

Enhancing Cloud Security and Privacy: Time for a New Approach?

Abstract—Achieving cloud security is not a trivial problem and developing and enforcing good cloud security controls is a fundamental requirement if this is to succeed. The very nature of cloud computing can add additional problem layers for cloud security to an already complex problem area. We discuss why this is such an issue, consider what desirable characteristics should be aimed for and propose a novel means of effectively and efficiently achieving these goals through the use of unikernel based systems. The main thrust of this position paper is to discuss the key issues which need to be addressed, noting which of those might be covered by our proposed approach. We discuss how our proposed approach may help better address the key security issues we have identified.

Index Terms—Cloud security and privacy; management control; compliance; complexity

I. INTRODUCTION

There are a great many barriers which must first be overcome before the goal of cloud security can be achieved. Much research has been conducted towards resolving this problem through technical means, but this presents a fundamental flaw. The business architecture of a company comprises people, process and technology [1] and any solution which will focus on a technological solution alone will be doomed to failure. People present a serious weakness to the security of a company [2], and while process may be very well documented within an organisation, often it is out of date due to the rapid pace of evolution of technology [3]. Technology can benefit companies due to the ever improving nature and sophistication of software, which is a good thing, but on the other hand, presents a much higher level of complexity for proper and secure implementation within company systems. However, the threat environment is also developing at a considerable pace [4].

Cloud brings a far higher level of complexity than is the case with traditional distributed systems, in terms of both the additional complexity of managing new relationships in cloud, and in the additional technical complexities involved in running systems within the cloud. It runs on other people's systems, and instances can be freely spooled up, and down, as needed. This leads to concerns over proper maintenance of an adequate audit trail [5], and forensic examination can range between difficult and impossible. We seek to address these problems by taking a more simplistic approach, with the goal of limiting the dependence of the company on the people who use the system, thus removing as many opportunities for human error as possible.

We are concerned with achieving both good security and good privacy. While it is possible to have security without

privacy, it is not possible to have privacy without security. Thus our approach in this position paper will be to first ensure a good level of security can be achieved, and to that end, we start by listing the specific security issues we seek to address and discuss how we propose to tackle them in Section II. This means we will leave addressing Privacy for a later paper. The remainder of the paper is organized as follows: in Section III, we discuss the outline of our proposed approach; in Section IV, we include some useful definitions; and in Section V, we discuss our conclusions.

II. THE SPECIFIC SECURITY ISSUES TO BE ADDRESSED

The fundamental concepts of information security are confidentiality, integrity, and availability (CIA), a concept developed when it was common practice for corporate management to run a company under agency theory. We have all seen how agency theory has failed to curb the excesses of corporate greed. The same is true for cloud security, which would suggest a different approach is needed [6]. The business environment is constantly changing [5], as are corporate governance rules and this would clearly imply changing security measures would be required to keep up to date. More emphasis is now being placed on responsibility and accountability [7], social conscience [8], sustainability [9][10], resilience [11] and ethics [12]. Responsibility and accountability are, in effect, mechanisms we can use to help achieve all the other security goals. Since social conscience and ethics are very closely related, we can expand the traditional CIA triad to include sustainability, resilience and ethics.

Ten key security issues have been identified, namely:

- The definition of security goals;
- Compliance with standards;
- Audit issues;
- Management approach;
- Technical complexity of cloud;
- Lack of responsibility and accountability;
- Measurement and monitoring;
- Management attitude to security;
- Security culture in the company;
- The threat environment.

A. The definition of security goals

Since many managers are unable, unwilling or unsure of how to define proper security goals [13] [14] [15], we seek to build this requirement into the system as a fundamental part of our approach. Thus, our basic approach will be “secure by design”[16].

B. Compliance with standards

Bearing in mind our comments on the issues with cloud security standards [17] [18], compliance audit mechanisms [19] [20], and the general consensus that standards need to shift from a rule based approach to a risk based approach [21] [22] [23] [24] [14] [25], we believe approaching this problem in a simple and robust manner will enable compliance with any relevant standard in future. One welcome change from the standards setting bodies is the move away from a strict rule based compliance approach to a more risk based approach, although the time taken from inception to agreement and implementation remains a concern.

C. Audit issues

There are three main purposes of audit [5], the most widely understood of which is the statutory requirement for financial statements to be audited by an independent external auditor, which has been a cornerstone of confidence in global financial systems since auditing was introduced [18]. The second purpose of audit is IT systems audit and the third is compliance, either with regulations, or more often with standards. We suggest that it is certainly necessary for management to understand better the purpose, and importance of audit [5] [26] [15] [27] [28]. It is also necessary to understand both the power and the weaknesses offered by the humble audit trail [6]. We will deal with this by providing everything that is necessary within our framework to ensure proper audit can be achieved, without complex configuration.

