Expressing and Enforcing Location Requirements in the Cloud Using Information Flow Control

Citation

Pasquier, Thomas F. J.-M., and Julia E. Powles. 2015. Expressing and Enforcing Location Requirements in the Cloud using Information Flow Control. In 2015 IEEE International Conference on Cloud Engineering (CLaw 2015), Tempe, AZ, March 9-13, 2015.

Published Version

10.1109/IC2E.2015.71

Permanent link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:37140356>

Terms of Use

This article was downloaded from Harvard University's DASH repository, and is made available under the terms and conditions applicable to Other Posted Material, as set forth at <http://nrs.harvard.edu/urn-3:HUL.InstRepos:dash.current.terms-of-use#LAA>

Share Your Story

The Harvard community has made this article openly available.
Please share how this access benefits you. [Submit a story](#).

[Accessibility](#)

Expressing and Enforcing Location Requirements in the Cloud using Information Flow Control

Thomas F. J.-M. Pasquier

Computer Laboratory, University of Cambridge
Email: tfjmp2@cam.ac.uk

Julia E. Powles

Faculty of Law, University of Cambridge
Email: jep50@cam.ac.uk

Abstract—The adoption of cloud computing is increasing and its use is becoming widespread in many sectors. As cloud service provision increases, legal and regulatory issues become more significant. In particular, the international nature of cloud provision raises concerns over the location of data and the laws to which they are subject. In this paper we investigate *Information Flow Control* (IFC) as a possible technical solution to expressing, enforcing and demonstrating compliance of cloud computing systems with policy requirements inspired by data protection and other laws. We focus on geographic location of data, since this is the paradigmatic concern of legal/regulatory requirements on cloud computing and, to date, has not been met with robust technical solutions and verifiable data flow audit trails.

I. INTRODUCTION

Cloud computing has been widely adopted as a means to provide always available, easily scalable computing resources. This has had an enormous impact on the IT industry, changing the way in which computing resources at all levels, including hardware, are designed and purchased. An innovative individual or company no longer requires a large amount of capital in order to deploy or maintain new internet-based services, and provisioning can scale (in cost and computing resources) to meet the demand for their solution. Furthermore, companies requiring timely processing of large amounts of data can buy computing time as needed, rather than maintaining expensive infrastructure that is not fully utilised. Finally, large companies can concentrate on their core business rather than on maintaining an IT infrastructure.

However, legal/regulatory requirements can create significant uncertainty around the use of cloud services [1], [2]. In this paper we investigate *Information Flow Control* (IFC) as a possible technical solution to expressing, enforcing and demonstrating compliance with policy requirements inspired by data protection and other laws. We take the geographic location of data as an example. The reconciliation of territorially-bound national law with international data flows has become a paradigmatic concern of cloud regulation. In addition to national law, cross-border regulations exist or are under discussion concerning the movement of data [3], and this is a subject of increasing political attention [4]. Other examples of policy drivers include the jurisdiction applicable to cloud-based commercial entities, requirements for encryption and anonymisation, ownership of derived data (metadata) including inferences from the data, persistence of data, etc.

A given cloud provider may offer services or process data in many different geographic locations; for example, Google has data centres in the US, EU, Chile, Singapore and Taiwan.

Data may be shipped from cloud to cloud for various purely technical reasons, such as load balancing or creating copies for reliability. In some cases, a third party cloud service, also in a different location, may be used by the cloud provider for certain specific services.

The delocalisation of data inherent in cloud propositions causes concerns over where the data of a country's citizens are stored and processed, as well as under which jurisdictions they might fall. We propose to investigate IFC as a technology for representing and enforcing such concerns, when there is sufficient clarity in the guiding law and principles that they can be expressed in policy terms. Specifically, we will investigate regulations and legal requirements with respect to *data location* and show how these policies might be expressed in IFC terms and enforced by IFC. We present briefly a cloud platform architecture based on our work on the CloudSafetyNet project¹. The examples we give are intended to illustrate the approach, rather than matching every possible nuance of the legal scenarios.

