



# Toward Accountability in the Cloud

Siani Pearson • HP Labs

Accountability is likely to become a core concept in both the cloud and in new mechanisms that help increase trust in cloud computing. These mechanisms must be applied in an intelligent way, taking context into account and avoiding a one-size-fits-all approach.

**T**he US National Institute of Standards and Technology defines cloud computing as “a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources (for example, networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.” In short, the cloud offers a huge potential both for efficiency and new business opportunities (especially in service composition), and is almost certain to deeply transform our IT. Not only will cost savings occur due to economies of scale on the service provider side and pay-as-you-go models, but business risk also decreases because organizations have less need to borrow money for upfront investment in infrastructure.

However, to help realize these benefits, we must address two primary barriers: lack of consumer trust and the complexity of compliance. Here, I argue that the concept of accountability is key to addressing these issues.

### Barriers to Cloud Adoption

Lack of consumer trust is commonly recognized as a key inhibitor to moving to software-as-a-service (SaaS) cloud models. People have increasing expectations that companies with which they share their data will protect it and handle it responsibly. Furthermore, compared to traditional server architectures, cloud consumers are more concerned about their data's integrity, security, and privacy as focus shifts from server health to data protection. However,

current terms of service push risk back on consumers and offer little remediation or assurance. Potential cloud customers perceive a lack of transparency and relatively less control than with traditional models, which is of particular concern in the context of sensitive information. Some cases have arisen in which cloud service providers (CSPs) have been forced by subpoena to hand over data stored in the cloud, and a fear persists that governments might get access to information stored in servers within their countries. Moreover, it isn't clear what would happen if things went wrong. Would providers notify users if a privacy breach occurred? Who would be at fault in such cases? Working out how victims could obtain redress is complex and hard to ascertain. It's also difficult to determine whether data has been properly destroyed (as it should be, for example, in the case of a CSP's bankruptcy or if a customer wishes to switch to a different CSP). So, people are concerned about weak trust relationships along the chain of service provision, especially as regards on-demand models in which users might have to find CSPs quickly; in such cases, trust won't necessarily be transitive along the chain.

A second barrier to cloud migration is the difficulty CSPs can have with compliance across geographic boundaries. Dataflows tend to be global and dynamic. Location matters from a legal viewpoint, leading to regulatory complexity. Complying with legislation can be difficult with regard to trans-border dataflow requirements and determining which laws apply and which courts should preside.

Issues such as unauthorized secondary data usage and inappropriate data retention are also difficult to address.

These two issues — trust and the complexity of compliance — are closely linked. CSPs have both legal and ethical obligations to ensure privacy and protect data and thereby demonstrate their services' trustworthy nature.

This higher risk to privacy and security in cloud computing is a magnification of issues faced in subcontracting and offshoring. Consumers aren't the only ones worried about privacy and security concerns in the cloud.<sup>1</sup> The European Network and Information Security Agency (ENISA)'s cloud computing risk assessment report states "loss of governance" as a top risk of cloud computing, especially for infrastructure as a service (IaaS).<sup>2</sup> "Data loss or leakages" is also one of the top seven threats the Cloud Security Alliance (CSA) lists in its *Top Threats to Cloud Computing* report.<sup>3</sup> The cloud's autonomous and virtualized aspects can bring new threats, such as cross-VM (virtual machine) side-channel attacks, or vulnerabilities due to data proliferation, dynamic provisioning, the difficulty in identifying physical servers' location, or a lack of standardization. Although service composition is easier in cloud computing, some services might have a malicious source. All these privacy and security risks might actually decrease, however, if users move from a traditional IT model to a cloud model with CSPs who have expertise in privacy and security.

Accountability can help us tackle these challenges in trust and complexity. It's especially helpful for protecting sensitive or confidential information, enhancing consumer trust, clarifying the legal situation in cloud computing, and facilitating cross-border data transfers. My focus here is on data-protection issues in the cloud. The term "data protection"

has more of a privacy focus in Europe but a broader data security context in the US. I focus primarily on privacy, but some of these issues transcend personal data handling and generalize to other types of data, beyond privacy concerns.

