once an intruder is in the system, and has escalated sufficient privileges, there is nothing to prevent them from deleting the forensic trail, all system logs and audit trails, thus hiding all evidence of their successful penetration and of the size and nature of their crime.

Under the forthcoming GDPR [1], any breached organisation must report the breach within 72 hours of discovery of the breach. They must also report how many relevant records have been compromised, whether by having been read, amended, deleted or ex-filtrated from the system. Given that many system logs are also not turned on by default [41], this means identifying which records have been compromised, whether by having been read, amended, deleted or ex-filtrated, will present a serious enough challenge in the first place, but since the intruder will likely have deleted all forensic trails in the system, the likelihood of an organisation being able to properly identify which records have been compromised may be impossible to determine.

This means not only non-compliance with the GDPR, triggering fines, but failure to tackle some elementary steps will then cause these fines to escalate following repeated events to the greater of €20million or 4% of global turnover. The size of the potential fines, along with the bad publicity will surely get the attention of organizations.

## III. Useful Techniques from the Accounting World

The process of accounting has been around for millennia, with the underlying standard approach of double entry bookkeeping in use for over 500 years, with the generally accepted story placing its creation in Italy. Accounting is primarily seen as a technique for collecting, measuring, processing and communicating financial information about the economic performance of entities, in order to provide decision useful information for interested parties, such as management, investors, creditors and regulators [30]. The International Accounting Standards Board (IASB) issued a similar, but more user-constrained definition in 2015, namely "The objective of general purpose financial reporting is to provide financial information about the reporting entity that is useful to existing and potential investors, lenders and other creditors in making decisions about providing resources to the entity. Those decisions involve buying, selling or holding equity and debt instruments, and providing or settling loans and other forms of credit." [31]

Auditing, too, has been around for millennia, as there has always been a need to provide assurance that accounts and financial statements present a "true and fair view" of the business under review. Naturally, many accounting and auditing techniques can also be applied to anything else that is measurable, and in this case, of particular interest to us is data. Hence, seeking to apply the more evolved and time tested techniques from accounting and auditing techniques to the management and governance of data in the cloud would seem logical.

A further extension of the processes of accounting and audit is forensic (OED [32] "pertaining to, connected with, or used in courts of law; suitable or analogous to pleadings in court") accounting, which as the definition suggests is the process of preparing evidence suitable for use in a court of law.

We can use these techniques, which have long been developed in the accounting world to good effect in helping us secure our cloud data. We can then liken the database system to an accounting system, whereby we collect, measure, process and communicate non-accounting information concerning a business to the people for whom it is intended or relevant.

In principal, we can then use cloud audit to provide assurance of the data provenance of all the data held in the database system, and in the event of a security breach, we can easily apply cloud forensic techniques, learning from the accounting world, in order to help us bring about a successful prosecution in the courts and be aware of the steps needed to

 ISBN: 978-1-61208-607-1 

improve security for the future. In practice, this will, of course, be far harder to achieve.

Of course, it is worth pointing out that for centuries, accountants have enjoyed the benefits of working with hard copy books, written with quill pen and ink. This medium presents the benefit of providing a hard ink trail to follow, something which we shall later see is no longer available with modern cloud systems.

## IV. Forensic Cloud Audit

An interesting distinction in definition between "forensic accounting" and "cloud computing forensic science" is the presence of that last word science. Hopwood et al., [33] give the following definition for forensic accounting:

Forensic accounting is the application of investigative and analytic skills for the purpose of resolving financial issues in a manner that meets standards required by courts of law. Notice that forensic accounting is not limited to the use of financial investigations that result in legal prosecution; however, if this is the purpose, the investigation and analysis must meet the standards required in the court of law that has jurisdiction. (page 3)

Whilst NIST [34] provides the following discussion and definition:

Many experts consider forensic science to be the application of a broad spectrum of sciences and technologies to the investigation and establishment of facts of interest in relation to criminal, civil law, or regulatory issues. However, the resulting techniques may also be used for purposes outside the scope of law to reconstruct an event that has occurred. Cloud computing forensic science is the application of scientific principles, technological practices and derived and proven methods to reconstruct past cloud computing events through identification, collection, preservation, examination, interpretation and reporting of digital evidence.

