

Modelling Cloud Computing Architecture Without Compromising Privacy: *A Privacy by Design Approach*



May 2010

NEC Company, Ltd.
and
Information and Privacy Commissioner,
Ontario, Canada

Acknowledgements

The authors are grateful to Min-Yu Hsueh, Toshikazu Fukushima, Jun Du, and Takahiro Sugiyama at NEC for their support in preparing this paper. Michelle Chibba and Vance Lockton, Policy Department staff at the Information and Privacy Commissioner's Office, Ontario, Canada are also acknowledged for their input to this paper. Thanks also to Hao Lei for valuable discussions during the development of the ideas herein presented.



**Information and Privacy Commissioner,
Ontario, Canada**

2 Bloor Street East
Suite 1400
Toronto, Ontario
M4W 1A8
Canada

416-326-3333
1-800-387-0073
Fax: 416-325-9195
TTY (Teletypewriter): 416-325-7539
Website: www.ipc.on.ca
Privacy by Design: www.privacybydesign.ca

Table of Contents

Foreword	1
1 Introduction	2
2 Cloud Computing	3
2.1 Cloud Computing Delivery Models	3
2.2 Deployment Models	4
2.3 Cloud Stack	5
3 Data in the Cloud – The Blurred Security Perimeter	5
4 Privacy by Design and the Positive-Sum Paradigm	6
4.1 The <i>Privacy by Design</i> Principles.....	7
5 PbD Cloud Computing	10
5.1 Protecting Privacy and Maintaining Appropriate Access.....	11
5.2 Protecting Privacy and Maintaining Data Integrity	14
6 Additional Privacy and Security Considerations	16
7 Conclusions	17
Bibliography	19
About the Authors	22

Foreword

As the Information and Privacy Commissioner of Ontario, Canada, one of my mandates is to raise awareness of privacy-related issues involved in emerging technologies or new programs that may impact one's privacy. I am pleased to have been approached by NEC to provide comment and input into this white paper regarding privacy-protective architecture design for Cloud computing, as my Office has undertaken significant research efforts into this area (see, for instance, our *Privacy in the Clouds* white paper). It is rewarding to see researchers at companies such as NEC considering privacy early on in their work.

My message has always been that respecting privacy should not present an impediment to ensuring the security of, and appropriate access to, data. I call this a 'positive-sum' paradigm – where optimized system functionality and privacy can be delivered in unison – a result that is achieved by incorporating privacy into the design phase of technologies, namely, by employing *Privacy by Design*. It is true that the broad range of potential uses of the Cloud come, as do many technologies, with potential threats to data privacy. These can be addressed, and overcome; however, it seems clear that these highly beneficial systems will only succeed if they are built with privacy in mind – delivering a positive-sum, win-win outcome.

I advance the view that *Privacy by Design* is the sine qua non or the essential element that must be embedded as we advance in technology, data management and Cloud computing. I encourage industry and academic researchers involved in furthering the use of Cloud computing to “think *Privacy by Design*” as they develop approaches to the various challenges ahead – I wish you every success in doing so.

Ann Cavoukian, Ph.D.

Information and Privacy Commissioner
Ontario, Canada

1 Introduction

Once the exclusive domain of networked computers at universities and other large organizations, home PCs and modem connections have opened the Internet, evolving it into the information superhighway that we know today. This commercial availability of Internet service, which has allowed individuals access to the vast and multi-faceted resources thereon, has radically changed the flow of information.

What the Internet allows, and indeed promotes, is connectivity and data flow. Traditionally, this data flow has been from point-A-to-point-B, with the intermediary network being passive except for data routing. Processing or other utilization of data was done at the communication's end-points. Ensuring the privacy and security of data in such a transaction required three principal measures: i) appropriate security protections (firewalls, etc.) at the end-point servers, ii) appropriate data collection and handling procedures at both points A and B, and iii) security protections (encryption, generally) for data in motion.

As the Internet has evolved, however, we have seen the emergence of "Cloud computing." Organizations have begun to leverage the connectivity created by the Internet to optimize the utility of computing. Ever-cheaper and more powerful processing and storage capabilities are allowing data centres to act as viable, large scale central computing hubs. Simultaneously, increasing network bandwidth and reliable yet flexible network connections make it possible for clients – both individual and enterprise – to utilize high quality services which reside solely on these remote central hubs. These services will often include data storage (and real time access) or processing (by remote software and computing resources). This possibility, however, forces clients to re-think the data protection schemes developed for the point-A-to-point-B data flow.

When using the processing or storage services offered by Cloud service providers, another variable is introduced into the A-to-B flow of data: the now-active intermediary, Cloud-based point C. This new actor is progressively being assigned more and more of the functions which formerly occurred within the internal (and thus, presumably, secure and trusted) servers and networks of points A or B. Securing information that enters the Cloud, and protecting the privacy associated therewith, thus requires a shift, moving protections deeper into the Cloud's infrastructure. As privacy issues are sure to be central to user concerns about the adoption of Cloud computing, building such protections into the design and operation of the Cloud is vital to the future success of this new networking paradigm.

In this paper, a brief introduction to Cloud computing is provided for context. The privacy issue being addressed is then introduced, by describing some of the unique factors to be considered when data enters the Cloud. Finally, a data protection scheme is outlined that will address a number of these factors, by providing a mechanism to allow for data to be encrypted in the Cloud without loss of accessibility or functionality for authorized parties. This scheme is not necessarily a replacement for traditional privacy and security measures for data, but rather an enhancement which allows users (again, at either the individual or enterprise level) a greater degree of confidence in the adoption of innovative, cost-saving Cloud computing technologies.

2 Cloud Computing

The ‘Cloud’ is a broad, loosely-defined construct that encompasses all resources made available through the Cloud computing paradigm. It refers both to services accessed via, and delivered through, the Internet and the hardware and systems software in remote datacenters that provide those services. Cloud computing changes the way we think about computing by decoupling data processing, data retention, and data presentation – in effect, divorcing components from location. The services made available by this paradigm are increasingly in demand - according to IDC’s analysis, the worldwide forecast for Cloud services in 2009 was in the order of \$17.4 billion, and is projected to rise to \$44.2 billion in 2013 (IDC, 2009).

While it is sometimes considered simply an alternative means of traditional server or website hosting, the Cloud is actually much more than that, offering many different layers and opportunities. Cloud computing enables on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management effort or service provider interaction. Of particular benefit is the flexible infrastructural platform that Cloud computing provides, which can change computing resources from a capital- and skill-intensive investment into an elastic-scale utility-model of allocation. Like the electricity or gas supply, users are offered a ‘pay-as-you-go’ model, for as much or as little of the resource as needed. Purchasing services from a Cloud environment may allow technology business decision makers to save money and allow companies to focus on their core business – an enticing proposition in the current economic climate.

