The Complexities of Auditing and Securing Systems in the Cloud — is there a Solution and will the GDPR move it up the Corporate Agenda?

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

Abstract—**It would seem that some companies have been slow or unable to secure their cloud activities or to be aware of breaches in a timely manner. The European Union (EU)s General Data Protection Regulation (GDPR) has been introduced with the intent of sufficient threat of meaningful fines that directors will now take cloud security seriously, even if they had not perceived it as a strategic priority before. However, just introducing such penal incentives does not mean that solutions are easy to implement. Whilst the perfect solution would always include stopping attackers from becoming intruders, once the attacker gets access the challenge is not just the immediate fiscal damage to the company or its trading partners, but also to the very records and integrity of the databases themselves. Once the intruder gains a foothold, they may then be able to grant themselves sufficient privileges to completely delete all trace of their incursion, possibly deleting far more records than they need to. They may remain undetected within the system, accessing, modifying, deleting or ex-filtrating data at will from the victim's system. This is referred to as the Cloud Forensic Problem. This, then, presents 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 experience and traditional techniques from the accounting world might be applied and adapted to mitigate this serious challenge.**

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. While this high level of complexity has been a fundamental part of cloud computing, we shall see that the capabilities of cloud computing have evolved considerably beyond what was first envisaged. The principal reason for the difficulty of this challenge is the so called "Cloud Forensic Problem".

The Cloud Forensic Problem arises when 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 ex-filtrating 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 through modifying or deleting entirely the forensic records of the system. In this paper, we consider how the use of forensic audit might help mitigate the impact of this problem based on our earlier work [1].

In addition to the cloud forensic problem, the EU General Data Protection Regulation (GDPR) [2] came into effect on 25th May 2018, and a principal requirement is the protection of any personally identifiable information of any EU resident 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 string of cyber breaches that are likely to ensue. Full compliance will inevitably pose a challenge for all organisations, but for those using cloud, due to the potential impact of the cloud forensic problem, the challenge will become so much more difficult.

It is too early to speculate on what approach the regulator might take towards setting penalties for breaches, but there is little doubt that where a company has an attitude problem towards proper compliance, or is complicit through poor internal security controls and provisions, then all these factors will be taken into account when gauging the level at which to set any potential fines. Equally, where a company can demonstrate that it has taken proper steps to mitigate the impact of the cloud forensic problem, it is clear that this will have the opposite effect, resulting in considerably lower levels of fines as a consequence of any breach.

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 address the importance of separation of duties. In Section V, we consider how we might develop some of these well established techniques to help us address this significant cloud security problem. In Section VII, we first consider some possible impediments to restoring the 'paper ink' trail. In Section VIII, we look at how we might use the immutable database as the core of this approach. In Section IX, we discuss the implications of this proposed solution. In Section X, we draw our conclusions and discuss our possible 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 [3] [4] and performance [5]–[7] it was not long before thoughts of security [8]–[10] and privacy [11] [12] started to surface.

While the US National Institute for Standards and Technology (NIST) were one of the first organisations to propose standard definitions [13] [14] interest in security [15]–[18] and privacy [19]–[21] started to grow.

Thoughts also started turning to accountability [9] [22]–[24] given the evolving complexities of cloud ecosystems. This ultimately led the EU to set up the Accountability for Cloud (A4Cloud) Project [25] to consider such important matters. The A4Cloud project drew much attention to the need for proper accountability in cloud systems and the contributors developed many useful mechanisms for ensuring proper levels of accountability could be monitored and achieved.

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 connected to the internet are subject to continuous and serious attack, and cloud systems are no exception. It would be realistic to state that no system is immune to attack, and this is particularly true for cloud systems. Attackers will always succeed in gaining entry to systems. The secret of success here is to be able to identify these occurrences the moment they happen, so that the attack can be shut down and the perpetrator removed from the system.

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 [26]. 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 is why it is so important for all companies to analyse their 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 set up a foothold in the system, thus allowing them to take steps to become an intruder.

The aim is not just to get in, and out, as quickly as possible, but to be able to develop a long term foothold, secreting themselves into corporate servers and other subsidiary systems which will allow them to return time and time again to help themselves to more information 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 [27].

If we look back five years ago, at previous cyber breach reports [28] 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 [29] 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 [30]. 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 GDPR [2] 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 [31] 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.

However, since the intruder will likely have worked hard to increase privileges to the point that they are able to modify or worse, delete all forensic trails in the system, the likelihood of an organisation being able to properly identify which records have been compromised may prove impossible to determine. Often, capturing adequate levels of forensic data does not happen due to many such features being turned off by default. It is bad enough when intruders delete forensic records, but it is inexcusable when an organisation fails to collect them in the first place.

The consequence of failure to detect such intrusions 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, their managers and all their stakeholders.

