Sticky Policies: An Approach for Privacy Management across Multiple Parties

Authors: Siani Pearson and Marco Casassa Mont

Cloud and Security Research Lab, Long Down Avenue, Stoke Gifford, Bristol. BS34 8QZ.

Phone: +44 117 3128438/3128794 FAX: +44 1173129250

Email: siani.pearson@hp.com; marco.casassa-mont@hp.com

Abstract

Organisations often have good privacy procedures in place for protection of personal, sensitive or confidential information within that organisation, but when this passes the organisational boundaries, the risk of inappropriate usage or exposure, whether deliberate or accidental, greatly increases. To mitigate this risk, technical mechanisms and solutions are needed, in addition to legal agreements and contracts, that enforce the wishes of end users and of organizations acting on behalf of customers or employees regarding the way in which that information is used. In this article we describe how 'sticky policies' can help provide a such technical solution by means of a user-centric control mechanism involving machine-readable policies (defining allowed usage and associated obligations) that are attached to data and travel with it as it is passed among multiple parties. This approach is suitable for enhancing privacy management in a broad range of domains, and particularly in sectors like healthcare, where sensitive information is involved, or to provide privacy protection in the cloud.

Keywords: sticky policies, consent, revocation, privacy management, EnCoRe

Introduction

The core problem addressed in this article is how to ensure fulfilment of privacy across organisational boundaries. The current mechanism for tackling this is by means of legal and business frameworks, including business contracts and Service Level Agreements. We believe that technical mechanisms can enhance and complement such approaches, in particular by supporting enforcement and auditing of the organisational obligations agreed in such frameworks. Before we consider how such a technical approach for privacy management across multiple parties may be achieved, let us consider what is meant by 'privacy management' and why it is an important issue that needs to be addressed.

Personally identifiable information (PII) can be defined as *information that can be traced to a particular individual.* It includes such things as name, address, phone number, social security number, national identity number, credit card numbers, email address, passwords and date of birth. It is also referred to as personal data or personal information, and the definition of this varies according to geography. Sensitive information – such as financial or medical information - can be considered as a subset of personal information, and because of its sensitive nature greater care must be taken in its handling.

In commercial contexts where consumers are involved, privacy is about the protection and careful use of the PII of customers, and about meeting the expectations of customers about the use of their PII. For corporations, privacy is about the application of laws, policies, standards and processes by which PII of individuals is managed. These ways in which organisations and individuals can control collection, usage and sharing of personal data, including sensitive information, is what *privacy management* is all about.

Privacy management is an important issue for organisations, as it can help foster trust with customers and data subjects, as well as there being regulatory requirements for compliance to data protection legislation. In a given context, there may be many different privacy-related regulatory requirements, including sector-specific laws, different national legislative requirements and transborder data flow restrictions [1]. Although assessing requirements in a given situation can be complex, there is a basic set of privacy principles that form the basis of most privacy legislation around the world [2].

To provide mechanisms for online privacy management, there has already been substantial research carried out related to anonymisation technologies, enforcement of privacy policies (for example, enforcement of privacy-enhanced access control policies, as considered for example in the Prime and PrimeLife EU projects; policy lifecycle management, etc.), and on compliance with global regulations relating to data protection (for example, tools for Governance, Risk and Compliance (GRC), modelling privacy regulations [1] and modelling organisational privacy policies and checking these down to an operational level [3]). However, major outstanding issues remain to be adequately addressed, including how to provide more control to end users, how to gather and manage end users' consent (and subsequent revocations) and how to make privacy management effective when information is transmitted across parties. In this article, we focus on these issues. Specifically, we introduce and discuss an approach based on sticky policies to tackle these problems. By the means we lay a foundation for an approach and related solutions that enable compliance with and enforcement of current requirements (e.g. Health Insurance Portability and Accountability Act (HIPAA) of 1996) along with future needs emerging by the adoption of new technologies and models (e.g. storage and processing of sensitive data in the cloud) .

Scenarios

As illustrated in Figure 1, we consider a scenario where users' confidential information, collected by an organisation, flows across organisational boundaries. For example, in a health care scenario, personal data and preferences could be disclosed to a GP via an online service. This information might need to be shared with hospital specialists, pharmaceutical companies and other third parties involved in the healthcare supply chain. A similar situation might apply for a travel agency service that needs to share data with various service providers (SPs) such as hotel booking, car rental, etc. More generally, these kinds of scenario will be increasingly common in a cloud computing environment, where users interact with front-end SPs which will have to share part of the information with other SPs in order to supply the required services.