2.1 Cloud Computing Delivery Models

There are three models by which Cloud computing services are delivered: Software as a Service (SaaS), Platform as a Service (PaaS), and Infrastructure as a Service (IaaS), each with different benefits and limitations. Understanding the relationship and dependencies between these models is important. IaaS is the foundation of all Cloud services (i.e. the bottom layer) and is overlaid with PaaS (the middle layer) and SaaS (the top layer), respectively.

Software as a Service (SaaS). This model of Cloud computing delivers applications to consumers (either individuals or enterprises) using a multitenant architecture. Consumers are able to utilize a service provider’s offerings, which are run on a Cloud infrastructure, through a thin client interface such as a web browser. The principal benefit to the user, here, is that no upfront investment in servers or software licensing is required. Although the services made available are constrained by the provider’s design and capabilities, the consumer does get predictability and pre-integration. The consumer, in this arrangement, does not manage or control the underlying Cloud infrastructure, including network, servers, operating systems, storage, or application capabilities. This model is generally the most public-facing, and is likely to be the most familiar to individual Cloud users who engage with applications such as the Google Docs suite. Salesforce.com is by far the best-known example among enterprise applications, but SaaS is also common for HR applications and has even worked its way up the food chain to enterprise-resource planning (ERP), with players such as Workday.

Platform as a Service (PaaS). This model of Cloud computing delivers development environments to consumers (typically at the enterprise level). Consumers are provided the capability to deploy onto the Cloud infrastructure applications they have created using programming languages and tools supported by the service provider. These applications can then be run on the provider's infrastructure and delivered to the consumer's users via the Internet from the provider's servers. As with SaaS, the consumer still does not manage or control the network, servers, operating systems, or storage, but does now have control over the deployed applications and hosting environment configurations. Prime examples include Coghead and the Google App Engine. For extremely lightweight development, Cloud-based mash-up platforms abound, such as Yahoo Pipes and Dapper.net.

Infrastructure as a Service (IaaS). This model of Cloud computing delivers resources such as servers, connections, and the related tools necessary to build an application environment from scratch. The consumer (again, typically at the enterprise-level) is able to rent processing, storage, networks, and other fundamental computing resources on which the consumer can then deploy and run arbitrary software, including operating system or application software. The consumer must deal with virtual machines, patches, and various other complex issues, but gains a high degree of flexibility. The consumer does not manage or control the underlying Cloud infrastructure, but does control operating systems, storage, deployed applications, and networking components (such as firewall and load balancer). Examples include virtual servers and storage available on demand from the likes of Amazon's Elastic Compute EC2 and Simple Storage Service S3.

2.2 Deployment Models

In addition to the above described delivery models, there are three deployment models for Cloud computing: public, private, and hybrid. "Public Cloud" describes Cloud computing in the traditional, off-site sense, while "private Cloud" emulates Cloud computing on private networks. However, the "hybrid Cloud" - a combination of public and private Cloud offerings - will be typical for most enterprises.

Public Cloud. This model of Cloud computing is provided by an off-site third-party service provider who shares resources in a multitenant operating environment, and bills on a utility computing basis. The physical infrastructure is generally owned and managed by the service provider.

Private Cloud. This model of Cloud computing is provided by an organization or its designated service provider and offers a single-tenant operating environment with all the benefits and functionality of elasticity and the accountability/utility model of Cloud computing. The physical infrastructure may be owned by and managed by the organization or the designated service provider with an extension of management and security control planes controlled by the organization.

Hybrid Cloud. This model of Cloud computing is a composition of two or more Clouds (public or private) that remain unique entities but are bound together by standardized or proprietary technology that enables data and application portability.

2.3 Cloud Stack

The term “Cloud stack” refers to one Cloud computing service provider relying on other Cloud computing providers, layering services to form a stack. The three delivery models of Cloud computing naturally form such a stack, where IaaS is the foundation of all Cloud services with the PaaS and SaaS layers atop. The on-demand model of Cloud provisioning coupled with high levels of virtualization and automation yields even more complex Cloud stacks in which one IaaS provider may call upon other IaaS providers to supply abstracted computing resources to a PaaS provider. In turn, one SaaS may utilize multiple PaaS providers to fulfill one application service, and finally mash-up SaaS may refine the outputs of several SaaS providers.

Finally, it should be noted that Cloud computing is an evolving paradigm. Consequently, its definitions and attributes will evolve and new classification of models will emerge over time.

3 Data in the Cloud – The Blurred Security Perimeter

In the traditional data security model, an imaginary information boundary, or formally, *security perimeter*, is established to divide the trusted from the untrusted. Inside of this perimeter was the local, self-controlled computing resources within which sensitive information resided, and by which it was processed. This was the case with the original Internet, which connected end hosts. Data storage and processing occurred within the secure resources of these end hosts, with the network simply providing transit. Thus, reasoning about data protection could largely involve privacy and security evaluations at the known end points of a data transaction, with appropriate security measures applied to protect the data in motion.

However, this is no longer the case when computing in the Cloud – particularly when utilizing a public or hybrid Cloud architecture. These new computing environments often have unclear boundaries as to where the processing or storage of data physically occurs or resides (see figure 1). An 3rd-party actor – the Cloud service provider – provides software, platform, and infrastructure resources to the consumer (an individual or an enterprise). Thus, an entity outside of an individual or organization’s trusted security perimeter will store or otherwise touch massive amounts of information – much of which the consumer might consider private, confidential, or otherwise sensitive.




Fig. 1 Left: Clear Distinction between the Trusted and the Untrusted; Right: Fuzzy Security Perimeter

Trust, however, must be extended – in some measure – to the Cloud service provider, in order to utilize the desired service. This ‘blurring’ of the security perimeter is a difficult negotiation. To the extent that consumers lack control of Cloud resources, they have little technical capacity to prevent secondary use of, or unauthorized access to, their data. Instead, they must rely on trust relationships or contracts to prevent data misuse, or on after-the-fact mechanisms to limit or compensate for damages, should they occur. Agreements between service providers and consumers regarding acceptable use of data must exist in any data system – particularly given the fact that many public Cloud-based service providers rely on secondary data use (to provide advertising, for example) as a major (if not sole) revenue stream. Ideally, of course, these agreements will be technologically enforceable, before-the-fact, by both parties. This enforceability is the enhancement to traditional usage agreements, and strengthening of the blurred security perimeter, that is provided by the system described in this paper.

4 Privacy by Design and the Positive-Sum Paradigm

Cloud computing brings benefits to all actors in the cyber supply chain through significant new functionalities, including the abstraction of infrastructure, elasticity, the utility model of consumption, and so on. At the same time, though, it brings forth new security and privacy challenges – in particular, those associated with the outsourcing of data beyond a clear and well-defined security perimeter.

Individuals and organizations who wish to utilize Cloud services thus seemingly face a dilemma regarding giving preference to privacy or functionality. The service provider may instinctively desire maximized functionality, while the consumer may require maximized privacy. The thought that only one of these requirements can be fulfilled, at the cost of the other, is referred to as ‘zero-sum.’ We believe, however, that this trade-off is a false one; if the best case for both parties occurs when both functionality and privacy are maximized, we must strive to achieve it. We call this the “positive-sum paradigm.”

