



Data Flow Management and Compliance in Cloud Computing

Citation

Thomas Pasquier. 2015. Data Flow Management and Compliance in Cloud Computing. Cloud Computing Magazine. IEEE Cloud Computing 2, no. 4: 24-32.

Published version

<https://doi.org/10.1109/mcc.2015.69>

Link

<http://nrs.harvard.edu/urn-3:HUL.InstRepos:42668751>

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 (LAA), as set forth at

<https://harvardwiki.atlassian.net/wiki/external/NGY5NDE4ZjgzNTc5NDQzMGIzZWZhMGFIOWI2M2EwYTg>

Accessibility

<https://accessibility.huit.harvard.edu/digital-accessibility-policy>

Share Your Story

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



Data Flow Management and Compliance in Cloud Computing

Jatinder Singh, Julia Powles, Thomas Pasquier, and Jean Bacon,
University of Cambridge

Cloud services are largely opaque. Information flow control aids transparency and compliance by enabling auditable, fine-grained control over data moving within and between cloud services.

As cloud computing becomes an increasingly dominant means of providing computing resources, the legal and regulatory issues associated with data in the cloud become more pronounced. These issues derive primarily from four areas: contract, data protection, law enforcement, and regulatory and common law protections for particularly sensitive domains such as health, finance, fiduciary relations, and intellectual property assets. From a technical perspective, these legal requirements all impose information management obligations on data sharing and transmission within cloud-hosted

applications and services. They might restrict how, when, where, and by whom data may flow and be accessed. These issues must be managed not only between applications, but also through the entire, potentially global, cloud supply chain.

Currently, cloud providers employ access controls to prevent unauthorized access to data and services, and containment mechanisms to prevent data leaking between tenants (those consuming cloud services) using a shared infrastructure. But these tend to be security rather than compliance focused and typically apply only at specific application, system, or user boundaries. Further, cloud services tend to be opaque and black box in nature. Despite some management tools (which depend on the service model/application), there's typically little scope for tenants to visualize, let alone specify, how data should be managed once within the cloud, or the precise circumstances in which data can be transferred.

Tenants and providers must ensure and demonstrate that they meet their legal and regulatory obligations. However, current technical mechanisms offer limited means for controlling data from afar, and insufficient tools for determining compliance and/or apportioning responsibility. This means that providers—and potentially their whole supply chain—must be trusted to act appropriately. This not only hinders accountability, but also represents a barrier to cloud adoption, particularly for personal data use and for industries such as healthcare and finance, where additional regulatory requirements pertain.

Clearly, more is required. We argue that one way forward is the development of flexible data-centric technical mechanisms that enable the visibility and control of data flows within and between cloud services. As an exemplar, we introduce our ongoing work on *information flow control* (IFC) to explore how greater technical controls over data flows can allow parties to better manage their legal obligations, improve accountability, and offer verifiable data trails for audit and compliance. Our focus here is on management and compliance with respect to civil, administrative, and criminal law obligations and responsibilities. Surreptitious actions, such as those by malicious parties and government agencies, are beyond the scope of this discussion, and require robust international policy efforts and domestic legal reforms.

Relationships and Responsibilities

Commercial cloud services are offered by *providers*, which may use third-party (*subprovider*) services as part of their supply chain. *Tenants* contract with the provider to leverage a cloud service, which they use to host applications and services for their *users*. Cloud services therefore involve a series of direct (typically contractual) relationships between users and tenants, tenants and cloud providers, and cloud providers and subproviders. Data flows typically correspond to these relationships.

Legal and regulatory considerations for data flows in the cloud revolve around four primary dimensions.

The first dimension is *contractual obligations*. Cloud services involve a number of contracts such as service-level agreements and privacy policies. These documents impose obligations for which it would be valuable to audit data flows and therefore verify compliance, detect breaches, and apportion responsibility. Although many cloud contracts appear nonnegotiable, in practice there might be room for negotiation—particularly for larger organizations.¹ A benefit of mechanisms that enable managed, auditable data flows is an increased likelihood of negotiable terms between parties, in addition to giving tenants and users greater capacity for control.