In all these situations, users need to reveal personal and even sensitive information in order to receive a service, but wish to control the way in which that information is used. We aim to enable users to directly control how their data should be processed, handled and shared by explicitly expressing their

preferences and data handling policies. We want these choices to be respected all along the service provision chain, including allowing the user to update these choices. We will show how this can be achieved by means of propagating these user choices to all the SPs and by deploying a number of mechanisms to ensure the policies are respected. Moreover, the user can be actively involved in the selection of multiple, interchangeable services that will track and audit the policy fulfilment.

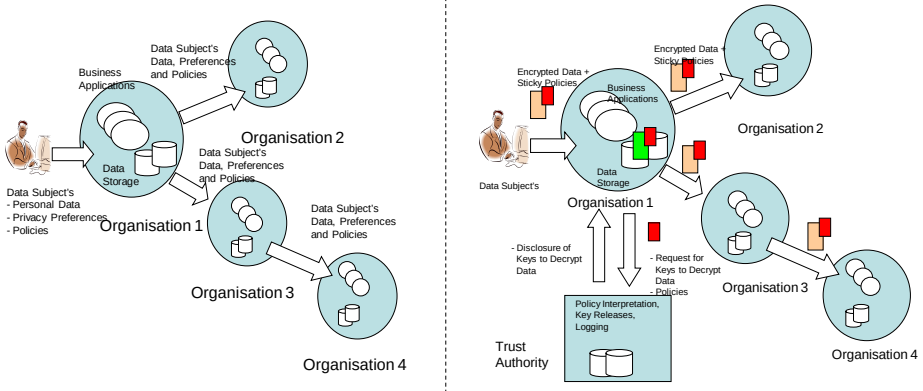

Figure 1: High level scenario and related management of sticky policies

Introducing Sticky Policies

Sticky policies are conditions and constraints attached ('stuck') to data that describe how that data should be treated. Depending on the degree of the policy stickiness, the involved data might be encrypted and access to their content in clear allowed only upon the satisfaction of these policies. Specifically, the policies govern the use of associated data, and may specify the following:

- Purposes of using data (e.g. for research, transaction processing, etc.).
- Data may only be used within a given set of platforms (with certain security characteristics), a given network or a subset of the enterprise
- Specific obligations and prohibitions (allowed third parties, people or processes; blacklists; notification of disclosure; deletion or minimisation of data after a certain time)
- List of Trusted Authorities (TAs) that will provide assurance and accountability in the process of granting access to the protected data (potentially the result of a negotiation process)

Figure 1 illustrates the core principles and mechanisms in place to handle sticky policies. In our approach, policies are strongly associated to the data by means of cryptographic mechanisms. There can

be different degrees of stickiness: we adopt a strong binding as it provides better accountability. Data are encrypted and only accessible upon the acceptance and satisfaction of constraints and duties imposed by the policies. TAs provide assurance by keeping track of promises made by the involved parties, in order to access data, along with controlling such access to data. The role of the TAs may be integrated with other functionality, such as being a consumer organisation, a certification authority, a well-known organisation or else it might be performed by a client-side software component or service that is under the control of end users or other parties, or it could be achieved using distributed components or a peer-to-peer mechanism.

The deployment of such a system is reasonably straightforward, as it does not require change from existing trusted third parties (apart from dealing with additional policy condition checks) nor from storage providers (if these are used to store the data and an authenticated reference is passed around instead of the data). However, SPs would need to handle the management of packaged sticky policies, or use an application to do this locally. This includes additional interactions with the TA and release of statements certifying their willingness to fulfill the policies. Hence this technique it is likely to be most suitable for service provision environments where the increased trust and protection would justify the additional expense, or alternatively to business partners of 'goodwilling' enterprises, who might encourage its use.

As illustrated in Figure 1, sticky policies are passed between organisations in order to capture obligations and other constraints that must be met by the receiving parties in order to access and use the associated personal data.

For example, if a healthcare record were to be passed from a hospital to a research institution and then a research team, it might be passed in a form where certain attributes (e.g. medical results, personal information such as name and address, etc.) are encrypted, with a sticky policy associated with this describing how parts of this could be used. Let's assume that the patient only wants this to be released to research teams, wants the information to be deleted after 3 years and wants to be notified every time their medical information is passed on. These constraints can be expressed in a number of different ways, for example just using a simple XML format.