D. Management approach

There is no doubt that management approach is an important consideration for management to be aware of in the myriad relationships involved in the cloud ecosystem [29]. While it is not necessary that all actors should utilise the same management approach, it is certainly helpful for management to be able to recognise the management approach used by each of the actors involved within their own cloud ecosystem. This will allow them to better identify the key risks they face and take appropriate steps to attempt to mitigate these risks. Management approach can have an enormous impact on the success of cloud security and privacy and [29] provide some useful background on this. We aim to reduce the impact of management approach by providing security and privacy by default, and by limiting management options, thus ensuring we can deliver a secure and private system.

E. Technical complexity of cloud

Yet another issue is the increasing complexity which new technology brings, and the ever increasing potential exposure to risk brought about by a failure to grasp the significance of risks arising as a result of this increase in complexity [30]. Traditional distributed information systems present a multiplicity of technical layers, each of which must interact with one or more other layers. Cloud introduces yet more layers, each of which can be operated by different actors. Cloud brokers may also be involved, leading to yet more

layers, yet more complexity, yet more risk. Cloud brings the possibility to quickly deploy, for example, a web server with a database back end, and users will often rely on default settings when setting up such an instance, which can lead to a number of weaknesses [5], as these default settings usually pay far more attention to usability than to security. The same is true for web server software, database software and many other complex pieces of business software. It is not often realised that the default settings of a database actually switch off audit logging, thus making the attacker's life much easier [6]. The additional technical complexities introduced by cloud need to be expressly addressed [31] [32] [33] to identify and deal with as many weaknesses as possible, particularly when considering the maintenance of a proper forensic trail [34].

F. Lack of responsibility and accountability

We must also consider the role that responsibility and accountability play in achieving this difficult objective. Monahan and Yearworth [35] observe that Service Level Agreements (SLAs) should be meaningful, both for customers and vendors as defined by some objective criteria, but evidence from procurement failures for large IT systems suggests otherwise. This observation has inspired an investigation into the possibility of offering alternative security SLAs that would be meaningful to both customers and vendors. Duncan and Whittington [15] provide some useful background on these issues in SLAs.

This can lead to issues with allocating proper responsibility to the right actors [36], personal data [37] and privacy [38]. Some [39] [38] [40] have long argued that responsibility and accountability should always be built in to the design of cloud systems. Much as we would like to have built these factors into our approach, the nature of these issues is such that they must be dealt with by a combination of management, measurement, and negotiation with CSPs in setting up more accountable SLAs.

G. Measurement and monitoring

We see a good deal of research into measurement of CSR, [41] [42] [43] [44] [45] [46] [47] [48], resilience [49] [50] [51] [52] [53] [54] [55] [56] and sustainability [57] [58] [59], yet there is still some way to go before effective measures are developed. Thus we can see that measurement is extremely important, yet can be very difficult to achieve. There exists a clear need to employ some method of continuous monitoring when it comes to security management. Reports from global security companies, which include both non-cloud and cloud data [1], [60], [61] suggest that over 85% of security breaches involve a low level of technical competence, often facilitated by lack of understanding, lack of competence, or poor configuration of victims' systems. Duncan and Whittington [27] provide useful background detail on this.

In the first key challenge, we note the need to define proper security goals. Once this is achieved, we then need to devise suitable measurements or metrics with which to determine whether these goals are being met. This should be achieved through constant monitoring [62] [63] [17] [27]. This is the

only way it will be possible to prove that the desired security objectives determined by management are being achieved [27].

H. Management attitude to security

Management attitude to security has long been a high priority [64]. In [65], 77% of security professionals have long recognised the need to set security attitudes from the top. According to a report [1] management attitude is high, if you listen to the executives, yet low when you listen to IT practitioners. Thus management need to be fully aware that it is not simply a technical issue to be passed down the line, rather it is a fundamental business process which needs to be driven right from the top of the organisation. Information security now presents one of the largest risks facing business today and needs to be given the proper attention and commitment it requires. We hope to ease this by taking a more simple approach.

I. Security culture in the company

One of the most important aspects of creating good security in a company is the development and maintenance of a good security culture within the organisation. The importance of security culture within the company has long been recognised [64] [65] [1]. This, of course, is dependant upon the attitude to security displayed by top management. This also needs to be coupled with proper staff training to ensure staff understand how to deal properly with security threats. It is estimated [1] that in 2012, only 26% of companies with a security policy believe their staff understand how to use them. This is again an issue for management to deal with. However, our simplified approach to the problem will help them to more easily create an effective security policy.

