

Isolation in Cloud Computing and Privacy-Enhancing Technologies

Suitability of Privacy-Enhancing Technologies for Separating Data Usage in Business Processes

Sustainability of Cloud Computing is assured for the uncritical services only. There is basically no company that entrusts critical data to a Cloud. Cloud Computing does not realize its potentials as far a cost reduction in IT is concerned. The reason is the lack of isolation. Isolation can be seen as a special sort of privacy, where the a service should not get in contact with other services, and the provider of the Cloud should not know what data are used in the service and for what purpose the service is used by the customer.

DOI 10.1007/s12599-011-0160-x

The Authors

Prof. Dr. Noboru Sonehara
Prof. Dr. Isao Echizen
Dr. Sven Wohlgemuth (✉)
 National Institute of Informatics
 2-1-2 Hitotsubashi, Chiyoda-ku
 Tokyo 101-8430
 Japan
sonehara@nii.ac.jp
iechizen@nii.ac.jp
wohlgemuth@nii.ac.jp

Received: 2010-07-01
 Accepted: 2011-03-29
 Accepted after four revisions by
 Prof. Dr. Buxmann.

Published online: 2011-05-05

This article is also available in German in print and via <http://www.wirtschaftsinformatik.de>: Sonehara N, Echizen I, Wohlgemuth S (2011) Isolation in Cloud-Computing und Mechanismen zum Schutz der Privatsphäre. Eignung von Mechanismen zum Schutz der Privatsphäre für die Trennung der Datenverarbeitung in Geschäftsprozessen. WIRTSCHAFTSINFORMATIK. doi: 10.1007/s11576-011-0274-2.

© Gabler Verlag 2011

1 The Three Layers of Cloud Services

Cloud Computing breaks open the access control domain of individuals' and companies' IT systems by processing their data within application frameworks and virtualized runtime environments of Cloud service providers. Cloud services are available as software services with standardized interfaces hiding their implementation. Hence, it is possible to make use of application services, application frameworks, and runtime environments of different service providers by combining them for the execution of a business process. Depending on the kind of Cloud service, it belongs to one of three abstraction layers: *Software as a Service (SaaS)*, *Platform as a Service (PaaS)*, and *Infrastructure as a Service (IaaS)* (Mather et al. 2009). *SaaS* stands for application services, e.g., office, e-mail, billing, and customer-relationship management applications. *SaaS* implements mainly function modules of business processes. The next abstraction layer *PaaS* provides the application framework for implementing and hosting Cloud application services. The third abstraction layer *IaaS* offers computing and data storage facilities as well as operating systems as a software stack. Thereby, providers of the Cloud services can differ and services can be executed in different Clouds. Due to the on-demand characteristics of Cloud Computing, the usage of Cloud services of given service providers should be possible for different orchestrations of business processes

of (competing) Cloud users (*availability of Cloud services*). Hence, a Cloud service of one service provider could provide the same function to instances of business processes, each belonging to a competing Cloud user. Thus, a flow of the Cloud users' data between different business processes via Cloud services – either of Cloud service providers or of other Cloud users – is possible.

This article provides a survey of the suitability of privacy-enhancing technologies for achieving isolation in Cloud Computing. Thereby, the Cloud user as owner of his data takes the role of an individual as the data subject of his data. Sustainability of Cloud Computing is assured for the uncritical services only. Basically no company entrusts critical data to a Cloud. Cloud Computing does not realize its potentials as far a cost reduction in IT is concerned due to the lack of isolation. Isolation can be seen as a special sort of privacy, where the service should not be making contact with other services, and the provider of the Cloud should not know what data are used in the service and for what purpose the service is used by the customer.

Today's security mechanism for isolation in Cloud Computing is virtualization (Armbrust et al. 2010). Virtualization encapsulates hardware, operating systems, or application frameworks respectively; it is their logical representation by software. Although virtualization can separate different service providers within a Cloud, firstly the virtualization mechanisms are controlled by the provider of the corresponding resource

and not by the Cloud user; and secondly this virtualization does not hold if Cloud services from different service providers are used horizontally on the same abstraction layer, such as *SaaS*, or vertically across different abstraction layers. However, this should be possible so that Cloud users can select on-demand the currently most suitable service for running their business process in a Cloud. This property corresponds to a liveness property meaning “let good things happen” (Alpern and Schneider 1985).

Privacy-enhancing technologies focus mainly on confidentiality of personal data and on avoiding a linking of an individual’s transactions to prevent non-authorized profiling. These properties derive from the privacy of informational self-determination which requires the ability of individuals to control the disclosure of their data (Westin 1967; Bundesverfassungsgericht 1983).

This state of the art contribution describes risks and thus stresses that security – as understood in computer science – is most likely impossible and not even desirable. The key term to sustainability is seen as the ability to isolate services from each other. Experience with exactly this property is available in the field of privacy. Cloud and its desire to orchestrate services requires that the service requestor has the possibility to formulate his own security policies, and to delegate rights and to control the usage of data and service results. These key technologies will be discussed in the following.