Sticky policies can help enable accountable management and disclosure of confidential data across boundaries, using the approach shown in Figure 1. Personal, private or confidential information that is stored and used is associated with machine-readable policies, in such a way that this cannot be compromised. When information is processed, this is done in such a way as to adhere to these constraints: as the data is replicated or shared in order to fulfil the service provision request, mechanisms will be in place to ensure that the customer's preferences are respected right along the chain. Specifically, keys to decrypt data need to be retrieved by TAs which will log all promises made by the requestors. This information can be used for forensic analysis if there are policy violations.

Figure 2 shows the basic mechanisms underpinning the management of sticky policies, using a Public Key Infrastructure (PKI) based approach. We provide more details in the following section about how PKI and other mechanisms may be achieved by using a variety of cryptographic techniques.

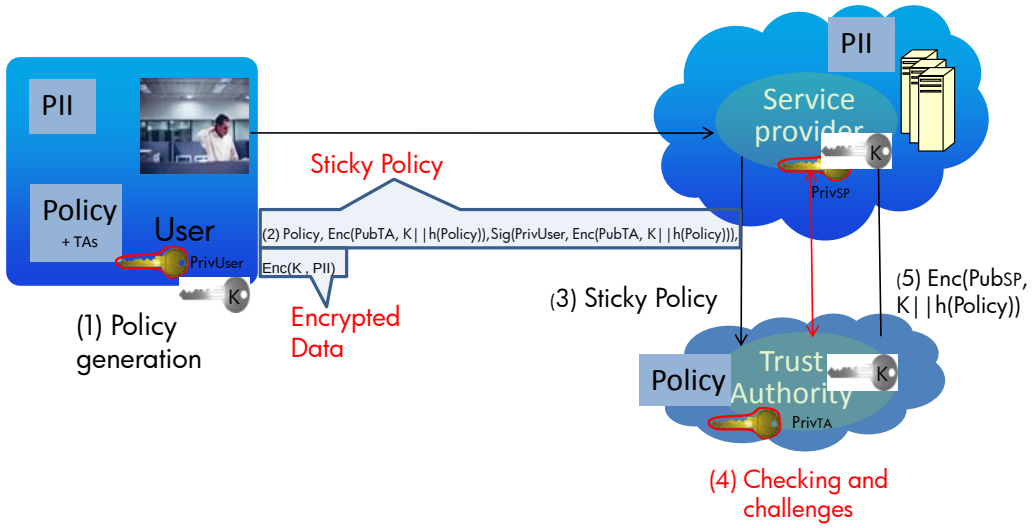

Figure 2: Macro Architecture

As shown in Figure 2, the following aspects are part of our solution:

- In order to be able to more easily interpret and enforce end user policies, instead of offering free expression of policies from end users, their preferences and policies are defined within a framework imposed by organisations. There are different ways of achieving this: one mechanism is that SPs publish a "Manifesto" containing the list of supported (macro) policies and TAs, where these policies relate to access control and obligation behaviours supported by the organisation;
- A user (customer) – interacting with a SP - can select the granularity of how policies apply to items or specific subsets of personal data to be disclosed (ranging from coarse grained to fine-grained) and customise related preferences (e.g. notification preferences, period of time after deletion, set of agreed purposes, list of parties not to interact with, etc.).
- The user selects a subset of TAs that are to be trusted
- Based on the above selections, a client-side component supports the creation of sticky policies and their association to data, i.e. the bundling of policies, preferences, data and TAs. In other words, the client-side component deals with the packaging of data along with selected parameterised policies and TAs.

- The user can select the option to refer to secured data (e.g. personally identifiable information (PII)) by another third party (this is a storage provider that stores the encrypted data) rather than passing the encrypted data directly to the SP.
- Encrypted data along with sticky policies is sent to the SP.
- In order to gain access to the data in clear, the SP needs to interact with one of the selected TAs (based on availability). During this interaction the SP has to assert its willingness to fulfil the customised, sticky policies, or – dependent upon the policy requirements – the TA may be able to check this independent of such signed statements, for example with reference to blacklists or reputation management systems that may be externally maintained, or by verifying system properties using mechanisms such as trusted attestation or remote software verification. This creates an audit trail that can be afterward used by the user and TA – in case of policy violations and misbehaviour.
- The SP will allow for a predefined period of time for connection with the TA. The solution supports the swapping between TAs based on needs.
- Only after satisfying all these requirements (and upon checking additional contextual information), can the TA decide to release the information (decryption keys) that would enable the decryption of data.
- The TA will be able to decrypt and access the data – either in the case where the data was directly disclosed or in the case where just a reference to it was provided (in this latter case, the SP will need to fetch the data).