J. The threat environment

It is necessary to recognise the magnitude of the threat environment. Companies are bound by legislation, sometimes regulation, the need to comply with standards, industry best practice, and are accountable for their actions. Criminals have no such constraints. They are completely free to bend every rule in the book, do whatever they want, manipulate, cajole, hack or whatever it takes to get to the money. They are constantly probing for the slightest weakness, which they are more than happy to exploit without mercy. It is clear that the threat environment is developing just as quickly as the technological changes faced by industry [18] [29] [28]. We need to be aware of this threat, and minimise the possible impact on our framework. While we have absolutely no control over attackers, we can help reduce the impact by removing as many of the “classical attack vectors” as possible, thus making their life far more difficult. The more difficult it becomes for them to get into the system, the more likely they will be to go and attack someone else’s system.

III. OUR PROPOSED SOLUTION

Again, by default, in the interests of usability, many more ports are open than may be needed to run a system. An

open port, especially one which is not needed, is another route in for the attacker. We also take the position that the probability of vulnerabilities being present in a system increases proportionally to the amount of executable code it contains. Having less executable code inside a given system will reduce the chances of a breach and also reduce the number of tools available for an attacker once inside. As Meireles [66] said in 2007 “... while you can sometimes attack what you can’t see, you can’t attack what is not there!”. We propose to address these issues by making the insides of the virtual machine simpler – but we should also address the outside. We also propose to tackle the audit issue by making configuration happen at build-time [67][68], and then making services be “immutable” after deployment, making “configuration lapses” (i.e. through conflicts caused by unnecessary updates to background services etc.) unlikely.

Given the success with which the threat environment continually attacks business globally, it is clear that many companies are falling down on many of the key issues we have highlighted in Section II. It is also clear that a sophisticated and complex solution is unlikely to work. Thus we must approach the problem from a more simple perspective.

A. Service isolation

A fundamental premise for cloud computing is the ability to share hardware. In private cloud systems, hardware resources are shared across a potentially large organization, while on public clouds, hardware is shared globally across multiple tenants. In both cases isolating one service from the other is an absolute requirement.

The simplest mechanism for service isolation is simply *process isolation* in classical kernels, relying on hardware supported virtual memory e.g. provided by the now pervasive x86 protected mode. While process isolation has been used successfully in mainframe setups for decades, access to terminals with limited user privileges has also been the context for classical attack vectors such as stack smashing, root-kits etc., the main problem being that a single kernel is being shared between several processes and that gaining root access from one terminal would give access to everything inside the system. As a result much work was done in the sixties and seventies to find ways to completely isolate a service without sharing a kernel. This work culminated with the seminal 1974 paper by Gerald J. Popek and Robert P. Goldberg [69] where they present a formal model describing the requirements for complete instruction level virtualization, i.e. *hardware virtualization*.

While hardware virtualization was in wide use on e.g. IBM mainframes since that time, it was not until 2005 that the leading commodity CPU manufacturers, Intel and AMD, introduced these facilities into their chips. In the meantime, paravirtualization had been re-introduced as a workaround to get virtual machines on these architectures, notably in [70]. While widely deployed and depended upon, the Xen project has recently been evolving its paravirtualization interface towards using hardware virtualization in e.g. PVH

[71] stating that “PVH means less code and fewer Interfaces in Linux/FreeBSD: consequently it has a smaller Trusted Computing Base (TCB) and attack surface, and thus fewer possible exploits” [72].

Another isolation mechanism is operating system-level virtualization with containers, e.g. Linux Containers (LXC) popularized in recent years by Docker, where each container represents a userspace operating environment for services that all share a kernel. The mechanism for isolating one container from another is classical process isolation, augmented with software controls such as cgroups and Linux namespaces. Containers do offer less overhead than classical virtual machines. An example where containers makes a lot of sense would be trusted in-house clouds, e.g. Google is using containers internally for most purposes [73]. We take the position that hardware virtualization is the simplest and most complete mechanism for service isolation with the best understood foundations, as formally described by Popek and Goldberg, and that this should be the preferred isolation mechanism for secure cloud computing.

B. Why Use Unikernels?

Using hardware virtualization as the preferred isolation mechanism, we take the view that there are three basic approaches we can use to deliver our requirements, namely the monolithic system/kernel approach, the microkernel approach and the unikernel approach. IaaS cloud providers will typically offer virtual machine images running Microsoft Windows or one or more flavours of Linux, possibly optimized for cloud by e.g. removing device drivers that are not needed. While specialized Linux distributions can greatly reduce the memory footprint and attack surface of a virtual machine, these are general purpose multi-process operating systems and will by design contain a large amount of functionality that is simply not needed by one single service. We take the position that virtual machines should be specialized to a high degree, each forming a single purpose micro service, to facilitate a resilient and fault tolerant system architecture which is also highly scalable.

