

Sticky Policies: An Approach for Managing Privacy across Multiple Parties

Siani Pearson and Marco Casassa Mont, *HP Labs Bristol*

Machine-readable policies can stick to data to define allowed usage and obligations as it travels across multiple parties, enabling users to improve control over their personal information. The EnCoRe project has developed such a technical solution for privacy management that is suitable for use in a broad range of domains.

Current mechanisms for ensuring privacy protection across organizational boundaries rely on legal and business frameworks, including contracts and service-level agreements. Technical mechanisms complement such approaches by supporting enforcement and auditing of the organizational obligations they outline.

Personally identifiable information (PII), also referred to as personal data or personal information, is data that can be traced to a particular individual—for example, a name, address, phone number, Social Security number, national identity number, credit card number, e-mail address, password, or date of birth. Because of its sensitive nature, greater care must be taken in the handling of the subset of PII that includes financial or medical data.

In commercial contexts, meeting customers' expectations regarding privacy requires the protection and careful use of PII. For corporations, privacy includes the application of laws, policies, standards, and processes for managing an individual's PII. Privacy management identi-

fies the ways in which organizations and individuals can control the collection, usage, and sharing of personal data, including sensitive information.

Privacy management and compliance with regulatory requirements for data protection can help organizations foster trust with their customers. In a given context, there may be many different privacy-related regulatory requirements, including sector-specific laws, national legislative requirements, and transborder dataflow restrictions.¹ Although assessing requirements in a given situation can be complex, an established set of principles forms the basis of most privacy legislation worldwide.²

To provide mechanisms for online privacy management, substantial research has been conducted related to

- anonymization technologies;
- enforcement of privacy-enhanced access-control policies, as in the Prime and PrimeLife EU projects (www.primelife.eu);
- policy life-cycle management;
- satisfying global regulations relating to data protection, including tools for governance, risk, and compliance (GRC);
- modeling privacy regulations; and
- modeling organizational privacy policies down to the operational level.³

However, major issues remain outstanding, 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.

An approach based on *sticky policies*—conditions and constraints attached to data that describe how it should be treated—enables compliance with and enforcement of current requirements such as the US Health Insurance Portability and Accountability Act (HIPAA) of 1996 along with future needs emerging from the adoption of new technologies and models, including the storage and processing of sensitive data in the cloud.

INFORMATION FLOW

In some scenarios, a user's confidential information flows across organizational boundaries. For example, a healthcare system could disclose personal data and preferences to a general practitioner via an online service provider (SP); the system also might need to share this information with hospital specialists, pharmaceutical companies, and other third parties involved in the healthcare supply chain. A similar situation might apply for a travel agency that needs to share data with various SPs such as hotel reservation brokers and car rental agencies.

More generally, these kinds of scenarios will be increasingly common in a cloud computing environment, where users interact with front-end SPs that will need to share part of the information with other SPs to supply the required services.

In all these situations, users must reveal personal and even sensitive information to receive a service, but they want to control how that information is used. They can directly control how their data should be processed, handled, and shared by explicitly expressing their preferences and data-handling policies. These choices must be respected all along the service provision chain, including allowing the user to update them. Achieving this objective requires propagating the user choices to all the SPs and deploying several mechanisms to ensure that the policies are respected. Moreover, the user can be actively involved in the selection of multiple, interchangeable services that will track and audit policy fulfillment.


CHARACTERISTICS OF STICKY POLICIES

Depending on the degree of a policy's stickiness, the data might be encrypted, with access to the content allowed only upon the satisfaction of these policies. Specifically, the policies govern the use of associated data, and could specify the following:

- proposed use of the data—for example, for research, transaction processing, and so on;
- use of the data only within a given set of platforms with certain security characteristics, a given network, or a subset of the enterprise;

- specific obligations and prohibitions such as allowed third parties, people, or processes;
- blacklists; notification of disclosure; and deletion or minimization of data after a certain time; and
- a 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 mechanisms for handling sticky policies. Our approach uses cryptographic mechanisms to strongly associate policies with the data. There can be different degrees of stickiness, but we adopt a strong binding as it provides better accountability. The data is encrypted and only accessible upon the acceptance and satisfaction of constraints and duties the policies impose.