Service providers are clearly capable of maximizing functionality. How, then, can privacy be integrated into this functional system? We argue that the most appropriate, and most successful, solution is to utilize *Privacy by Design (PbD)*. *PbD* makes privacy a design requirement – a consideration that must be integrated into every phase of a technology’s design and development. By considering privacy at every step of the development cycle, the functionality and usability of a technology need not be compromised by after-the-fact, ‘bolted-on’ privacy measures.

Cloud computing is no exception; with proper privacy protections designed-in from the very beginning of the system lifecycle, and integrated at every system layer, the “positive-sum” paradigm is achievable. Both the need for protection of personal information and the need for functionality can be satisfied simultaneously.

4.1 The *Privacy by Design* Principles

A set of best practices for the development of a privacy-respecting Cloud computing architecture can be found in the *Privacy by Design* Principles, as developed by Dr. Ann Cavoukian¹. Below, these principles are described, along with the ways in which they be might applied in a Cloud environment.

1) *Proactive* not Reactive; *Preventative* not Remedial

“The *Privacy by Design* approach is characterized by proactive rather than reactive measures. It anticipates and prevents privacy invasive events *before* they happen. *PbD* does not wait for privacy risks to materialize, nor does it offer remedies for resolving privacy infractions once they have occurred — it aims to *prevent* them from occurring. In short, *Privacy by Design* comes before-the-fact, not after.”

When designing privacy measures for the Cloud, proactivity and prevention are the only solutions. Once data has been breached, it cannot be recovered – and the associated loss of consumer trust will be difficult, if not impossible, to remedy. Breaches must be stopped long before they happen.

For example, a Cloud consumer (individual or enterprise) may choose to encrypt all personal or otherwise sensitive data both while the data is stored on a Cloud service provider’s servers (at rest) and while being transmitted to end users (in motion)² – along with, of course, appropriate protections while the data is in-use. Industry-standard (at minimum) encryption will address some of the privacy risks associated with accidental loss or malicious exposure of information, as without the

1 See <http://www.ipc.on.ca/images/Resources/7foundationalprinciples.pdf>

2 This practice mirrors the IPC’s recommendation that all personal data be encrypted while in-transit (between ‘secure’ facilities) in the ‘offline’ world – for instance, while on a mobile device or USB key.

corresponding key (or a method to break the encryption algorithm³), that information is considered unusable. This principle is reflected in laws and regulations around the world, particularly privacy requirements. For instance, Safe Harbor provisions in laws and regulations consider lost encrypted data to be non-personal in nature. Encryption is at the heart of the architecture proposed in this paper (to be described in Section 5), along with systems to ensure appropriate access to data is not reduced – maintaining the positive-sum paradigm.

2) Privacy as the *Default*

“We can all be certain of one thing — the default rules! *Privacy by Design* seeks to deliver the maximum degree of privacy by ensuring that personal data are automatically protected in any given IT system or business practice. If an individual does nothing, their privacy still remains intact. No action is required on the part of the individual to protect their privacy — it is built into the system, *by default*.”

It should not be assumed that a user has either the capability or capacity to oversee and control access rights to the entirety of his or her data. Thus, privacy must be protected by default. Choice, of course, should be maintained; however, the user’s decision should be whether or not to allow *more* access to data, not whether they should enforce restrictions.

For instance, as described above, the ability for a data subject and/or holder to classify (and thus encrypt) data with regards to sensitivity may be a highly beneficial offering. However, to be truly privacy-respecting, this differentiation must take place as an ‘opt-out’ – that is, data must be protected by default until an authorized party allows expanded sharing. Access control mechanisms should similarly default to the most restrictive appropriate subgroup, to be expanded at the desire of the consumer. Such a system would allow the individual to take an active role in the determination of appropriate collection, use and disclosure policies for their data (instead of having these terms dictated to them for acceptance or denial). It would also benefit data handlers, as data subjects would, in effect, have the opportunity to ‘pre-consent’ to data uses – reducing the handler’s burden to seeking consent only if they desire to deviate from the defined terms. However, most importantly, it would maintain all appropriate restrictions on data until such time as the data owner chose to expand its associated access rights.

3) Privacy *Embedded* into Design

“*Privacy by Design* is embedded into the design and architecture of IT systems and business practices. It is not bolted on as an add-on, after the fact. The result is that privacy becomes an essential component of the core functionality being delivered. Privacy is integral to the system, without diminishing functionality.”

Privacy considerations must be taken into account during the design phase of a technology – this is the heart of *Privacy by Design*. The Cloud’s blurred security perimeter makes this principle particularly important to meet. Not only are reactive or remedial solutions poor business practice, but as the Cloud consumer does not have control over Cloud-based resources, they may not be possible. This inability to fix a privacy breach is yet one more reason to prevent one from happening.

3 Of course, for encryption to offer sufficient protection it must meet all relevant industry standards for strength and non-reversability.

Privacy in the Cloud will best be protected if measures are designed directly into the architecture – this is the rationale behind the solution that is presented in this paper.

4) Full Functionality – Positive-Sum, not Zero-Sum

“Privacy by Design seeks to accommodate all legitimate interests and objectives in a positive-sum “win-win” manner, not through a dated, zero-sum approach, where unnecessary trade-offs are made. Privacy by Design avoids the pretense of false dichotomies, such as privacy vs. security, demonstrating that it is possible to have both.”

Creating a privacy-protective Cloud architecture is, of course, meaningless if it is not also functional. The privacy of data should be assured, but so should appropriate access. The integrity of stored data should also be assured. A zero-sum approach forces trade-offs for the *user*, as well, who must determine which factors they value above others (privacy, security, functionality, etc.). For a paradigm shift as large as Cloud computing, the positive-sum must be realized in order to meet the technology’s full potential.

5) End-to-End Lifecycle Protection

“Privacy by Design, having been embedded into the system prior to the first element of information being collected, extends securely throughout the entire lifecycle of the data involved, from start to finish. This ensures that at the end of the process, all data are securely destroyed, in a timely fashion. Thus, Privacy by Design ensures cradle to grave, lifecycle management of information, end-to-end.”

The entire lifespan of data should be evaluated, and privacy threats identified and addressed. For instance, while the encryption of data addresses many of the privacy and security concerns of the data subject and/or data holder, significant meta-data may still be generated based on patterns of data access (including the identities of the data requestors), which might belie the contents of the data itself. To further enhance data protections, a Cloud architecture may thus wish to provide the capability for users to access, upload, modify or otherwise control data in an anonymous or pseudonymous fashion.

One must also remember that privacy attack countermeasures on a single layer of the Cloud computing system stack are not sufficient, as privacy is a cross-layer issue. For example, an application layer privacy enhancing technology that perfectly conceals user identity could be easily defeated if this application sits on a network where unique addressing is in place. Thus, technologies such as Trusted Computing or mix networks/onion routing may also be required to fully address the challenge of privacy protection.