In our next paper, we will discuss five security observations about various unikernel operating systems: choice of service isolation mechanism; use of a single address space, shared between service and kernel; no shell by default and the impact on debugging and forensics; the concept of reduced attack surface; and micro services architecture and immutable infrastructure. We shall argue that the unikernel approach offers the potential to meet all our needs, while delivering a much reduced attack surface, yet providing exactly the performance we require. An added bonus will be the reduced operating footprint, meaning a more green approach is delivered at the same time.

C. How Does This Compare to a Conventional System?

Looking at what Frederick P. Brooks Jr. suggests in [74] “Because ease of use is the purpose, this ratio of function

to conceptual complexity is the ultimate test of system design. Neither function nor simplicity alone defines a good design”, we can see where modern software systems are missing the point. The more complex a system becomes, the more overhead is introduced, leading to greater complexity and ultimately unnecessary bloat, draining performance, and exposing vulnerabilities.

Conventional cloud systems tend to be over-complicated, unnecessarily bloated, and therefore expensive to scale. Unikernels, on the other hand in [75] “Unikernels are specialized, single-address-space machine images constructed by using library operating systems”, meaning they are exactly the right size to carry out their given task — no larger, and no smaller. Our proposed approach, using unikernels, limits/enforces the software architect to use a given pattern (event-based computing using the single-responsibility-principle, service-oriented architectures, separation of data and processing, and modularity) — which is very good from a software design point of view. We are trying to get people to use “best-of-breed” patterns and thus develop better software through this limitation.

IV. DEFINITIONS

For the sake of clarity, we provide a number of useful definitions below.

Definition IV.1. Popek-Goldberg virtual machine. An environment created by the virtual machine monitor, which is functionally equivalent to the physical machine, as defined in [69]

Definition IV.2. Software library A collection of compiled program objects, $\{O_1, \dots, O_n\}$ providing symbol resolutions, i.e. linkable objects, for the symbols $\{S_1, \dots, S_n\}$

Definition IV.3. Library operating system. For a software service SW where the objects $\{O_1, \dots, O_n\}$ form the set of object files necessary and sufficient for the service to run on a given hardware platform, a library operating system is a software library that can provide exactly $\{O_1, \dots, O_n\}$

In the context of virtual machines and cloud computing it makes sense to describe the whole virtual machine as a unikernel [67], as there is in fact no classical boundary between kernel- and user space, and also because any combination of objects that can be pulled from the library operating system individually can be combined with a piece of software to form a unique whole. This piece of completely linked software will have full access to hardware on the same level as a classical kernel. However, in the context of classical operating system kernels, the library operating system designed to produce unikernels may also be called a unikernel [68] in reference to *microkernel*, *nanokernel*, *monolithic kernel* etc. It could be argued that if all the contents of a virtual machine were to be considered a unikernel, there wouldn’t really be any point in using the word “kernel”.

Definition IV.4. unikernel Given a library operating system OS, a unikernel U is defined as $U \subseteq OS$ such that U is

sufficient and necessary to provide complete linkage to some service S for a given hardware platform.

Definition IV.5. unikernel service A unikernel consisting of the minimal functionality needed to implement the contract of a given service definition.

Definition IV.6. Popek-Goldberg unikernel A single address-space popek-goldberg virtual machine, running a unikernel service

V. CONCLUSIONS

We have identified in Section 2, some ten key security issues which need to be addressed in order to address potential barriers to successful implementation of good cloud security and privacy. We have provided a range of definitions for both conventional systems and unikernels in Section 3, and have gone on to provide a comparison of these with unikernel based systems in order to identify how successful each might be in the context of dealing with security and privacy issues. We have identified how our unikernel approach might offer a better solution. In addition, due to the reduced size of the unikernel, we can benefit from reduced resource consumption.

Having now established that our proposed approach can offer significant potential benefits in a cloud scenario, we develop a detailed technical description and analysis in our second paper which can demonstrate both how our proposed system might work, and how well it might compare to conventional systems. Our third paper will consider the potential threats posed by malicious actors, and will compare how robust our proposed approach might be as compared against conventional systems.