We envisage the deployment, within organisations, of privacy-management components that complement Identity and Access Management solutions, as experimented and tested in the context of the EnCoRe collaborative project [4]. Specifically these components will complement organisations' middleware solutions, in the space of Identity and Access Management to provide: privacy-aware access control, obligation management, data tracking, processing of sticky policies and interactions with TAs. The role of these TAs is not just to release keys but also to provide accountability by means of logging and auditing, and subsequently supporting forensic analysis.

Mechanisms for Stickiness

The original 'sticky policy' paradigm was espoused by Karjoth *et al.* [5], and specifies that privacy preferences should flow with personal data to make sure that they can always be enforced. However, no method for strong stickiness of policies to data was suggested. Initial work illustrating how policies can be strongly associated to data was espoused by the authors of this article [11].

The following common central approach applies to all the variations of mechanism considered below. Customers allow SPs to have access to specific data based on agreed policies and by forcing interactions with interchangeable independent third parties (the TAs). The access to data can be as fine-grained as necessary, based on policy definitions, underlying encryption mechanisms (supporting the stickiness of policies to the data) and a related key management approach that allows (sets of) data attribute(s) to be encrypted specifically based on the policy. Access to data is mediated by a TA that checks for compliance to policies in order to release decryption keys, in such a way that it is not necessarily the case that

checking for compliance is only about having the SP assert willingness to comply. By these means users can be provided with fine-grained control over access and usage of their data. This is the case even in public cloud models.

We describe below a variety of techniques that can be used to provide sticky policy-based protection of data from data subjects to SPs, based on different underlying encryption mechanisms. In each case, this technique may be extended to cover the propagation of data along the service provision chain. It is an analogous process to user to SP protocols, in which the first SP may add additional policy constraints to form a superset of the previous policy constraints. The proposed techniques can be refined and deployed in a variety of mechanisms currently used to exchange information, including web technologies and protocols (e.g. http/s, SOAP, etc.), document formatting and protection (e.g. Adobe, DRM techniques, etc.) and various messaging tools (e.g. emails, IM, etc.).

These protocols not only apply to users representing people, but also more broadly to interactions from machine to machine or from service to SP.

Sticky Policies using Public Key Encryption Techniques

In this case, we assume that all the stakeholders have, at least, certified public/private key pairs from trusted certification authorities (CAs). An approach to provide a strong binding that enhances integrity is to bind policies to data by encrypting the data under a symmetric key, conditionally shared by sender and receiver (i.e. based on fulfilment of policies), and sticking the data to the policy using public key enveloping techniques similar to Public Key Cryptography Standard (PKCS) 7. An example of this process is shown in Figure 2, where the labelled stages are as follows:

1. Generation of the policy by the sender, together with a symmetric key K used to encrypt the data (for efficiency, a symmetric key is used rather than an asymmetric key). If desired, this process may be generalized to allow different attributes to be encrypted separately (i.e. using different symmetric keys generated at this stage), and hence only part of the information revealed when an attribute is decrypted.
2. Generation of the message from the sender to the SP. One part is that the data is encrypted with the symmetric key K. The other part is a ‘sticky policy’, where first the symmetric key K appended to a hash of the policy is encrypted with the TA’s public key and then this is signed using the user’s private key. (This makes it possible to verify the source and integrity of the policy as well as binding the symmetric key K to the data and the policy.) The resultant ‘sticky policy’ is sent together with the encrypted data to the SP.
3. Generation of the message from SP to TA, which involves passing on just the sticky policy and encrypted shared keys.
4. The TA carries out policy checking, potentially including challenges to the SP. The SP may have to provide signed statements about its policies.
5. If all checks are fulfilled, the TA will release the shared key: it generates a message from the TA to SP, which involves encrypting the symmetric key K appended to the hash of the policy with the SP’s public key. By these means SP can get access to K and can check the integrity of the policy and then can decrypt the PII.

Further detail about this approach (for a cloud context) is given in [6].

Sticky Policies using Identifier-Based Encryption

Identifier-based encryption (IBE) [7] is a cryptographic schema where any kind of string (including a name, role, terms and conditions, etc.) can be used as a public encryption key. The generation of the corresponding IBE decryption key can be postponed until later. A TA can generate this decryption key on the fly, under specific circumstances.

Conceptually similar to the PKI approach described above, we adapt the IBE approach by mapping a 'sticky policy' to an IBE encryption key. The role of the TA is expanded to perform a policy checking role (it checks for the integrity and trustworthiness of the requestor's credentials and their IT environment before releasing the decryption key) and an auditing role (it logs and audits disclosures of confidential data).