6) Visibility and Transparency

“Privacy by Design seeks to assure all stakeholders that whatever the business practice or technology involved, it is in fact, operating according to the stated promises and objectives, subject to independent verification. Its component parts and operations remain visible and transparent, to users and providers alike. Remember, trust but verify.”

In general, users of Cloud services will have little direct knowledge, involvement or control over computing resources or processes. However, this is not to imply that Cloud service providers should operate in a wholly ‘black-boxed’ environment. In order to allow the consumer to make informed decisions with regard to their relationship with a company that operates in the Cloud, all systems must be clearly explained, and processes surrounding the handling of the data formalized. This level of control and awareness allows the consumer (individual or enterprise) to decide on her level of engagement with Cloud computing much more effectively. Audit mechanisms may be beneficial, in this situation, to provide the user with information that they cannot derive on their own.

It is also important to have open standards that allow interoperation across systems, to prevent the formation of ‘islands’ within the Cloud. The Open Cloud Manifesto⁴ establishes a core set of principles to ensure that organizations will have freedom of choice, flexibility, and openness as they take advantage of Cloud computing.

7) Respect for User Privacy

“Above all, *Privacy by Design* requires architects and operators to keep the interests of the individual uppermost by offering such measures as strong privacy defaults, appropriate notice, and empowering user-friendly options. Keep it user-centric.”

Respect for user privacy is the overarching theme of *Privacy by Design*. Privacy should not be a box-checking exercise; it should instead be integrated into an organization’s culture. The reason for ensuring privacy – respect for the user – should be kept in awareness when designing systems. This may be seen in appropriate defaults, data minimization, strong security measures, or the ability for the user to exercise control over his or her own data – but this user-centricity should always be present.

To achieve this, organizations should to continuously raise the education standards with respect to privacy – working to instill a culture of privacy within core operating procedures. To that end, the non-profit Cloud Security Alliance⁵ has formed to provide education on the uses of, and promote the best practices for providing privacy and security assurance within, Cloud computing. This is but one example of how to bring forth the importance of privacy – and *Privacy by Design*.

5 PbD Cloud Computing

It is clear that a privacy-protective Cloud computing architecture both can and must address the *Privacy by Design* principles described above. In this section, potential architectural elements that would achieve the ‘positive-sum’ of ensuring data privacy while maintaining system functionality are introduced and described. These elements will look to address two primary problems:

- Protecting data that enters the Cloud and maintaining appropriate access to this protected data, and
- Ensuring the integrity of protected data, without losing privacy.

⁴ Available at: <http://www.opencloudmanifesto.org/>

⁵ See: <http://www.cloudsecurityalliance.org/>

5.1 Protecting Privacy and Maintaining Appropriate Access

Ensuring data privacy without negatively impacting usability (i.e., ensuring appropriate access) within Cloud computing requires the protection of data throughout its lifecycle. The nature of the Cloud – in which it is common for the data subject or holder not to control the computing back end (including hardware, firmware or software) – implies that it would be prudent to assume that all communications and all stored data may be visible to arbitrary outsiders, even if such visibility may occur only through an error by some other party. That is, plaintext data in the Cloud should be considered ‘exposed’ to, at best, the service provider that controls the computing resources, or, at worst, the world at large in case of a failure in those resources. Thus, in order to achieve an appropriate level of privacy and/or security, data subjects and/or holders should encrypt all private or otherwise sensitive data prior to releasing it to a third-party in the Cloud.

Of course, this data must remain available for access and/or use by authorized parties. Figure 2 shows a potential minimalist Cloud computing architecture that preserves privacy and usability when data is encrypted and outsourced into the Cloud. This minimalist architecture is designed to solve one challenging problem – ensuring that organizations that make legitimate requests are granted access to encrypted data. Again, a positive-sum solution is sought with this architecture – obtaining privacy without an associated loss of functionality.

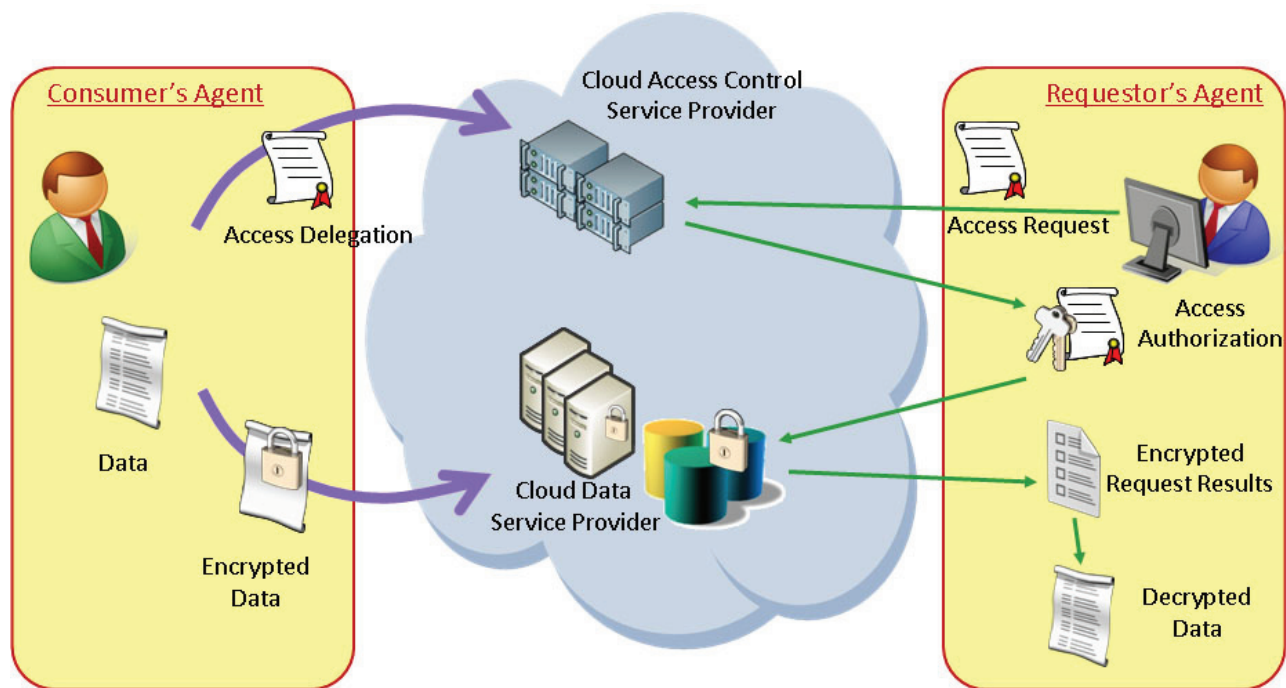


Fig. 2 Minimalist Cloud Computing Architecture for Privacy-Preserving and Usable Data Outsourcing

This architecture requires collaboration between two agents – the consumer’s agent and the requestor’s agent – and two service providers – the Cloud access control service provider (ACSP) and the Cloud data service provider (DSP). The consumer’s agent encrypts data prior to sending it to the Cloud DSP, and issues access delegation to the Cloud ACSP that will handle data utilization requests from the requestor. The requestor could be any party that has a personal or business relationship with the consumer who has outsourced data to the Cloud. The consumer could act as a requestor as well.