In §III we discuss the legal motivation for controlling the location of data. §II introduces key technical terms used throughout the paper. §IV introduces IFC and defines aspects of an IFC model. §V presents briefly our proof of concept platform for IFC enforcement. §VIII lists open challenges. §VI discusses how IFC could be used to define and enforce location policies. §VII outlines some related work and §IX summarises and concludes.

II. KEY TERMINOLOGY

In this section we introduce terminology that will be used in the paper. The vocabulary we use is from computer science; for example, *processing* data is understood as actively performing computation over data, as opposed to 'processing' in a data protection context. *Storage* indicates persistent storage on disc or other non-volatile device and *caching* of data indicates the temporary storage of data in main memory.

Infrastructure as a Service (IaaS): cloud service providers are responsible for the management of the network, hardware and hypervisor.

Platform as a Service (PaaS): cloud service providers are responsible for the management of all the previous categories and, in addition, the operating system (OS) and application environment.

Software as a Service (SaaS): cloud service providers manage everything on behalf of the tenants, including the applications.

¹<http://www.cl.cam.ac.uk/research/srg/opera/projects/csn/>

Cloud Provider: the company that provides some component of cloud computing, typically IaaS, PaaS or SaaS, to tenants.

Tenant: the individual/company using the service of a cloud provider to provide cloud applications to end-users, in situations where the cloud provider is not doing so directly (e.g. some Google services).

Cloud Application: software provided by a tenant and running over a cloud platform.

End-user: an individual using a cloud application provided by a tenant or a cloud provider directly.

III. LEGAL MOTIVATION FOR LOCATION CONTROL

Regulations concerning the geographic location of data are familiar in data protection law, particularly in the EU [3]. There are various reasons for a *cloud provider* to ensure the geographic location of its *tenants'* data, and reliably represent the origin of any data that may be processed in a given cloud. In this section we propose to briefly introduce some of these aspects. More expert and precise discussion can be found in [4], [5]. We highlight, in a generalised fashion, the legal and regulatory requirements (the *policies*) that might be specified in IFC terms. In later sections we describe technical enforcement mechanisms for controlling and auditing the flow of information.

The overall motivation for both data location requirements and their technical responses is to help establish a degree of certainty regarding applicable law, and therefore minimise compliance and litigation risk. Through the use of IFC, we demonstrate one way in which technical mechanisms can assist in addressing well-specified policies. This might be in conjunction with, or as a reinforcement of, contract and certification-based mechanisms [6].

As introduced in §I, cloud providers currently move data between different countries, and may use third-party services, for purely technical reasons. A prime example is the Heroku² PaaS. Heroku offer a PaaS platform for web applications, and workers³ for processing those applications' data. The Heroku PaaS runs over Amazon IaaS and provides additional services (such as databases, key-value stores, logging) through third party offerings. Furthermore, some SaaS offerings may run over Heroku. In such a scenario, it may be extremely complicated for a customer to understand where their data is located as there is no guarantee that the workers, storage or logging system are situated within the same data centre, operated by the same company, and the physical hardware may even belong to different companies.

Different aspects of regulations with implications for data location can be summarised as follows:

1. Some countries restrain the processing, storage and caching of data originating in that country to certain well-defined locations. For example, the European Union states that sensitive and confidential information should stay within its borders or certain Safe Harbor destinations that are party to the EU Data Protection Directive 95/46/EC. This has been extended, in the wake of the Edward Snowden revelations, to discussions over

a potential 'EU cloud', with analogues in BRICS and other countries [4], [5].

2. Some countries explicitly state their right to access and intercept foreign data within their borders in order to preserve their security, economic or scientific interests [7]–[9]. Often this is a matter of executive discretion, based on fluid concepts of national security and interest, and may not be subject to rigorous democratic oversight and judicial safeguards. In order to maintain the trust and custom of companies/individuals, greater certainty regarding where data are processed, stored and cached is desirable.

3. Increasingly, nations are claiming the right to prosecute and investigate foreign companies that process the data of their citizens beyond their borders [10]. However, it should be possible for a company to be able to isolate information belonging to, for example, US customers from that belonging to EU customers. If the company is not able to clearly separate the two it may be forced to release both.