Users can directly control how their data should be processed, handled, and shared by explicitly expressing their preferences and data handling policies.

TAs provide assurance by keeping track of promises the involved parties make to access data, along with controlling access to such data. The TAs' role may be integrated with other functionality, such as being a consumer organization, a certification authority (CA), or a well-known organization. The TA's role also can be performed by a client-side software component or service that is under the control of end users or other parties, or it can 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 except for dealing with additional policy condition checks or from storage providers if they are used to store the data and an authenticated reference is passed around instead of the data. However, SPs would either need to manage 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 is likely to be most suitable for service provision environments in which the increased trust and protection would justify the additional expense. Alternatively, business partners of goodwill enterprises that are trying to employ best practices might encourage its use.

Sticky policies are passed between organizations to capture obligations and other constraints that the receive-

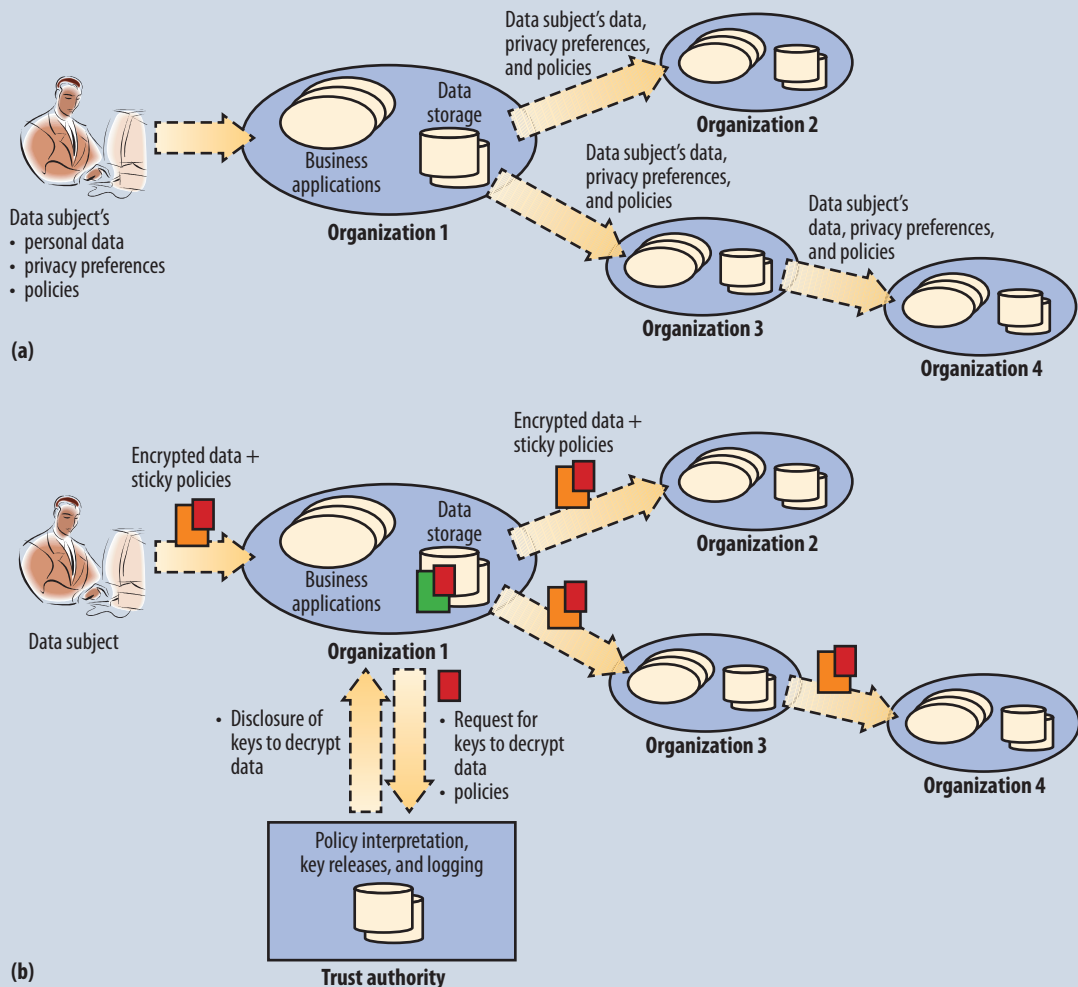


Figure 1. High-level scenario and related management of sticky policies. (a) High-level scenario involving data disclosures across organizations. (b) Overview of sticky policy approach.