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. It can be argued that accounting and the reporting of accounting numbers has had two overriding purposes that are in tension with each other. The first, that had dominance in earlier centuries [32] is stewardship, though the term itself has seemed to evolve over time [33] (pp. 264) from being the careful, honest and accurate recording of transactions to efficient use of resources to finally ensuring an appropriate return for shareholders.

This progression is dependent on the trust that the earlier definition can now be taken as given due to improvements in recording mechanisms and the outside eye of an auditor. The confidence in the recording mechanism requires a complete history of transactions that means the accounts can be checked and even re-built if necessary whether in the mythical "shoe-box" of receipts for the small business or the sophisticated computerised ledgers of a multinational.

The integrity of the items recorded and the potential value of the detail highlights another concern that the data could be useful to people for whom it was not intended (competitors and fraudsters). Hence the need to lock up the accounting ledgers (or their computerised descendants) to keep them from being corrupted or seen by those who have no right of access.

From the 1950s onwards a more developed theory of accounting and reporting evolved with the focus being on accounting 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 [34].

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." [35]

"Decision usefulness", particularly for investors, became the central determinant of "good" or "bad" accounting methods, again one could argue, because of a confidence (sometimes misplaced) that stewardship, the basic recording, could be taken for granted.

In the above story of accounting development, we already needed to introduce the term "auditing". 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", or some similar phrase, of the business under review. Audit, the checking of conformity or of being fit for purpose, takes place in many fields, each of which develop over time and may (or may not) learn appropriate lessons from audit practices that have been honed over decades or centuries in more mature situations or professions .

Hence, not only accounting but also financial auditing techniques can also be applied to any other sphere where there is a need for recording, safety or trust and where there are records and some element of measurement, in this case, of course, we are particular interested in data. Hence, seeking to apply the more evolved and time tested techniques from accounting and auditing to the management and governance of data — and specifically data in the cloud would seem logical.

A further extension of the processes of accounting and audit is forensic (OED [36] "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, though such approaches are often used without a courtroom on the horizon. Forensic accounting is tuned to expose fraud and manipulation.

We can potentially 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 any database system to an accounting system, whereby we collect, measure, process and communicate data and the information gleaned from it concerning a business to the people for whom it is intended or relevant. Of course, the reliability, and even completeness, of data is a prerequisite for assessing any organisational efficiency level or for decision making.

We can see that the completeness of recording, the trust in the methods of processing the transactions and the ability of an auditor to interrogate the raw transactions are key building blocks for any effective data management system — whether accounting or otherwise focused.

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. This trail of records written in ledgers and of pieces of paper with signatures, comments and account codes provides for even the smallest business a trail of evidence for the accountant or auditor to follow through. The occasional missing item can usually be determined through "incomplete records analysis" as there is a surety concerning the other data and the bank statements. A larger business would think through more streamlined and consistent approaches to record keeping which then evolved into some of the earliest computer records, where (with known hard drives and no internet) anything entered would stay entered with the identity of the person undertaking the transaction, a time stamp and the matching double entry.

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 should then be able to easily apply cloud forensic techniques, learning from the accounting world, in order to help us bring about a successful prosecution in the courts and to become aware of the steps needed to improve security for the future. In practice, this, of course, will 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. We can learn lessons from the accounting world, specifically in the area of the audit trail, as used with accounting systems for centuries. A further relevant step in business and accounting risk mitigation in the accounting process is separation of duties, and we will now discuss this more fully in the next section.

IV. The Importance of Separation of Duties

For many decades, a key part of the structure of departments and of businesses overall is that of "separation (or segregation) of duties." This is a simple but straight forward security measure that could be employed by all but the smallest businesses. The logic is to carefully separate out the tasks in a business process so that no one person can have input or control into steps that might give them the opportunity, and temptation, to commit fraud or to effect theft. The smallest business would struggle to achieve this as different employees will be required to be responsible for specific tasks.

Ashton, [37] used a questionnaire to ask auditors a series of questions with the intent of being able to weigh their consistency as they inspected accounts and applied judgement. The first questions in his questionnaire addressed the segregation of tasks —

- Are the tasks of both timekeeping and payment of employees adequately separated from the task of payroll preparation?
- Are the tasks of both payroll preparation and payment of employees adequately separated from the task of payroll bank account?

It is not hard to see the result of a negative answer to either question. In both cases, an employee would be faced with the chance to change numbers in order to benefit themselves or, applying a little cunning, someone else. Involving two or more people may not be perceived to be enough, a further good safety feature would be to site the wages and salaries staff away from most of the workforce, reducing the chance of collaboration on a fraudulent scheme. A further gain from segregation, even when all employees are honest, is the opportunity to spot mistakes — a second person being required to take up the next stage of a process will mean either a clearly defined check or at least a "reasonableness" check on the work done to date.