Note that the forensic accounting definition does not include the word science, despite the area (see for example two textbooks Taylor [35] and Hopwood et al., [33]) including scientific methods. Taylor [35], as a more introductory text, focuses initially and at some length on the need to understand background and environmental issues, using this as a backdrop before moving on to, again, a largely discursive review of the wide range of relevant criminal activities that might require the attention of the forensic accountant. He also addresses risk management issues in relation to IT systems, briefly including the cloud, and the process of investigation. Hopwood et al., [33] have a similar structure but give a little greater weight to forensic science and computer forensics.

From the computer science camp, Choo and Dehghantanha [36], a more scholarly work, reflects a greater weight placed on technical issues, the tools and techniques needed, for forensic cloud investigations. Almulla et al., [37] review the cloud forensic literature and find some discursive, though many technical papers.

Issues requiring computer forensic audit are likely to involve the stealing of money, the stealing of monetizable data or the misrepresentation of data to personal or group advantage. These are areas which accountants have strived to address over decades in less technical and complex settings. It would seem logical that their group learning over time would have some relevance and currency to the new cloud situation.

Like most professions, accountants have well organised professional exams. The Association of Chartered Certified Accountants (ACCA), an international professional body with over 200,000 members [38], has an exam at its professional stage, Advanced Auditing and Assurance [39], that includes a section on forensic audit though it should be noted that it is only a small part. It would seem that qualified accountants are ill-prepared for the complexities of the cloud environment, both in terms of understanding the environmental issues, though there is accessible material for them to pick up some of this (see Taylor [35] and Hopwood et al., [33]), as well as comprehending the technical ones, which would be a far more complex and difficult step. Whilst there are a few small organisations focusing on forensic accounting and audit, these appear peripheral (for example..), it does not seem that many qualified accountants have moved into this more rarefied space by adding years of further learning to their accounting badge.

From the other direction, computer specialists clearly have an understanding of the technology and some understanding of the softer environmental, legal and behavioural issues (see Choo and Dehghantanha [36]) though little if any accounting awareness.

So, it would seem, that apart from a few exceptional, motivated, highly skilled individuals there is not yet a significant body that balances the three areas in the venn diagram below. The diagram is, of course, highly simplistic intending to just give a broad view of the difficulties in bringing the wide range of knowledge and experience required for forensic cloud investigation.

Whilst there are many audit tools, the computing literature already uses the "audit trail" [37] when discussing evidence integrity, however in previous work [40]–[43], we have questioned the level of development of these audit trails and whether all the lessons from the rich accounting history in this area have been taken on board. One stark difference between the accounting approach and the computing one is that of redundancy. To the accountant, there is an expectation of keeping more rather than less, with computer scientists having a focus on efficiency and minimising costs. Another is some level of agreement on what should be in an audit trail. For example, Bernstein [44] sees the trail including: events, logs, and the analysis of these, whilst Chaula [45] gives a longer, more detailed list: raw data, analysis notes, preliminary development and analysis information, processes notes, and so on. Pearson et al. [9], as far back as 2010, accept that attaining consistent, meaningful cloud audit trails is a goal rather than reality. More worryingly, Ko et al. [21] point out that it is possible to delete the audit trail along with a cloud instance, meaning there is no record then remaining. Ko [46] details the requirements for accountability.

## V. The Special Skills Mix Needed for Cloud Forensic Audit

As we mentioned earlier, with modern cloud systems, we no longer are able to enjoy the benefits of the permanent ink trail. While reasonable alternatives can be available with conventional distributed network systems, this is not the case for cloud systems. We discussed the Cloud Forensic Problem earlier, and it is this security weakness inherent in cloud

 ISBN: 978-1-61208-607-1 

systems that makes this job significantly harder to accomplish effectively.