ing parties must meet to access and use the associated personal data. For example, if the system passes a health-care record from a hospital to a research institution and then to a research team, the information might be in a form in which certain attributes such as medical results or personal information such as name and address are encrypted, with an associated sticky policy describing how parts of this could be used. For example, a patient wants this information to be released only to research teams, requests that it be deleted after three years, and asks to be notified every time the medical information is passed on. These constraints can be expressed in several ways, including using a simple XML format.

Sticky policies can help enable accountable management and disclosure of confidential data across boundaries. In the approach shown in Figure 1, personal, private, or confidential information is associated with machine-readable policies in a way that can't be compromised. The system

processes the information in a way that adheres to these constraints. As it replicates the data or fulfills the service provision request, mechanisms will be in place to ensure that the customer's preferences are respected all along the chain. Specifically, TAs need to retrieve keys to decrypt data and 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, which can be achieved using various cryptographic techniques, including public-key infrastructure-based and other approaches.

Our solution includes the following aspects:

- To more easily interpret and enforce policies, organizations impose a framework that defines their preferences and policies. In one approach to achieving

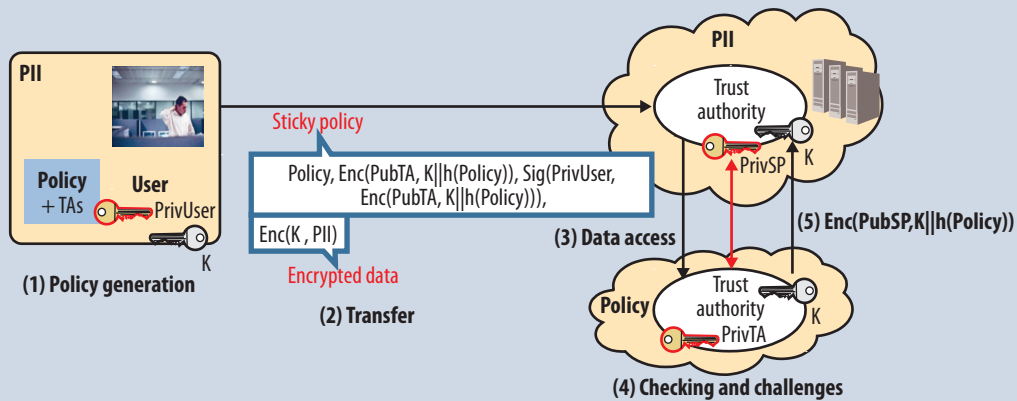


Figure 2. Core mechanisms underpinning the management of sticky policies. (1) Creation of sticky policies at the user side. (2) Sending sticky policies and data to the service provider. (3) Sending sticky policies to the agreed trust authority to get access to data. (4) Service provider interacts with trust authority to satisfy sticky policy constraints. (5) Getting cryptographic keys for use in accessing the data.

this, SPs publish a “manifesto” containing the list of supported (macro) policies and TAs and defining how these policies relate to access control and obligation behaviors the organization supports.

- A user (customer) can interact with an SP to select the granularity (ranging from coarse-grained to fine-grained) of applying policies to items or specific subsets of personal data to be disclosed and customize related preferences such as notification, period of time after deletion, set of agreed purposes, and the list of parties not to interact with.
- The user selects a subset of TAs that are to be trusted.
- Based on these selections, a client-side component supports the creation of sticky policies and their association to data—the bundling of policies, preferences, data, and TAs. In other words, the client-side component manages the packaging of data along with selected parameterized policies and TAs.
- Rather than passing the encrypted data directly to the SP, the user can select the option to refer to PII secured by a third party—a storage provider that stores the encrypted data.
- The system sends the encrypted data along with sticky policies to the SP.
- To gain access to the data, the SP needs to interact with one of the selected TAs (based on availability). During this interaction, the SP must assert its willingness to fulfill the customized sticky policies. Alternatively, depending on the policy requirements, the TA might be able to check this independently of such signed statements—for example, with reference to externally maintained blacklists or reputation management systems, or by verifying system properties using mechanisms such as trusted attestation or remote software verification. This creates an audit

trail available to the user and TA afterward in case of policy violations or misbehavior.