2 Risks for Confidentiality of Cloud Users’ Data

According to the abstract system model of Pretschner et al. (2006), participants in an information system take the role of a data provider and data consumer. If a participant discloses (personal) data d to another participant, the sending participant acts as a *data provider* whereas the receiving participant acts as a *data consumer*. Roles can change dynamically. If the receiving participant discloses the data d or a modified/extended data set d' , then it changes its role from a *data consumer* to a *data provider*. Figure 1 shows the instantiation of this model on Cloud Computing. Services $SaaS_2$ and $SaaS_3$ of the abstraction layer *SaaS* use the same service $PaaS_2$ of the abstraction layer *PaaS*.

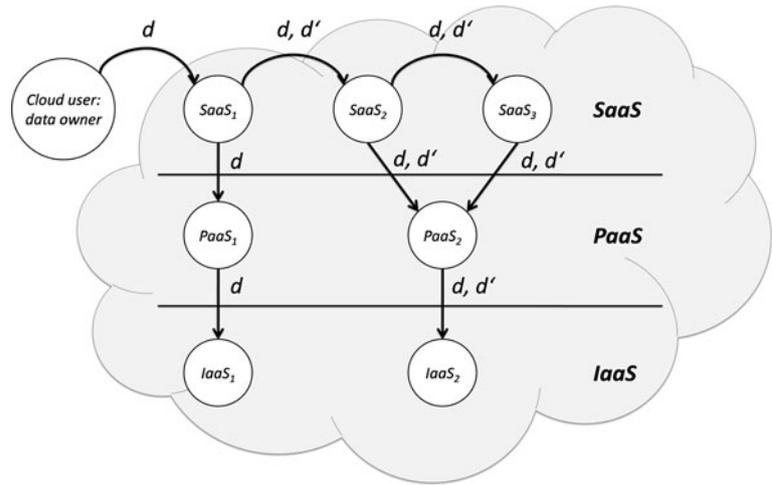


Fig. 1 Data flow in Cloud Computing according to the system model of Pretschner et al. (2006)

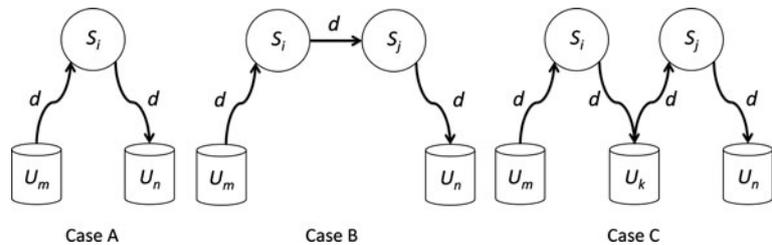


Fig. 2 Possible data flows of Cloud users’ data via Cloud services

Due to the on-demand characteristics of Cloud Computing, the usage of the same Cloud services should be possible for different orchestrations of business processes of (competing) Cloud users (*availability of Cloud services*). Hence, a Cloud service of one service provider could provide the same function to at least two instances of a business process belonging to competing Cloud users. Thereby, a (non-authorized) flow of these Cloud users’ data between the business processes via the Cloud services is possible and has to be prevented to achieve isolation. Figure 2 shows the possible data flows of data d from the orchestration of a business process of Cloud user U_m to the orchestration of a business process of a (competing) Cloud user U_n by the Cloud service S_i and the Cloud services S_i and S_j , respectively. It is assumed that these services are not vulnerable against attacks, such as guest-hopping attacks. These data flows are carried out without authorization of the Cloud user U_m . They violate the IT-Security protection goal of confidentiality.

- Case A: The Cloud service S_i reads d from the dataset of U_m and writes d into the dataset of U_n .

- Case B: The Cloud service S_i reads d from the dataset of U_m and sends d to the Cloud service S_j . The Cloud service S_j writes d into the dataset of U_n .

- Case C: The Cloud service S_i reads d from the dataset of U_m and writes d into the dataset of U_k . The Cloud user U_k is not in competition with the Cloud users U_m and U_n . The Cloud service S_j has access to the dataset of U_k , reads d from this dataset and writes it into the dataset of U_n .

The exemplary case study “Telemedicine” in which medical services and companies make use of Cloud services and personal data (x-ray images) of patients illustrates these possible data flows. A commercial provider of an Electronic Health Record (EHR) data service collects health data from its users (patients) with their agreement for the purpose of sharing the data among others with clinics, health insurance agencies, and pharmaceutical companies. Existing systems comply with the US American Health Insurance Portability and Accountability Act (US Department of Health & Human Services 1996) by letting users decide on the usage and disclosure of their medical

data. However, they do not offer mechanisms in order to enforce the rules of a privacy policy (Haas et al. 2010). Let us assume that a patient needs medical treatment abroad. A clinic in the homeland has shot an x-ray image of the patient and has disclosed it to a data centre. Patients have shown their digital identity to the first service provider, i.e., the clinic in the homeland, and have agreed on obligations for disclosing their medical data via an EHR data service to a hospital and to a clinic abroad. Additional disclosures of these data are not permitted. If in general at least one participant makes use of a Cloud service for the usage of the patients' health data, non-authorized disclosure of health data to a third party might take place, if data are disclosed by the service itself (cases A and B), e.g., by the EHR service, or via the used Cloud services of the lower layer (IaaS or Paas).