Second, *data protection laws*, adopted in many countries, place obligations and responsibilities on tenants and providers for the management of personal data. The fundamental premise of data protection is that all uses of information identifiable to an individual should be strictly regulated and controlled, with various audit mechanisms, flow and purpose restrictions, and penalties (at least, in theory) for noncompliance. Building on the Organization for Economic Cooperation and Development's 1981 Guidelines on the Protection of Privacy and Transborder Flows of Personal Data, the most influential set of laws are those concerning the European Union's Directive 95/46/EC and the new EU regulation currently under negotiation. In some jurisdictions, such as the United States, data protection regulates only some industries and types of processing, rather than providing a general schema for all data processing. In this article, we draw particularly on the EU data protection authorities' 2012 guidelines for cloud operators,² which set out rigorous requirements for technical and organizational measures that ensure transparency, purpose

specification and limitation, and data erasure. Note, however, that the technical concepts presented are general, and therefore can operate in other data protection regimes.

The third dimension is *law enforcement access* for crime/national security. For global businesses with international clients, there is increasing pressure to report government demands for data. In this article, we draw particularly on the recent high-profile Microsoft-Ireland case concerning the potential for IFC to track unauthorized (direct requests to cloud providers) and authorized (via warrant and mutual legal assistance treaties, possibly only to an accredited public institution) data flows out of a specified territory or agreed location.

Finally, *regulatory and common law protections* exist for particularly sensitive domains such as health, finance, doctor-patient, and lawyer-client relations, as well as, in a commercial context, protection of trade secrets and other intellectual property.

Within each of these areas, the physical (geo) location of data, storage, processing, and equipment are particularly pertinent considerations,^{3,4} as these are all relevant factors in determining legal jurisdiction to legislate, adjudicate, and enforce obligations on otherwise delocalized cloud operations.

Different types/instances of data can come with different obligations and responsibilities, meaning that tenants might have a number of data management constraints, and therefore require flexible technical controls.

The critical aspect that seems to be missing from existing legal and regulatory obligations, and the way that providers have responded to date, are guarantees that responsibilities are being met: demonstrable compliance through technical means. This article furthers the proposition that proper steps and policies should be in place to show explicitly how, where, when, and by whom data is accessed. This not only provides assurances to cloud consumers and to authorities, but offers a verifiable audit trail so that evidence is available if something goes wrong or responsibility needs to be apportioned. This can consolidate and be reinforced by external, general audits to move beyond higher-level reactive checks and into proactive, data-centric, and context-aware compliance.

Data Management Mechanisms: The Status Quo

These legal and regulatory issues concern data. It follows that the technical mechanisms for data management directly affect the ability of a party to meet, demonstrably, its obligations. Currently, the

well-established and commercially deployed control mechanisms tend to focus more on security. Compliance and accountability, despite being crucial cloud considerations, receive considerably less attention.⁵

The type of cloud service offering determines the capacity for management. Cloud offerings tend to be described in terms of a *service model*, which reflects the parts of the cloud stack that are managed by the provider. That is, the type of service model relates to the degree of control a tenant has over the service. Generally, tenants have limited (if any) means to influence or view those aspects managed by the provider.

For *infrastructure as a service* (IaaS), the cloud provider manages the low-level aspects (hardware, software-hardware interfaces, and isolation mechanisms), giving tenants the freedom to determine the operating system and software stack, and of course, to deploy and manage their own applications. (Note that IaaS tenants could use preconfigured virtual machine images, and often leverage other provider-managed services such as storage.) *Software as a service* (SaaS) is at the opposite end of the spectrum, where the entire application is offered and managed by the provider. This might be a university email service run by a large webmail provider, for example. SaaS tenants have far less freedom, because any management is determined by the configuration functionality offered by the application. *Platform as a service* (PaaS) offerings run tenant-provided applications on top of the cloud provider's software and service stack.

For all service models, the tenant lacks some control and visibility over various aspects of service. For users, the situation is similar to the SaaS provider-tenant scenario, in that the functionality of the (tenant) application determines any controls to manage their data. This is analogous to technical control over one's Facebook or Amazon profile, which is determined by those services' privacy-setting functionality. Further, cloud services can be composed. For instance, Dropbox (a SaaS offering for users) runs over Amazon (an IaaS offering), and Salesforce Heroku PaaS runs over Amazon IaaS. Thus, it might not be clear to users, or even tenants, which offerings comprise the service's supply chain and therefore where responsibility lies.