The implications of judging that the answer to either of these two questions is "no" are obvious — an opportunity and a temptation arises for an individual to manipulate the payroll to their advantage. Clearly if it were possible to locate the payroll department away from the main work location and be confident that no one in payroll knew anyone in the rest of the company, then confidence would be increased yet further. Such separation not only makes fraud difficult, but also means unintentional errors are more likely to be spotted.

According to Gelinas et al. [38] there are four areas in a business or accounting process that need to be separated: authorising transactions, executing transactions, recording transactions and safeguarding resources subsequent to the transactions being completed. Vaassen et al. [39] add a further need for separation with — "authorisation; custody; recording; checking and execution". We move to more direct relevance to our key concern with the work of Hall [40] who addresses segregation when computerised accounting has been implemented. Hall sees that further concerns now need to be added, including "Is the logic of the computer program correct? Has anyone tampered with the application since it was last tested? Have changes been made to the programme that could have caused an undisclosed error?" (page 208). Even such apparently obvious questions need to be asked if the integrity of the system is to be assured..

The Sarbanes-Oxley Act [41] introduced new disclosure requirements for senior directors to commit personally to the quality of their numbers, and by implication the systems that produced those numbers. Ge and McVay [42] found 261 firms in 2002-2004 that admitted weaknesses in control and of these 45 included a reference to separation of duties. A further concern for us is that companies from the computer sector were a noticeably high proportion of the 261 problem companies in Ge and McVays study.

Taking these examples and then applying the same logic to both programming in general to software use is straightforward, though one might still question whether knowledge of this is sufficiently carried through into practice. If practice is indeed good, we still need a record of activity in order for audit and to investigate when things go awry (i.e., the audit trail).

This may all sound like unnecessary work and detail. However, when one of the authors ran a large purchase ledger department, a ledger clerk became confused when processing a number of very similar invoices from one large supplier that totalled to 2 million. She had entered invoices, then entered credit notes to cancel them and then repeated this a number of times before coming to him in some distress. The problem was sorted, but, as the author expected, an auditor spotted this unusual activity some time later and asked for an explanation. The event log and records showed each step, who had carried out which transactions and how the issues were corrected.

V. 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. [43] 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 [44] 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 [45] and Hopwood et al. [43]) including scientific methods. Taylor [45] 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.

[43] have a similar structure but give a little greater weight to forensic science and computer forensics.

From the computer science camp, Choo and Dehghantanha [46] a more scholarly work, reflects a greater weight placed on technical issues, as well as the tools and techniques needed, for forensic cloud investigations. Almulla et al. [47] review the cloud forensic literature and find some discursive, though mostly technical papers. Some of our previous research [48]–[50] has focused on the critical nature of understanding human frailties and interactions as well as what seems the more technically demanding elements of computer science.

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. There are many accounting associations, many with long histories and experience in exam setting. The syllabi of these bodies depends somewhat on the countries in which they operate and are revised over time to reflect new priorities and changes in the world -both technical and social. The Association of Chartered Certified Accountants (ACCA), is a significant international accounting professional body and we will take their professional examination content as an exemplar of others. The ACCA has over 200,000 members [51] and has an exam at its professional stage, Advanced Auditing and Assurance [52] that includes — a section on forensic audit though it should be noted that — it is only a very small part of the content of that exam.

Each of the professional bodies faces a dilemma when revising their syllabi for a changing and ever more complex world. In many countries, the market to attract accounting students is competitive, hence a more complex world cannot lead to more exams and a longer route to qualification without the body facing a competitive disadvantage. It is also very difficult to decide to drop traditional content to make appropriate space for newer material or issues.

It would seem that qualifying accountants are ill-prepared by their professional bodies for the complexities of the cloud environment. This is both in terms of understanding the environmental issues, though there is accessible material for them to pick up some of this (see Taylor [45] and Hopwood et al. [43]), 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 and 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. The large accounting “firms”, commonly referred to as the “Big 4” (KPMG, PWC, EY and Deloitte), who audit nearly all the worlds big companies and collectively employ about a million people (2017), do offer forensic services along with a broad range of consultancy services (see [53] for example).

The ever-widening scope of the Big 4, making far more money from consultancy than audit, is contentious in some countries. The reliance on an oligopolistic audit industry with seemingly conflicted aims of professionalism and commercial gain, along with what many see as questionable competence in their core audit activity, is building a crescendo for change Marriage2018 and Marriage2018a highlight some audit quality issues and FRC2018 also includes an example of the consultancy dilemma at paragraph 34). Whilst these firms have undoubtedly built up some expertise in the area, there are some significant issues with their continuing range of activities. Two discussed options, in the UK at least, are either splitting up the firms or to separate the audit arm from consultancy.

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 [46]) 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 Figure 1 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. One logical conclusion would be the need to build multi-disciplinary teams, though the development of sufficient common understanding, shared technical language, never mind recognition of mutual professional credibility and importance should not be taken as insignificant.