3 Privacy and Legal Requirements

In view of the spread of electronic data processing, the information flow of personal data was included in the term of privacy by Alan Westin: “*Privacy is the claim of individuals, groups, or institutions to determine for themselves when, how, and to what extent information about them is communicated to others*” (Westin 1967). The protection of personal data for its collection and use under the concept of informational self-determination was specified in 1983 for the first time by the German Federal Constitutional Court in the so-called “census judgment” (Bundesverfassungsgericht 1983). This judgment limits the right to informational self-determination, if there is a substantial general interest for the limitation.

The named technical development and resulting threats for privacy led to the European data protection acts (Roßnagel 2005). The general basic principles are transcribed in the European Data Privacy Directive 95/46/EC (European Commission 1995) which was extended to a general electronic communication by the European Data Privacy Directive 2002/58/EC (European Commission 2002):

- Transparency of data processing through briefing and notification of the person concerned,
- Necessity of the data collected for a certain purpose,

- Restriction of data processing to a certain purpose,
- Correction rights of the person concerned regarding the required data and the processing phases,
- Data avoidance and data economy,
- Data protection through technology, and
- Implementation control through a data protection representative.

The minimal principles for privacy protection found in the named Data Privacy Directive of the European Union originate from the *Fair Information Practices*. These principles were first published in the *United States Departments for Health Education and Welfare* (HEW) report and were incorporated in the *US Privacy Act of 1974* (Smith 1993). The five principles of *Fair Information Practices* were accepted by the *Organization for Economic Cooperation and Development* (OECD) and standardized in the form of eight principles for the protection of privacy in cross-frontier data exchange in the *OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data* (OECD 1980).

Privacy according to data protection legislations can be seen as processing personal data for a specific usage by a specific party and with the consent of the corresponding individual. Hence, by their security principles data protection legislations define access to data according to the principle of least privilege. If a Cloud service violates such a security principle, the privacy of the corresponding individual has been violated.

4 Privacy-Enhancing Technologies and Isolation

Privacy-enhancing technologies (PET) aim at supporting individuals in controlling the disclosure of their data according to the agreed-upon privacy policy between the corresponding individual and service provider. A privacy policy consists of rules regarding the collection, processing, and storage of personal data as well as regarding their disclosure to third parties by the data collecting party. Privacy policies are formalized by privacy policy languages in a machine-readable form enabling automatic decisions on requests to access personal data. PET should enable individuals to enforce these privacy policies or at least to check whether participating service providers follow the agreed-upon privacy policy.

PET can be grouped into (a) privacy policy languages, (b) encryption schemes, and (c) anonymity and pseudonymity. This section investigates the state of the art of PET with regard to their suitability for isolation of data usage in Cloud Computing.

4.1 Privacy Policy Languages

P3P (Platform for Privacy Preferences) for formalizing privacy policies has been standardized in 2002 by the *W3C (World Wide Web Consortium)* (Cranor et al. 2002). A P3P policy consists of rules published by a service provider on his website. A web browser with a P3P implementation automatically reads this P3P policy and compares it with the privacy preference of the corresponding individual. An individual gives his agreement to the collection and processing of the given personal data by agreeing to a P3P policy. P3P supports a disclosure of personal data to third parties by a P3P attribute for the recipient of the data and the kind of data processing. However, this attribute refers only to those data which have been directly collected from the individual. In addition, even though P3P supports transparency of data processing by enabling machine-readable privacy policies, its vocabulary is restricted to contact details of individuals. It does not support an extension of the vocabulary. Furthermore, rules for the disclosure of personal data to third parties apply only to those data which have been directly collected from the individual. Hence, P3P does not consider a chain of data disclosures to third parties and therefore is not suitable for the formalization of isolating data usage in Cloud Computing.

EPAL (Enterprise Privacy Authorization Language) considers the formulation of privacy policies within a company (Ashley et al. 2003). The aim is that service providers can substantiate that they have followed their privacy policy, e.g. published as a P3P policy (Karjoth et al. 2002). Therefore, an EPAL policy is linked to the corresponding personal data so that access requests to these data can be decided based upon this EPAL policy. This is called *sticky policy* (Karjoth et al. 2003). EPAL does not provide a vocabulary but a list of hierarchies of data-categories, user-categories, purposes, set of actions, conditions, and obligations. EPAL considers a successive disclosure of personal data to third parties by rules regarding the data consumers and their

processing of the corresponding personal data. According to the concept of *sticky policies*, the EPAL policy is passed on together with the corresponding personal data. However, EPAL does not support an interaction with the individuals/data owners.

4.2 Encryption Schemes

An implementation of sticky policies for disclosure of personal data to third parties is the **Adaptive PMS (Adaptive Privacy Management System)** of Hewlett-Packard (Casassa and Pearson 2005). Adaptive PMS links sticky policies to certain personal data at the time of their collection by means of an encryption scheme. A data center stores encrypted personal data of several individuals for the purpose of their disclosure to a third party. A data consumer will obtain the decryption key from a TTP (Trusted Third Party), if he is authorized by the sticky policy. That means that the data center obtains the encrypted individual's data. Hence, it cannot conduct additional services on the collected data. In addition, data consumers can further disclose the decrypted personal data without any control by the individual or the TTP.