Isolation

Given the cloud's shared nature, a key focus has been on isolating tenants (data and processing) to prevent information leakage. Few mechanisms exist for tenants and providers to determine if or when data has been leaked as a result of some misconfiguration or software bug, or due to a security issue.

A common approach involves isolating tenants by allocating virtual machines to them, meaning that tenants share only the provider-managed hardware and hypervisor (although there might be other shared infrastructure, such as storage services). More recently, containers have enabled strong isolation of tenants over a shared operating system. The goal of isolation is to segregate tenants, protecting their data and computation, and to limit a tenant's (direct) knowledge of others.

Although strong isolation is clearly important, many applications and services will require data sharing across and outside of isolation boundaries. This might occur in order to span a range of applications or to directly access other service components and resources, storage, or billing services. Such interactions are managed through access controls.

Access Controls

Access controls regulate the actions that a *principal*—such as a human user, application, software, or process—may perform, such as read or write data, reconfigure a system, use a particular service, and so on. These actions typically relate to data.

Applications and services use access controls to manage data they hold, through authentication (identification) of the principal (“you are who you say you are”) and authorizing the actions the principal attempts to take.

Authorization involves applying a policy at a particular policy enforcement point within the application or service, considering the principals directly involved. This determines whether the action is allowed. As a simple illustration, users can log in to a social media platform (authentication), where authorization rules ensure that they can view a profile's detail only if they're “friends” with the profile owner. This relationship would be evaluated on an attempt to view a profile.

In a cloud context, this means that access controls generally govern the user–tenant, tenant–cloud provider, and provider–subprovider interactions at the interface between them. These mechanisms typically don't, of themselves, offer control beyond that point; for example, they wouldn't regulate indirect interactions between a user and subprovider. Further, the application-centric nature of many access controls renders it difficult to have a consistent management policy that can apply across the range of applications and services.

Encryption

As a data management tool, encryption provides an orthogonal form of protection to the access controls

just described. Encryption doesn't restrict physical access to data, but rather it affects its usability by making it unintelligible. Access to the (intelligible) data is regulated through the distribution of keys that enable decryption. In a cloud context, this means that if a user places encrypted data in the cloud, this data won't be accessible by the provider, or anyone else, unless they hold the requisite keys.

It is worth noting that key management is hard. Keys must be distributed to the relevant parties, and revoked (and reallocated) when conditions change. This comes with overhead, and quickly becomes unmanageable in dynamic and distributed environments. Further, there are no means for determining when, where, and by whom data was decrypted. This makes detecting leaks difficult and hinders accountability. We argue that even the distribution of encrypted data should be carefully managed, because a broken encryption scheme or compromised key at any time in the future places the data at risk.⁴

In a cloud context, encryption can protect a communication channel from eavesdropping, and protect data items that are transferred outside their boundary of control. Regarding the former, Transport Layer Security (TLS/SSL) transmission is commonplace, particularly in the post-Snowden era where large providers use encrypted communication channels, even within their datacenters.³

The practice of encrypting data before upload might be appropriate for storage services, and more generally applied to protect against surreptitious access. However, many cloud service offerings entail data processing. This generally requires the provider to have access to the customer's intelligible data (and/or the customer's keys if data is encrypted) to provide the service. There is ongoing work on homomorphic encryption, which allows operations to be performed on encrypted data without revealing the plaintext,⁶ but the current state of the art is not yet practical for use at scale.

The Need for More

These controls clearly have their place: containment prevents data leaking between tenants on a shared infrastructure, and access controls regulate the circumstances in which particular applications, services, and data may be accessed. However, these mechanisms weren't designed to demonstrate compliance with respect to contractual, legal, and regulatory obligations, nor to account for the fact that applications and the provision of cloud services necessarily entail data sharing.