When considering cloud forensic issues, it is now clear that we can no longer afford to rely on conventional discipline boundaries when trying to address these issues, as it is now likely that all the disciplines affected are likely to suffer from a potentially significant knowledge gap. Clearly, the cloud environment is considerably different from conventional distributed network models under the sole control of a company. There are also a great many actors involved in such an environment, each potentially with its own agenda. Legal and regulatory issues are a lot less clearly defined for cloud environments, with the increased likelihood of multiple companies and jurisdictions.

We also have to contend with a number of uninvited actors — namely, the attackers and the intruders, with the latter presenting the greater challenge.

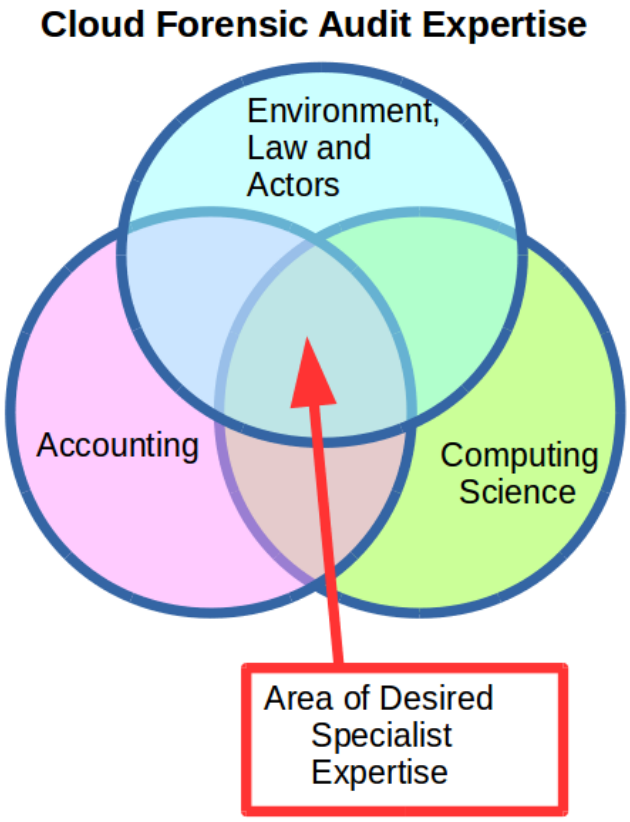


**Figure 1.** The Area of Desired Expertise

This means we can no longer consider addressing cloud forensic audit from an insular perspective, since accountants, computer scientists and legal, regulatory and other actors within the cloud environment will all suffer from incomplete knowledge, which rather suggests there are likely to be weaknesses in their various individual approaches. Equally, in the absence of the solid 'ink trail', this increases the complexity of the task exponentially. In Figure fig:venn, we show the overlapping Area of Desired Expertise that is needed for all three disciplines to fully understand where this knowledge gap needs to be addressed.

Currently, when it comes to Cloud, intruders have it all their own way. Once they are in the system, it is merely a matter of time until they have sufficient privileges to delete the forensic trail, thus allowing them to bed down for the long run. Since they are likely to be deleting all audit and forensic trails as they go, this also means an increase in difficulty, verging on the impossible, for data controllers to safely keep the organisation fully compliant with all regulatory and legislative requirements they must adhere to in order to achieve compliance, security and privacy.

There are, therefore, two major goals that must be dealt with. First, we need to restore the permanent 'ink trail' so that we have something to follow, and this is where the immutable audit trail process comes in. Second, we have to fill the various knowledge gaps to ensure that all parties involved in Cloud Forensic Audit are fully up to speed. This will come down to a combination of collaboration and proper training. This latter is outside the scope of this paper, but the first is very much a part of it, and we discuss this further in the next section.