### **What Is Accountability?**

For several years, computer science has used the term accountability to refer to a narrow and imprecise requirement that's met by reporting and auditing mechanisms. Here, however, I use the term in the context of corporate data governance. Accountability (for complying with measures that give effect to practices articulated in given guidelines) has been present in many core frameworks

demonstrate that they are, and will be, compliant with requirements that EU Data Protection Authorities (DPAs) have defined for transferring data outside the EU. More recently, several groups have highlighted accountability's significance and utility in introducing innovations to the current legal framework in response to globalization and new technologies (see "The Future of Privacy," from the Article 29 Working Party,<sup>7</sup> its opinion of July 2010,<sup>8</sup> and the Madrid Resolution's global data protection standards, which the International Conference of Data Protection and Privacy Commissioners adopted in October 2009).

The Galway project started by privacy regulators and privacy

---

## **Organizations must employ responsible decision making and report, explain, and be answerable for decisions they've made.**

---

for privacy protection, most notably the Organization for Economic Cooperation and Development (OECD)'s privacy guidelines (1980),<sup>4</sup> Canada's Personal Information Protection and Electronic Documents Act (2000),<sup>5</sup> and Asia Pacific Economic Cooperation (APEC)'s Privacy Framework (2005).<sup>6</sup>

More recently, region block governance models are evolving to incorporate accountability and responsible information use, and regulators are increasingly requiring that companies prove they're accountable. In particular, legislative authorities are developing frameworks such as the EU's Binding Corporate Rules (BCRs) and APEC's Cross Border Privacy Rules to provide a cohesive and more practical approach to data protection across disparate regulatory systems. For example, BCRs require that organizations

professionals defines accountability in the context of these latest regulations:

Accountability is the obligation to act as a responsible steward of the personal information of others, to take responsibility for the protection and appropriate use of that information beyond mere legal requirements, and to be accountable for any misuse of that information.<sup>9</sup>

Central components of this notion are transparency, responsibility, assurance, and remediation. With regard to responsibility, organizations must demonstrate that they've acknowledged and assumed responsibility, in terms of both having appropriate policies and procedures in place and promoting good practices that include correction and remediation for failure and misconduct. Such organizations must employ responsible decision making and, in particular,

report, explain, and be answerable for the consequences of decisions they've made with regard to data protection.

### Retrospective vs. Prospective Accountability

Some have argued that to provide accountability, we must shift from hiding information to ensuring that only appropriate uses occur.<sup>10</sup> Information usage should be transparent so that we can determine whether a use is appropriate under a given set of rules. CSPs can maintain a history of data manipulation and inferences (providing transparency) that can then be checked against the policies that govern them. This provides retrospective accountability – that is, if actor A performs action B, then

static, and in which the data controller must conduct an ongoing assessment of harm and a privacy review process throughout the contractual or service provision chain.

Broadly speaking, an accountability approach in accordance with current regulatory thinking requires organizations to

- commit to accountability and establish policies consistent with recognized external criteria;
- provide transparency and mechanisms for individual participation, including sharing these policies with stakeholders and soliciting feedback;
- use mechanisms to implement these policies, including clear documentation and communication

must take into account the entire life cycle of personal data usage, including deletion. Companies must think about not only what data they'll collect and how they plan to use it but also what potential harm the proposed use of that data could cause to individuals. Without going into the intricacies of legal ownership, the data subject is normally, in a fundamental sense, the real owner of his or her data and is ultimately the person harmed in the event of a privacy breach; this person should be empowered and supported. For example, if you're tracking someone's behavior online, under an accountability approach you might provide clear notice that tracking is happening, an explanation of how you plan to use the data, and a mechanism for individuals to opt out of tracking and request that you delete previous tracking data about them.

## Accountability places a legal responsibility on an organization to ensure that the contracted partners to whom it supplies data are compliant.

we can review B against a predetermined policy to decide if A has done something wrong and so hold A accountable.

We must extend this approach to include prospective effects because, for instance, the environment might change – for instance, new risks might arise for data subjects because the service provisioning chain alters, the location of the physical servers storing or processing data changes, a CSP has new ownership, or a new type of attack occurs. Reducing the risk of disproportionate harm to data subjects thereby reduces negative consequences for data controllers. To do this, we must build in processes and reinforce good practices such that liability doesn't arise in the first place.<sup>11</sup> This is a reflexive privacy process that isn't

(encompassing an organization's ethical code), support from all levels within the organizational structure, tools, training, education, ongoing analysis, and updating;

- allow validation – that is, provide means for external enforcement, monitoring, and auditing; and
- provide mechanisms for remediation, which should include event management (such as dealing with data breaches) and complaint handling.