Should the requestor want to access the consumer’s data in the Cloud, requests would not go directly to the Cloud DSP, which, for the sake of privacy protection, does not hold the encryption key for the data it holds. Instead, this architecture would mandate that the requestor’s agent must contact the Cloud ACSP for access authorization. Upon authentication of the requestor, and satisfaction of any criteria set out in the access delegation, the Cloud ACSP would issue an access authorization to the requestor.

This proposed authorization message would consist of three components, each with a different effect. First, it would indicate to the Cloud DSP that the requestor had been authenticated, and was permitted to access the consumer’s data. Second, the Cloud ACSP would include in the message any available information regarding the subset of data to be released to the requestor, with the goal of restricting requestor access to be only the minimum required for its stated purposes. Finally, the authorization message would also contain a decryption key for the released data, engineered so as to only allow the requestor decrypt capabilities.

Should the requestor be able to circumvent this system, contacting the Cloud DSP directly and managing to succeed in retrieving data, the absence of the appropriate decryption key implies that all that is retrieved is meaningless ciphertext. Similarly, if the Cloud DSP were compromised or actively colluded with the requestor, again, only ciphertext could be obtained.

In this architecture design, consumer data is protected by encrypting it prior to outsourcing, and by enforcing access control even after the data has resided with a Cloud service provider. Access is also only granted to the smallest subset necessary to meet the requestor’s specified purposes, which implies both increased consumer privacy and fast data transfer for the requestor. Further, the party who utilizes the data sees no change in the data access process, as the requestor’s agent would be designed to function without user intervention, excepting the definition of the original request. There is no practical limitation on who can be a requestor; it could be a business partner, a family member, an advertiser, or the Cloud service provider itself. Depending on the consumer’s specifications, however, any set of requestors (including all, or none) may be granted access or use permissions. The consumer thus retains tight access control on data that has been released to, stored on, or otherwise made accessible to one or more entities in the Cloud. It is therefore clear that everyone involved in this architecture wins, as both privacy and usability are enhanced.

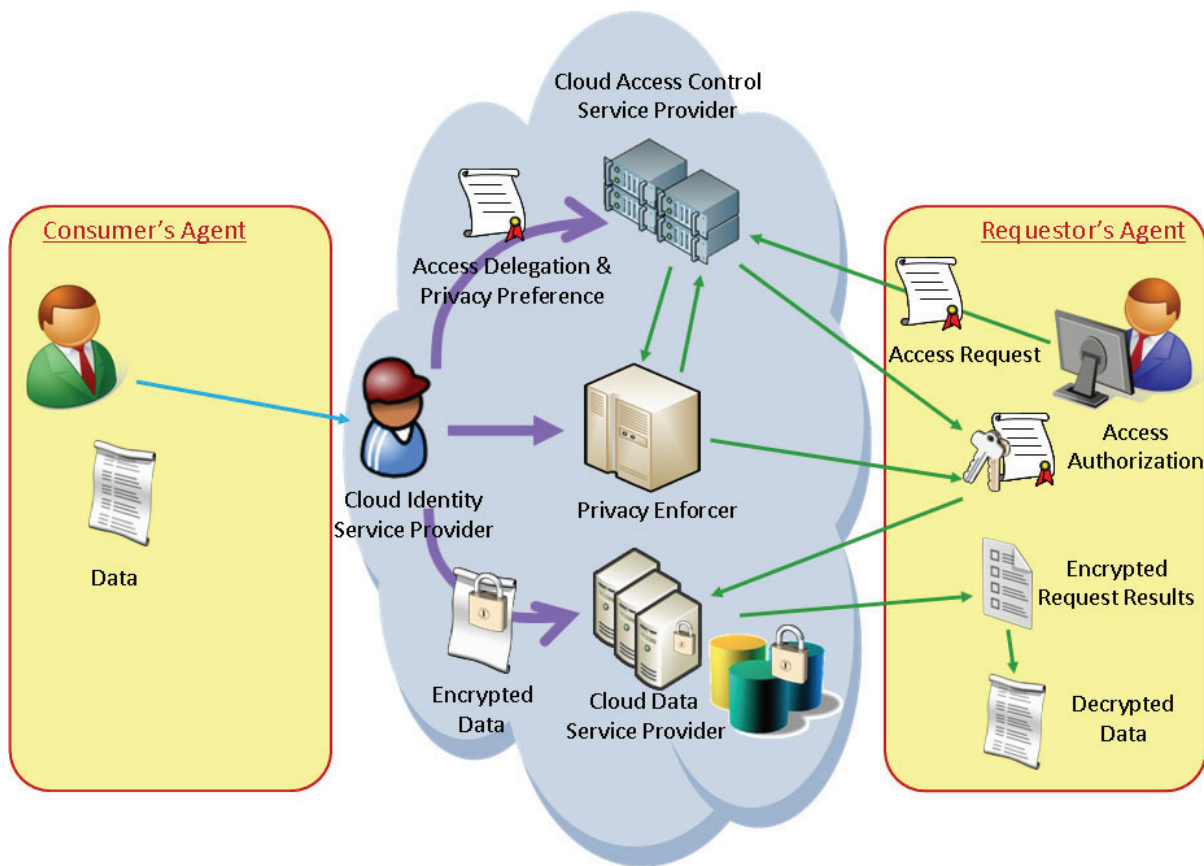


Fig. 3 Better Cloud Computing Architecture for Privacy-Preserving and Usable Data Outsourcing

To further the privacy protections available through this model – for instance, preventing service discrimination – additional architectural components, such as those shown in Figure 3, should be considered. In this enhanced design, a Cloud identity service provider will help the consumer in identity management, allowing the consumer the option of interacting with the Cloud service providers under the protection of secure and manageable pseudo identities. This identity service provider could allow, similarly, data requestors to obtain pseudo identities. This added element would help to ensure that information is not leaked into the Cloud through metadata – specifically, through data access patterns.

Additionally, as an increased security measure, a Privacy Enforcer could be introduced to prevent the Cloud access control service provider from unilaterally issuing access authorizations. Following the principle of duty separation, this privacy enforcer would match the requestor’s stated purposes against the consumer’s privacy preferences. It is then only when the requestor satisfies both the Cloud access control service provider and the privacy enforcer that an access authorization will be issued.

Finally, it is worth noting that in this architecture, scalable key management and effective auditing are necessary ingredients for the success of this architecture, although their corresponding components may or may not reside in the Cloud.