There's a clear role for technical mechanisms that enable data to be managed beyond application

and system boundaries, within and between applications and cloud services, and throughout the cloud supply chain. This is particularly relevant as the cloud becomes part of wider architectures, such as for the Internet of Things.⁷ Mechanisms that facilitate visibility are also needed to determine when and where data flows, to help identify the occurrence of any leakage and/or other data obligation failures. This provides evidence indicating who might (or might not) be responsible. Finally, mechanisms are needed to flexibly deal with the subtleties of data management requirements, which might be contextual or apply only to certain data items.

IFC: Managing the Flow

Current cloud mechanisms depend heavily on trust between parties that data will be properly managed after it's transferred. As such we're developing means for managing information flows within and between cloud services and cloud-hosted applications. The aim of IFC is to complement other mechanisms by enabling visibility and control across isolation and application and system boundaries, and by providing evidence for compliance purposes.

Conceptually, the approach involves coupling data with a known management policy that's continuously enforced wherever the data flows. This is implemented by tagging data, where lightweight *tags*—representing particular management aspects/concerns—are tightly linked with data. The policy is enforced when tagged data flows (attempts to flow) between system components, within or between machines. Policy enforcement might involve allowing or preventing information flows, or transforming and/or retagging data so it will be allowed to flow. Importantly, this approach means that the data management policy need not be encoded within the software logic (that is, application-centric), but rather enables separation of that policy, which can therefore apply across different infrastructures and be managed over time.

Specifically, data and system components are linked with secrecy and integrity *labels*. A label is a set of tags, where the *secrecy* label encapsulates tags relating to confidentiality/privacy/sensitivity concerns, and the *integrity* label data quality/provenance considerations. The labels' state represents a *security context*, and a flow is permitted only if the security context of the data and components agree. This is achieved through tag-matching rules (evaluated by subset relationships) that ensure that outgoing flows that would violate the secrecy constraints are prevented and that incoming data has certain properties (integrity).⁸ More simply, secrecy con-

straints concern the decision to release data, whereas integrity constraints concern the decision to receive it. For example, data tagged as *private-bob* in its secrecy label can flow only to a component containing a *private-bob* secrecy tag in its secrecy label, which can ensure that data produced by Bob's mobile phone is processed only by a cloud application and associated services running on his behalf. Similarly, a storage component tagged as *validated-data* in its integrity label can only input data that has gone through some system validation or sanitization process and has been assigned a *validated-data* tag in its integrity label.

Policy is continuously enforced on every flow in the system, which is made possible by the fact that each component runs in a security context, defined by its labels. This facilitates a more complete audit, because the policy decision for every flow throughout the entire supply chain can be recorded, as can the security contexts for the components involved.

Importantly, IFC provides the flexibility for fine-grained, data-centric control, and is inherently dynamic, meaning that flows can be managed in accordance with changes in context. *Decentralized IFC (DIFC)*⁹ allows IFC tags to be managed in a decentralized fashion, meaning that policy can be expressed and tailored to infrastructure, application, or even the needs of an individual user.

IFC in the Cloud

IFC, although well established, has only recently been considered for cloud computing.¹⁰

We've developed a full DIFC prototype for PaaS offerings that provides a lightweight IFC policy enforcement regime to protect every OS-mediated I/O operation (within a virtual machine), with an integrated middleware responsible for policy-compliant intermachine data exchange. Technical details of this work are provided elsewhere.⁸

IFC can help ensure that data is being properly protected by acting as a safety net around the isolation mechanisms for all service models. Further, it operates continuously, beyond specific enforcement points, thus securing information flows across various parties and aspects of the infrastructure.

Other approaches to wide-scale policy enforcement, such as those leveraging "sticky" policies, are complementary because they tend to be at higher levels of abstraction, where policy definition and interpretation are comparatively heavyweight.¹¹

As a general approach, IFC offers flexibility over who manages policy. For instance, users could express IFC policy for their data, which is respected by the tenant applications and cloud infrastructure. Al-