I argue that we can extend the third item in this list to encompass both preemptive approaches (to assess risk and avoid privacy harm) and reactive approaches that provide transparency and auditing. These privacy policies and mechanisms

### Data Stewardship

A closely related notion to accountability is *data stewardship*.<sup>12</sup> In a cloud model, many different cloud providers in an ecosystem consume IT. Understanding such ecosystems can be challenging and we must make a paradigm shift in our thinking. Security and privacy management evolves into an information stewardship problem – that is, how organizations can properly look after and protect information (in a broader sense than just personal data) on behalf of the data owners, subjects and third parties. In the cloud, establishing risks and obligations, implementing appropriate operational responses, and dealing with regulatory requirements will be more difficult than with traditional server architectures. The notions of transparency and assurance are more relevant and data controllers and CSPs must ensure “chains of accountability.” Accountability places a legal responsibility on an organization that uses personal information to ensure that the contracted

partners to whom it supplies this data are compliant, wherever they might reside worldwide. So, the communities responsible for data stewardship (who are typically organizational IT security, legal, operations, and compliance staff) place responsibilities and constraints on other individuals or on how systems operate, and these constraints are met along the chain of provision.

### Intelligent Accountability

Baroness O'Neill first proposed the idea of "intelligent accountability" as a means to provide greater accountability without damaging professional performance in her 2002 Reith Lectures on "A Question of Trust" ([www.bbc.co.uk/radio4/reith2002/](http://www.bbc.co.uk/radio4/reith2002/)). She argued that much of what individuals and organizations must account for isn't easily measured and can't be reduced to a set of stock performance indicators. O'Neill said that intelligent accountability "requires more attention to good governance and fewer fantasies about total control" and that "good governance is possible only if institutions are allowed some margin for self-governance of a form appropriate to their particular tasks."

We must introduce accountability in an intelligent way, or trust won't increase and the overall effect could be quite negative with regard to the increased administrative burden. As relates to the cloud, intelligent accountability could involve

- moving away from "box checking" and static privacy mechanisms;
- assessing potential harms to data subjects before exposing data to risks; this would be part of ongoing risk assessment and mitigation, for which *privacy impact assessments* (PIAs) are one important tool;
- allowing organizations more flexibility in how they provide data protection so that they can

use internal mechanisms and controls that make the most sense for their business situation, rather than a one-size-fits-all prescriptive set of rules;

- employing various degrees of accountability; it might be that more stringent standards and tests for accountability could facilitate proof of CSPs' readiness to engage in certain activities (such as those that involve processing highly sensitive data) or even relieve them of certain administrative burdens (such as renotification of minor changes in processing); and
- developing clever, automated analysis, automated internal policy enforcement, and other technologies to enhance enforcement and avoid increasing the human burden.

As an integral part of an intelligent accountability approach, organizations will need to spend time and resources analyzing what it means to them and gaining management support for implementing necessary changes.

### How to Provide Accountability in the Cloud

Accountability promotes the implementation of practical mechanisms whereby legal requirements and guidance are translated into effective data protection. Legislation and policies tend to apply at the data level, but mechanisms for accountability can exist at various levels, including system and data levels. Solution builders could provide data controllers with a toolbox of measures to enable the construction of custom-built solutions whereby controllers could tailor measures to their context (taking into account the systems involved, the type of data, dataflows, and so on).

We can codesign legal mechanisms, procedures, and technical

measures to support this approach. We might integrate design elements to support

- prospective (and proactive) accountability, using preventive controls and
- retrospective (and reactive) accountability, using detective controls.

Preventive controls can help mitigate whether an action continues or takes place at all (for example, an access list that governs who can read or modify a file or database, or network and host firewalls that block all but allowable activity). The cloud is a special example of how businesses must assess and manage risk better.<sup>13</sup> Preventive controls for the cloud include risk analysis and decision support tools, policy enforcement (for example, machine-readable policies, privacy-enhanced access control, and obligations), trust assessment, obfuscation techniques, and identity management.