- The SP allows a predefined period of time for connection with the TA. The solution supports swapping data between TAs based on needs.
- Only after satisfying all these requirements and checking additional contextual information will the TA decide to release the keys for decrypting data.
- The TA will be able to decrypt and access the data regardless of whether it was directly disclosed or if only a reference to it was provided. In the latter case, the SP would need to fetch the data.

We envision the deployment within organizations of privacy-management components that complement identity and access management solutions, as tested in the context of the EnCoRe collaborative project (www.encore-project.info). Specifically, these components will complement organizations’ 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.


CREATING STICKY POLICIES

The original sticky policy paradigm specified that privacy preferences should flow with personal data to make sure that they can always be enforced.⁴ Subsequent research suggested a method for creating strong stickiness of policies to data.⁵

In a common central approach, customers allow SPs to have access to specific data based on agreed policies in interactions with interchangeable independent third par-

ties (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 specifically encrypts data based on the policy. A TA mediates access to data, checking for compliance to policies to release decryption keys, so that checking for compliance requires more than having the SP assert its willingness to do so. This provides users with fine-grained control over access and usage of their data, even in public cloud models.

Various techniques using different underlying encryption mechanisms can provide sticky-policy protection of data. In each case, the system can extend the selected technique to cover the propagation of data along the service provision chain. The process is analogous to user-to-SP protocols, in which the first SP can add policy constraints to form a superset of previous policy constraints. Multiple mechanisms currently used to exchange information can refine and deploy the proposed techniques, including Web technologies and protocols such as http/s, SOAP, and so on; document formatting and protection techniques such as Adobe and DRM; and various messaging tools including e-mail and instant messaging.



The access to data can be as fine-grained as necessary, based on policy definitions, underlying encryption mechanisms, and a related key-management approach that specifically encrypts data based on the policy.

These protocols apply not only to human users but also more broadly to machine-to-machine or service-to-SP interactions.

Using public-key encryption techniques

When using public-key encryption, we assume that all the stakeholders have certified public or private key pairs from trusted CAs. An approach that enhances integrity binds policies to data by encrypting the data under a symmetric key that a sender and receiver conditionally share based on fulfillment of policies, and sticking the data to the policy using public-key enveloping techniques similar to the Public-Key Cryptography Standard (PKCS) 7. Figure 2 shows an example of this process, in which the labeled stages are as follows:

1. The sender generates the policy, 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 can be generalized to allow encrypting different attributes separately—that is, using

different symmetric keys generated at this stage—revealing only part of the information when an attribute is decrypted.

2. The sender generates a message to the SP. One part of the message is the data encrypted with K . The other part is a sticky policy, in which K , appended to the policy's hash, is encrypted with the TA's public key and then is signed using the user's private key. This makes it possible to verify the policy's source and integrity and binds K to the data and the policy. The system sends the resultant sticky policy together with the encrypted data to the SP.
3. The SP generates a message to the TA, which involves passing on just the sticky policy and encrypted shared keys.
4. The TA checks policies, potentially including challenges to the SP. The SP might need to provide signed statements about its policies.
5. If all checks are fulfilled, the TA releases the shared key. This generates a message from the TA to the SP, which involves encrypting K appended to the policy's hash with the SP's public key. The SP can get access to K to check the policy's integrity and then decrypt the PII.⁶

Using identifier-based encryption

An identifier-based encryption (IBE) cryptographic schema can use any kind of string as a public encryption key, including a name, role, terms, or conditions.⁷ The generation of the corresponding IBE decryption key can be postponed. A TA can generate this decryption key on the fly, under specific circumstances.

While it is conceptually similar to the PKI approach, we adapt IBE by mapping a sticky policy to an IBE encryption key. The TA's role is expanded to check the integrity and trustworthiness of the requestor's credentials and its IT environment before releasing the decryption key. It also logs and audits disclosures of confidential data.⁸

VARIATIONS ON ENCRYPTION

We can potentially use any encryption mechanism to associate policies with data. For example, Voltage and Navajos provide format-preserving encryption and search-enabled encryption, respectively. If the operation involves indexing, it would still be possible to search and index encrypted attributes.