Homomorphic encryption achieves the confidentiality of the individual's data while at the same time enabling the use of these data for statistics, e.g. for benchmarking. It is an instance of multi-party computation meaning that several parties participate in a protocol and keep their input values confidential (Goldreich et al. 1987; Kerschbaum 2008; Bogetoft et al. 2009). By using homomorphic encryption schemes, personal data can be kept confidential for a computation, however they are not suitable if data in cleartext are required.

4.3 Anonymity and Pseudonymity Mechanisms

Anonymity and pseudonymity mechanisms aim at non-linkability of individual's transactions. Both focus on the collection of personal data. Whereas **anonymity mechanisms** are not suitable for personalized service since they either give unlimited access to personal data or prohibit any access, pseudonymity mechanisms allow a controlled disclosure of a subset of an individual's personal data. In 1985, David Chaum introduced the concept of **identity management** for non-linkability by using

pseudonyms for authentication. The aim is both to protect an individual against a non-authorized data collection and profiling and to assure accountability of their transactions (Chaum 1985). The aim of current **identity management systems** is twofold: to preserve informational self-determination and to simplify authentication in the sense of Single Sign-On (SSO). Our investigation embraces the authentication protocols of the following identity management systems and protocols: Shibboleth (Erdos and Cantor 2004), Liberty Alliance (Wason 2004), iManager (Wohlgemuth et al. 2004), and IBM idemix (Camenisch and van Herreweghen 2002).

Shibboleth and **Liberty Alliance** are identity management systems with an identity provider. The role of an identity provider is to certify an individual's identity and to manage his attributes. This implies two different trust models. Firstly, an identity provider is a certification authority (CA), and service providers as well as individuals trust a CA to certify identities according to its certification policy. This is the usual trust model as it is used in public-key infrastructures (Ford and Baum 1997). Secondly, since an identity provider manages an individual's attributes such as his pseudonyms, an identity provider is able to trace this individual. It follows that a user has to trust his identity provider to keep the user's attributes in confidence according to the agreed privacy policy. This is the difference to a CA, so we call an identity provider with these trust relationships a privacy CA.

The identity manager **iManager** focuses on the usability of identity management for security novices by managing the identity of its user by means of partial identities and self-signed credentials. **IBM idemix** is an anonymous credential system and part of the PRIME identity management system (Camenisch et al. 2005). Its main property is that an individual can use the same anonymous credentials with several pseudonyms. Hence, the individual's transactions are non-linkable and accountable. Even the credential issuing party, e.g., a CA, cannot link these transactions with additional information.

However, if these identity management systems are applied to business processes requiring disclosure of personal data to third parties, individuals will lose control of the usage of their credentials. If these

credentials were used to access the individual's data which are stored at another service provider, this individual would lose control on the disclosure of these data. This is due to the fact that identity management systems follow the all-or-nothing delegation principle (Wohlgemuth and Müller 2006).

5 Usage Control and Enforcement of Services

In principle PET realizes access control on personal data, which is controlled by the respective individual. However, PET assume that (a) the user is aware of the data collection and (b) pseudonymized or encrypted data are sufficient to be used by service providers. An unconscious data collection undermines the protection by PET (Sackmann et al. 2006) and business processes for personalized services require personal data in clear text (Wohlgemuth 2008).

The concept of usage control extends the access control by means of rules which have to be followed after access has taken place (Park and Sandhu 2004). These rules are called obligations. They describe acceptable states of for the usage of personal data without restricting their access in advance. Together with the rules for obtaining access to data – provisions – obligations are part of a privacy policy. While provisions define known safety properties and are therefore decidable, obligations are in general not decidable from a start but can become so during the execution of a process. If obligations specify a restricted timeline for their enforcement, they are observable. However, non-observable obligations can be transformed into more strict and observable obligations. Depending on their type, obligations can be enforced by usage control mechanisms (a) before, (b) during, and (c) after an execution of a business process.

The enforcement of obligations is linked to their properties of controllability and observability (Hilty et al. 2005). Regarding controllability, a data provider can ensure that a data consumer enforces the agreed-upon obligations, e.g., by extending his access control domain to the data consumer with mechanisms of Digital Rights Management or by verifying the data consumer's process for using the disclosed data. However, this is not always practicable. On the other hand, observability means that

a data provider can check whether a data consumer has followed the agreed-upon obligations without extending his access control domain. Observable obligations are rules that can be enforced directly by the data provider's reference monitor, e.g., "re-access the data within the next k days". Non-observable delegations are access requests which are outside of the data provider's reference monitor's domain, e.g., "do not disclose data to given third parties", "use data only for the agreed-upon purposes", and "delete data after k days at the latest". To enforce non-observable obligations, a reference monitor should be enhanced by taking signals of the data consumer into account or to decrease the data quality (Pretschner et al. 2006; Sackmann 2007). A reference monitor is a trusted component intercepting each and every request to the system. That means a reference monitor mediates the complete access to the system and its resources, and it may not be possible to alter it (Anderson 1972).