4. There is a growing concern in Europe over third party use of data, particularly regarding advertisement or recommendation systems. The developments in relation to the so-called 'right to be forgotten' have led to European data protection authorities requiring that US-based companies enforce EU law across global services [11]. This could potentially mean that data should be processed differently depending on their origin. For example, we could imagine that in a not too distant future, data used in conjunction with advertising-based services in Europe must go through a differential privacy algorithm before being used.

At present, such concerns are not enforced continuously and systematically by technical means. Therefore, data mismanagement practices and scandals tend only to be revealed after the fact, and presumably there are others which never reach the public eye. It therefore seems that it would be beneficial to end-users, cloud providers, tenants and regulators if the location of data were controlled through technical means, to provide transparency, compliance and assurance that current and emerging regulatory concerns are being addressed.

IV. INFORMATION FLOW CONTROL

Standard data access control mechanisms only enforce control at certain points within a system, called *Policy Enforcement Points (PEP)*. For example, at a given PEP, it might be ensured that users are authenticated, and that their access to resources such as files and databases pass appropriate authorisation checks. Access control policy is invariably principal-specific (relating to the authenticated user) and the context of the access is not included. Moreover, once access has been granted, no further controls are possible on the subsequent use of the data [12]. What IFC offers, by contrast, is a way of enforcing and tracking where data goes and where it is received, throughout its life cycle, and not only at PEPs. IFC can also offer some guarantee about the context in which data are used. This would be a welcome development in a regulation-intensive environment, and in a situation where bugs, malicious acts or even business practices can lead to information being leaked, i.e. to data being used in an unenvisioned context, different from their original purpose.

²<https://heroku.com>

³for example, to perform batch processing or data analysis.

Though solutions are under investigation to encrypt all cloud data and to perform operations over encrypted data [13], such *homomorphic encryption* is not yet broadly applicable, and the performance cost is likely to continue to be significant and unsupportable. A further issue is that encrypting data before transfer may be an inadequate protection mechanism, because it is possible that decryption keys can be demanded by law from a provider, once data has flowed [14]. With IFC, we seek solutions in an environment where it is assumed that computations are performed over unencrypted data.

We propose the use of IFC to control data flows in the cloud and offer some guarantee about their usage over time. In this section we describe the fundamental principles of IFC necessary to understand our proposed usage. For a more in-depth treatment of implementations at the OS level see [15], [16]; and within a distributed system, see [17], [18].

A. Safe Information Exchange

In an IFC system, each entity is associated with a *security context* comprising two labels representing the *secrecy* (S) and the *integrity* (I) of the information contained within this entity. Entities can be processes, files, sockets, pipes [15], [16], database entries [19], etc. Transfer of information between entities is controlled in order to guarantee the secrecy and the integrity of the information.

The *secrecy* concern is intended to limit the propagation of sensitive information following Bell and LaPadula's [20] *no-read up, no-write down* rule. This comes from traditional, military, system-wide security classifications such as top-secret, secret, ..., unclassified. In more general terms, the *secrecy* concern represents **to** where an entity is authorised to send data. For example, secret information can only be sent to a recipient with secret or top-secret *secrecy* label.

The *integrity* concern is intended to be used to guarantee the quality of data or the authority of their source, following Biba's [21] *no-read down, no-write up* rule. In more general terms, the *integrity* concern represents **from** where an entity is authorised to receive data.

Secrecy limits where the data can be sent. For example, a process running with a location-EU secrecy label can only send information to a process cleared (and appropriately labelled) to handle EU information; while a process handling non-sensitive data which has no particular requirement would be able to send information to any location.

Integrity limits the provenance of the data a process can receive. For example, a process for which the integrity label is location-EU can only receive information from a process whose label also contains location-EU; while a process with no integrity constraints is able to receive information from anywhere.

A process with both its integrity and secrecy labels set to location-EU would only be able to receive and send information within the EU, guaranteeing geographic isolation. However, mechanisms need to exist to allow the transfer of information between processes whose labels indicate different locations: these mechanisms are known as *declassification* and *endorsement*.