An alternative solution permits binding of privacy preferences to data and conveying the individual's consent as well.⁹ However, this solution does not avoid the unauthorized use of data.

This approach can be adapted to support multiple verification and control functions. Instead of having individual certificates, each entity could be provided with a key component, called a "share." An option such as Shamir's threshold-based secret-sharing scheme¹⁰ could be used

to require l of m shares for the cloud service provider to recover K and decrypt the PII, while still providing some redundancy among TAs. Secret-sharing schemes form a particular group of multiparty key establishment protocols that enable distribution of control or trust in 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 (the 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.

Trusted computing group integrity-checking mechanisms can verify that the receiver's platform is trusted, its software state is conformant with the disclosure policies, and it correctly implements defined privacy-management mechanisms.

Furthermore, there are several variations on these approaches 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.

In addition, the protocols themselves can be amended. In the PKI approach, for example, the user can bind the policy to the data within a signing operation rather than within the encryption. Other options include using the signcryption algorithm specified in ISO/IEC 29150.2 (www.iso.org/iso/iso_catalogue/catalogue_ics/catalogue_detail_ics.htm?ics1=35&ics2=040&ics3=&csnumber=45173), performing a single operation and separately encrypting the data (or reference to the data). An alternative is to encrypt attributes with different keys, enveloping the sensitive data and passing it on without revealing the key from the TA while revealing different attributes to different entities in the chain.

CASE STUDY: ENCORE PROJECT

The processes and components designed for privacy management within the EnCoRe project demonstrate the feasibility of sticky policies. EnCoRe is a collaborative research effort undertaken by UK academic and industrial partners that uses consent and revocation management to give individuals more control over their personal information. In this context, revocation essentially means change of consent, potentially in a fine-grained way.


The project provides mechanisms for users to define consent policies and to change them. EnCoRe uses sticky policies to represent and enforce the consent and revocation preferences of end users. In general, EnCoRe supports the following:

- *Explicit management of consent and revocation.* Negotiating, setting, changing, and enforcing sticky policies are integrated with the management of security and

privacy policies. Compliance checking and auditing are integrated capabilities.

- *Bridging the disconnect between high-level and lower-level policies.* This includes mapping legal, business, social, and security requirements to high-level policies. We define an intermediate conceptual framework to model policies and reason on top of them. We then map these concepts into monitorable and enforceable policies driven by users' preferences.

Our solution is applicable in a variety of business contexts, and it is especially valuable where sensitive information is involved—for example, in healthcare scenarios such as biobanks and assisted-living facilities, providing third-party access to employee data, government scenarios, and cloud computing.



The EnCoRe project provides mechanisms for users to define consent policies and to change them, as well as for enforcement of these policies.

In the EnCoRe project, we have developed a flexible toolbox solution that can be customized and deployed consistently within the business processes of each involved SP. EnCoRe-compliant capabilities provide assurance about a given SP's privacy management practices and related management of consent and revocation.

Figure 3 illustrates the overall set of functionalities and capabilities that EnCoRe provides. The system can provide these components as a set of services in the cloud or it can deploy them as an overall stand-alone infrastructural solution. The components include the following:

- *Personal consent and revocation assistant.* This component provides user-side capabilities to help people express their consent by making privacy choices such as opt-in/opt-out, identifying preferences, and so on, and submitting revocation requests, along with the explanation of privacy practices that organizations provide. A Web browser plug-in can trigger this function during data-disclosure processes. The system can embed these privacy choices into sticky policies to ensure that third parties receiving the data will fulfill them.
- *Virtual data registry.* This repository—or an aggregation of synchronized repositories—keeps track of where each known individual's data has been stored within and outside the organization and identifies which type of data has been disclosed and to whom, along with any relevant associated sticky policies.