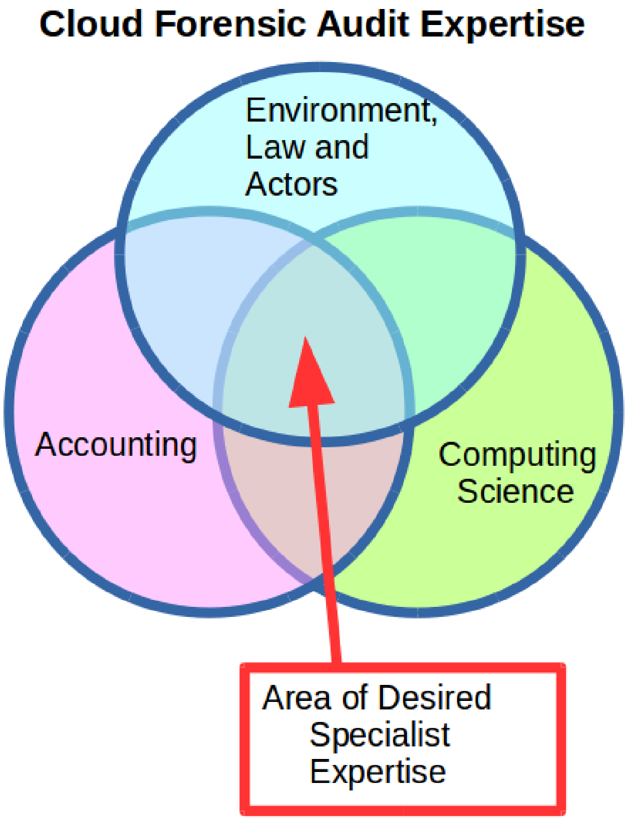
Cloud Forensic Audit Expertise
Environment, Law and Actors
Accounting
Computing Science
Area of Desired Specialist Expertise

Figure 1: The Area of Desired Expertise

Whilst there are many audit tools used in accounting, the computing literature already uses the “audit trail” [47] when discussing evidence integrity, however in previous work [31], [54]–[56] 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 understood and then 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 — indeed the whole concept of double entry is to record every transaction twice. Computer scientists, on the other hand, have a focus on efficiency and minimising costs, using such terms as “redundancy” for seemingly unnecessary or duplicate recording.

An audit trail needs to be developed and be fit for purpose — it may require some thought and planning to decide firstly on the purpose(s) of the trail and then logically what data needs to be safely and securely recorded. For example, Bernstein [57] sees the trail including: events, logs, and the analysis of these, whilst Chaula [58] gives a longer, more detailed list: raw data, analysis notes, preliminary development and analysis information, processes notes, and so on. Pearson et al. [10] as far back as 2010, accept that attaining consistent, meaningful cloud audit trails is a goal rather than reality. More worryingly, Ko et al. [22] point out that it is possible to delete the audit trail along with a cloud instance, meaning there is no record then remaining. In the traditional accounting external audit, the external accountant appears at the end of the year and would need to access all the records they might need in order to satisfy themselves that everything is in order — an ephemeral audit trail would not be fit for purpose. Ko [59] also details the requirements for accountability.

VI. The Special Skills Mix Needed for Cloud Forensic Audit

As we mentioned earlier, with modern cloud systems, we are no longer 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 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 potentially significant knowledge gaps. Clearly, the cloud environment is considerably different from conventional distributed network models under the sole control of a company. There are now a great many actors involved in such an environment, each potentially with their own agenda. Legal and regulatory issues are also a lot less clearly defined for cloud environments, with the increased likelihood of multiple companies and jurisdictions.

We also have to contend not only with the invited actors but also with the potential of a number of uninvited actors too. The list of the invited players is longer than we might first think. Company employees and managers may not be as competent or trustworthy as we would wish. Outside the company itself, there will be many others, including the software provider, the cloud service provider, the auditor and, in a modern business-to-business environment, the suppliers and the customers. This is a complex mix of actors with disparate agendas and, frustratingly, it cannot even be taken as given that these legitimate actors will be willing to co-operate fully with each other if a problem arises. Of course, there are also the potential uninvited guests — namely attackers and intruders, with the latter presenting the greater challenge.

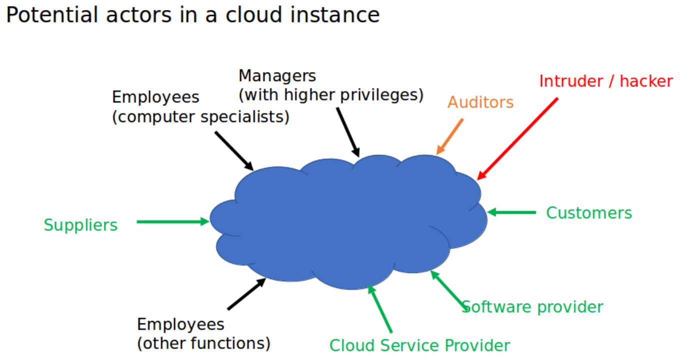

Figure 2: Who is in your cloud?