5.1 Enforcement of Obligations

Obligations can be enforced before access to data, e.g., by integrating organizational control instruments into a business process and by verifying before its execution whether the process includes such controls. At the moment the data is accessed, this can be done e.g. by a monitor with a history of events and knowledge of the probable future process executions or by encrypting personal data before disclosing it. After an access obligations can be enforced by an audit regarding the granted access requests by the data consumer (Müller et al. 2010). In order to enforce obligations before access to personal data has taken place, patterns for control instruments can be integrated into the business process to be executed (Namiri and Stojanovic 2007). The drawbacks are that a business process cannot be changed afterwards. If the personal data have to be used in clear text, encryption protects only one disclosure of personal data to a third party. At the time of access, reference monitors could decide on the access request based upon a likelihood calculation regarding a violation of an obligation if access is granted. After access has been granted and the data have been decrypted, a data provider has no control of the usage of the data in clear text. Ex post enforcement of obligations is achieved by an audit. An audit takes the

granted accesses to personal data into account. Thereby, a usage of personal data by the data consumer and a further disclosure of personal data to third parties can be checked by an auditor, e.g. a data protection officer. Privacy evidences are a result of an audit and indicate possible violations of an obligation (Sackmann 2007). EXAMINA (Accorsi 2008) inverts policy rules in order to identify violations by one service by means of falsification. The APPLE framework defines a logic for an ex post enforcement of policies on trust management (Etalle and Winsborough 2007). This logic framework using logs to verify that actions taken by the system were authorized, assumes a secure system. This is in contradiction to our assumption that Cloud services may violate the agreed-upon privacy policy. The Optimistic Security Model of Povey (1999) presents a formal model to ensure that a rollback of an information system from executed privileges/policy breaches is possible. The protection goal however is integrity of the system and not confidentiality. If data were no longer confidential, a rollback of the system would not restore the confidentiality of data.

5.2 Delegation of Rights

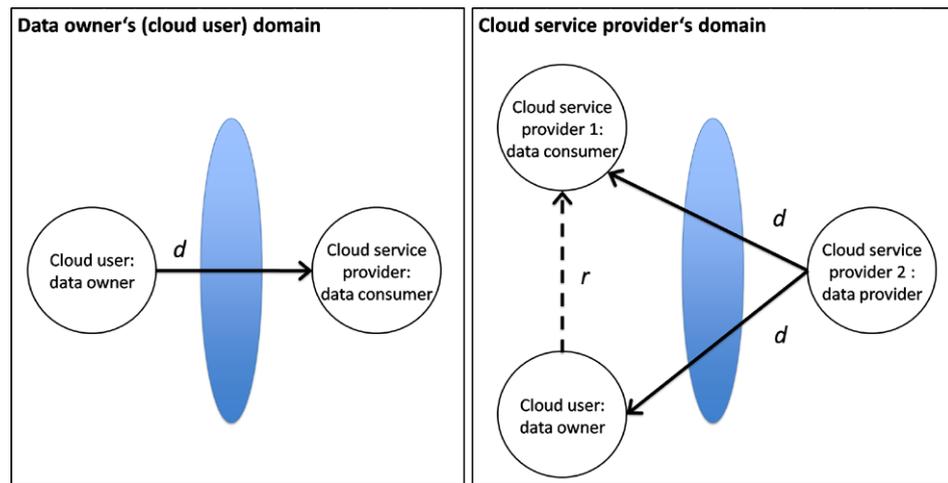
Our proposal is to establish isolation by defining the authorized cooperation relationship between Cloud users and Cloud services, i.e. defining the authorized data flows. Since a cooperation relationship depends on the instance of a business process in a Cloud, the privacy policy (authorizations) must be modifiable according to the business process and revocable if a cooperation relationship does no longer exist. Our approach makes use of distributed trust management of credentials (Blaze et al. 1996), since it supports a case-by-case issue and revocation of authorizations by binding them to credentials. Credentials are a representation of access rights, which are delegated to service providers to obtain access to personal data in agreement with the individual in question. After having granted access to personal data, they are no longer in the access control domain of the data provider. The access control model consists of two access control domains each possessing a reference monitor for access decisions: the individual's domain and a service provider's domain. The individual (in the role of the data subject) specifies these access decisions by delegating the access rights together with obligations

for using these access rights to the requesting service provider (in the role of a data consumer). To obtain a user's agreement for each disclosure of his personal data to a third party, it should be possible to delegate rights and to revoke them for a single disclosure. Thereby, a delegation of rights defines a collaboration between service providers including the exchange of given user's personal data between them. **Figure 3** shows the usage control model with the delegation of rights. Ellipses represent a reference monitor. The dashed line shows a delegation of a right r from a data subject to a subject in the role of a data consumer. The continuous lines represent granted access to the personal data d of the individual. In the data subject' domain, the user controls access to his personal data; in the service provider's domain, the service provider 2 controls access to the user's data according to the enforcement of the delegated right r . The service provider 1 shows his authorization to obtain access to the personal data d by showing the right r to the service provider 2.