## VI. The Immutable Database

We can see that compliance with the GDPR is not a readily achievable goal that can be easily met by any organisation using Cloud services, due to the difficulties associated with the Cloud Forensic Problem. Thus, we must ensure we create and maintain both a secure forensic and audit trail in order to have any chance of making this happen.

We need to consider very carefully exactly what we need to log to ensure we can achieve compliance with the GDPR. This means we need to monitor our Cloud instances, we need to monitor who is accessing our systems and we need to monitor what is happening with our database systems.

We start by considering our Cloud instances. As Duncan and Whittington have shown in [41] [47] [42], a working solution can be found using an immutable database at its core to record all the relevant information we would require. This means we must first consider carefully exactly what that information should be.

We would want to log all significant events as they transpire during the life cycle of each Cloud instance, with the first significant event being the creation of the Cloud instance, and the last being the shutting down of that instance. Under normal circumstances, these, and all other lifetime events, would be logged on the instance itself. This, as we know from Ko et al. [21], is a dangerous thing to do; thus our first step will be to ensure this data is logged additionally onto an external secure immutable database to ensure it achieves full persistence.

This external database must run on a dedicated secure server, protected by an Intrusion Detection System (IDS), and the database must be immutable, i.e., append only. This secure server will also use dedicated software agents to police the activities being logged, so that the occurrence of any significant event (such as the shutting down of a Cloud instance) will be instantly identified and reported for approval/further investigation.

Turning to the question of who is using our systems, we want to understand who is logging in to our systems, where they come from and what they do once they have successfully logged in. Thus we must capture the relevant detail from the access logs. The detail of how this may be carried out

 ISBN: 978-1-61208-607-1 

will depend on the systems architecture deployed, the type of access control credentials used and means of controlling access to the various systems available to specific users. A multi-factor authentication approach is always better than access by password. Proper logging of each step in the process is also always preferable.

Once a user gains access to any system, we still want to know where the user came from, and we certainly want to know what the user is doing with the system after they gain access. Thus we should be logging all the steps that the user takes, regardless of whether access is via physical presence or via remote login. In other words, we need to log every single query made or instruction given to the system. We might wish to consider whether we want to record what the result of that query would be, since this might generate inordinate amounts of data in the case of a database query. Whatever we decide is required, we must ensure a separate copy of the queries recorded are stored into our dedicated secure immutable database. It is clear that redundancy can be a good thing.

## VII. Discussion

Having devloped a workable solution to this problem, we may well have some questions, such as:

- How easy is it to implement?
- How quickly and how well will interested parties adhere to such a solution?
- In the event of a breach, who will be responsible and what might the consequences be?

The answer to the first question is that we take the view that this approach needs to be simple to implement and simple to maintain. It is as simple as switching on the necessary forensic and audit trail logging, then writing a chron job to forward the resulting logs to the immutable database. Wherever possible, such programmes should be set to immutable to make it difficult for attackers and intruders to delete them. A regular check on the configuration files would also be a useful thing to do.

For the second question, it is likely that the easier something is to implement, the more likelihood that it will be implemented. It is not challenging to implement, nor to maintain, and the cnsequences of failing to do so could have a huge adverse impact, so there is a considerable incentive to both implement and maintain this approach.

As to the third question, it is not a question of 'in the event of a breach', but rather a case of accepting there will be breaches, and these are likely to be a continuous feature. As soon as a breach occurs, a forensic trail will be generated and stored both within the Cloud instance , as well as in the off-site immutable database. Under normal circumstances, the attacker will now attempt to dig deep, escalate privileges and delete the forensic trail. The longer the intruder remains inside the system, the more likelihood that a successful deletion of the audit trail will take place. However, with a covert copy of the forensic and audit trail data available, this will allow some potentially fruitful investigative work to take place.