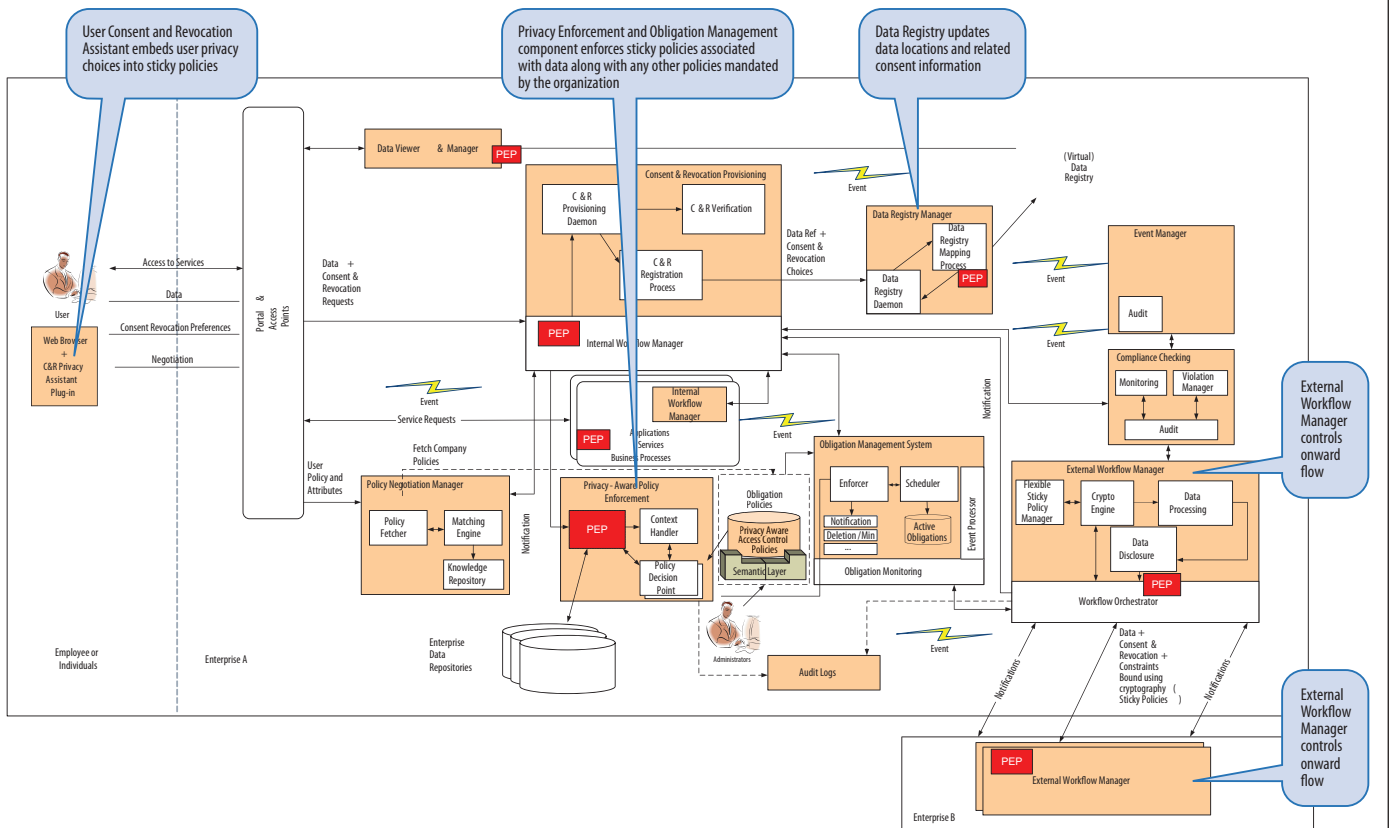


Figure 3. EnCoRe architecture. EnCoRe components process personal data and enforce preferences. Users disclose their personal data with privacy preferences; the EnCoRe privacy-aware access control and obligation components enforce the preferences when third parties access the data; the data registry tracks the data's location; the external workflow manager creates and attaches sticky policies to data before the system discloses it to third parties. The system applies this approach recursively across chains of organizations.

- *Consent and revocation provisioning.* This component automatically updates the data registry every time there is a new expression of consent and revocation. It uses internal workflows to update an individual's preferences and identify constraints that affect the enforcement of access control and obligation policies.
- *Privacy-aware policy enforcement and obligation management.* Driven by consent, this component deals with access control over data and obligations. It enforces sticky policies associated with the data along with any other policies the organization mandates.
- *External workflow manager.* This component intercepts and tracks the flow of personal data, both within and between organizations, and propagates the associated consent information. Sticky policies ensure degrees of compliance with agreed policies and data subject's expressed preferences. Applications and services might need to be instrumented with agents that communicate with this component.
- *Auditing.* This component logs and tracks what happens to data, consent, and revocation during operational and administrative activities, includ-

ing the flow of personal data within and beyond the organization.