Our approach consists of the following participants:

- **Cloud user:** A Cloud user acts as a data provider to Cloud services by disclosing his data to them. According to the requirement of data protection legislations authorizing the processing of personal data on the premise of the corresponding individual's agreement, the Cloud user acts as the data owner and delegates the authorizations to the cooperating Cloud services.
- **Cloud services:** A Cloud service acts as data consumer under the identity of his service provider when processing the data of a Cloud user. According to Pretschner et al. (2006) a Cloud service acts as a data provider if it discloses the data to another Cloud service which then acts as a data consumer.
- **Auditor:** An auditor checks whether an authorization of a privacy policy has been violated. If a privacy policy has been violated, the auditor should identify the originator of this privacy policy violation, i.e. identify the Cloud service which has disclosed the Cloud user's data to another Cloud service without the authorization of the Cloud user as owner of the corresponding data.
- **Certification authority:** A CA certifies the identities of the participants and issues authorizations to Cloud services on behalf of the Cloud user as the data

Fig. 3 Access control domains of delegation of rights



owner. In addition, a certification authority ensures the availability of the delegated authorizations to an auditor. The assumptions are as follows:

- The CA has certified the identity of every participant in the system.
- Every participant protects access to its data storage by an access control system.
- A participant can only access data of another participant if the latter has successfully shown his certified identity to the access control system of the Cloud service.
- Access rights are of the set of rules {*read, write, delete*}.

The requirements for isolation of data usage in Cloud Computing by delegation of rights are:

- Access to a Cloud user's data is only granted if the requesting Cloud service possesses the corresponding valid authorization. Thereby it is possible for a Cloud service to obtain access to data of different Cloud users as long as they have authorized the Cloud service to access the data. If a Cloud service requests data without showing a valid authorization, the access request must be denied.
- A CA only issues an authorization to a service provider for access to data of a user (data owner) stored at another service provider, if the corresponding Cloud user (data owner) has explicitly agreed to the requested data access. An authorization issued by a CA proves to an auditor that the corresponding user has given his agreement to the service provider to obtain access to his data within a business process.
- Only the participating services should get access to those Cloud user's data which are necessary to run the business

process. Additional data of the Cloud user should not be disclosed or gained by the delegation of the authorizations.

- A Cloud service S_i is allowed to write data d of Cloud user U_m into the dataset Cloud user U_n , if
 - Cloud user U_m has given the Cloud service S_i the authorization to write d into the dataset of U_n or if the data d have been anonymized AND
 - Cloud user U_n has given the Cloud service S_i the authorization to write d into its dataset.
- A Cloud service S_j is allowed to send d of Cloud user U_m to the Cloud service S_j , if Cloud user U_m has given the Cloud service S_j the authorization to read d from Cloud service S_j or the data d have been anonymized.
- If a data owner has removed the authorization to access data d for the Cloud service S_i , then Cloud service S_j has to delete the corresponding data d . This has to be done recursively.
- To enable an auditor to identify the originator of a non-authorized disclosure of a Cloud user's data, the history of the Cloud user's data flow has to be documented with the data. This ensures that the auditor can re-construct the data flow and compare it with the delegated authorizations of the corresponding Cloud user.

5.3 DREISAM—An Experimental System for Usage Control

The DREISAM protocols extend identity management to a non-linkable delegation and revocation of rights. Their novelty is (a) to grant access to given personal data without disclosing any further identifying data about the individual and (b) to keep the personal data in clear text

so that the data center can use them for its service according to the agreed-upon privacy policy. Therefore, these protocols combine the mechanisms for delegation of rights by credentials with mechanisms for enforcing non-linkability when using credentials. These are the building blocks of the DREISAM credential system.

The non-linkability mechanisms for credentials are cryptographic modules. Anonymous credentials, as they are used by IBM idemix, make use of a cryptographic commitment scheme for binding authorizations to a cryptographic key and of zero-knowledge proofs for showing this relationship without revealing any identifying data, but with the mentioned disadvantage of the all-or-nothing delegation. An identity manager of the individual's client software uses the anonymous credentials of the user and executes its delegation decisions.

A proxy credential replaces the sharing of an individual's master identity and cryptographic key k . For the CA it represents the individual's delegation request for a certain right to a service provider. If the service provider obtains a proxy credential, he has the individual's authorization to get the requested access right by means of an anonymous credential.

The CA logs requests from users and service providers with the issued proxy and anonymous credentials in the delegation list. This list is a realization of the access control matrix (Harrison et al. 1976). The CA uses this list to check the service providers' requests for anonymous credentials and to resolve disputes between users and service providers. An entry of this list refers to the delegation of an access right of the user to service providers. To allow the CA to resolve disputes, the pertinent credential of

the user is stored together with the transcript of the showing protocol's run. To enforce the maximum number of access rights uses, the CA records the number of issuances of anonymous credentials to service providers and issues only anonymous one-show credentials concerning a delegation.