Further detail about how IBE-based sticky policies can be used to protect content may be found in [8].

Variations on this Approach

We can potentially use any encryption mechanism in order to associate policies with data. For example, Voltage and Navajos provide format-preserving encryption and search-enabled encryption respectively; if the operation involves indexing, this could be used when encrypting the different attributes and thereby it would still be possible to do searching and indexing on the encrypted attributes.

An alternative solution, that permits binding of privacy preferences to data and conveying the consent of the individual as well, has been proposed by Pöhls [9]. However, the solution does not avoid the non-consented use of data.

This approach can be adapted to support multiple verification and control, and instead of each entity having a certificate they could be provided with a key component (called a 'share'). A secret sharing scheme such as Shamir's threshold-based secret sharing scheme [10] could be used to require *l* of *m* shares for the CSP to recover K and decrypt the PII, while still providing for some redundancy among TAs. Secret sharing schemes form a particular group of multi-party key establishment protocols that enable distribution of control or trust in many critical activities. The central idea of such a (*l*,*m*) threshold scheme is that a key (in our case, the key used to encrypt the data) would be divided into *m* pieces (called shares) such that any *l* of them can be used to reconstruct the whole original key, but using any number of shares less than *l* will not help to reconstruct the key.

TCG integrity checking mechanisms can be used to check that the receiver's platform is a trusted platform, that the software state of this platform is conformant with the disclosure policies and that the platform correctly implements defined privacy management mechanisms.

Furthermore, there are a number of variations on the approaches described above, including in terms of policy definition, the degrees of stickiness and the fine-grained nature of the encryption that occurs. The mechanisms are independent of the particular representation used for the policies. The protocols themselves may be amended: in the PKI approach, the binding of the policy to the data may be achieved

by the user within a signing operation rather than within the encryption, or by using 'signcryption'(e.g. signcryption algorithm specified in ISO/IEC 29150), so that the user need only perform this single operation and separately encrypt the data (or reference to the data); attributes can be encrypted with different keys, and there are different ways of exploiting that, including enveloping the sensitive data and passing it on without revealing the key from the TA, and with different attributes revealed to different entities in the chain.

Case Study: EnCoRe Project

In this section we show the feasibility of sticky policies by highlighting how this can be achieved by means of processes and components designed for privacy management, within the EnCoRe project [4]. EnCoRe is a collaborative research project being undertaken by UK academic and industrial partners that aims to give individuals more control over their personal information, by means of consent and revocation management. Here, revocation essentially means change of consent, potentially in a fine-grained way.

Mechanisms are provided for users to define consent policies and to be able to change these. As we shall see, sticky policies can be used to represent and enforce the consent and revocation preferences of end users. In general the EnCoRe system supports:

- *Explicit management of consent and revocation*: over a lifecycle that includes policy negotiation, setting, changing, enforcement, involvement in sticky policies, etc. This is achieved in an integrated way with the management of security and privacy policies. Compliance checking and auditing are integrated capabilities.

- *Bridging the disconnection between high and lower level policies.* This is done by mapping legal, business, social and security requirements into high-level policies; we define an intermediate conceptual framework to model policies and reason on top of them. Finally, we map these conceptual policies into monitorable and enforceable policies driven by users' preferences.

Our solution is applicable in a variety of business contexts, and is especially valuable where sensitive information is involved: for example, healthcare scenarios (including biobanks), employee data (and access to this by third parties), government scenarios, assisted living and cloud computing.

In the EnCoRe project we have developed a flexible toolbox solution that can be potentially customized and deployed in the context of each involved SP, consistently with their business processes. The presence of EnCoRe compliant capabilities – in a given service provider – will provide assurance about its privacy management practices and related management of consent and revocation. Assurance is further supported by using "sticky policy" mechanisms.