B. Declassification and Endorsement

The *security context* of an entity is defined by its current labels (S , I). It cannot be modified by any external factor. The modification of a security context is strictly controlled by the cloud platform managed by the cloud provider independently of the application managed by the tenant. The secrecy and integrity of information can only be modified through well specified means; that is, label modification is part of trusted IFC management, enforced by the OS. This allows a higher level of confidence regarding how information flows within the cloud.

By changing the security context it is possible to perform two types of operation: *declassification* and *endorsement*.

Declassification involves removing secrecy constraints (tags) so that data can flow more freely within the system. This might happen if sensitive data is fully anonymised, and therefore becomes compliant with local law. Another example could be data flowing freely to various data stores for replication, after being encrypted.

Endorsement is the operation that makes some piece of information validated as trustworthy. For example, it could represent that consent has been given for processing for specified purposes, or it could be used to indicate that the data has been transformed to comply with some legal requirement, such as encrypting it before doing a cross-border data transfer, or validating that it came from an authorised source.

V. AN IFC-ENABLED CLOUD PLATFORM

As the use of cloud computing increases, so do regulatory requirements [22]. There are already strong contractual requirements regarding the quality of service (QoS) of cloud computing [23] and governments are starting to strongly regulate the use of cloud computing in strategic sectors such as the medical sector (HIPAA⁴) or government sector (e.g. UK G-Cloud⁵ or US FedRAMP⁶). There is a wealth of research on respecting performance guarantees and demonstrating that contractual requirements are met, regarding latency under high demand, reliability, etc [24]. Surprisingly, to our knowledge, little attention is being given to demonstrating that cloud providers comply with law and regulation. Cloud service providers can be certified—a manual process that captures compliance at a point in time but needs to be redone after any update to hardware or software. We expect that, in the future, greater transparency in demonstrating compliance with contracts, laws and regulations will be seen as a priority for cloud providers. Providing IFC for data location is one aspect of how providers may pursue responsible compliance.

A. IFC-Enabled Cloud Platform Overview

We illustrate the overall architecture of our system enforcing IFC in the cloud in Fig. 1. The idea behind our platform is to provide a PaaS service where the end-users of applications are provided transparent and audited assurances that the cloud provider is enforcing a specified security policy.

⁴<http://www.hhs.gov/ocr/privacy/>

⁵<https://www.gov.uk/government/publications/g-cloud-security-accreditation-application>

⁶<http://cloud.cio.gov/fedramp>

IFC is enforced within the OS by the PaaS provider. The different components of the platform are outlined below; further detail can be found in the publications indicated.

IFC OS Enforcement Mechanism [16]: enforces the IFC constraints based on labels associated with processes, pipes, sockets and files. This enforcement mechanism allows unmodified cloud applications to run under IFC constraints.

The Context Manager: is responsible for storing processes' security contexts and privileges (i.e. the privilege to change security context in order to declassify or endorse) and making them persist across machines and applications' life cycles. The context manager can also be hooked with end-users' interfaces to allow them direct control over the security context within which their cloud application instances run.

The Messaging Middleware [18]: is responsible for inter-process communication within and between (virtual) machines. Interaction with persistent storage is also done through the messaging middleware. IFC can be integrated directly in the data stores, as in [19] or through some trusted process interposing between the application and the data store, as in [15], [16].

The Audit System: IFC enforcement, in addition to providing strong assurances that policy is being enforced, also provides a data-centric log. Cloud logging systems are generally based on legacy (OS, web-server, database etc) logging systems that either fail to capture the needed information or are extremely complicated to interpret in an useful manner [25]. However, IFC logs, as provided by our platform, allow us to capture exactly the relevant information—in the case of geographic location, as described here, that information means the source of data and where this data has flowed through the cloud infrastructure. We are actively investigating processing tools to analyse the logs generated by our IFC system, to allow tenants, end-users or mutually-trusted third parties a simple way to verify compliance.

B. Running Web Applications

In previous work, we demonstrated how unmodified standard web applications can be run over an IFC-enabled platform [16]. The framework provided 'workers' (application instances) running within different security contexts and routed end-users' requests to an appropriate worker, based on the security context specified by them.