REFERENCES

- [1] PWC, "UK Information Security Breaches Survey - Technical Report 2012," London, Tech. Rep. April, 2012. [Online]. Available: www.pwc.com/uk/bis/gov.uk
- [2] R. E. Crossler, A. C. Johnston, P. B. Lowry, Q. Hu, M. Warkentin, and R. Baskerville, "Future directions for behavioral information security research," *Comput. Secur.*, vol. 32, pp. 90–101, 2013. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S0167404812001460>
- [3] G. T. Willingmyre, "Standards at the Crossroads," *StandardView*, vol. 5, no. 4, pp. 190–194, 1997.
- [4] Cisco, "2013 Cisco Annual Security Report," Tech. Rep., 2013.
- [5] B. Duncan and M. Whittington, "Enhancing Cloud Security and Privacy: The Cloud Audit Problem," in *Cloud Comput. 2016*. Rome: IEEE, 2016, pp. 119–124.
- [6] —, "Enhancing Cloud Security and Privacy: The Power and the Weakness of the Audit Trail," in *Cloud Comput. 2016*. Rome: IEEE, 2016, pp. 125–130.
- [7] M. Huse, "Accountability and Creating Accountability: a Framework for Exploring Behavioural Perspectives of Corporate Governance," *Br. J. Manag.*, vol. 16, no. S1, pp. S65–S79, mar 2005. [Online]. Available: <http://doi.wiley.com/10.1111/j.1467-8551.2005.00448.x>
- [8] A. Gill, "Corporate Governance as Social Responsibility: A Research Agenda," *Berkeley J. Int'l L.*, vol. 26, no. 2, pp. 452–478, 2008.
- [9] C. Ioannidis, D. Pym, and J. Williams, "Sustainability in Information Stewardship: Time Preferences, Externalities and Social Co-Ordination," in *Weis 2013*, 2013, pp. 1–24.
- [10] A. Kolk, "Sustainability, accountability and corporate governance: Exploring multinationals' reporting practices," *Bus. Strateg. Environ.*, vol. 17, no. 1, pp. 1–15, 2008. [Online]. Available: <http://search.ebscohost.com/login.aspx?direct=true&db=psych&AN=2008-02303-001&site=ehost-live&kolk@uva.nl>
- [11] F. S. Chapin, G. P. Kofinas, and C. Folke, *Principles of Ecosystem Stewardship: Resilience-Based Natural Resource Management in a Changing World*. Springer, 2009. [Online]. Available: [http://books.google.co.uk/books?id=98\\$backslash\\$0Tgfd25kC](http://books.google.co.uk/books?id=98$backslash$0Tgfd25kC)
- [12] S. Arjoon, "Corporate Governance: An Ethical Perspective," *J. Bus. Ethics*, vol. 61, no. 4, pp. 343–352, nov 2012. [Online]. Available: <http://www.bravo-mag.com/25-nigerian-ceos-in-fraud-scandal/>
- [13] 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*, pp. 1–4, 2011.
- [14] A. Baldwin, D. Pym, and S. Shiu, "Enterprise Information Risk Management: Dealing with Cloud Computing," *Abdn.Ac.Uk*, pp. 257–291, 2013. [Online]. Available: http://link.springer.com/content/pdf/10.1007/978-1-4471-4189-1_8.pdf
- [15] B. Duncan and M. Whittington, "Enhancing Cloud Security and Privacy: Broadening the Service Level Agreement," in *14th IEEE Int. Conf. Trust. Secur. Priv. Comput. Commun. (IEEE Trust., Helsinki, Finland, 2015*, pp. 1088–1093.
- [16] Wikipedia, "Secure by Design." [Online]. Available: https://en.wikipedia.org/wiki/Secure_by_Design