A detailed description of the DREISAM protocols and their proof-of-concept implementation for the example of customer relationship management is given in Wohlgemuth and Müller (2006).

6 Outlook

Access control and security policies are not sufficient to assure isolation in Cloud processing on all three layers of Cloud usage. The reason is that granting access is just the first step when meeting the challenge of isolation in Clouds. The data and information flows, as they can be seen in the actual service usage, are of economic interest. They should be determined on a contractual basis. The key to this is the specification of obligations and delegations of rights as well as the enforcement of these characteristics. With the DREISAM protocols for a non-linkable delegation of rights, one aspect has been shown. The definition of data disclosure to third parties without disclosing the individual's additional data is one important aspect among others. Especially the enforcement of delegated rights remains an open issue.

Acknowledgements

This work was funded by the FIT-NII-Postdoctoral-Program of the German Academic Exchange Service (DAAD) and is a result of the Japanese-European Institute for Security (JEISec) at the National Institute of Informatics (Japan). The authors also gratefully acknowledge the helpful comments and suggestions of the reviewers, which have improved the article.

References

Accorsi A (2008) Automated privacy audits to complement the notion of control for identity management. In: Fischer-Hübner S, Tseng JC, Borking J (eds) Proc of first IFIP conference on policies and research in identity management (IDMAN'07), Rotterdam

- Alpern B, Schneider F (1985) Defining liveness. *Inf Process Lett* 21(4):181–185
- Anderson JP (1972) Computer security technology planning study. Technical report ESD-TR-73-51, Electronic system division/AFSC, Bedford, MA
- Armbrust M, Fox A, Griffith R, Joseph A, Katz R, Konwinski A, Lee G, Patterson D, Rabkin A, Stoica I, Zaharia M (2010) A view of cloud computing. *Commun ACM* 53(4):50–58
- Ashley P, Hada S, Karjoth G, Powers C, Schunter M (2003) Enterprise privacy authorization language (EPAL 1.2). <http://www.w3.org/Submission/EPAL/>. Accessed 2011-02-10
- Blaze M, Feigenbaum J, Lacy J (1996) Decentralized trust management. In: Symposium on security and privacy, Oakland
- Bogetoft P, Christensen DL, Damgard I, Geisler M, Jakobsen T, Krogaard M, Nielsen JD, Nielsen JB, Nielsen K, Pagter J, Schwartzbach M, Toft T (2009) Secure multiparty computation goes live. In: Dingledine R, Golle P (eds) *Financial cryptography and data security*, Barbados
- Bundesverfassungsgericht (1983) Volkszählungsurteil. In: *Entscheidungen des Bundesverfassungsgerichts*. Urteil vom 1983-12-15. Az.: 1 BvR 209/83; NJW 84, 419
- Buneman P, Khanna S, Tan WC (2001) Why and where: a characterization of data provenance. In: 8th int conf on database theory, London
- Camenisch J, van Herreweghen E (2002) Design and implementation of the idemix anonymous credential system. In: Proc of the 9th ACM conf on computer and communications security, Washington, DC
- Camenisch J, Shelat A, Sommer D, Fischer-Hübner S, Hansen M, Krasemann H, Lacoste G, Leenes R, Tseng J (2005) Privacy and identity management for everyone. In: Proc of the 2005 workshop on digital identity management, DIM 05, Fairfax, VA
- Casassa MM, Pearson S (2005) An adaptive privacy management system for data repositories. In: Katsikas SK, Lopez J, Pernul G (eds) *TrustBus 2005*, Copenhagen
- Chaum D (1985) Security without identification: transaction systems to make big brother obsolete. *Commun ACM* 28(10):1030–1077
- Cox IJ, Miller ML, Bloom JA, Fridrich J, Kalker T (2008) Digital watermarking and steganography. Morgan Kaufmann, Los Altos
- Cranor L, Langheinrich M, Marchiori M, Presler-Marshall M, Reagle J (2002) The platform for privacy preferences 1.0 (P3P1.0) specification. <http://www.w3.org/TR/P3P>. Accessed 2011-02-10
- Ellison G (ed) (2005) Liberty. ID-WSF security mechanisms version: 1.2. Liberty alliance project. <http://www.projectliberty.org/specs/liberty-idwsf-security-mechanisms-v1.2.pdf>. Accessed 2011-02-10
- Erdos M, Cantor S (2004) Shibolet-Architecture DRAFT v05. <http://shibolet.internet2.edu/docs/draft-internet2-shibolet-arch-v05.pdf>. Accessed 2011-02-10
- Etalle S, Winsborough WH (2007) A posteriori compliance control. In: ACM SACMAT'07, Nice-Sophia Antipolis
- European Commission (1995) Directive 95/46/EC of the European parliament and of the council of 24 October 1995 on the protection of individuals with regard to the processing of personal data and on the free movement of such data. *Official Journal of the European Communities*, L 281(395L0046):31–50

Abstract

Noboru Sonehara, Isao Echizen, Sven Wohlgemuth

Isolation in Cloud Computing and Privacy-Enhancing Technologies

Suitability of Privacy-Enhancing Technologies for Separating Data Usage in Business Processes

Cloud Computing lifts the borders between the access control domain of individuals' and companies' IT systems by processing their data within the application frameworks and virtualized runtime environments of Cloud service providers. A deployment of traditional security policies for enforcing confidentiality of Cloud users' data would lead to a conflict with the availability of the Cloud's software services: confidentiality of data would be assured but Cloud services would not be available for every user of a Cloud. This state-of-the-art contribution shows the analogy of the confidentiality of external data processing by Cloud services with mechanisms known and applied in privacy. Sustainability in Cloud is a matter of privacy, which in Cloud is called "isolation".