- *Compliance checking and risk assurance.* The enterprise's privacy administrators use this key offline component to assess current risks and provide indications of compliance.

The sticky policies that the EnCoRe system sends to other organizations specify the purposes of using the data and any obligations and prohibitions, including notification and deletion after a certain time, that the user has specified in the consent and revocation preferences associated with that data. The TA is distributed in the sense that the EnCoRe external workflow manager 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 perform some additional checks if the external workflow manager cannot perform them directly.

If the receiving party is EnCoRe-enabled, the system translates the high-level requirements expressed in the sticky policies into fine-grained access and obligation poli-

cies to be enforced along with the original privacy choices. To achieve this, mapping capabilities systematically translate high-level constraints (defined in the policy manifesto) into enforceable ones. If the receiving parties do not have EnCoRe-compliant systems, the external workflow manager assesses the extent to which the data can be released for a given purpose, sanitizing it before release if needed. EnCoRe administrators predefine the criteria for sanitizing data—for example, omitting some details or providing statistical information. The criteria for releasing data include evaluating the purpose for which the data was required and the outcome of risk assessment carried out on the receiving parties—for example, their ability to deliver the required privacy controls on specific data items.

To revoke consent, users edit their consent preferences through Web-based UIs. EnCoRe batches and automatically propagates these preferences throughout the system as well as beyond it to the other organizations involved, leveraging the information stored in the data registry. Organizations can apply this approach recursively to disclose information to one another.

FUTURE DIRECTIONS

We have developed the core mechanisms for managing 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 envision using the EnCoRe system to add information to the sticky policy regarding technical and process control mechanisms or boundaries that the receiving entity should have in place for it to be considered trustworthy or that it is EnCoRe-compliant. We are also researching better ways of propagating consent and revocation changes along the chain within which data is shared, including the external workflow managers of the other entities that periodically check, update, and trigger enforcement of relevant user preference options stored elsewhere.

Open issues that we are currently researching include stronger enforcement and trying to prevent SPs from cheating by breaking promises to TAs. A logical binding can easily be unbound, but even with a cryptographic binding, after the personal data has been decrypted, the binding is broken in the sense that the users' data is then fully available to the authorized party and subsequent actions could be taken that contravene the policy. The solution needs to protect data after it has been decrypted,¹¹ but current options result in stronger protection at the cost of poor scalability or unrealistic expectations regarding the hardware or operating system environment the SPs use.

Trusted computing also might be used to ensure that receivers act according to associated policies and con-

straints. However, the digital signature only proves the authenticity of a binding the data subject established in the past. If encryption is applied only to text files that adhere to a predefined structure, it can be relatively easy to corrupt policies; thus, a skilled hacker could tamper with the file and make the policy illegible. Watermarking schemes¹² and obfuscation techniques¹³ also can provide content protection, but they do not ensure policy enforcement or offer protection for the data after access.

Sticky policies offer a promising approach for privacy management within and across organizational boundaries that can be leveraged in various contexts, including in the cloud. The user defines sticky policies when disclosing data to an organization. These policies dictate the preference conditions and ensure that appropriate constraints will be audited and degrees of assurance provided.

Using sticky policies allows tracing and auditing via TAs and enforcement of user preferences by SPs. In addition to advancing the state of the art by providing an end-to-end data management solution, the approach is scalable, provides different options to drive the interaction process between the SPs and TAs, and allows optional involvement of storage service providers.

Privacy advisors or client applications will mediate user interactions to mitigate the complexity of creating sticky policies and binding them to data. This solution could be used in several business areas, but would be particularly appropriate where sector-specific legislation or user concerns are strongest—for example, in domains relating to healthcare, finance, or defense.