- [17] B. Duncan, D. J. Pym, and M. Whittington, "Developing a Conceptual Framework for Cloud Security Assurance," in *Cloud Comput. Technol. Sci. (CloudCom), 2013 IEEE 5th Int. Conf. (Volume 2)*. Bristol: IEEE, 2013, pp. 120–125.
- [18] B. Duncan and M. Whittington, "Compliance with Standards, Assurance and Audit: Does this Equal Security?" in *Proc. 7th Int. Conf. Secur. Inf. Networks*. Glasgow: ACM, 2014, pp. 77–84.
- [19] S. K. Asare and A. Wright, "The Effectiveness of Alternative Risk Assessment and Program Planning Tools in a Fraud Setting," *Contemp. Account. Res.*, vol. 21, no. 2, pp. 325 – 352, 2004.
- [20] B. Duncan and M. Whittington, "Reflecting on whether checklists can tick the box for cloud security," in *Proc. Int. Conf. Cloud Comput. Technol. Sci. CloudCom*, vol. 2015-Febru, no. February. Singapore: IEEE, 2015, pp. 805–810.
- [21] E. Humphreys, "Information security management standards: Compliance, governance and risk management," *Inf. Secur. Tech. Rep.*, vol. 13, no. 4, pp. 247–255, nov 2008. [Online]. Available: <http://dx.doi.org/10.1016/j.istr.2008.10.010>
- [22] F. Albersmeier, H. Schulze, G. Jahn, and A. Spiller, "The reliability of third-party certification in the food chain: From checklists to risk-oriented auditing," *Food Control*, vol. 20, no. 10, pp. 927–935, 2009. [Online]. Available: <http://dx.doi.org/10.1016/j.foodcont.2009.01.010>
- [23] K. Prislán and I. Bernik, "Risk Management with ISO 27000 standards in Information Security," *Inf. Secur.*, pp. 58–63, 2010.
- [24] IsecT, "Information Security Frameworks from "Audit" to "Zachman"," Tech. Rep. March, 2011.
- [25] Order, "Executive Order 13636: Improving Critical Infrastructure Cybersecurity," pp. 1–8, 2013.
- [26] T. Sang, "A Log-based Approach to Make Digital Forensics Easier on Cloud Computing," *Proc. 2013 3rd Int. Conf. Intell. Syst. Des. Eng. Appl. ISDEA 2013*, pp. 91–94, 2013.
- [27] B. Duncan and M. Whittington, "The Importance of Proper Measurement for a Cloud Security Assurance Model," in *2015 IEEE 7th Int. Conf. Cloud Comput. Technol. Sci.*, Vancouver, 2015, pp. 1–6.
- [28] —, "Information Security in the Cloud: Should We be Using a Different Approach?" in *2015 IEEE 7th Int. Conf. Cloud Comput. Technol. Sci.*, Vancouver, 2015, pp. 1–6.
- [29] —, "Company Management Approaches — Stewardship or Agency: Which Promotes Better Security in Cloud Ecosystems?" in *Cloud Comput. 2015*. Nice: IEEE, 2015, pp. 154–159.
- [30] E. Zio, "Reliability engineering: Old problems and new challenges," *Reliab. Eng. Syst. Saf.*, vol. 94, no. 2, pp. 125–141, feb 2009. [Online]. Available: <http://linkinghub.elsevier.com/retrieve/pii/S0951832008001749>
- [31] M. Almorsy, J. Grundy, and A. S. Ibrahim, "Collaboration-Based Cloud Computing Security Management Framework," in *Proc. - 2011 IEEE 4th Int. Conf. Cloud Comput. CLOUD 2011*, 2011, pp. 364–371.
- [32] 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," *Proc. - 2011 IEEE World Congr. Serv. Serv. 2011*, pp. 584–588, 2011.
- [33] S. Thorpe, T. Grandison, A. Campbell, J. Williams, K. Burrell, and I. Ray, "Towards a Forensic-based Service Oriented Architecture Frame-