Figure 2 shows a little of this complexity with the internal actors in black, other companies in green, the auditor in orange and the intruder in red. Whilst one might hope that the authorized participants will play by the rules, any intruder will make up their own, hence gaining access via a customer, supplier or even the auditor is a reasonable option for them.

This level of complexity means we should 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 both incomplete knowledge and skill sets. These gaps are not just missing pieces of a jigsaw that someone else from a different discipline can potentially add, but the filling of the gap may also redefine some of the other problems and issues in other disciplinary domains within the complex situation.

Further, in the absence of the solid ‘ink trail’, this increases the complexity of the task exponentially. In Figure 1 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. Assuming individuals are approaching the centre from an initial disciplinary perspective, the need for each to move well beyond their professional comfort zone is obvious and challenging.

Currently, when it comes to Cloud, intruders can potentially have it all their own way. Once they are in the system, it can be merely a matter of time before they have built up sufficient privileges to delete the forensic trail of their activity, thus allowing them to either bed down for the long run, or to withdraw without leaving fingerprints or “steps in the mud”. The deleting of all audit and forensic trails as they proceed, means that there is an significant 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 find a way to restore in some way the permanent 'ink trail' so that we have something to fall back on and to enable us to re-trace and re-build if necessary, this is where an immutable audit trail process needs to be conceived, designed and implemented. Second, we need to recognise all the experience and skill sets required and then enable the knowledge gaps to be bridged, ensuring that all the parties who need to be 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 following sections.

VII. Some Impediments to Restoring the Permanent 'Ink Trail'

Before discussing how we might resolve the matter of the permanent 'ink trail', we should consider some impediments that are often inadvertently placed by companies on themselves when using cloud systems. Companies should not rush this decision, but should prepare properly ahead of time. They should not assume it will be easy. Instead, they should think it through, understand the costs properly, and ensure they have the right service package rather than continuing to use the first one that came along [60]. Companies frequently wore cost blinkers when choosing cloud provisioning. It is vital to factor in risks and potential exposure too [61] not just looking at the short term, but also taking the long view too.

Many companies have failed to prepare a proper disaster recovery plan [62]. They must always expect the unexpected, and plan for it. It is vital to be aware of precisely which data needs to go to cloud, who should be able to see it, and this data needs to be protected with proper access control. Companies must understand where their data is stored [63] and how they can get their data back, if required. They must understand who all can gain access to their data. As we discussed in Section V cloud systems will not necessarily only be exposed to CSP personnel, but also other sub-contracted organisations [64] whose security and privacy approach might not be as robust as that of the CSP. Companies often fail to account for data privacy risks, which presents a really good incentive for using encryption for their data.

As far as cloud security and privacy is concerned, there is no single solution [54]. It is a mistake to assume the CSP's security is good. CSPs have a huge incentive not to disclose full details of previous security and privacy breaches, as they do not wish to lose future sales. Companies themselves should not use the wrong privacy approach. Wherever possibly, they should try to align security with their own business goals [65]. No matter which approach is used, it should be cloud-friendly. For GDPR compliance, companies should always consider encryption [66] preferably with split encryption keys. Companies often sign up to cloud accepting the standard SLA, which can be a big mistake. Many standard contracts are extremely vague about security and privacy, or do not even mention it. This lack of accountability on the part of the CSP can only help attackers breach company systems more easily.

All companies must understand the true threat against their employees, customers, suppliers and ultimately, their data. Their information security plan must be cutting-edge and comprehensive. They should view security not just as an "IT problem", but rather as a "business problem" that also includes IT. Many who implement security as an IT problem have ended up with a strong IT implementation of data security controls but limited (if any) attention paid to the required security controls such as physical security, security policies and procedures, training, and other administrative and environmental controls. People are frequently the weakest link in the security chain, meaning this is why special attention must be paid to their proper training in all security matters. This is precisely the reason why security mirrors the business architecture of any company, namely it is a combination of people, process and technology [67] not technology alone.

There have been many interesting approaches to trying to resolve some of the obvious issues in cloud security. One such area is how to ensure data integrity in the cloud. There have been a number of interesting proposals, such as [68] [66] [69] [70] [71] which seek to provide data integrity assurance to users through various forms of audit, which as a rule do work quite well. Others, such as [72] [73] [74] [75] have suggested trust computing might be the way to solve these problems. Again, these often work well, but despite establishing trust between providers and users, nevertheless, the fact remains that the work is being performed on someone else's systems, thus risks will always remain. Yet others, such as [76] [77] [78] [79] believe provable data possession could help address this problem, whereas others believe that timeline entanglement, such as [80] [81] [82] is the best solution.