We are working to extend and broaden this approach to achieve accountability by using contractual assurances along the service provision chain from SPs to organizations, enhanced on the technical side by enforcement of corresponding machine-readable policies propagated with data, integrated risk assessment, assurance, and auditing.¹⁴ Thus, organizations can ensure that all who process data observe their obligations to protect it, regardless of where that processing occurs. **■**

Acknowledgments

Several parties provided helpful input and feedback about this research, most notably Liqun Chen, Gina Kounga, and Archie Reed.

References

1. S. Pearson, T. Sander, and R. Sharma, "Privacy Management for Global Organisations," *Data Privacy Management and Autonomous Spontaneous Security*, LNCS 5939, Springer, 2009, pp. 9-17.

2. Organization for Economic Cooperation and Development, "OECD Guidelines on the Protection of Privacy and Transborder Flows of Personal Data," 1980; www.oecd.org/document/18/0,3343,en_2649_34255_1815186_1_1_1_1,00.html.
3. S. Pearson and D. Allison, "Privacy Compliance Checking Using a Model-Based Approach," *E-Business Applications for Product Development and Competitive Growth: Emerging Technologies*, IGI Global, 2011, pp. 199-220.
4. G. Karjoth, M. Schunter, and M. Waidner, "Platform for Enterprise Privacy Practices: Privacy-Enabled Management of Customer Data," *Proc. 2nd Workshop Privacy Enhancing Technologies (PET 02)*, LNCS 2482, Springer, 2002, pp. 69-84.
5. M. Casassa Mont, S. Pearson, and P. Bramhall, "Towards Accountable Management of Identity and Privacy: Sticky Policies and Enforceable Tracing Services", 2003; www.hpl.hp.com/techreports/2003/HPL-2003-49.pdf.
6. S. Pearson, M. Casassa Mont, and G. Kounga, "Enhancing Accountability in the Cloud via Sticky Policies," *Secure and Trust Computing, Data Management and Applications*, vol. 187, Springer, 2011, pp. 146-155.
7. D. Boneh and M.K. Franklin, "Identity-Based Encryption from the Weil Pairing," *SIAM J. Computing*, vol. 32, no. 3, 2003, pp. 586-615.
8. M. Casassa Mont, S. Pearson, and P. Bramhall, "Towards User Control and Accountable Management of Privacy and Identity Information," *Proc. 8th European Symp. Research in Computer Security (ESORICS 03)*, LNCS 2808, Springer, 2003, pp. 146-161.
9. H.C. Pöhls, "Verifiable and Revocable Expression of Consent to Processing of Aggregated Personal Data," *Proc. 10th Int'l Conf. Information and Communications Security (ICICS 08)*, LNCS 5308, Springer, 2008, pp. 279-293.
10. A. Shamir, "How to Share a Secret," *Comm. ACM*, vol. 22, no. 11, 1979, pp. 612-613.
11. Y. Zuo and T. Keefe, "Post-Release Information Privacy Protection: A Framework and Next-Generation Privacy-Enhanced Operating System," *Information Systems Frontiers*, vol. 9, no 5, pp. 451-467.
12. L. Perez-Freire et al., "Watermarking Security: A Survey," *Trans. Data Hiding and Multimedia Security*, LNCS 4300, Springer, 2006, pp. 41-72.
13. R. Bayardo and R. Agrawal, "Data Privacy through Optimal k-Anonymisation," *Proc. Int'l Conf. Data Engineering (ICDE 05)*, IEEE CS Press, 2005, pp. 217-228.
14. S. Pearson, "Toward Accountability in the Cloud," *IEEE Internet Computing*, July/Aug. 2011, pp. 64-69.

Siani Pearson is a senior researcher in the Cloud and Security Research Lab at HP Labs Bristol. Her research focuses on privacy-enhancing technologies, accountability, and the cloud. She received a PhD in artificial intelligence from the University of Edinburgh, UK. Pearson is a fellow of the British Computer Society, a senior member of IEEE, and a Certified Information Privacy/Information Technology Professional. Contact her at siani.pearson@hp.com.

Marco Casassa Mont is a senior research scientist in the Cloud and Security Research Lab at HP Labs Bristol. His research interests include strategic aspects of risk management, security and privacy, and technologies applied to business contexts and emerging scenarios, including the cloud. 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.



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

computing now

ACCESS | DISCOVER | ENGAGE

Let us bring technology news to you.



<http://computingnow.computer.org>
Subscribe to our daily newsfeed