- work for Auditing of Cloud Logs,” in *Proc. - 2013 IEEE 9th World Congr. Serv. Serv. 2013*, 2013, pp. 75–83.
- [34] K. Ruan, J. Carthy, T. Kechadi, and I. Baggili, “Cloud forensics definitions and critical criteria for cloud forensic capability: An overview of survey results,” *Digit. Investig.*, vol. 10, no. 1, pp. 34–43, 2013. [Online]. Available: <http://dx.doi.org/10.1016/j.diin.2013.02.004>
- [35] B. Monahan and M. Yearworth, “Meaningful Security SLAs,” HP Labs, Bristol, Tech. Rep., 2008. [Online]. Available: <http://www.hpl.hp.com/techreports/2005/HPL-2005-218R1.pdf>
- [36] M. Theoharidou, N. Papanikolaou, S. Pearson, and D. Gritzalis, “Privacy Risk, Security, Accountability in the Cloud,” in *IEEE Int. Conf. Cloud Comput. Technol. Sci. Priv.*, 2013, pp. 177–184.
- [37] C. Millard, I. Walden, and W. K. Hon, “Who is Responsible for ‘Personal Data’ in Cloud Computing? The Cloud of Unknowing, Part 2,” *Leg. Stud.*, vol. 27, no. 77, pp. 1–31, 2012. [Online]. Available: <http://ssrn.com/abstract=1794130>
- [38] S. Pearson, “Taking account of privacy when designing cloud computing services,” *Proc. 2009 ICSE Work. Softw. Eng. Challenges Cloud Comput. CLOUD 2009*, pp. 44–52, 2009. [Online]. Available: <http://ieeexplore.ieee.org/xpls/abs/all.jsp?arnumber=5071532>
- [39] S. Pearson and A. Charlesworth, “Accountability as a way forward for privacy protection in the cloud,” *Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics)*, vol. 5931 LNCS, no. December, pp. 131–144, 2009. [Online]. Available: <http://citeseerx.ist.psu.edu/viewdoc/summary?doi=10.1.1.406.4389>
- [40] CSA, “Security research alliance to promote network security,” Cloud Security Alliance, Tech. Rep. 2, 1999. [Online]. Available: <http://scholar.google.com/scholar?hl=en&btnG=Search&q=intitle:Security+Guidance+Critical+Areas+of+Focus+for+0>
- [41] T. Hahn, F. Figge, J. Pinkse, and L. Preuss, “Editorial Trade-Offs in Corporate Sustainability: You Can’t Have Your Cake and Eat It,” *Bus. Strateg. Environ.*, vol. 19, no. 4, pp. 217–229, 2010.
- [42] A. Lindgreen and V. Swaen, “Corporate Social Responsibility,” *Int. J. Manag. Rev.*, vol. 12, no. 1, pp. 1–7, 2010. [Online]. Available: <http://doi.wiley.com/10.1111/j.1468-2370.2009.00277.x>
- [43] D. J. Wood, “Measuring Corporate Social Performance: A Review,” *Int. J. Manag. Rev.*, vol. 12, no. 1, pp. 50–84, 2010. [Online]. Available: <http://www.gotolisi.com/WOS:000273687300005>
- [44] T. Green and J. Pelozo, “How does corporate social responsibility create value for consumers?” *J. Consum. Mark.*, vol. 28, no. 1, pp. 48–56, 2011.
- [45] A. Christofi, P. Christofi, and S. Sisaye, “Corporate sustainability: historical development and reporting practices,” *Manag. Res. Rev.*, vol. 35, no. 2, pp. 157–172, 2012. [Online]. Available: <http://www.emeraldinsight.com/doi/abs/10.1108/01409171211195170>
- [46] N. Rahman and C. Post, “Measurement Issues in Environmental Corporate Social Responsibility (ECSR): Toward a Transparent, Reliable, and Construct Valid Instrument,” *J. Bus. Ethics*, vol. 105, no. 3, pp. 307–319, 2012.
- [47] M. A. Delmas, D. Etzion, and N. Nairn-Birch, “Triangulating Environmental Performance: What Do Corporate Social Responsibility Ratings Really Capture?” *Acad. Manag. Perspect.*, vol. 27, no. 3, pp. 255–267, 2013. [Online]. Available: <http://amp.aom.org/cgi/doi/10.5465/amp.2012.0123>
- [48] I. Montiel and J. Delgado-Ceballos, “Defining and Measuring Corporate Sustainability: Are We There Yet?” *Organ. Environ.*, vol. Advance on, pp. 1–27, 2014. [Online]. Available: <http://oae.sagepub.com/cgi/doi/10.1177/1086026614526413>
- [49] D. Bodeau, R. Graubart, L. LaPadula, P. Kertzner, A. Rosenthal, and J. Brennan, “Cyber Resiliency Metrics,” *MITRE Rep. MP 120053 Rev 1.*, no. April, pp. 1–40, 2012.
- [50] H. Carvalho, S. G. Azevedo, and V. Cruz-Machado, “Agile and resilient approaches to supply chain management: influence on performance and competitiveness,” *Logist. Res.*, vol. 4, no. 1–2, pp. 49–62, 2012. [Online]. Available: <http://link.springer.com/10.1007/s12159-012-0064-2>
- [51] M. Vieira, H. Madeira, K. Sachs, and S. Kounev, “Resilience Benchmarking,” *Resil. Assess. Eval. Comput. Syst.*, pp. 283–301, 2012. [Online]. Available: http://link.springer.com/chapter/10.1007/978-3-642-29032-9_14
- [52] A. V. Lee, J. Vargo, and E. Seville, “Developing a Tool to Measure and Compare Organizations’ Resilience,” *Nat. Hazards Rev.*, no. February, pp. 29–41, 2013.
