Towards Accountable Management of Identity and Privacy: Sticky Policies and Enforceable Tracing Services

Marco Casassa Mont, Siani Pearson, Pete Bramhall
Trusted Systems Laboratory
HP Laboratories Bristol
HPL-2003-49
March 19th , 2003*

E-mail: marco_casassa-mont@hp.com, siani_pearson@hp.com, pete_bramhall@hp.com

identity management, privacy, accountability, IBE, TCPA, audit, tracing, sticky policies

Digital identities and profiles are precious assets. On one hand they enable users to engage in transactions and interactions on the Internet. On the other hand, abuses and leakages of this information could violate the privacy of their owners, sometimes with serious consequences.

Nowadays, most of the people have limited understanding of security and privacy polices when applied to their confidential information. In addition, people have little control over the destiny of this information once it has been disclosed to third parties. In most cases this is a matter of trust.

This document describes an innovative approach and related mechanisms to enforce users' privacy by putting users in control and making organizations more accountable.

We introduce a technical solution based on sticky policies and tracing services that leverages Identifier-based Encryption (IBE) and TCPA technologies. Work is in progress to build a full working prototype.

* Internal Accession Date Only
Approved for External Publication

1. Introduction

Digital identities and profiles are more and more relevant to enable Internet transactions and interactions among citizens, service providers, enterprises and government institutions. Confidential information (including personal data, financial details, business data) needs to be disclosed in order to enable these interactions. Particularly interesting is the case where Internet interactions span across multiple parties (in B2C, B2B and government scenarios) due to sub-contracting, outsourcing and integration of services supplied by multiple providers. In this case the disclosure of personal identity and profile information can be used to enable single-sign on, reduce the overall complexity and simplify users' experiences.

Personal identity and profile information is precious and valuable to organisations: it can be used to improve and customise services, to provide statistical, strategic and marketing information or it can be sold to third parties. On the other hand, misuses and unauthorised leakages of this information can violate users' privacy, cause frauds and encourage spamming.

People perceive and address the related security and privacy issues in different ways, ranging from completely ignoring them (and indiscriminately disclosing their personal data) to being so concerned to prevent them from using any Internet and web-based applications. Situations commonly occur where users do not bother to read long lists of terms and conditions concerning privacy and confidentiality because they cannot understand them or they have no time. Often users are asked to grant to web sites the authorization to electronically manage their information, in order to carry on their transactions.

Identity and privacy management solutions are going to play a key role in protecting identities and profiles, enforcing good management practices and helping to detect criminal activities and support forensic analysis.

These solutions need to simplify users' experience so that people can feel they are in control of their confidential data and that this data is managed in an accountable way. If people are not willing to be involved in the active protection and management of their digital assets, trusted third parties could do this on their behalf and could provide people with easy-to-use tools to monitor and keep the situation under control.

2. Addressed Problem and Related Work

In this paper we address the problem of providing people with more control over their personal information and enforce accountable management of such information.

In order to describe some of the aspects involved by the problem, we refer to an e-commerce scenario. In no way are the issues and aspects we highlight limited to this sector, as they are common to financial, government and enterprise areas.

Figure 1 shows a scenario where users deal with electronic transactions that span across multiple e-commerce sites:

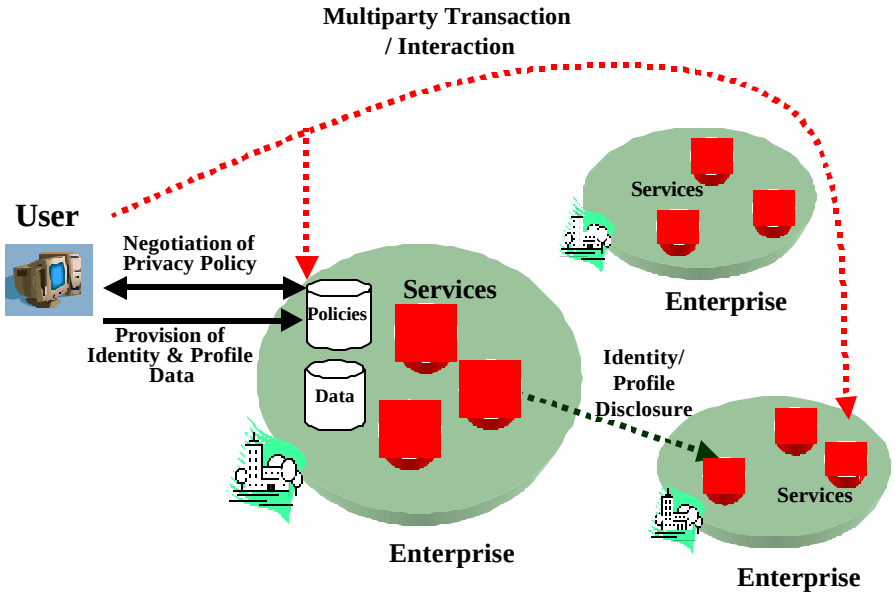

Figure 1: A Multiparty Transaction

In this scenario a person initially provides their digital identity and profile information to an e-commerce site in order to access their services, possibly after negotiations about which privacy policies need to be applied (description of such a negotiation process is beyond the scope of this paper). Then the user logs in and interacts with these services: it might happen that in so doing he/she needs to involve other web sites or organisations. The user might be conscious of this or this might happen behind the scenes, for example due to the fact that the e-commerce site interacts with partners and suppliers.

The e-commerce site might need to disclose personal data to third parties (such as suppliers, information providers, government and financial institutions, etc.) in order to fulfil the specific transaction. The involved e-commerce sites do not necessarily have prior agreements or belong to the same web of trust.

The above scenario highlights a few key issues: how to fulfill users' privacy rights and make users be in control of their information. At the same time users' interactions need to be simple and intuitive.

In general, users have little understanding or knowledge of the privacy laws and legislation that regulate the management of their information and their implications. Privacy and data protection laws that regulate this area do exist but it is hard to enforce or monitor them, especially when private information spread across organisations and nations' boundaries. In addition, further complexity arises due to the fact that privacy laws can differ quite substantially depending on national and geographical aspects. For example in US privacy laws restrict what the government can do with personal data but they introduce few restrictions on trading of personally identifiable information by private enterprises. In Europe (EU) people can consent to have their personally identifiable information used for commercial purposes but the default is to protect that information and not allow it to be used indiscriminately for marketing purposes.

Little has been done so far to directly involve users (or entities acting on their behalf) in the explicit management and enforcement of privacy policies, especially in a context of multiparty interactions. Users have lack of control over their personal information, especially after the initial disclosures. In addition third parties (such as delegates, e-commerce sites or enterprises) have lack of control over the confidential information they manage on behalf of their customers, in particular when they disclose it to other organisations, during transactions or interactions. In most cases it is a matter of trust.

Mechanisms such as W3C's Platform for Privacy Preferences (P3P) [16] allow users to define simple privacy policies but only for point-to-point interactions.

Liberty Alliance [10] and Microsoft Passport [11] efforts in federated identity management are (for the time being) based on a closed web of trust. Identity providers must be part of trusted clubs and be compliant with predefined privacy policies. This approach limits scalability and flexibility of the allowed interactions and transactions.

Seminal work towards a more fine-grained control over the privacy of personal information has been described by [8, 9]. In paper [8] the authors defines a privacy control language that includes user consent, obligations and distributed administration. They introduce the core elements of privacy policies and their formalisation. In paper [9] the authors describe a platform for enterprise privacy practices (E-P3P). They introduce the "sticky policy" paradigm and mechanisms for enterprise privacy enforcement.

Particularly interesting is the concept of "sticky policy": when submitting data to an enterprise, a user consents to the applicable privacy policies along with selected opt-in and opt-out choices. Sticky policies are strictly associated to users' data and drive access control decisions and privacy enforcement.

Papers [8] and [9] do not describe how the strong associations between policies and confidential data is enforced, especially across enterprise boundaries. Users still need to trust the enterprise when disclosing their data. Leakage of personal and confidential information might happen, despite data protection laws and privacy policies, because of lack of security, dishonesty of some of the involved intermediaries and the complexity of the overall systems.

In this paper we extend the work done in [8, 9] by suggesting mechanisms to strongly associate disclosure policies to personal data and increase the accountability of the involved parties.