In the event that an attack against the Cloud instance is successful, where will liability sit? The GDPR regulation is quite clear. In the event of a breach, the Data Controller has a legal obligation to notify the Supervisory Authority within 72 hours of becoming aware of a breach. Individuals must also be notified in the event that encryption is not used. Clearly the use of encryption would be a prudent approach to minimise the impact of the breach, as well as the amount of any possible fine.

## VIII. Conclusion and Future Work

We have seen that compliance with the EU GDPR for all Cloud users is likely to present a significant challenge. Without special measures being taken, it is likely that compliance will prove impossible to achieve. This is likely to expose such Cloud users to the full force of the penalties of this regulation, which are significant.

It is clear that a minimal requirement will be to generate both a secure forensic trail and an audit trail, in order to have the basic requirements to be able to consider fulfilling the regulatory requirements in the event of a breach. Without this in place, it is likely to be impossible to comply with the legislation, thus rendering the organisation liable to some serious penalties.

In this paper, we have identified the particular issues that companies who are Cloud users and are liable to be GDPR compliant must be able to deal with. There is no point in relying on Cloud service providers to take care of this matter. The company data controller is accountable to the regulator for ensuring the company is compliant, and without both a forensic trail and a full audit trail for the PII held on behalf of EU residents, then compliance will not be possible. This will lead to potentially massive fines being applied — a situation that is potentially avoidable.

We are in the process of building a miniature real life Cloud system on which to test our ideas. The server will run a full Cloud management system, which will be used to run a number of independent Cloud instances, all of which will run web servers with database back ends to replicate the approach of many Cloud users. This will be subject to rigorous attack, with the view to discover whether the immutable database approach can allow Cloud users to be GDPR compliant.

We have a range of permutations to test, and we seek to find the optimum solution providing the right balance between usability, performance, cost and ease of dealing with breaches. We shall be reporting our results later this year, and we will be working towards delivering a workable solution to keep Cloud users compliant.

## References

[1] EU, "EU General Data Protection Regulation (GDPR)," 2017. [Online]. Available: http://www.eugdpr.org/ [Retrieved: December 2017]

[2] T. Takahashi, Y. Kadobayashi, and H. Fujiwara, "Ontological Approach toward Cybersecurity in Cloud Computing Categories and Subject Descriptors," in Science And Technology, 2010, pp. 100–109.

[3] L. M. Vaquero, L. Rodero-Merino, J. Caceres, and M. Lindner, "A Break in the Clouds: Towards a Cloud Definition," ACM SIGCOMM Computer Communication Review, vol. 39, no. 1, 2008, p. 50.

[4] M. Alhamad, T. Dillon, C. Wu, and E. Chang, "Response Time for Cloud Computing Providers," Response, 2010, pp. 8–10.

[5] D. Durkee, "Why Cloud Computing Will Never Be Free," Communications of the ACM, vol. 53, no. 5, may 2010, p. 62.

[6] S. Fraser et al., "Cloud Computing Beyond Objects: Seeding the Cloud," Communications, 2009, pp. 847–850.

 ISBN: 978-1-61208-607-1 

[7] A. Haeberlen, "A Case for the Accountable Cloud," ACM SIGOPS Operating Systems Review, vol. 44, no. 2, 2010, p. 52.

[8] S. Pearson, "Towards Accountability in the Cloud," IEEE Internet Computing, vol. 15, no. 4, jul 2011, pp. 64–69.

[9] S. Pearson and A. Benameur, "Privacy, Security and Trust Issues Arising from Cloud Computing," in 2010 IEEE Second International Conference on Cloud Computing Technology and Science, no. December. Ieee, nov 2010, pp. 693–702.

[10] R. Buyya, C. S. Yeo, S. Venugopal, J. Broberg, and I. Brandic, "Cloud computing and emerging IT platforms: Vision, hype, and reality for delivering computing as the 5th utility," Future Generation Computer Systems, vol. 25, no. 6, 2009, p. 17.