5.2 Protecting Privacy and Maintaining Data Integrity

Along with best privacy practices – such as the *Privacy by Design* principles listed above – maintaining the integrity of data plays a vital role in the establishment of trust between data subject and service provider. As with any data storage and processing medium, data integrity vulnerabilities exist within the Cloud computing paradigm. For Cloud services, these threats can be put into three classes: latent faults (e.g., those caused by a bit error in the storage medium), correlated faults (e.g., those caused by a lack of geographic location diversity), and recovery faults (e.g., those caused by improperly debugged procedures). In making the decision to outsource particular data, consumers – at either the individual or enterprise level – must be able to evaluate the risk of that data being corrupted or otherwise lost to use. Thus, as an additional architectural component for Cloud computing, an auditor to whom the consumers could delegate the task of checking data integrity is introduced. This auditor would periodically check the integrity of all data stored with Cloud service providers and release, for instance, monthly audit reports. Based on these reports, Cloud consumers could evaluate the risks associated with any particular Cloud service provider before they decide to rely on its service. The audit report may also be beneficial to the Cloud service provider: in addition to serving as a promotional tool, a positive audit report from a third party may assist the service provider in obtaining a favorable insurance rate, based on the measured stability of their primary asset (data).

Of course, the benefits of introducing an auditor into the Cloud computing environment should not be offset by the privacy-invasiveness of allowing the auditor to actually read consumers' outsourced data. For instance, a clear breach occurs if data must be decrypted for an audit to occur. A more subtle problem, though, can take place even if the data remains encrypted. If the auditor is able to see metadata, such as an increase or decrease in the size of the ciphertext, personal information might be inferred. For example, an increase in the size of a health record may reveal frequent visits to a clinic center, which could serve as a basis for discrimination if the consumer is in need of medical insurance. Thus, the audit architecture will need to be designed to allow sufficient audit capabilities without loss of user privacy.

Cloud consumers and service providers alike have a need for data integrity assurance. The basic requirement to be met is that when the consumer wishes to retrieve data from a service provider, that data should be readily available and in the same state as originally outsourced, with not one single bit error. To that end, Figure 4 describes an architecture capable of providing audit capabilities in the Cloud – assuring the Cloud consumer that outsourced data has not been altered through error, neglect, or malice, and allowing the Cloud provider to demonstrate to its consumers that the data stored with it is indeed the same data that the consumer placed into the Cloud.




Fig. 4 Minimalist Cloud Computing Architecture for Privacy-Preserving and Trustworthy Data Outsourcing

This architecture requires collaboration between two agents – the consumer’s agent and the auditor’s agent – and a (re-engineered) data service provider (DSP). Again, the consumer’s agent outsources encrypted data to the Cloud DSP, while now additionally contracting an auditor to handle integrity audits. This auditor could either be internal or external to the consumer. Once contracted, the auditor is given the capacity to send audit requests to the Cloud DSP. In turn, the Cloud DSP, which has been re-engineered to handle such requests, would reply with an audit result. The auditor, after further processing the audit results, would then release an audit report on data integrity. What is crucial in this architecture design is that the consumer never need disclose her encryption key. The auditor must be designed to fulfill its duty using only the public key of the consumer. Therefore, even if in the case that both the Cloud DSP and the auditor’s agent are compromised, the consumer’s data, and thus privacy, remains safe. The parties involved in this process should learn no more than the true/false result of the integrity test – no other information should be generated or released to any party. Through these added architectural elements, the requirement for auditing the integrity of outsourced data is satisfied, and importantly, the consumer’s privacy remains protected.

In the Cloud, to achieve a higher degree of data security and/or availability, the Cloud DSP may opt to use other Cloud DSPs for backup purposes; this use of redundant data repositories is, in fact, relatively common. However, once the Cloud DSP uses other DSPs for enhanced data availability, all the Cloud DSPs involved need to be auditable for data integrity. To better cope with this scenario, again additional architectural elements should be considered, yielding the more complex audit architecture shown in Figure 5. Under this new architecture, a Cloud integrity service broker will help both the consumer and the Cloud data service provider in contracting auditors. The Cloud integrity service broker may also be the entity which relays the consumer’s public key to both the auditor and the Cloud data service provider. This relaying of keys further enforces a crucial point of the architecture design – that the auditor must be able to fulfill his duty using only the public key of the consumer, with the consumer’s private encryption key never being disclosed.

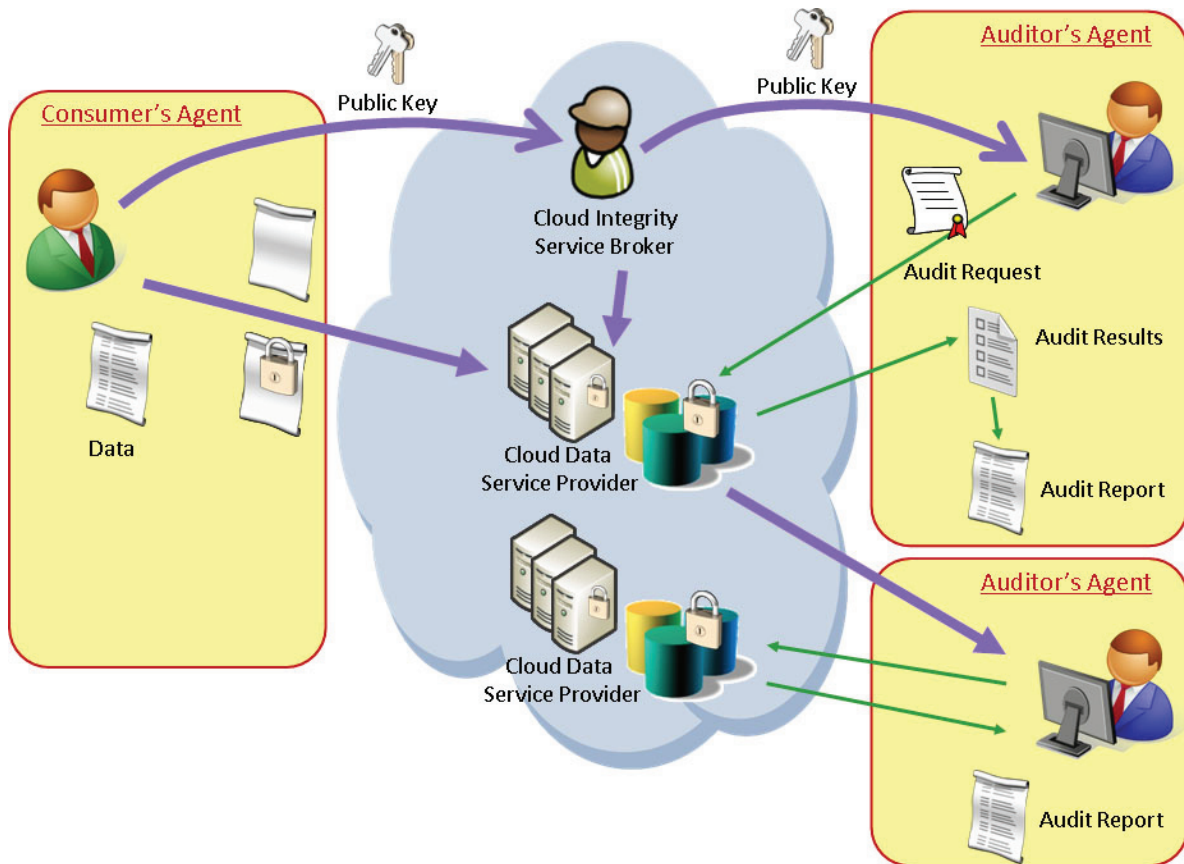


Fig. 5 Cloud Computing Architecture for Privacy-Preserving, Trustworthy, and Available Data Outsourcing