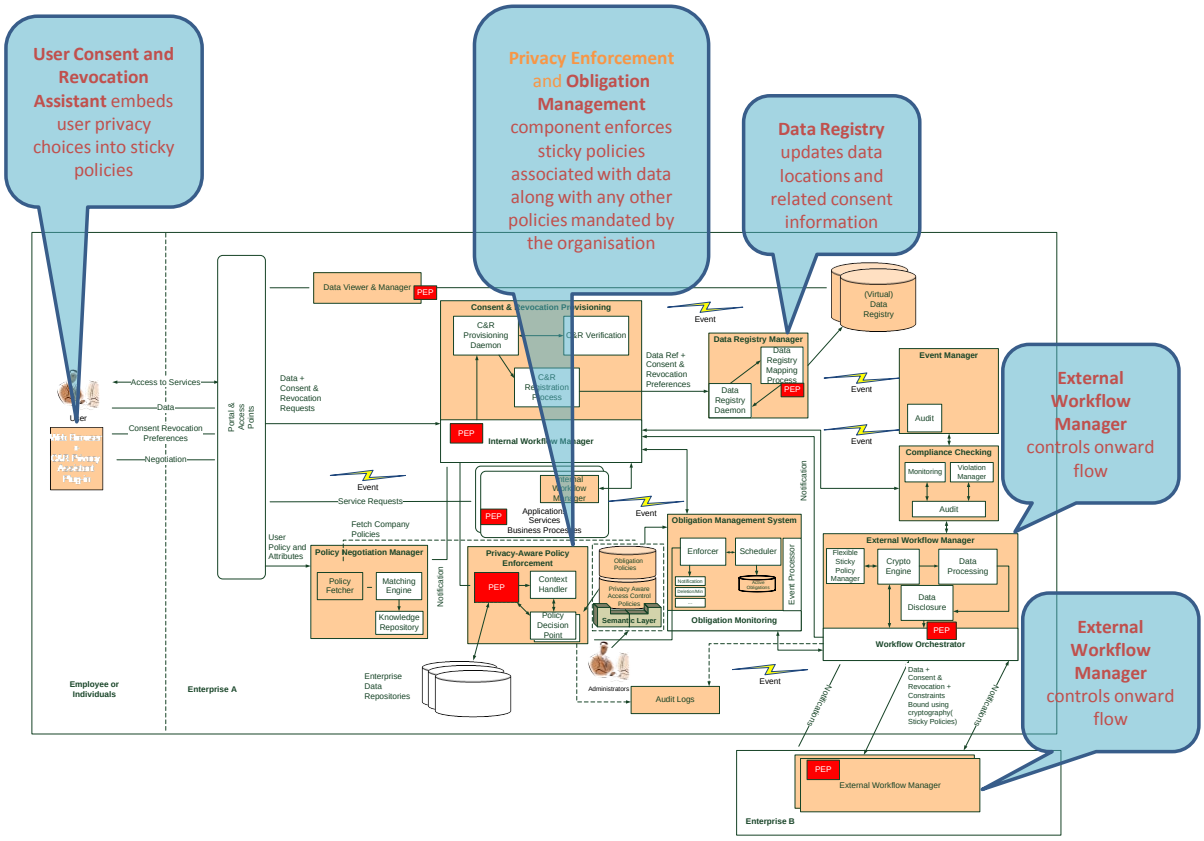

Figure 3: *High-level EnCoRe Architecture*

Figure 3 illustrates the overall set of functionalities and capabilities provided by the "EnCoRe toolbox"; these functionalities could be provided themselves as a set of services in the cloud or be deployed as an overall standalone infrastructural solution. These include:

- **Personal Consent & Revocation Assistant:** this component assists and provides user-side capabilities to help ordinary people express their consent by means of privacy choices (opt-in/opt-out choices, privacy preferences, etc.) and revocation requests, along with explanation of privacy practices provided by the organisations. It is triggered (e.g. via a plug-in in a web browser) during data disclosure processes. Part of these privacy choices can be embedded into sticky policies to ensure that they will be fulfilled by third parties receiving the data.

- **(Virtual) Data Registry:** this repository (or aggregation of synchronised repositories) keeps track, for each known individual, where their data has been stored within and outside the organisation, which type of data has been disclosed and to whom, along with any relevant associated sticky policies.

- **Consent and Revocation Provisioning:** this component is in charge of automatically updating the data registry every time there is a new expression of consent and revocation. It makes use of internal workflows to update individuals' preferences and constraints that affect the enforcement of access control and obligation policies.

- **Privacy-aware Policy Enforcement and Obligation Management**: this component deals with access control over data and obligations, driven by consent. It is in charge of enforcing sticky policies associated to the data along with any other policies mandated by the organisation.

- **External Workflow Manager**: this component intercepts and tracks flows of personal data, both within and between organisations, and propagates the associated consent information. In our scenarios, sticky policies are created (or derived from previous data subjects' policies), associated to the data to be disclosed and managed to ensure degrees of compliance to agreed policies and data subject's preferences. Applications and services might need to be instrumented with agents that communicate with this component.

- **Audit:** this is an essential component to log and track what happens to data, consent and revocation during operational and administrative activities, including flows of personal data within and beyond the organisation.