[11] L. M. Kaufman, "Data security in the world of cloud computing," IEEE Security and Privacy, vol. 7, no. 4, jul 2009, pp. 61–64.

[12] P. Mell and T. Grance, "Effectively and Securely Using the Cloud Computing Paradigm," NIST, Information Technology Laboratory, vol 2, no 8, pp 304-311.

[13] P. Mell and T. Grance, " The NIST definition of cloud computing.,"Tech. Rep., 2011. [Online]. https://doi.org/10.6028/NIST.SP.800-145 [Retrieved: December 2017]

[14] S. Bradshaw, C. Millard, and I. Walden, "Contracts for Clouds: Comparison and Analysis of the Terms and Conditions of Cloud Computing Services," International Journal of Law and Information Technology, vol. 19, no. 3, 2011, pp. 187–223.

[15] K. Hon, C. Millard, and I. Walden, "The Problem of 'Personal Data' in Cloud Computing - What Information is Regulated ?" 2011. [Online]. Available: http://ssrn.com/abstract=1562461 [Retrieved: December 2017]

[16] M. Iansiti and G. L. Richards, "Economic Impact of Cloud Computing," Economics of Innovation and New Technology, vol. 7, no. 2000, 2010, pp. 1–42.

[17] N. Papanikolaou, S. Pearson, M. C. Mont, and R. K. L. Ko, "Towards Greater Accountability in Cloud Computing through Natural-Language Analysis and Automated Policy Enforcement," Engineering, 2011, pp. 1–4.

[18] Data Protection Working Party, "Opinion 05/2012 on Cloud Computing," 2012. [Online]. Available: http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2012/wp196_en.pdf [Retrieved: December 2017]

[19] W. Jansen and T. Grance, "Guidelines on Security and Privacy in Public Cloud Computing," NIST, Tech. Rep. 7, 2011. [Online]. Available: http://csrc.nist.gov/publications/nistpubs/800-144/SP800-144.pdf [Retrieved: December 2017]

[20] N. Papanikolaou, S. Pearson, and M. C. Mont, "Towards Natural-Language Understanding and Automated Enforcement of Privacy Rules and Regulations in the Cloud: Survey and Bibliography," Analysis, 2011, pp. 1–9.

[21] R. K. L. Ko et al, "TrustCloud: A framework for accountability and trust in cloud computing," Proceedings - 2011 IEEE World Congress on Services, SERVICES 2011, 2011, pp. 584–588.

[22] R. K. L. Ko, B. S. Lee, and S. Pearson, "Towards achieving accountability, auditability and trust in cloud computing," Communications in Computer and Information Science, vol. 193 CCIS, no. PART 4, 2011, pp. 432–444.

[23] S. Pearson and A. Charlesworth, "Accountability as a way forward for privacy protection in the cloud," Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics), vol. 5931 LNCS, no. December, 2009, pp. 131–144.

[24] M. Hammock, "A Review of the Economics of Information Security Literature," pp. 1–38, 2010. [Online]. Available: http://ssrn.com/abstract=1625853 [Retrieved: December 2017]

[25] A. M. Froomkin, "Government Data Breaches," Berkeley Technology Law Journal, 2009, pp. 1–42.

[26] Verizon, "2012 Data Breach Investigations Report," Verizon, Tech. Rep., 2012.

[27] Verizon, "2016 Verizon Data Breach Report," Tech. Rep., 2016.

[28] D. Kossmann, T. Kraska, and S. Loesing, "An evaluation of alternative architectures for transaction processing in the cloud," in Proceedings of the 2010 International Conference on Management of Data. Indianapolis, Indiana: ACM Press, 2010, pp. 579–590.

[29] B. Duncan and M. Whittington, "Enhancing Cloud Security and Privacy: The Cloud Audit Problem," in Submitted to Cloud Computing 2016, Rome, 2016, pp. 1–6.