All of these systems, while proving capable of delivering what they promise, share a common flaw. They provide an excellent means of achieving their objectives, but not a means to deal with what happens after a serious security breach. In such cases, the intruders often act in a brutal and indiscriminate way, modifying or deleting multiple records. Any user who does not understand the true purpose of an audit trail may quickly discover that they no longer have access to the necessary data with which to restore the modified or deleted data to its original state.

Thus warned, we will now turn to looking at the immutable database in the next section.

VIII. 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. Three fundamental weaknesses exist and need to be addressed.

First, failure to use adequate default logging options will result in a reduced level of required audit trail data being collected. Second, there is often a lack of understanding that audit trail data can be accessed by a malicious user gaining root privileges. This can allow the malicious user remove key data which would otherwise have provided evidence of who compromised the system, and what actions they performed once they took control. Third, log data must be properly collected in permanent storage such that there can be no loss of audit trail data, either when an instance is shut down, or when it is compromised. The obvious answer would be to store this data away from the running cloud instance, in a secure environment using an immutable database which will allow "append-only" transactions to be made.

Starting with the first point, most database software offers a considerable range of audit trail options that can be used to keep proper records of what is happening within the system. However, by default, logging is set to "off"! Since many organisations rely on default installation settings, it is clear that they will be at an immediate disadvantage unless the logging options are fully explored and activated. An obvious, yet simple point missed by many.

Looking at the second point, as Anderson [83] states, the audit trail should only be capable of being read by users. In a cloud setting, this presents a problem, as the software being used is usually running on someone else's hardware, with the output being stored there as well. There is always a risk of compromise from any outside user with malicious intent. There may also be a risk of compromise by some malicious actor working for the CSP. The CSP may very well take vetting of staff seriously, but there may be situations that arise where a temporary contract worker, subject to lesser scrutiny, is engaged at short notice.

Turning to the third point, where database logging is actually switched on, this data is logged to the currently running instance. Thus, this data remains accessible to any intruder who is able to successfully breach the system. This will afford them the opportunity to cover their own tracks by modifying or deleting any entries relating to their intrusion of the system. Equally, they may simply delete the entire audit trail files. Finally, when the instance is shut down, all the data would disappear anyway.

These three points are generally not much thought about, yet they present a serious weakness to the success of maintaining the audit trail. Equally, these are relatively trivial to address. All too often, management and IT staff will take the view "so what?. We don't need to collect redundant data". This entirely misses the point that this data is the only source of proof of what intruders have done whist inside the system. Without these records, it will not be possible to comply with GDPR compliance procedures.

Of course, we need to consider very carefully exactly what data we should log to ensure we can achieve compliance with the GDPR. First, we need to monitor our Cloud instances. We need to understand exactly who is accessing our systems, whether authorized or not and we need to monitor what is happening with our database systems to understand what these users are doing with them.

Looking at our Cloud instances, as Duncan and Whittington have shown in [31] [84] [55] 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. [22] 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 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 did with the system after they gained 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.

IX. Discussion

It is clear that without the assistance of the humble audit trail, compliance with the GDPR while using cloud is likely to prove an unattainable goal. Of course, not being breached would also provide a solution, but based on events to date, there is no guarantee that such a situation would be readily achievable, let alone sustainable in the long run.

Having developed an effective, yet simple and workable solution to this problem, we may well have some further 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 consequences 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. It is also the case that some practical measure should also be taken, such as ensuring that the encryption and decryption keys are not stored on the cloud instance they are designed to protect.

Clearly, doing nothing is not an option. Without a means of being able to tell which records have been accessed, modified or deleted, compliance with the GDPR will not be possible, and that will potentially carry a very high price tag indeed.

X. 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 article, 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 have built a miniature real life Cloud system on which to test our ideas. The server runs 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 system will run on a closed network where it will be subject to rigorous attack, with the view to discover whether the immutable database approach can succeed in allowing Cloud users to be GDPR compliant.

We have developed a range of scenarios 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 on our results next year, and we will be working towards delivering a workable solution to keep Cloud users compliant.

References

[1] B. Duncan and M. Whittington, "Can Forensic Audit Help Address the Cloud Forensic Problem in Light of the Requirements of the Forthcoming European Union General Data Protection Regulation?" in Cloud Computing 2018: The Ninth International Conference on Cloud Computing, GRIDs, and Virtualization, no. February. Barcelona, Spain: IARIA, 2018, pp. 84–89.

[2] EU, "EU General Data Protection Regulation (GDPR)," 2017. [Online]. Available: http://www.eugdpr.org/ [Last accessed: August 2018]

[3] 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.

[4] 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.

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

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

[7] S. Fraser, R. Biddle, S. Jordan, K. Keahey, B. Marcus, E. M. Maximilien, and D. Thomas, "Cloud Computing Beyond Objects: Seeding the Cloud," Communications, 2009, pp. 847–850.

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

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

[10] 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.

[11] 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.

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

[13] P. Mell and T. Grance, "Effectively and Securely Using the Cloud Computing Paradigm," NIST, Information Technology Laboratory, vol. 2, no. 8, 2009, pp. 304–311.