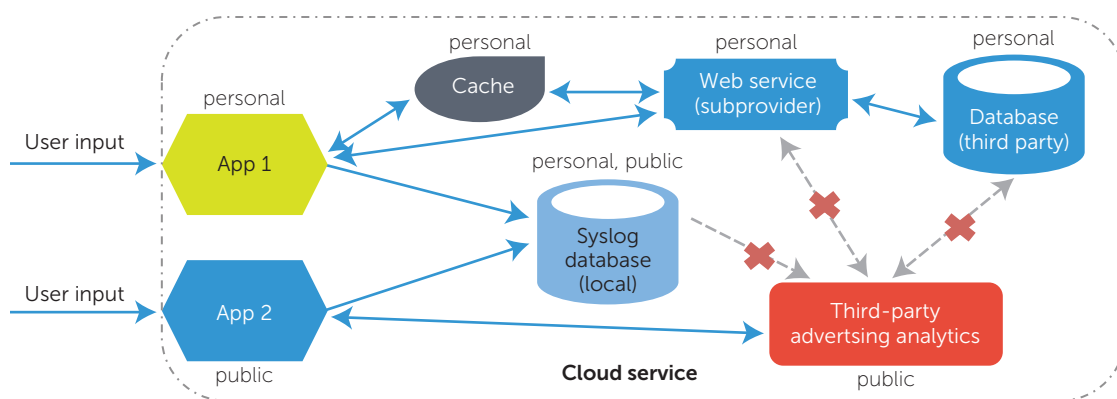


FIGURE 1. Two tenant applications. Because App 1 is dealing with personal information, the data from it can flow only to entities marked as properly dealing with personal information (labelled “personal”), and thus not to the advertising service, which isn’t certified to process personal data. Information flow control (IFC) ensures that this constraint is adhered to as data flows through the cloud, including to subproviders and other third parties. This is visible through audit, as each flow is recorded.

ternatively, IFC policy can be managed by the cloud provider without tenant involvement to guarantee, for example, data-location requirements.¹² A mixture of these is also possible.

IFC provides guarantees only above the level at which it’s enforced. For instance, our implementation enforces IFC at the operating system kernel level, and thus protects user-space flows but not lower-level aspects concerning the hardware or the hypervisor. Therefore, we assume that a cloud provider that implements IFC doesn’t actively try to circumvent its enforcement. This isn’t unreasonable, since a degree of trust in the cloud provider is implied by its use. This is reinforced by contractual relationships between tenants and cloud providers, and the role of regulators that operate in many jurisdictions. There is work on using trusted hardware components to enforce geolocation guarantees for cloud instances.¹³ A similar approach could be taken to give stronger guarantees regarding IFC policy enforcement.

Flow Controls and Legal Case Studies

IFC offers an additional security mechanism for cloud services, allowing for greater control over data flows and helping parties to meet their data management responsibilities. It also improves accountability by providing detailed logs to demonstrate compliance and/or failures. The following examples illustrate the potential of these capabilities within a legal context.

Compliance with Contractual and Data Protection Obligations

IFC is an effective mechanism for enforcing policy requirements set out in contracts and data pro-

tection regulations. We take as an example the stringent standards set out in the 2012 opinion by European data protection authorities on cloud computing, which requires cloud providers and tenants to implement technical and organizational measures towards the following aspects.²

Transparency. There must be transparency concerning all subproviders contributing to the provision of the cloud service, as well as the physical locations of all datacenters in which personal data can be processed (including storage, caching, and computation). Both of these must be auditable by the tenant, the user, or a certified third party. As Figure 1 illustrates, IFC can help improve audits by providing fine-grained, data-centric constraints over data flows within the cloud, increasing visibility to those involved, rather than relying on external physical/process audits and certifications that might go out of date, for example, because of a simple configuration change. Auditable data flows provide the means for identifying (and bounding the fallout from) misconfigurations or compliance failures. Although Figure 1 focuses on subprovider relations, location constraints can also be considered, as we’ve explored in earlier work¹² and revisit in the Microsoft-Ireland example later in this article.