Organizations can use detective controls to identify privacy or security risks that go against policies and procedures (for example, intrusion-detection systems, policy-aware transaction logs, language frameworks, and reasoning tools). Detective controls for the cloud include auditing, tracking, reporting, and monitoring. In addition, corrective controls are necessary (such as an incident management plan or dispute resolution) that can help fix an undesired outcome that's already occurred. These controls complement each other: a combination would ideally be required for accountability.

Provision of accountability wouldn't occur only via procedural means, especially for the cloud, which is an automated and dynamic environment: technology can play an important role in enhancing solutions by enforcing policies and providing decision support, assurance, security, and so on.



Procedural measures for accountability include determining CSPs' capabilities before selecting one, negotiating contracts and service level agreements (SLAs), restricting the transfer of confidential data to CSPs, and buying insurance. Organizations should also appoint a data-protection officer, regularly perform privacy impact assessments on new products and services, and put mechanisms in place to allow quick response to data subject access and deletion requests.

Technical measures for accountability can include encryption for data security mitigation, privacy intermediaries, and agents to help increase trust. We must also be able to rely on infrastructure to maintain appropriate separations, enforce policies, and report information accurately. At HP Labs, we're investigating how to build and exploit trusted virtualized platforms with precisely these properties.

Another mechanism we're researching is the use of sticky policies, in which machine-readable policies (defining allowed usage and associated obligations) are attached to data within the cloud and travel with it. Other mechanisms include risk assessment, decision support, obfuscation in the cloud, and policy translation from higher-level policies to machine-readable ones that are enforced and audited. We don't have the space here to describe all this work, so I'll just briefly outline three examples of our research.

First, we've worked with HP Privacy Office to develop and deploy a tool called the HP Privacy Advisor that takes employees through a series of dynamically generated contextual questions and outputs the risk for privacy compliance in any new product, service, or program. It encodes HP's privacy rulebook and other sources and provides privacy by design guidance. An associated workflow with privacy managers ensures that employees address the

suggested actions mitigating these risks.

The Cloud Stewardship Economics project is defining mathematical and economic models of the cloud ecosystem and the different choices cloud stakeholders face. The goal is to help cloud consumers, providers, regulators, and other stakeholders explore and predict the consequences of different policies, assurance mechanisms, or even ways of regulating accountability. This can facilitate consumer choice; as chains of providers become more complex, the models can highlight how and why evidence sharing is likely to provide necessary assurance.

Finally, we're working to achieve accountability using contractual assurances along the service provision chain from CSPs to accountable organizations, enhanced on the technical side by enforcement of corresponding machine-readable policies propagated with (references to) data through the cloud, integrated risk assessment, assurance, and auditing. By these means, the accountable organizations can ensure that all who process data observe their obligations to protect it, irrespective of where that processing occurs.

### Moving Forward


Current regulatory structure places too much emphasis on recovering and not enough on trying to get organizations to proactively reduce privacy and security risks. New data governance models for accountability can provide a basis for providing data protection when people use cloud computing. Accountability is becoming more integrated into our self-regulatory programs as well as future privacy and data protection frameworks globally. If CSPs don't think beyond mere compliance and demonstrate a capacity for accountability, regulations will likely develop that could be difficult to follow and might stifle innovation; a

backlash might also arise from data subjects.

Strengthening an accountability approach and making it more workable by developing intelligent ways to apply accountability and information stewardship is a growing challenge. It goes beyond traditional approaches to protect data (such as security and the avoidance of liability) in that it includes complying with and upholding values and obligations, and enhancing trust. Hewlett-Packard is actively working in this area to produce practical solutions, both on the policy (HP Privacy Office) and technical fronts (HP Labs).

At present we're just starting to see some technical work emerging from other parties in this area. The CSA – a non-profit organization formed to promote the use of best practices for providing security assurance within cloud computing – has a Governance, Risk Management, and Compliance (GRC) stack that includes two very relevant activities: CloudAudit, which aims to provide a technical foundation to enable transparency and trust in private and public cloud systems, and the Trusted Cloud Initiative, which is working toward certifying "trusted clouds." HyTrust Appliance is a hypervisor consolidated log report and policy-enforcement tool that logs from a system perspective. The Commonwealth Scientific and Industrial Research Organisation (CSIRO) has produced a prototype in which CSPs are accountable for faulty services. The Computer Sciences Corporation (CSC) is developing a CloudTrust protocol that will promote CSP transparency.