3. Proposed Model and Technical Solution

This section introduces a high-level model and a related technical solution that allows users to enforce their privacy polices and, at the same time, makes organisations more accountable whilst dealing with users' data.

3.1 Model

The proposed model extends [8, 9] by including the following key aspects:

- Obfuscation of (any aggregation of) personal information before it leaves users' premises [1], in order to protect its content;

[1] Note that this stage could be extended using methods that allow users to release only minimised/selected information about themselves appropriate to the circumstance, such as via self-profiling [14], in which trusted profiles can be formed and released and trusted hardware can directly

- Association of "tamper resistant" sticky policies defined by users (or trusted third parties, acting on their behalf – agent technology is particularly useful here) to the obfuscated data, to explicitly declare the relevant disclosure constraints.
- Disclosure of data subject to the fulfilment of the sticky policies' constraints.
- Enforced tracing and auditing of disclosures of confidential data, to increase data receivers' accountability.

Figure 2 graphically shows how this model fits in the e-commerce scenario described in the previous section:

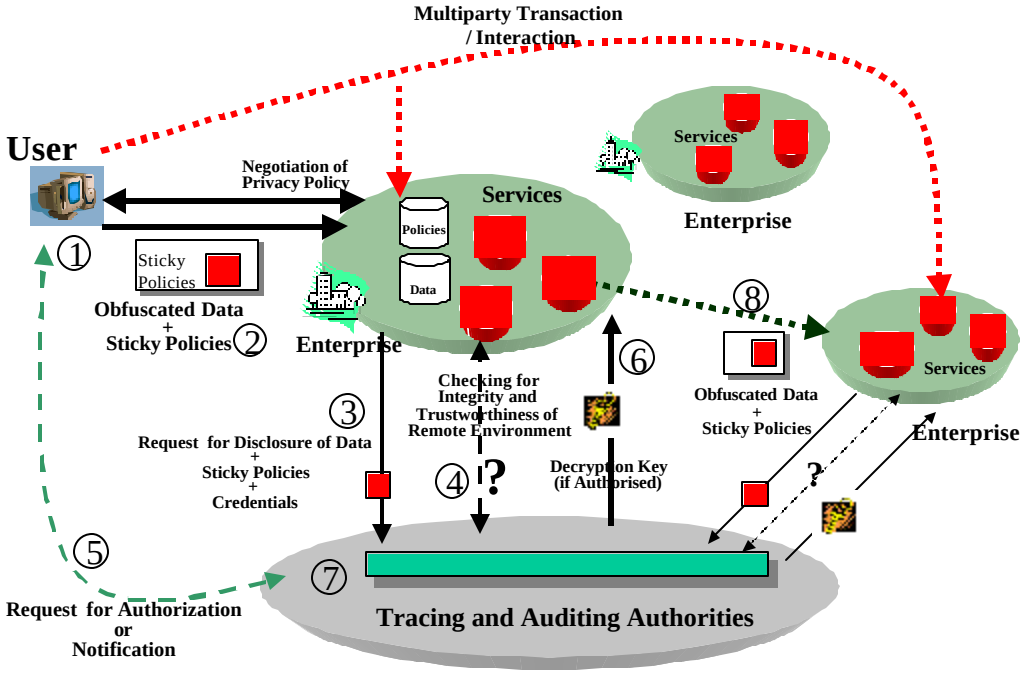

Figure 2: Proposed Privacy Model

In this model people use graphical tools (1) to:

- Locally author their disclosure policies (i.e. sticky polices) in a fine-grained way;
- Obfuscate their confidential data by directly using these disclosure polices;
- Associate these policies to the obfuscated data.

Some of the above activities can be automated by using predefined policy templates and scripts.

Digital packages (2) containing obfuscated data along with their sticky polices can be provided to requestors, for example e-commerce sites. These digital packages might contain a

---

certify (parts of) these if required). This can be done pseudonymously if desired. Furthermore, the extent and type (e.g. whether it is generalised) of the information released can be dependent upon the trust decision as to the recipient, including whether the recipient's platform is in a trustworthy state and has an appropriate enforcement system installed, etc.

superset of the required information, to reduce the number of users' interactions. Selective disclosure of (part of) their contents will be authorised, depending on needs.