[14] P. Mell, T. Grance, and Others, "The NIST Definition of Cloud Computing," National Institute of Standards and Technology, Tech. Rep., 2011.

[15] 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.

[16] 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 [Last accessed: August 2018]

[17] 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.

[18] 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.

[19] 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 [Last accessed: August 2018]

[20] 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 [Last accessed: August 2018]

[21] 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.

[22] R. K. L. Ko, P. Jagadpramana, M. Mowbray, S. Pearson, M. Kirchberg, Q. Liang, and B. S. Lee, "TrustCloud: A framework for accountability and trust in cloud computing," Proceedings - 2011 IEEE World Congress on Services, SERVICES 2011, 2011, pp. 584–588.

[23] 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.

[24] 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.

[25] EU, "Accountability for Cloud (A4Cloud)," 2018. [Online]. Available: http://a4cloud.eu/ [Last accessed: August 2018]

[26] M. Hammock, "A Review of the Economics of Information Security Literature," pp. 1–38, 2010. [Online]. Available: http://ssrn.com/abstract=1625853 [Last accessed: August 2018]

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

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

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

[30] 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.

[31] 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.

[32] R. M. Skinner and J. Milburn, Accounting standards in evolution, 2nd ed. Toronto: Prentice Hall, 2001.

[33] S. A. Zeff, "The objectives of financial reporting: a historical survey and analysis," Accounting and Business Research, vol. 43, no. 4, 2013, pp. 262–327.

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

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

[36] OED, "Oxford English Dictionary," 2017. [Online]. Available: http://www.oed.com [Last accessed: August 2018]

[37] R. H. Ashton, "An experimental study of internal control judgements," Journal of Accounting Research, 1974, pp. 143–157.

[38] S. S. G. Gelinas U.J. and A. E. Oram, Accounting Information Systems (4th edition). South-Western College Publishing, Cincinnati, Ohio, US., 1999.

[39] E. Vaassen, R. Meuwissen, and C. Schelleman, Accounting information systems and internal control. Wiley Publishing, 2009.

[40] J. A. Hall, Accounting Information Systems (3rd edition). South-Western College Publishing, Cincinnati, Ohio, US., 2001.

[41] Sox, "Sarbanes-Oxley Act of 2002," p. 66, 2002. [Online]. Available: news.findlaw.com/hdocs/docs/gwbush/sarbanesoxley072302.pdf [Last accessed: August 2018]

[42] W. Ge and S. McVay, "The disclosure of material weaknesses in internal control after the Sarbanes-Oxley Act," Accounting Horizons, vol. 19, no. 3, 2005, pp. 137–158.

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

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

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

[46] 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.

[47] 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.

[48] B. Duncan and M. Whittington, "Information Security in the Cloud: Should We be Using a Different Approach?" in 2015 IEEE 7th International Conference on Cloud Computing Technology and Science (CloudCom), Vancouver, 2015, pp. 523–528.

[49] B. Duncan and M. Whittington, "Company Management Approaches Stewardship or Agency: Which Promotes Better Security in Cloud Ecosystems?" in Cloud Computing 2015: The Sixth International Conference on Cloud Computing, GRIDs, and Virtualization, IARIA, Ed. Nice, France: IEEE, 2015, pp. 154–159.

[50] B. Duncan and M. Whittington, "Reflecting on whether checklists can tick the box for cloud security," in Proceedings of the International Conference on Cloud Computing Technology and Science, CloudCom, vol. 2015-Febru, no. February. Singapore: IEEE, 2015, pp. 805–810.

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

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

[53] KPMG, "Forensics," 2018. [Online]. Available: https://home.kpmg.com/uk/en/home/services/advisory/risk-consulting/forensic-landing.html [Last accessed: August 2018]

[54] 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.

[55] 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.

[56] 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.

[57] 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.

[58] 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 [Last accessed: August 2018]

[59] R. K. L. Ko, "Data Accountability in Cloud Systems," in Security, Privacy and Trust in Cloud Systems. Springer, 2014, pp. 211–238. [Online]. Available: http://link.springer.com/10.1007/978-3-642-38586-5 [Last accessed: August 2018]

[60] Q. Zhang, L. Cheng, and R. Boutaba, "Cloud computing: State-of-the-art and research challenges," Journal of Internet Services and Applications, vol. 1, no. 1, 2010, pp. 7–18.

[61] K. Dahbur, B. Mohammad, and A. B. Tarakji, "A Survey of Risks, Threats and Vulnerabilities in Cloud Computing," Computing, 2011, pp. 1–6.

[62] Z. Chen and J. Yoon, "IT Auditing to Assure a Secure Cloud Computing," in Proceedings - 2010 6th World Congress on Services, Services-1 2010, 2010, pp. 253–259.