[30] A. A. A. C. to Prepare a Statement of Basic Accounting Theory, A statement of basic accounting theory. American Accounting Association, 1966.

[31] IASB, "IASB ED/2015/3 - Exposure Draft Conceptual Framework for Financial Reporting Comments," IASB, Tech. Rep., 2015.

[32] "forensic, adj. and n." OED Online. Oxford University Press, December 2017. [Retrieved: December 2017.]

[33] W. S. Hopwood, J. J. Leiner, and D. G. R. Young, Forensic accounting and fraud examination. McGraw-Hill, 2012.

[34] NIST, "NIST Cloud Computing Forensic Science Challenges," 2014, p. 51.

[35] J. Taylor, Forensic accounting. Financial Times Prentice Hall, 2011.

[36] K.-K. Choo and A. Dehghantanha, "Contemporary Digital Forensics Investigations of Cloud and Mobile Applications," in Contemporary Digital Forensic Investigations of Cloud and Mobile Applications. Elsevier, 2017, pp. 1–6.

[37] S. A. Almulla, Y. Iraqi, and A. Jones, "A State-of-the-Art Review of Cloud Forensics," Journal of Digital Forensics, Security and Law, vol. 9, no. 4, 2014, pp. 7–28.

[38] ACCA, "ACCA celebrates hitting 200,000 members worldwide with a global tour to honour each and every one," London, 2018.

[39] ACCA, "Advanced Audit and Assurance: Syllabus and Study Guide September 2018 to September 2019," 2017.

[40] B. Duncan and M. Whittington, "Compliance with standards, assurance and audit: does this equal security?" in Proceedings of the 7th International Conference on Security of Information and Networks - SIN '14. Glasgow: ACM, 2014, pp. 77–84.

[41] B. Duncan and M. Whittington, "Enhancing Cloud Security and Privacy: The Power and the Weakness of the Audit Trail," in Cloud Computing 2016: The Seventh International Conference on Cloud Computing, GRIDs, and Virtualization, no. April. Rome: IEEE, 2016, pp. 125–130.

[42] B. Duncan and M. Whittington, "Creating an Immutable Database for Secure Cloud Audit Trail and System Logging," in Cloud Computing 2017: The Eighth International Conference on Cloud Computing, GRIDs, and Virtualization. Athens, Greece: IARIA, ISBN: 978-1-61208-529-6, 2017, pp. 54–59.

[43] B. Duncan and M. Whittington, "Creating and Configuring an Immutable Database for Secure Cloud Audit Trail and System Logging," International Journal On Advances in Security, vol. 10, no. 3 & 4, 2017, pp. 155–166. [Online]. Available: ISSN: 1942-2636, http://www.iariajournals.org/security/index.html [Retrieved: December 2017]

[44] D. Bernstein, E. Ludvigson, K. Sankar, S. Diamond, and M. Morrow, "Blueprint for the intercloud - Protocols and formats for cloud computing interoperability," in Proceedings of the 2009 4th International Conference on Internet and Web Applications and Services, ICIW 2009, 2009, pp. 328–336.

[45] J. A. Chaula, "A Socio-Technical Analysis of Information Systems Security Assurance: A Case Study for Effective Assurance," Ph.D. dissertation, 2006. [Online]. Available: http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:A+Socio-Technical+Analysis+of+Information+Systems+Security+Assurance+A+Case+Study+for+Effective+Assurance#1 [Retrieved: December 2017]

[46] R. K. L. Ko, "Data Accountability in Cloud Systems," in Security, Privacy and Trust in Cloud Systems. Springer, 2014, pp. 211–238.

[47] B. Duncan and M. Whittington, "Cloud cyber-security: Empowering the audit trail," International Journal on Advances in Security, vol. 9, no. 3 & 4, 2016, pp. 169–183. [Online]. Available: https://www.iariajournals.org/security/tocv9n34.html [Retrieved: December 2017]