**A**t HP Labs, our broader vision is to deliver seamless, secure, context-aware experiences for a connected world. The richness, choice and convenience of how we interact with our

devices and a pervasive computing environment will be enhanced. At the same time, we want this to be safe and ultimately controlled by end users. We've been introducing and will continue to research new innovative techniques to uphold HP's ethics and values internally and demonstrate this to our stakeholders and customers. 

## References

1. R. Gellman, "Privacy in the Clouds: Risks to Privacy and Confidentiality from Cloud Computing," *World Privacy Forum*, 2009; [www.worldprivacyforum.org/pdf/WPF\\_Cloud\\_Privacy\\_Report.pdf](http://www.worldprivacyforum.org/pdf/WPF_Cloud_Privacy_Report.pdf).
2. *Cloud Computing: Benefits, Risks and Recommendations for Information Security*, D. Catteddu and G. Hogben, eds., ENISA, Nov. 2009; [www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment/at\\_download/fullReport](http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment/at_download/fullReport).
3. *Top Threats to Cloud Computing*, version 1.0, tech. report, Cloud Security Alliance, Mar. 2010; <https://cloudsecurityalliance.org/topthreats/csathreats.v1.0.pdf>.
4. *Guidelines Governing the Protection of Privacy and Transborder Flow of Personal Data*, Organization for Economic Cooperation and Development (OECD), 1980.
5. *Personal Information Protection and Electronic Documents Act (PIPEDA)*, Canada, schedule 1, principle 1, 2000.
6. *APEC Privacy Framework*, Asia-Pacific Economic Cooperation, 2005; [www.apec.org/Groups/Committee-on-Trade-and-Investment/~media/Files/Groups/ECSG/05\\_ecsg\\_privacyframewk.ashx](http://www.apec.org/Groups/Committee-on-Trade-and-Investment/~media/Files/Groups/ECSG/05_ecsg_privacyframewk.ashx).
7. "The Future of Privacy: Joint Contribution to the Consultation of the European Commission on the Legal Framework for the Fundamental Right to Protection of Personal Data," EU Article 29 Working Party, WP168, Dec. 2009; [http://ec.europa.eu/justice/policies/privacy/docs/wpdocs/2009/wp168\\_en.pdf](http://ec.europa.eu/justice/policies/privacy/docs/wpdocs/2009/wp168_en.pdf).
8. "Opinion 3/2010 on the Principle of Accountability," EU Article 29 Working Party, WP173, July 2010; [http://ec.europa.eu/justice/policies/privacy/docs/wpdocs/2010/wp173\\_en.pdf](http://ec.europa.eu/justice/policies/privacy/docs/wpdocs/2010/wp173_en.pdf).
9. *Galway Project Plenary Session Introduction*, Galway Project, 28 Apr. 2009, p 5.
10. D. Weitzner et al., "Information Accountability," *Comm. ACM*, vol. 51, no. 6, 2008, pp. 82–87.
11. S. Pearson and A. Charlesworth, "Accountability as a Way Forward for Privacy Protection in the Cloud," *Proc. 1st Int'l Conf. Cloud Computing*, LNCS 5931, M.G. Jaatun, G. Zhao, and C. Rong, eds., 2009, pp. 131–144.
12. D. Pym and M. Sadler, "Information Stewardship in Cloud Computing," *Int'l J. Service Science, Management, Engineering and Technology*, vol. 1, no. 1, 2010, pp. 50–67.
13. A. Baldwin and S. Shiu, *Managing Digital Risk: Trends, Issues, and Implications for Business*, tech. report, Lloyds 360 Risk Insight, 2010.

**Siani Pearson** is a senior researcher in the Cloud and Security Research Lab at HP Labs Bristol. Her current research focus is on privacy enhancing technologies, accountability and the cloud. Pearson has a PhD in artificial intelligence from the University of Edinburgh. She's a technical lead on regulatory compliance projects with the HP Privacy Office and HP Enterprise Services and on the collaborative TSB-funded Ensuring Consent and Revocation project. Contact her at [siani.pearson@hp.com](mailto:siani.pearson@hp.com).

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