- [53] I. Linkov, D. A. Eisenberg, M. E. Bates, D. Chang, M. Convertino, J. H. Allen, S. E. Flynn, and T. P. Seager, “Measurable Resilience for Actionable Policy,” *Environ. Sci. Technol.*, vol. 47, no. ii, p. 130903081548008, 2013. [Online]. Available: <http://pubs.acs.org/doi/abs/10.1021/es403443n>
- [54] I. Linkov, D. A. Eisenberg, K. Plourde, T. P. Seager, J. Allen, and A. Kott, “Resilience metrics for cyber systems,” *J. Risk Decis.*, vol. 33, no. 4, pp. 471–476, 2013. [Online]. Available: <http://link.springer.com/10.1007/s10669-013-9485-y>
- [55] T. Prior and J. Hagmann, “Measuring resilience: methodological and political challenges of a trend security concept,” *J. Risk Res.*, vol. 17, no. 3, pp. 281–298, 2014. [Online]. Available: <http://www.tandfonline.com/doi/abs/10.1080/13669877.2013.808686>
- [56] C. Ioannidis, D. Pym, J. Williams, and I. Gheyas, “Resilience in Information Stewardship,” in *Weis 2014*, vol. 2014, no. June, 2014, pp. 1–33.
- [57] SASB, “Sustainability Accounting Standards Board,” 2015. [Online]. Available: <http://www.sasb.org/>
- [58] R. Eccles, K. Perkins, and G. Serafeim, “How to Become a Sustainable Company,” *MIT Sloan Manag. Rev.*, vol. 53, no. 4, pp. 43–50, 2012. [Online]. Available: http://www.mapfre.com/portal/CFormacion/bol4/Docs/012013/7_{_}sustainablecompany.pdf
- [59] R. G. Eccles, I. Ioannou, and G. Serafeim, “The Impact of Corporate Sustainability on Organizational Processes and Performance,” *Manag. Sci.*, vol. 60, no. 11, pp. 2835–2857, 2014. [Online]. Available: <http://pubsonline.informs.org/doi/abs/10.1287/mnsc.2014.1984>
- [60] Trend, “2012 Annual Security Roundup: Evolved Threats in a ‘Post-PC’ World,” Trend Micro, Tech. Rep., 2012.
- [61] Verizon, N. High, T. Crime, I. Reporting, and I. S. Service, “2012 Data Breach Investigations Report,” Verizon, Tech. Rep., 2012.
- [62] R. Taylor and C. Tofts, “Business as a Control System - the Essence of an Intelligent Enterprise,” *HP Lab. Tech. Report, HPL-2003-247*, pp. 1–9, 2003.
- [63] L. Badger, D. Bernstein, R. Bohn, F. de Vaulx, M. Hogan, M. Iorga, J. Mao, J. Messina, K. Mills, E. Simmon, A. Sokol, J. Tong, F. Whiteside, and D. Leaf, “US Government Cloud Computing Technology Roadmap,” Tech. Rep., 2014. [Online]. Available: http://www.nist.gov/manuscript-publication-search.cfm?pub_{_}id=915112_{_}backslashnhttp://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.500-293.pdf
- [64] ISACA, “An Introduction to the Business Model for Information Security,” Tech. Rep., 2009.
- [65] PWC, “Information Security Breaches Survey 2010 Technical Report,” pp. 1–22, 2010. [Online]. Available: <http://www.pwc.co.uk/>
- [66] P. Meireles, “Narkive Mailinglist Archive,” 2007. [Online]. Available: <http://m0n0wall.m0n0.narkive.com/OI4NbHQQ/m0n0wall-virtualization>
- [67] A. Madhavapeddy, R. Mortier, C. Rotsos, D. Scott, B. Singh, T. Gazagnaire, S. Smith, S. Hand, and J. Crowcroft, “Unikernels: Library Operating Systems for the Cloud,” *ASPLOS ’13 Proc. eighteenth Int. Conf. Archit. Support Program. Lang. Oper. Syst.*, vol. 48, pp. 461–472, 2013. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2499368.2451167>
- [68] A. Bratterud, A.-A. Walla, H. Haugerud, P. E. Engelstad, and K. Begnum, “IncludeOS: A Minimal, Resource Efficient Unikernel for Cloud Services,” *2015 IEEE 7th Int. Conf. Cloud Comput. Technol. Sci.*, pp. 250–257, 2015. [Online]. Available: <http://ieeexplore.ieee.org/lpdocs/epic03/wrapper.htm?arnumber=7396164>
- [69] G. J. Popek and R. P. Goldberg, “Formal Requirements for Virtualizable Third Generation Architectures,” *Commun. ACM*, vol. 17, no. 7, pp. 412–421, 1974.
- [70] P. Barham, B. Dragovic, K. Fraser, S. Hand, T. Harris, A. Ho, R. Neugebauer, I. Pratt, and A. Warfield, “Xen and the Art of Virtualization,” *AGM SIGOPS Oper. Syst. Rev.*, vol. 5, no. 5, p. 164, 2003.
- [71] D. Chisnall, “Xen pvh: Bringing hardware to paravirtualization,” *Inf. IT*, 2014.
- [72] X. Project, “Xen project software overview,” 2015.
- [73] A. Verma, L. Pedrosa, M. Korupolu, D. Oppenheimer, E. Tune, and J. Wilkes, “Large-scale cluster management at Google with Borg,” *Proc. Tenth Eur. Conf. Comput. Syst. - EuroSys ’15*, pp. 1–17, 2015. [Online]. Available: <http://dl.acm.org/citation.cfm?doid=2741948.2741964>
- [74] F. P. Brooks Jr, *The Mythical Man-Month: Essays on Software Engineering, Anniversary Edition, 2/E*. Pearson Education India, 1995.
- [75] Wikipedia, “Unikernels,” 2015. [Online]. Available: <https://en.wikipedia.org/wiki/Unikernel>