Keywords: Cloud computing, Disclosure of personal data to third parties, Isolation, Privacy, Usage control, Delegation of rights

- European Commission (2002) Directive 2002/58/EC of the European parliament and of the council of 12 July 2002 concerning the protection of personal data and the protection of privacy in the electronic communications sector (directive on privacy and electronic communications). Official Journal of the European Commission L201:37–47
- Ford W, Baum M (1997) Secure electronic commerce. Prentice-Hall, New York
- Goldreich O, Micali S, Wigderson A (1987) How to play ANY mental game. In: Aho AV (ed) Proc of the 19th annual ACM symposium on theory of computing (STOC'87), New York
- Haas S, Wohlgemuth S, Echizen I, Sonohara N, Müller G (2010) Aspects of privacy for electronic health records. International Journal of Medical Informatics for its special issue on security. <http://dx.doi.org/10.1016/j.ijmedinf.2010.10.001>
- Harrison MA, Ruzzo WL, Ullman JD (1976) Protection in operating systems. Commun ACM 19(8):461–471
- Hilty M, Basin D, Pretschner A (2005) On obligations. In: European symp on research in computer security (ESORICS 2005), Milan
- Karjoth G, Schunter M, Waidner M (2002) Privacy-enabled services for enterprises. In: 13th int workshop on database and expert systems applications, Aix-En-Provence
- Karjoth G, Schunter M, Waidner M (2003) Platform for enterprise privacy practices: privacy-enabled management of customer data. In: 2nd workshop on privacy enhancing technologies (PET 2002), San Francisco
- Kerschbaum F (2008) Building a privacy-preserving benchmarking enterprise system. Enterprise Information Systems 2(4):421–441
- Namiri K, Stojanovic N (2007) Using control patterns in business processes compliance. In: Int conf on web information systems engineering (WISE), New York
- Mather T, Kumaraswamy S, Latif S (2009) Cloud security and privacy: an enterprise perspective on risks and compliance. O'Reilly, Sebastopol
- Müller G, Accorsi R, Höhn S, Sackmann S (2010) Sichere Nutzungskontrolle für mehr Transparenz in Finanzmärkten. Informatik-Spektrum 33(1):3–14
- Organisation for Economic Co-operation and Development (1980) OECD guidelines on the protection of privacy and transborder flows of personal data. http://www.oecd.org/document/18/0,2340,en_2649_34255_1815186_1_1_1_1,00.html. Accessed 2011-02-10
- Park J, Sandhu R (2004) The UCON_{ABC} usage control model. 24th ACM Transactions on Information and System Security 7(1):128–174
- Povey D (1999) Optimistic security: a new access control paradigm. In: ACM new security paradigm workshop'99, Caledon Hills
- Pretschner A, Hilty M, Basin D (2006) Distributed usage control. Commun ACM 49(9):39–44
- Roßnagel A (2005) Modernisierung des Datenschutzrechts für eine Welt allgegenwärtiger Datenverarbeitung Multimedia und Recht 8(2)
- Sackmann S, Strüker J, Accorsi R (2006) Personalization in privacy-aware highly dynamic systems. Commun ACM 49(9):32–38
- Sackmann S (2007) Personalization and privacy in ubiquitous computing – resolving the conflict by legally binding commitments. In: IEEE conference on E-commerce technology (CEC'07), Tokyo
- US Department of health & human services (1996) Health insurance portability and accountability act of 1996 privacy rule. <http://www.cms.hhs.gov/HIPAAGenInfo>. Accessed 2011-02-10
- Smith RE (1993) The law of privacy in a nutshell. Privacy Journal 19(6):50–51
- Wason T (ed) (2004) Liberty ID-FF architecture overview version: 1.2. Liberty alliance project. <http://www.projectliberty.org/specs/liberty-idff-arch-overview-v1.2.pdf>. Accessed 2011-02-10
- Westin A (1967) Privacy and freedom. Atheneum, New York
- Wohlgemuth S (2008) Privatsphäre durch die Delegation von Rechten. Vieweg+Teubner, Wiesbaden
- Wohlgemuth S, Jendricke U, Gerd tom Markotten D, Dorner F, Müller G (2004) Sicherheit und Benutzbarkeit durch Identitätsmanagement. In: Spath D, Haasis K (eds) Tagungsband zum doIT Software-Forschungstag 2003, Stuttgart
- Wohlgemuth S, Müller G, (2006) Privacy with delegation of rights by identity management. In: Emerging trends in information and communication security (ETRICS 2006), Freiburg i.Br.