- **Compliance Checking and Risk Assurance**: this is a key offline component that is used by the enterprise's privacy administrators to assess current risks and provide indications of compliance.

The sticky policies sent out from the EnCoRe system to other organizations specify the purposes of using the data and any obligations and prohibitions (including notification and deletion after a certain time), that have been specified by the user in their consent and revocation preferences associated with that data. The TA functionality described above is distributed in the sense that the EnCoRe external workflow manager component controls sharing of the information associated with the sticky policies, and the data registry records how it has been distributed. Optionally, an external TA can also be involved to perform some additional checks if the external workflow manager is not able to make those directly.

At the receiving party side, if EnCoRe enabled, there is a translation of the high-level requirements expressed in the sticky policies into fine-grained access and obligation policies into local access control and obligations policies to be enforced, along with the original data subjects' privacy choices. This is achieved by means of mapping capabilities that systematically translate high-level constraints (defined in the Policy Manifesto) into enforceable ones. If the receiving parties do not have EnCoRe compliant systems, then the external workflow manager assesses the extent to which the data may be released for a given purpose and controls release of the data accordingly, potentially sanitizing it if needed. EnCoRe administrators pre-define the criteria by which data should be sanitized and how (e.g. by omitting some details or providing statistical information) depending on the purpose for which data was required and the outcome of risk assessment carried out on the receiving parties (e.g. on their capability to deliver the required privacy controls on specific data items).

In order to revoke consent, the users use the same mechanisms to edit their consent preferences as those they used to set them in the first place, i.e. via web-based UIs: these preferences are automatically propagated throughout the EnCoRe system as well as beyond it, in a batched manner, to the other organizations involved, by leveraging the information stored in the Data Registry.

This approach can be applied recursively, for a chain of organizations disclosing information between them.

We have already developed the core mechanisms for the management of sticky policies within the EnCoRe project along with a PKI-based implementation of the required mechanisms. Next steps are to deploy them in a case study with a customer and provide advanced implementations of the protocols, including multiple verification and control capabilities.

In the longer term, we envisage that additional information may be added to the sticky policy by the EnCoRe system e.g. about technical and process control mechanisms or boundaries that should be in place by the receiving entity (in order for it to be considered trustworthy) or else that it is EnCoRe-compliant, checks whether that organization itself is not blacklisted or is considered trustworthy as an organisation, etc. We are also researching better ways of propagating consent and revocation changes along the chain within which data is shared, including trusted components (e.g. the external workflow managers of the other entities, if they are EnCoRe-compliant) that periodically check, update and trigger enforcement of relevant user preference options stored elsewhere.

Open research issues that we are currently researching include stronger enforcement and trying to prevent SPs cheating by breaking promises to TAs. A logical binding can be easily unbound, but even with a cryptographic binding, after the personal data have been decrypted, the binding is broken, in the sense that the users' data is then fully available to the authorized party and subsequently actions may be taken that contravene the policy. Hence a solution such as the one proposed by Zuo *et al.* [12] needs to be used in combination, to protect data after it has been decrypted, but current options result in stronger protection at the cost of poor scalability, or unrealistic expectations as to the hardware or operating system environment used by SPs. Trusted computing might also be used to try to ensure that receivers act according to associated policies and constraints. Furthermore, the digital signature only proves the authenticity of a binding established in the past by the data subject. If encryption is applied only to text files that adhere to a predefined structure, policies can be relatively easy to corrupt and a skilled hacker can tamper with the file and make the policy illegible. Watermarking schemes [13] and obfuscation techniques [14] can also be used for content protection, but do not provide policy enforcement or support for protecting the data after access.

Conclusions

Sticky policies provide a promising approach for privacy management within and across organisational boundaries that can be leveraged in a variety of contexts, including in the cloud. Sticky policies are defined and used at the time the user disclosed data to an organisation. These sticky policies dictate the preference conditions and ensure that policy constraints will be audited and degrees of assurance provided.

The solutions we have described allow tracing and auditing via TAs and enforcement of user preferences by SPs. This approach progresses the state of the art in that it provides an end to end data management solution, is scalable, provides different options to drive the interaction process between the SPs and TAs and allows optional involvement of storage service providers.

User interactions will be mediated by "Privacy Advisors" and/or client applications to mitigate the complexity of creating sticky policies and binding them to data. This solution could be used in a number of different business areas, but would be particularly appropriate where sector-specific legislation or user concerns are strongest – for example in domains relating to defence, healthcare or finance.