Finally, it should be noted that for an enhanced level of privacy protection, a Cloud identity service provider may be included in the above audit architecture, which would allow the consumer to find an auditor anonymously or pseudonymously. This would prevent a malicious auditor from obtaining any advantage towards privacy invasion simply by being contracted to audit a consumer's data.

6 Additional Privacy and Security Considerations

In this paper, two portions of an architectural design that meets security, usability, data integrity and privacy requirements simultaneously – the ‘positive-sum’ – have been proposed. However, although the right architecture is a crucial step toward the end goal of positive-sum Cloud computing, there remain further considerations that must be addressed. A selection of these considerations are described below:

- Assessment with respect to the system and data being considered for Cloud resources should be conducted to ensure all risks are identified, understood, and accounted for. There are many tools available, such as the Privacy Impact Assessment (PIA) and Federated Privacy Impact Assessment (F-PIA) with which an organization or organizations can demonstrate privacy requirements at different phases of design and delineate their data protection efforts.

- In a Cloud environment there are new and different threats (both to privacy and security), and evaluators need different guidance on how to test Cloud delivery mechanisms against these threats. Regulators and standards bodies should thus update and modify their existing, computing-related regulations and standards to account for the differences in architecture and operation in a Cloud environments. These new standards and regulations should be principles-based and media-neutral, in order to accommodate the required technical flexibility. Further, Cloud-based organizations may wish to engage these groups in order to assist in the development of appropriate principles – such input will almost certainly be crucial to the success of these new models.
- Organizations must rethink their established software development, validation, certification, and accreditation processes in response to the need to push or pull applications in the Cloud. They may thus need to re-design their Software Development Life Cycle (SDLC) – building privacy in and looking to solutions or evaluatory techniques that extend beyond the trusted perimeter. Similar to security, designing privacy into a system is the best way to achieve effective protections. Early and comprehensive integration of privacy and security into the software design phase should be the focus.

7 Conclusions

Individuals and organizations could view Cloud computing as a disruptive innovation which challenges norms and forces out-of-the-box thinking – or, in much the same way as other innovations, it can be embraced as an opportunity to re-evaluate the approach to computing resources and to develop new efficiencies and strategic responses. However, as of this writing, the possibility remains for Cloud computing to not reach its full potential, as individual or enterprise-level consumers shy away due to data security and/or privacy concerns. Cloud-based mechanisms are required to ensure data security and privacy, and to fulfill the regulatory and audit requirements of enterprises. Without a profound increase in the understanding and assessment of the risks accompanying Cloud computing, and without proper countermeasures being deployed and evaluated, the Cloud may become the exclusive domain of non-sensitive data – a status far below its anticipated potential. The joint issues of privacy and security must be addressed, if Cloud computing service providers intend to attain a status in the computing domain similar to “utility providers” in the general world – as trusted, reliable purveyors of pay-per-use access to fundamental (in this case, computing) resources.

The IPC has addressed this issue in the past, in its *Privacy in the Clouds* white paper. In that paper, 4 technological approaches are presented to the problem of “collectively [assuring] confidence and trust in the privacy of our personally identifiable information, when our [data] is held by others and we are not directly involved in data transactions in the Cloud.” The architecture described above address the 4th of these approaches: “Trust intermediary [service] providers to behave.” Such an approach states that:

“... trusted actors will increasingly act on our behalf, disclosing our [data] for the purposes we define in advance, and under specific conditions. They must find credible technological mechanisms for assuring us that they are behaving in an open and accountable manner, and that our privacy is in fact being protected. [This] might include automated audit and enforcement tools that can also convey up-to-the-minute privacy and security status reports to users, regulators and other trusted third parties.”

The architectural elements proposed in this paper contain the means of addressing a number of these requirements: automated protection, privacy assurance, audit and reporting capabilities, etc. Better yet, though, it does so without negatively impacting usability or data availability. This is the positive-sum outcome that it is so crucial to obtain. It is our belief that this combination of security, privacy and usability will be a key differentiator in the Cloud computing market – and as such, the proposed architectural elements would be of great benefit to both users and service providers.

There is, of course, further work to be done in the research and engineering disciplines. Required advances include, for example, the development of privacy preserving data provenance, encrypted processing, privacy preserving forensics, resource isolation, security as a Cloud service, and so forth. As such we call to action the research and engineering domains for provision of security and privacy-enhancing technologies, and those in the operational domain to deploy these technologies. Let positive-sum be the goal, and let a trustworthy and fully functional Cloud be the future.

Bibliography

Altman, I. (1997) Privacy Regulation: Culturally Universal or Culturally Specific. *Journal of Social Issues*, 33:3, p. 66-84.

Armknect, F. et al. (2007) Crosslayer Privacy Enhancement and Non-Repudiation in Vehicular Communication. *4th Workshop on Mobile Ad-Hoc Networks (WMAN)*, Bern Switzerland, March 2007.

Bayardo, R.J., and Srikant, R. (2003) Technological Solutions for Protecting Privacy. *IEEE Computer*, 36(9), p. 115-118.

Bertino, E. (2009) Privacy-preserving Digital Identity Management for Cloud Computing. *IEEE Data Engineering Bulletin*, 32, p. 21-27.

Brodkin, J. (2008) Seven Cloud-Computing Security Risks. *InfoWorld*. Available online at: <http://www.infoworld.com/d/security-central/gartner-seven-cloud-computing-security-risks-853>

Buyya, R., Yeol, C.S., and Venugopal, S. (2008) Market-Oriented Cloud Computing: Vision, Hype, and Reality for Delivering IT Services as Computing Utilities. In *Proc. of 10th IEEE International Conference on High Performance Computing and Communications (HPCC'08)*, p.5-13.