[63] D. J. Abadi, "Data management in the cloud: limitations and opportunities," IEEE Data Eng. Bull., vol. 32, no. 1, 2009, pp. 3–12.

[64] R. Chow, P. Golle, M. Jakobsson, E. Shi, J. Staddon, R. Masuoka, and J. Molina, "Controlling data in the cloud," in Proceedings of the 2009 ACM workshop on Cloud computing security - CCSW '09, 2009, p. 85.

[65] K. Popovic and Z. Hocenski, "Cloud computing security issues and challenges," MIPRO, 2010 Proceedings of the 33rd International Convention, 2010, pp. 344–349.

[66] S. Subashini and V. Kavitha, "A Survey on Security Issues in Service Delivery Models of Cloud Computing," Journal of Network and Computer Applications, vol. 34, no. 1, 2011, pp. 1–11.

[67] PWC, "UK Information Security Breaches Survey - Technical Report 2012," London, Tech. Rep. April, 2012. [Online]. Available: www.pwc.com www.bis.gov.uk [Last accessed: August 2018]

[68] Z. Hao, S. Zhong, and N. Yu, "A Privacy-Preserving Remote Data Integrity Checking Protocol with Data Dynamics and Public Verifiability," IEEE Transactions on Knowledge and Data Engineering, vol. 23, no. 9, 2011, pp. 1432–1437.

[69] D. Chen and H. Zhao, "Data Security and Privacy Protection Issues in Cloud Computing," 2012 International Conference on Computer Science and Electronics Engineering (ICCSEE), vol. 1, no. 973, 2012, pp. 647–651.

[70] K. Yang and X. Jia, "An Efficient and Secure Dynamic Auditing Protocol for Data Storage in Cloud Computing," IEEE Transactions on Parallel and Distributed Systems, vol. I, no. 9, 2014, pp. 2–5.

[71] L. Wei, H. Zhu, Z. Cao, X. Dong, W. Jia, Y. Chen, and A. V. Vasilakos, "Security and privacy for storage and computation in cloud computing," Information Sciences, vol. 258, 2014, pp. 371–386.

[72] N. Santos, K. P. Gummadi, and R. Rodrigues, "Towards Trusted Cloud Computing," Information Security Technical Report, vol. 10, no. 2, 2005, p. 5. [Online]. Available: http://www.sciencedirect.com/science/article/pii/S1363412705000221 [Last accessed: August 2018]

[73] Z. Shen, L. Li, F. Yan, and X. Wu, "Cloud Computing System Based on Trusted Computing Platform," 2010 International Conference on Intelligent Computation Technology and Automation, 2010, pp. 942–945.

[74] S. Sengupta, V. Kaulgud, and V. S. Sharma, "Cloud Computing Security - Trends and Research Directions," 2011 IEEE World Congress on Services, no. October, 2011, pp. 524–531.

[75] Z. Xiao and Y. Xiao, "Security and Privacy in Cloud Computing," Communications Surveys Tutorials, IEEE, vol. 15, no. 2, 2013, pp. 843–859.

[76] Y. Zhu, H. Wang, Z. Hu, G.-j. Ahn, H. Hu, S. S. Yau, H. I. Storage, and R. Information, "Efficient Provable Data Possession for Hybrid Clouds," in Proceedings of the 17th ACM Conference on Computer and communications security, 2010, pp. 756–758.

[77] G. Ateniese, R. Burns, R. Curtmola, J. Herring, O. Khan, L. Kissner, Z. Peterson, and D. Song, "Remote data checking using provable data possession," ACM Transactions on Information and System Security, vol. 14, no. 1, 2011, pp. 1–34.

[78] Q. Wang, C. Wang, K. Ren, W. Lou, and J. Li, "Enabling Public Verifiability and Data Dynamics for Storage Security in Cloud Computing," in IEEE Transactions on Parallel and Distributed Systems, vol. 22, no. 5, 2011, pp. 847–859.

[79] Y. Zhu, H. Hu, G. J. Ahn, and S. S. Yau, "Efficient audit service outsourcing for data integrity in clouds," Journal of Systems and Software, vol. 85, no. 5, 2012, pp. 1083–1095.

[80] H. T. T. Truong, C.-L. Ignat, and P. Molli, "Authenticating Operation-based History in Collaborative Systems," Proceedings of the 17Th Acm International Conference on Supporting Group Work, 2012, pp. 131–140.

[81] M. Mizan, M. L. Rahman, R. Khan, M. Haque, and R. Hasan, "Accountable proof of ownership for data using timing element in cloud services," Proceedings of the 2013 International Conference on High Performance Computing and Simulation, HPCS 2013, 2013, pp. 57–64.

[82] S. L. Reed, "Bitcoin Cooperative Proof of Stake," 2014, pp. 1–16.

[83] R. J. Anderson, Security Engineering: A Guide to Building Dependable Distributed Systems, C. A. Long, Ed. Wiley, 2008, vol. 50, no. 5.

[84] 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.