We allowed the end-user to manage its security context independently of the application. End-users may rely on the IFC-enabled platform to enforce the desired policy, regardless of the application itself. This is based on our assumption that cloud providers tend to be better resourced, more technically competent and less vulnerable to security breaches than tenants, whose applications may contain bugs and other issues. Furthermore, hardware technology could be leveraged to increase the trust placed in the provider (see §VIII). Applications may give no clear indication to end-users that third parties might use their data, so it is desirable to enforce policy at a higher level. This might be attractive to providers seeking competitive advantage in a post-Snowden era, as well as trying to minimise their exposure to compliance risk for data mishandling.

Fig. 1. Outline Architecture of an IFC-Enabled Cloud Platform.

VI. LOCATION CONTROL BY IFC

Now consider the legal concerns presented in §III. At a basic level, they can be separated into two broad categories: (1) concerns about where data are authorised to flow to; and (2) concerns about where data comes from. These two concerns are explicitly and accurately captured by IFC policies. Where data can go to is represented by *secrecy* labels and where data can come from is represented by *integrity* labels.

A. Safe Information Exchange

Suppose an international, US-headquartered cloud provider wishes both to comply with European law and to protect the privacy of its US customers' information. One of the steps towards achieving this is by enforcing a policy that EU customers' personal information should not leave the boundaries of the EU. This is achieved in our IFC system by labelling all EU data, and all entities within the EU that process or store EU data, with the *secrecy* label $S = [\text{location-EU}]$. Such labelled information cannot leave the EU (assuming correct labelling) and, with a proper audit log, the provider could demonstrate its compliance by showing data flows over time.

One of the cloud provider's US users might raise the concern that, if the US user becomes legally implicated in Europe, the provider may be forced to disclose information that includes this US user. However, the US user does not operate outside the US, and it is only through their use of the international cloud service that such data is even potentially vulnerable to European authorities. To avoid exposing the US user to European authorities, the cloud provider decides that all information from this client will be marked as $S = [\text{location-US}]$ and processed within the US by processes labelled in the same way. Not only does this guarantee the US user that such processes are located in the US, but also that they will only process US data, therefore reducing the risk that a foreign authority will be able to request such data in the absence of overreaching extraterritorial application of law, or by warrant. Through IFC, the provider will be able to demonstrate through an audit log that data has been exclusively dealt with in US territory, and that there has been no processing, storage or caching in Europe.

Similarly, *integrity* constraints can be used to specify the location from which a service is willing to accept data for processing. We have seen that there is a risk that processing data from a certain location may expose cloud providers to scrutiny by the corresponding jurisdiction of *all* the data processed by the service. In the example, the US provider, wishing to limit the exposure of its non-US users, may enforce an integrity constraint to ensure that its processing of European data, for example, is not in contact with data originating in the US, and that the outputs of its European processing do not include US data input.

By setting both *secrecy* and *integrity* tags on a given application to reflect the geographic location under which a service operates, or under which data has been generated, we can provide transparency and assurances about geographic location of data.

B. Declassification and Endorsement

Declassification and endorsement within IFC policy are used to guarantee that a certain path (and transformations) are followed before information reaches its destination. For example, a cloud service could provide storage and make no assumptions about the data being stored there. Such a store will be labelled $S = \emptyset$. European information labelled $S = [\text{location-EU}]$ could not flow there. A declassifier based in the EU could potentially provide the *declassification* from $S = [\text{location-EU}]$ to $S = \emptyset$, by encrypting the information, for example. If this was implemented, the IFC policy would ensure that unless the data has been encrypted, it cannot reach this location.

Similarly *endorsement* can be used as an attempt to limit a service exposing itself to other jurisdictions. Indeed, an endorsement from $I = [\text{location-EU}]$ to $I = [\text{location-US}]$, through the anonymisation of the data set, would allow the US service to reduce exposure to liability for manipulating EU personal information, assuming that anonymisation is sufficient in legal terms.

By using declassification and endorsement it is possible to limit the type of data that can be used to perform some task according to local regulation. For example, it is possible to build an IFC policy such that in a given country, advertising is only done over anonymised data and there are assurances provided in relation to this.