Cavoukian, A. (2009) The New Federated Privacy Impact Assessment (F-PIA): Building Privacy and Trust-Enabled Federation, and Privacy by Design. *Information and Privacy Commissioner of Ontario*. Available online at: <http://www.ipc.on.ca/english/Resources/Discussion-Papers/Discussion-Papers-Summary/?id=836>

Cavoukian, A. (2008) Privacy in the Clouds – A White Paper on Privacy and Digital Identity: Implications for the Internet. *Information and Privacy Commissioner of Ontario*. Available online at: <http://www.ipc.on.ca/english/Resources/Discussion-Papers/Discussion-Papers-Summary/?id=748>

Cavoukian, A. (2006) 7 Laws of Identity: The Case for Privacy-Embedded Laws of Identity in the Digital Age. *Information and Privacy Commissioner of Ontario*. Available online at: <http://www.ipc.on.ca/english/Resources/Discussion-Papers/Discussion-Papers-Summary/?id=470>

Cavoukian, A. and Hamilton, T. (2002) *The Privacy Payoff*. McGraw-Hill.

Caloyannides, M. (2003) Privacy vs. Information Technology. *IEEE Security and Privacy*, 1:1, p. 100-103.

Chaum, D. (1985) Security without Identification: Transaction Systems to Make the Big Brother Obsolete. *Communications of the ACM*, 28:10, p. 1030-1044.

Chong, S. et al. (2007) Secure web applications via automatic partitioning. In *Proc. of 21th ACM Symposium on Operating Systems Principles*, pp. 31-44.

Cloud Computing Use Case Discussion Group. (2009) Cloud Computing Uses Cases White Paper. Available online at: <http://bit.ly/9aylix>

Cloud Security Alliance (2009) Security Guidance for Critical Areas of Focus in Cloud Computing. Available online at: <http://www.cloudsecurityalliance.org/guidance/>

ENISA (2009) Cloud Computing Risk Assessment. Available online at: <http://www.enisa.europa.eu/act/rm/files/deliverables/cloud-computing-risk-assessment>

Geelan, J. (2009) The Future of Cloud Computing. *Cloud Computing Journal*. Available online at: <http://cloudcomputing.sys-con.com/node/771947>

Guha, S., Tang, K., and Francis, P. (2008) NOYB: Privacy in Online Social Networks. In *Proc. of the 1st ACM SIGCOMM Workshop on Online Social Networks (WOSN)*, p.49-54

IDC (2009) IDC's New IT Cloud Forecast: 2009-2013. IDC eXchange. Available online at: <http://blogs.idc.com/ie/?p=543>

Imam, A. (2009) Cloud Computing: Prospects and Challenges. Newcastle University Business School. Available online at http://docs.google.com/View?id=dcf9h45v_2076mjw5dc

Kavis, M. (2009) Secure Hybrid Cloud Architectures. *Kavis Technology Consulting*. Available online at: <http://www.kavistechnology.com/blog/?p=957>

Microsoft (2009) Securing Microsoft's Cloud Infrastructure. Available online at: <http://www.globalfoundationservices.com/security/documents/SecuringtheMSCloudMay09.pdf>

Reed, M., Syverson, P., and Goldschlag, D. (1998) Anonymous connections and Onion Routing, *IEEE J. Selected Areas in Communication*, 16:4, p. 482-494.

Scheier, R.L. (2009) Busting the Nine Myths of Cloud Computing. *CIO.com*. Available online at: http://www.cio.com/article/495523/Busting_the_Nine_Myths_of_Cloud_Computing

Syverson, P., Goldschlag, D., and Reed, M. (1997) Anonymous Connections and Onion Routing. In *Proceedings of the IEEE Symposium on Security and Privacy*, p. 44-54.

Narten, T. and Draves, R. (2001). RFC 3041 Privacy Extensions for Stateless Address Autoconfiguration in IPv6.

Tian, Y. et al. (2009) A Fast Search Method for Encrypted Medical Data. *1st International Workshop on Medical Applications Networking*, MAN 2009, Dresden, Germany.

Wang, C.X. (2009) A Close Look At Cloud Computing Security Issues. CSO Online. Available online at: <http://www.csoonline.com/article/496388/>

Westin, A. (1987) *Privacy and Freedom*. Bodley Head.

White, S. (2005) A Brief History of Computing. Available online at: <http://trillian.randomstuff.org.uk/~stephen/history/>

Wildstrom, S.H. (2009) Cloud Computing: Understand the Risks. *Business Week*. Available online at: http://www.businessweek.com/magazine/content/09_14/b4125000676483.htm

Wuchner, A. (2009) Privacy of Information: Do we fully understand the issue? *IT Risk Space*. Available online at: <http://itriskspace.com/2009/06/18/1245307200000.html>

Zeng, K. (2006) Pseudonymous PKI for Ubiquitous Computing. *LNCS 4043*, p. 207-222, EuroPKI'06.

Zeng, K. (2008) Publicly Verifiable Remote Data Integrity. *LNCS 5308*, p. 419-434, ICICS'08.

About the Authors

Ke Zeng, Ph.D., Researcher, NEC Laboratories China

Ke Zeng has been with NEC Laboratories China since 2003, where he plays key leadership roles in research and development teams examining security, privacy, and performance issues in building large-scale distributed service-oriented systems. He has been working on various aspects of distributed data intensive-systems, ranging from distributed authentication and authorization for ubiquitous computing, high-privacy security for inter-vehicle communication, privacy-enhanced information retrieval, and secure Cloud computing.

He is always interested in and working on innovative solutions through novel algorithm and architecture design together with smart coding and system implementation. He is an advocator and practitioner of secure SDLC management and secure programming.

Ke Zeng received Bachelor, Master, and Ph.D. degrees in EE from Tsinghua University. In 2006 and 2007, he was with the Computer Science Department of Stanford University as Visiting Scholar where he studied cutting-edge technologies in online search, GPGPU computing, and Cloud computing. He has 20+ peer-reviewed papers published and 10+ U.S. patents in application. He was among the winners of the Science & Technology Innovation Award and Best Paper Award. He has held CISSP designation since 2008.

Ann Cavoukian, Ph.D., Information and Privacy Commissioner of Ontario, Canada

Dr. Ann Cavoukian is recognized as one of the leading privacy experts in the world. Noted for her seminal work on Privacy Enhancing Technologies (PETs) in 1995, her concept of Privacy by Design seeks to embed privacy into the design specifications of technology, thereby achieving the strongest protection. An avowed believer in the role that technology can play in protecting privacy, Dr. Cavoukian's leadership has seen her office develop a number of tools and procedures to ensure that privacy is strongly protected, not only in Canada, but around the world. She has been involved in a number of international committees focused on privacy, security, technology and business, and endeavours to focus on strengthening trust and confidence in emerging technological applications. Dr. Cavoukian also serves as the Chair of the Identity, Privacy and Security Institute at the University of Toronto, Canada and is a member of the Future of Privacy Advisory Board. Recently reappointed as Commissioner for an unprecedented third term, Dr. Cavoukian intends to grow Privacy by Design and hopes to make it go "viral."



The information contained herein is subject to change without notice. NEC Company, Ltd. and IPC shall not be liable for technical or editorial errors or omissions contained herein.

May 2010

<http://www.nec.com> | <http://www.privacybydesign.ca>