Purpose specification. Personal data must not be used in ways incompatible with the purposes for which it was originally provided. IFC provides information on the paths data has taken, which can be used to indicate proper usage at a fine-grained, data-centric level. Further, if the purpose can be

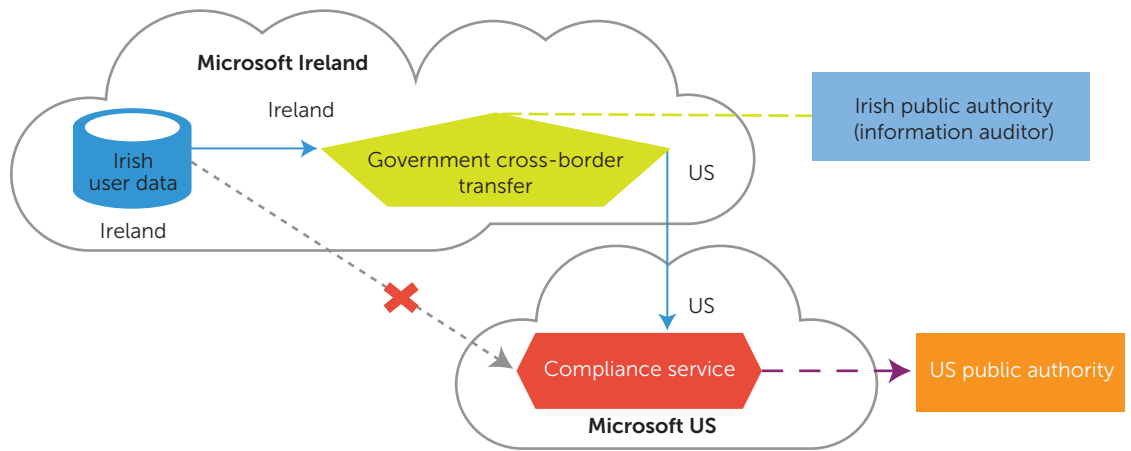


FIGURE 2. Illustration of the Microsoft-Ireland situation, where IFC could regulate the flow of information to the US for law enforcement by forcing it through a cross-border transfer process.

encapsulated within labels, it could proactively ensure that purpose constraints are met.

Data erasure. Guarantees must be in place to ensure that when personal data is no longer required, it's erased or truly anonymized. If this data can't be erased because of legal retention rules (such as tax regulations), access to it should be blocked. Because personal data can be kept in various locations, each instance (and fragment) must be erased irretrievably, including from backup, caches, and potentially log files. IFC assists with erasure concerns. First, as data flows are audited, IFC can determine where data has gone, and thus can ensure (verifiably) that the deletion requests are directed to all relevant entities. Further, if the erasure operation is a defined, encapsulated process, IFC could provide evidence that data was moved through a deletion operation (recording that it was sent to "trash").

Location and Law Enforcement

The ongoing Microsoft-Ireland legal dispute involves a search warrant issued in a US drug trial that has been claimed to extend to accessing a Microsoft customer's email data held exclusively in Ireland. Microsoft and several interested parties are appealing the decision on the basis that it involves evading the EU-US Mutual Legal Assistance Treaty (MLAT), which ordinarily governs cross-border law enforcement requests to access personal data.

Here, IFC offers the potential to regulate and audit data flow across jurisdictions by ensuring that any transfer of Irish user data to US authorities aligns to a visible, auditable, and traceable process,

as Figure 2 illustrates. The data residing in Ireland is labeled as such. For data to move across jurisdictions, the integrity tags require successful passage through a privileged government cross-border transfer (GCBT) process, which, if appropriate, changes the data's tags to allow it to flow to the US.

This example illustrates several benefits of an IFC approach. First, it shows the containment of data by location according to its integrity tags. Data can't simply flow to the US (or other jurisdictions); it must travel a certain path. Second, all flows are recorded, including the flows in and out of the GCBT process. This means the data transferred can be audited—for example, by Irish public authorities—to ensure that only data processed according to an MLAT-approved request is transferred. Individuals could also potentially verify that their information hasn't been transferred abroad. However, because this isn't always appropriate for cases of law enforcement, access to the audit log would likely be protected through access controls.

The approach we illustrate provides a stronger safety net than the current situation, where both MLAT and non-MLAT-processed requests are invisible and trust-based, with few (technical) guarantees.

Strict Processing Constraints

Personal health data is intrinsically sensitive, but it can also be useful for medical research and public health. For personal data to be used for medical research purposes, there are generally strict requirements about informed consent, anonymization processes, and appropriate ethics and governance frameworks. IFC can help ensure and provide evidence that these constraints and requirements are respected.