We have shown that IFC allows the isolation of data within geographic location as needed. The requirement of declassification/endorsement for cross-location data flow ensures that decisions requiring such exchanges be explicit. All cross-location data exchanges become either intentional or are prevented from occurring at all. This imposes an obligation on the application designer (i.e. the tenant) to assess the risks and the implications of such exchanges. Greater certainty about the location of data and how it is treated may also assist in determining complex jurisdictional issues about the applicable law and attendant obligations.

VII. RELATED WORK

Jarayam et al. [26] propose to provide GPS-enabled trusted platform modules (TPM) in the cloud. Their focus is on IaaS and hybrid cloud (tenant-hosted, with outsourcing to cloud providers when scaling is required), ensuring that data is not able to cross geographic borders unless it has been properly encrypted. This offers useful complementarities with the IFC approach described here, which embodies a broader approach to labelling data sent to and from applications.

Henze et al. [27] describe the need to be able to specify policy related to data pushed into the cloud. However, the policy is not strictly tied to the data but rather, different layers in the cloud stack (IaaS, PaaS, SaaS) negotiate with each other to find an offer matching the end-user's requirement.

We believe that ultimately, such data handling policy would usefully be bound to data through IFC.

Mundada et al. [28] presented an IFC-like scheme at the *virtual machine (VM) layer*, where IaaS tenants could label their data to control their dissemination within the infrastructure. In our proposed CSN platform architecture, see Fig. 1, we implement IFC at the finer granularity of the process rather than the whole VM, which is more adapted to a PaaS environment.

VIII. FUTURE WORK

Building a Trusted Platform: Recent developments in hardware technologies [29] enable new levels of hardware-rooted trust [30]. This can be leveraged to increase the level of trust that tenants have in the provider; for instance, by enabling data integrity and confidentiality to be guaranteed regardless of the platform on which the data is processed [31], or to provide assurances and logs concerning the physical location of data [26]. Furthermore, there are techniques to ensure that generated audit logs are tamper proof [32], [33] Such technology is yet to be implemented into IFC systems.

Global Tag Naming: So far we have assumed that agreement exists to use tags such as $S = [\text{location-EU}]$, $S = [\text{location-US}]$. This assumes an international naming scheme, akin to DNS (Domain Name System). If tag naming schemes were to evolve independently, negotiation and translation between domains would be needed.

IFC Implementation Layer: In this paper we presented an example based on our work [16] which aims to provide IFC enforcement to PaaS applications. However, IFC has been implemented at different layers of the software stack: application layer [34], Java VM [35], OS [15], hypervisor [28] and network [36]. Interaction between different layers of enforcement has been proposed (between Java VM and OS) [37].

We believe that some legal requirements, beyond that of location, can be implemented more effectively within different layers of the cloud stack. Future research may focus on formalising interactions between different layers of policy enforcement, and understanding at which layer a given policy is meaningful (for example, policy concerning data exchange between cloud tenants may not be meaningful within an application, but should rather be an OS enforcement concern). Understanding and formalising the interaction between different layers becomes critical, when each layer may be managed by different legal entities (as discussed in §III).

IX. CONCLUSION

Our CloudSafetyNet project is investigating IFC as a mechanism for expressing and enforcing policy within cloud computing systems. We anticipate that in the near future, only a small number of large companies will be offering IaaS services. PaaS services will be built on these IaaS offerings and in turn offer hosting of SaaS services. End-users will typically interact with cloud-hosted web applications. Cloud service provision will be increasingly regulated and it will be important to providers, users and regulators of services that there is a high degree of proactive compliance with law and regulation.

The international nature of cloud service provision presents many challenges for the framing, expression and enforcement of policies, law and regulation. In this paper we have taken a minimal example to show how data flows can be restricted to particular geographic locations. This allows transparency and accountability about the jurisdiction in which data is processed and stored, as well as enabling proactive policy enforcement adapted to the requirements of providers and users. This is a significant improvement on the present situation, where there are no existing internal mechanisms for cloud services to enforce policy or reliably demonstrate compliance. We believe that IFC is a promising technology for these purposes, worthy of further investigation.