We are working to extend and broaden this approach in order to achieve accountability by using contractual assurances along the service provision chain from SPs to accountable organizations, enhanced on the technical side by enforcement of corresponding machine-readable policies propagated with (references to) data, integrated risk assessment, assurance, and auditing [15]. By these means, the accountable organizations can ensure that all who process data observe their obligations to protect it, irrespective of where that processing occurs.

Acknowledgements

We would like to acknowledge helpful input and feedback about this research from a number of parties, and most notably Liqun Chen, Gina Kounga and Archie Reed.

References

1. Siani Pearson and Tomas Sander, "A decision support system for privacy compliance", Strategic and Practical Approaches for Information Security Governance: Technologies and Applied Solutions, IGI Global, 2011.
2. *Guidelines Governing the Protection of Privacy and Transborder Flow of Personal Data*, Organization for Economic Cooperation and Development (OECD), 1980.
3. Siani Pearson and Damien Allison, "Privacy Compliance Checking using a Model-Based Approach", E-Business Applications for Product Development and Competitive Growth: Emerging Technologies, e. I. Lee, Business Science Reference, New York, IGI Global, pp. 199-220, 2011.
4. EnCoRe project, Enforcing Consent and Revocation, www.encore-project.info
5. Karjoth, G., Schunter, M., Waidner, M.: Platform for Enterprise Privacy Practices: Privacy-enabled management of customer data, PET, LNCS, Springer Verlag, 2002.
6. Siani Pearson, Marco Casassa Mont and Gina Kounga, "Enhancing Accountability in the Cloud via Sticky Policies", Proc. STAVE, Springer, June 2011.
7. D. Boneh and M. Franklin, Identity-based Encryption from the Weil Pairing. Crypto 2001, 2001.
8. Marco Casassa-Mont, Siani Pearson and Pete Bramhall, "Towards User Control and Accountable Management of Privacy and Identity Information", *Proc. ESORICS*, pp. 146-161, LNCS 2808, Springer, 2003.
9. Pöhls, H. C.: Verifiable and Revocable Expression of Consent to Processing of Aggregated Personal Data, ICICS, 2008.
10. Shamir, A. "How to Share *a Secret", C*ommunications of ACM, 22, pp. 612-613, 1979.
11. Marco Casassa Mont, Siani Pearson, Pete Bramhall: Towards Accountable Management of Identity and Privacy: Sticky Policies and Enforceable Tracing Services, IEEE TrustBus 2003, Prague, 2003. Available via http://www.hpl.hp.com/techreports/2003/HPL-2003-49.pdf
12. Zuo, Y, O'Keefe, T.: Post-release information privacy protection: A framework and next-generation priacy-enhanced operating system, ISF, vol 9, no 5, pp. 451-467, Springer http://www.springerlink.com/content/03718003288553u5/ 2007.
13. Perez-Freire, L., Comesana, P., Troncoso-Pastoriza, J. and Perez-Gonzalez, F.: Watermarking security: a survey, Transactions on Data Hiding and Multimedia Security, LNCS, 2006.
14. Bayardo, R., Agrawal, R.: Data Privacy through Optimal k-Anonymisation, International Conference on Data Engineering, pp. 217-228, 2005.

15. Siani Pearson, “Toward Accountability in the Cloud”, View from the Cloud, IEEE Internet Computing, IEEE Computer Society, July/August issue, vol. 15, no. 4, 2011.

Siani Pearson is a senior researcher in the Cloud and Security Research Lab at HP Labs Bristol. Her current research focus is on privacy enhancing technologies, accountability and the cloud. She has a PhD in artificial intelligence from the University of Edinburgh. She’s a technical lead on regulatory compliance projects with the HP Privacy Office and HP Enterprise Services and on the collaborative TSB-funded Ensuring Consent and Revocation (EnCoRe) project. She is a fellow of the British Computer Society, senior member of IEEE and a Certified Information Privacy Professional/Information Technology (CIPP/IT). Contact her at siani.pearson@hp.com.

Marco Casassa Mont is a senior research scientist at the HP Labs in Bristol, Cloud and Security Lab. His research interests include strategic aspects of risk management, security, privacy and technologies applied to business contexts and emerging scenarios, including the Cloud. He has been involved in various technology transfers with HP business groups. He is currently the coordinator and the technology lead of the UK collaborative (TSB-funded) EnCoRe Project. Casassa Mont received an MSc in computer science from the University of Turin, Italy. He is a senior member of IEEE and a member of the UK Institute of Information Security Professionals. Contact him at marco.casassa-mont@hp.com