Figure 3 shows a simplified scenario in which tags ensure that the only means for personal data to flow to researchers is through a particular (that is, designated and approved) anonymization process, with consent a prerequisite. The medical research database is labeled such that it will only receive—or be willing to accept, for reasons of liability and responsibility—data that has passed through a particular anonymizer, again where consent is given. Individual researchers and projects are bound by the same IFC constraints as the medical research database. In transacting with the medical research database, additional controls—such as *differential privacy* techniques¹⁴ could also apply.

This example shows how IFC can assist in managing particularly sensitive data and where additional regulatory requirements pertain. Recording all flows facilitates audit. This allows consenting patients to see whether their data was actually staged for research. More complex constraints are also possible, such as consenting only to particular types of research.

Data Processing Regimes

The examples demonstrate the effect of the flow controls to essentially isolate (or, more accurately, ensure noninterference of) data transmission and processing. IFC effectively sets up a data processing regime, potentially down to the data item level, in which the flows are transparent and, importantly, the flow of data into and out of that regime can be tightly managed. If the management concerns can be specified in tags, they can be technically enforced.

Cloud providers could leverage such mechanisms to offer services based on the processing regime, rather than on infrastructure aspects such as the service model. Sector-specific clouds are feasible—for example, a financial services authority-compliant cloud in which data is guaranteed to flow only to processes certified as compliant with particular requirements. It's also feasible that tenants could define their own regimes. The extra control and transparency that IFC brings should improve levels of trust in cloud services,⁵ and therefore encourage cloud uptake. Further, effecting such controls imposes comparatively little effort on the cloud provider relative to segregated infrastructure offerings.

Discussions concerning the current conflation of legal jurisdiction and physical location are ongoing. IFC is relevant to these discussions, because it can ensure particular data management aspects based on and/or irrespective of the physical location of technical infrastructure.³

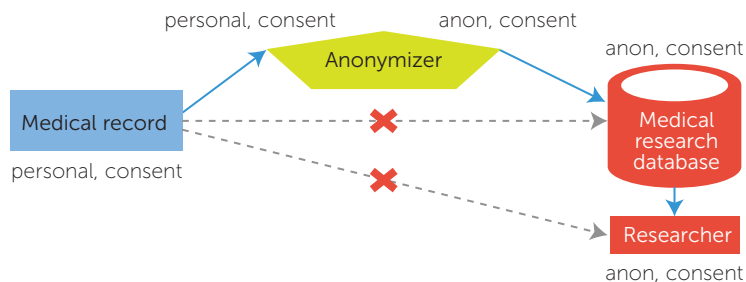


FIGURE 3. Illustration of the health data example, in which a personal medical record is transferred to a medical research database through an anonymization process. Only the pertinent labels are shown.

Although IFC offers much potential for the legal and regulatory dimensions of cloud computing, it represents ongoing research. To date, there are no commercial IFC deployments by cloud services. To make IFC mainstream, further work is needed on issues of trusted enforcement, global naming schemes, policy authoring mechanisms, tag sensitivity and management, to name a few.⁸ All of the examples discussed in this article have nevertheless been implemented in our prototype, demonstrating the approach's feasibility. Again, our vision for IFC is to complement other management technologies, be they well-established or the subject of ongoing research.

Note the research directions in cloud computing are toward smaller clouds, reducing the size and scope of cloud deployments.^{15,16} Such work is particularly relevant to the emerging Internet of Things.⁷ Smaller clouds can be dynamically and temporarily migrated—for example, from a phone (collecting data while mobile) to a larger provider (for more complex processing). They also tend toward specific purposes—for example, one for an individual's fitness data, another for that person's home appliances. As clouds become smaller, their functions become more tailored, obvious, and explicit. This might facilitate more focused and informed policy decisions. As more data flows occur to, from, and within and between these clouds, mechanisms for consistent and continuous enforcement of data management policy across applications and infrastructure become crucial.

Acknowledgments

Those of us from the Computer Laboratory are supported through the UK EPSRC (grant EP/K011510), and the Microsoft Cloud Computing Research Centre. We thank our colleagues at the Centre for Commercial Law Studies, Queen Mary University of London, for their input.