Acknowledgement

This work was supported by the UK Engineering and Physical Sciences Research Council under grant EP/K011510 “Cloud-SafetyNet: End-to-End Application Security in the Cloud”. We acknowledge the support of Microsoft for work on Cloud Law and contributions of colleagues in the QMUL Faculty of Law. We thank our colleagues Dr Jatinder Singh and Prof Jean Bacon for sustained advice.

REFERENCES

- [1] J. Bacon, D. Eyers, T. Pasquier, J. Singh, I. Papagiannis, and P. Pietzuch, “Information Flow Control for Secure Cloud Computing,” *IEEE TNSM SI Cloud Service Management*, vol. 11, no. 1, pp. 76–89, March 2014.
- [2] C. J. Millard, Ed., *Cloud Computing Law*. OUP, 2013.
- [3] “European Commission: Proposal for a General Data Protection Regulation, 2012/0011(COD), C7-0025/12, Brussels COM(2012) 11 final,” 2012.
- [4] K. Hon, C. Millard, C. Reed, J. Singh, I. Walden, and J. Crowcroft, “Policy, Legal and Regulatory Implications of a Europe-Only Cloud,” Queen Mary University of London, School of Law, Tech. Rep., 2014. [Online]. Available: http://papers.ssrn.com/sol3/Delivery.cfm/SSRN_ID2527951_code1577160.pdf
- [5] J. Singh, J. Bacon, J. Crowcroft, A. Madhavapeddy, T. Pasquier, W. K. Hon, and C. Millard, “Regional Clouds: Technical Considerations,” University of Cambridge, Computer Laboratory, Tech. Rep. UCAM-CL-TR-863, 2014. [Online]. Available: <http://www.cl.cam.ac.uk/techreports/UCAM-CL-TR-863.pdf>
- [6] Article 29 Data Protection Working Party, “Opinion 05/2012 on Cloud Computing,” 2012.
- [7] “French Military Programming Law 2013-1168,” 2013.
- [8] “USA PATRIOT Act,” 2001.
- [9] “UK Data Retention and Investigatory Powers Act,” 2014.
- [10] “Microsoft Corp. v. United States, No. 14-2985, am. notice of appeal 2nd Cir.” 10 Sept. 2014.
- [11] Article 29 Data Protection Working Party, “Guidelines on Implementation of the Court of Justice of the European Union Judgement on “Google Spain and INC V. Agencia Espanola de proteccion de datos (AEPD) and Mario Costeja Gonzalez” C-131/12,” 2014.
- [12] R. S. Sandhu, “Lattice-based Access Control Models,” *Computer*, vol. 26, no. 11, pp. 9–19, 1993.
- [13] K.-M. Chung, Y. Kalai, and S. Vadhan, “Improved Delegation of Computation Using Fully Homomorphic Encryption,” in *Advances in Cryptology—CRYPTO 2010*. Springer, 2010, pp. 483–501.
- [14] “UK Regulation of Investigatory Powers Act,” 2000.
- [15] M. Krohn, A. Yip, M. Brodsky, N. Cliffer, M. F. Kaashoek, E. Kohler, and R. Morris, “Information Flow Control for Standard OS Abstractions,” in *21st ACM Symposium on Operating Systems Principles*, 2007, pp. 321–334.
- [16] T. F. J.-M. Pasquier, J. Bacon, and D. Eyers, “FlowK: Information Flow Control for the Cloud,” in *6th International Conference on Cloud Computing Technology and Science (CloudCom)*. IEEE, Dec 2014.
- [17] N. Zeldovich, S. Boyd-Wickizer, and D. Mazières, “Securing Distributed Systems with Information Flow Control,” in *5th USENIX Symposium on Networked System Design and Implementation*, 2008, pp. 293–308.
- [18] J. Singh, T. Pasquier, J. Bacon, and D. Eyers, “Integrating Middleware with Information Flow Control,” in *International Conference on Cloud Engineering (IC2E)*. IEEE, 2015.
- [19] D. Schultz and B. Liskov, “IFDB: Decentralized Information Flow Control for Databases,” in *8th ACM European Conference on Computer Systems (Eurosys)*. ACM, 2013, pp. 43–56.
- [20] D. E. Bell and L. J. LaPadula, “Secure Computer Systems: Mathematical Foundations and Model,” The MITRE Corp., Bedford MA, Tech. Rep. M74-244, May 1973.
- [21] K. J. Biba, “Integrity Considerations for Secure Computer Systems,” MITRE Corp., Tech. Rep. ESD-TR 76-372, 1977.
- [22] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. Katz, A. Konwinski, G. Lee, D. Patterson, A. Rabkin, I. Stoica *et al.*, “A View of Cloud Computing,” *Communications of the ACM*, vol. 53, no. 4, pp. 50–58, 2010.
- [23] R. Buyya, C. S. Yeo, and S. Venugopal, “Market-oriented Cloud Computing: Vision, Hype, and Reality for Delivering IT Services as Computing Utilities,” in *10th IEEE International Conference on High Performance Computing and Communications*. IEEE, 2008, pp. 5–13.
- [24] D. Armstrong and K. Djemame, “Towards Quality of Service in the Cloud,” in *25th UK Performance Engineering Workshop*, 2009.
- [25] R. K. Ko, M. Kirchberg, and B. S. Lee, “From System-centric to Data-centric Logging-accountability, Trust & Security in Cloud Computing,” in *Defense Science Research Conference and Expo (DSR), 2011*. IEEE, 2011, pp. 1–4.
- [26] K. R. Jayaram, D. Safford, U. Sharma, V. Naik, D. Pendarakis, and S. Tao, “Trustworthy Geographically Fenced Hybrid Clouds,” in *ACM/IFIP/USENIX Middleware*. ACM, 2014.
- [27] M. Henze, R. Hummen, and K. Wehrle, “The Cloud Needs Cross-Layer Data Handling Annotations,” in *Security and Privacy Workshops (SPW)*. IEEE, 2013, pp. 18–22.
- [28] Y. Mundada, A. Ramachandran, and N. Feamster, “Silverline: Data and network isolation for cloud services,” *Proc. of HotCloud*, 2011.
- [29] “Software Guard Extensions Programming Reference,” Intel, Tech. Rep. 329298-001US, 2013. [Online]. Available: <https://software.intel.com/sites/default/files/329298-001.pdf>
- [30] J. S. Dwoskin and R. B. Lee, “Hardware-rooted Trust for Secure Key Management and Transient Trust,” in *14th ACM conference on Computer and communications security*. ACM, 2007, pp. 389–400.
- [31] A. Baumann, M. Peinado, and G. Hunt, “Shielding Applications from an Untrusted Cloud with Haven,” in *Proceedings of the 11th USENIX conference on Operating Systems Design and Implementation*. USENIX Association, 2014, pp. 267–283.
- [32] S. A. Crosby and D. S. Wallach, “Efficient Data Structures For Tamper-Evident Logging,” in *USENIX Security Symposium*, 2009, pp. 317–334.
- [33] C. N. Chong, Z. Peng, and P. H. Hartel, *Secure Audit Logging with Tamper-resistant Hardware*. Springer, 2003.
- [34] T. F. J.-M. Pasquier, J. Bacon, and B. Shand, “FlowR: Aspect Oriented Programming for Information Flow Control in Ruby,” in *13th International Conference on Modularity*. ACM, April 2014.
- [35] W. Cheng, D. R. K. Ports, D. Schultz, V. Popic, A. Blankstein, J. Cowling, D. Curtis, L. Shriram, and B. Liskov, “Abstractions for Usable Information Flow Control in Aeolus,” in *Proc. USENIX Annual Technical Conference*, Boston, 2012.
- [36] A. Alghothami and F. Kammüller, “Network Information Flow Control: Proof of Concept,” in *Systems, Man, and Cybernetics (SMC), 2013 IEEE International Conference on*, Oct 2013, pp. 2957–2962.
- [37] I. Roy, D. E. Porter, M. D. Bond, K. S. McKinley, and E. Witchel, “Laminar: Practical Fine-grained Decentralized Information Flow Control,” *SIGPLAN Not.*, vol. 44, no. 6, pp. 63–74, Jun. 2009.