References

1. C.J. Millard, ed., *Cloud Computing Law*, Oxford Univ. Press, 2013.
2. European Commission, “Opinion 05/2012 on Cloud Computing,” 2012; http://ec.europa.eu/justice/data-protection/article-29/documentation/opinion-recommendation/files/2012/wp196_en.pdf.
3. K. Hon et al., *Policy, Legal and Regulatory Implications of a Europe-Only Cloud*, tech. report, School of Law, Queen Mary Univ. of London, 2014; http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2527951.
4. J. Singh et al., *Regional Clouds: Technical Considerations*, tech. report UCAM-CL-TR-863, Computer Laboratory, Univ. of Cambridge, 2014; www.cl.cam.ac.uk/techreports/UCAM-CL-TR-863.pdf.
5. Cloud Security Alliance, “Security Guidance for Critical Areas of Focus in Cloud Computing v3.0,” 2011; <https://cloudsecurityalliance.org/research/security-guidance>.
6. M. Naehrig, K. Lauter, and V. Vaikuntanathan, “Can Homomorphic Encryption Be Practical?” *Proc. 3rd ACM Workshop on Cloud Computing Security (CCSW)*, 2011, pp. 113–124.
7. J. Singh et al., “Twenty Security Considerations for Cloud-Supported Internet of Things,” to be published in *IEEE Internet of Things J.*
8. T. Pasquier et al., “CamFlow: Managed Data-Sharing for Cloud Services,” arXiv:1506.04391, 2015; <http://arxiv.org/abs/1506.04391>.
9. A.C. Myers and B. Liskov, “A Decentralized Model for Information Flow Control,” *Proc. 17th Symp. Operating Systems Principles (SOSP)*, 1997, pp. 129–142.
10. J. Bacon, “Information Flow Control for Secure Cloud Computing,” *IEEE Trans. Network and Service Management*, vol. 11, no. 1, 2014, pp. 76–89.
11. S. Pearson and M.C. Mont, “Sticky Policies: An Approach for Managing Privacy across Multiple Parties,” *Computer*, vol. 44, no. 9, 2011, pp. 60–68.
12. T. Pasquier and J. Powles, “Expressing and Enforcing Location Requirements in the Cloud Using Information Flow Control,” *Proc. Int’l Conf. Cloud Engineering (IC2E 15)*, 2015, pp. 410–415.
13. K.R. Jayaram et al., “Trustworthy Geographically Fenced Hybrid Clouds,” *Proc. 15th Int’l Middleware Conf.*, 2014, pp. 37–48.
14. C. Dwork, “Differential Privacy,” *Automata, Languages and Programming*, Springer, 2006, pp. 1–12.
15. M. Satyanarayanan et al., “The Case for VM-Based Cloudlets in Mobile Computing,” *IEEE Pervasive Computing*, vol. 8, no. 4, 2009, pp. 14–23.
16. J. Crowcroft et al., “Unclouded Vision,” *Distributed Computing and Networking*, LNCS 6522, Springer, 2011, pp. 29–40.

JATINDER SINGH is a senior research associate in the Computer Laboratory at the University of Cambridge. His research interests concern management control in distributed systems, particularly regarding cloud and the Internet of Things. Singh has a PhD in computer science from the University of Cambridge. Contact him at Jatinder.Singh@cl.cam.ac.uk.

JULIA POWLES is a lawyer and PhD student at the Centre for Intellectual Property and Information Law at the University of Cambridge. Her research interests include the law and politics of information-based assets, from data privacy to patent law. Powles has a master's degree in law from the University of Oxford. Contact her at jep50@cam.ac.uk.

THOMAS PASQUIER is a PhD student and research assistant in the Computer Laboratory at the University of Cambridge. His research interests include identity and data management in distributed systems, particularly the cloud. Pasquier has an MPhil in computer science from the University of Cambridge. Contact him at tfjmp2@cam.ac.uk.

JEAN BACON is a professor of distributed systems at the University of Cambridge, where she leads the Opera research group. Her research interests include open, large-scale, secure, widely-distributed systems. Bacon's most recent degree is an honorary doctorate in computer science from the Open University. Contact her at Jean.Bacon@cl.cam.ac.uk.



Selected CS articles and columns are also available for free at <http://ComputingNow.computer.org>.