A requestor (3) has to demonstrate to the Tracing Authority that he/she understands the involved terms and conditions. A Tracing Authority checks for the integrity and trustworthiness of the requestor's credentials and their IT environment (4), accordingly to the disclosure policies.

The owner of the confidential information can be actively involved in the disclosure process (5) by asking for his authorizations or by notifications, according to the agreed disclosure policies. In our model nothing prevents the owner of the confidential information from running a Tracing Authority.

The actual disclosure (6) of any obfuscated data to a requestor (for example the e-commerce site) only happens after the requestor demonstrates to a trusted third party – i.e. the "Tracing Authority" - that it can satisfy the associated sticky policies.

Disclosures of confidential data are logged and audited by the Tracing Authority (7). This increases the accountability of the requestors by creating evidence about their knowledge of users' confidential data. In particular this apply when confidential information is indiscriminately disclosed to third parties, as this evidence can be used for forensic analysis. In case a requestor sends the obfuscated data package to a third party (8), the same process, described above, applies.

Multiple trusted third parties (Tracing Authorities) can be used in the above process in order to minimise the risks involved in the management of trust, for example having to rely only on one entity.

3.2 Technical Solution

This section describes a technical solution that implements the above model by leveraging two key technologies:

- Identifier-based Encryption (IBE) [2, 4, 5]: an emerging cryptographic schema where any kinds of string (including a name, a role, terms and conditions, etc.) can be used as encryption keys (public keys). The generation of the corresponding IBE decryption key can be postponed until later. A Trust Authority (TA) (a type of trusted third party) can generate this decryption key on the fly, under specific circumstances. Appendix A provides more details about core IBE principles.

- Trusted Computing Platform Alliance (TCPA) technology [15]: this provides mechanisms and tools to check the integrity of computer platforms and their installed software. Appendix B provides more details about this technology.

In our technical solution a "sticky policy" is mapped to an "IBE encryption key". The "Tracing Authority" is a "Trust Authority".

IBE encryption keys can be modelled to define any kind of constraints or terms and conditions. At the very base an IBE encryption key is a string: it is self-explanatory and is directly used to encrypt confidential data.

An IBE encryption key does stick with the encrypted data. Any alteration or tampering of this string will make impossible to the Trust Authority to generate the correct IBE decryption key.

No secret needs to be generated and exchanged between users and the receivers of confidential information. The Trust Authority (TA) will generate the IBE decryption key on the fly, when required.

After describing the high-level architecture for our solution, we will move on to consider how:

- sticky policies are used;
- policies can be enforced;
- multiple Trust Authorities can be involved;
- non-compliance can be tracked;
- information owners can themselves act as a Trust Authority.

Figure 3 shows the architecture and components of a distributed system implementing our model:

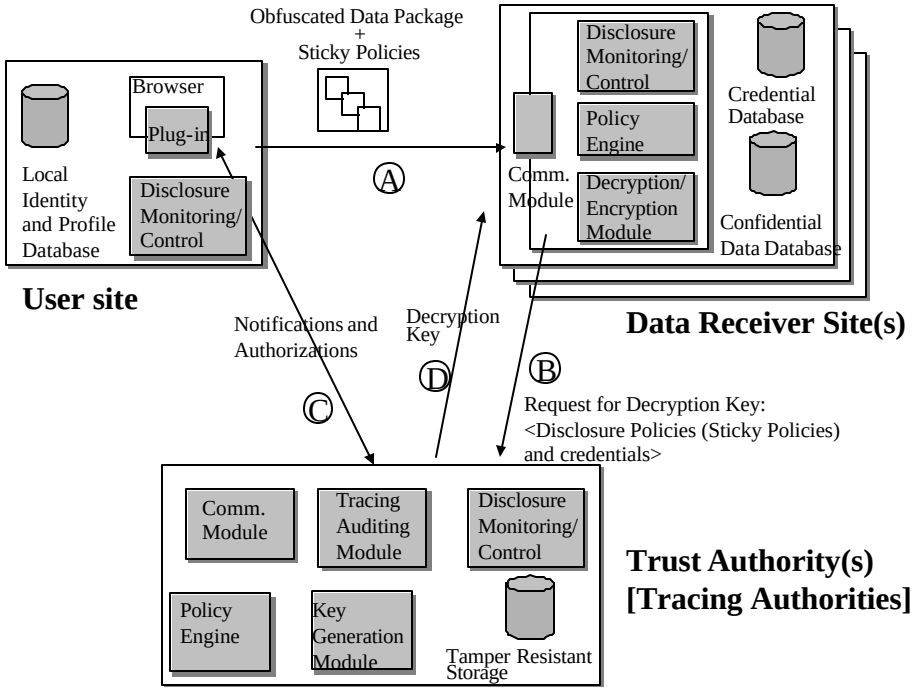

Figure 3: High Level Architecture

Messaging protocols (A)-(D) are carried out in order, and involve transfer of the information indicated in the directions shown by the arrows. In general, this process works as follows:

Identity or profile information is protected by encryption with sticky policies before its disclosure to third parties (A), by means of convenient plug-ins or trusted applications. These policies are used as IBE encryption keys (public keys) and might include:

- References to logical names of identity and profile attribute(s);
- Disclosure constraints;
- Actions (i.e. notification of the owner in case of multiparty disclosure);
- Lifetime, etc.

To obtain a valid IBE decryption key (B), the receiver needs to interact with TAs and provide information (including authentication credentials, business related information, company/individual policy related to data disclosure, usage and storage, software state,

platform configuration etc.) as required by the disclosure policies. In doing this, the receiver is explicitly aware of (and understands) these policies.

As part of this process, an extension of the TCPA integrity checking mechanisms [15] can be used to check that the receiver's platform is a trusted computing platform, that the software state of this platform is conformant with the disclosure policies and that the platform correctly implements defined privacy management mechanisms.

A TA will issue a decryption key (D) if it acknowledges the compliance with the disclosure policies. Before doing this it might interact with the information owner (C) to ask for his/her authorization or notification. The TA traces and stores all the information exchanged during these interactions in audit-trails, as evidence for future contentions or forensic analysis.

The remaining part of this section provides more details about some of the key aspects of our technical solution.

3.2.1 Sticky Policies

Users' identity and profile information is exchanged by means of data packages and the associated sticky policies. An example of a data package, containing obfuscated data along with its sticky policies, is as follows:

```
<data package>
    <data component>                                    // Identity and profile - attribute 1
     <sticky policy>                                    // disclosure policy – IBE public key

       <Trusted Authority>
         address and location of the Trusted Authority
       </Trusted Authority>
       <owner>
         <reference name> pseudonym1 </reference name> //reference name – IBE public key
         <owner's details>
              encrypted call back address
         <owner's details>                              //encrypted call back address
                                                        //by using the user's reference name
       </owner>
       <target>
            name of the identity or profile attribute
       </target>
       <validity>                                       //validity
            expiration date
       </validity>
       <constraint>                                     //constraints
            require_strong_X.509_authentication
       </constraint>
       <constraint>
             allow_sharing_of_data
       </constraint>
       <action>                                         //actions
            notify_owner
       </action>
     </sticky policy>
     <encrypted data>
            encrypted attribute value, using the above policy as IBE public key
     </encrypted data>
    </data component>
  </data package>
```

In the above example the data package relates to only one confidential attribute (i.e. piece of data, for example a credit card number), for simplicity. The associated sticky policy contains:

- *An encrypted "identifier" of the owner*. This can be any type of information, including the owner's e-mail address, URL, etc. Note that a "reference name" (a pseudonym, for example) has been used as an IBE encryption key to encrypt this information. Only the competent Trust Authority will be able to retrieve the owner's identifier (and use it, for example, to notify the owner of a disclosure or ask for an authorization).

- The *name of the attached confidential attribute.*

- An *expiration date*: date after which the Trust Authority will not issue anymore the decryption key.

- *Constraints and actions*: these constrain the requestor to strongly authenticate to the Trust Authority (for example by using PKI-based X.509 identity certificates [7]) and specify the usage of the attribute. An additional constraint is to notify the user of a disclosure.

Sticky policies (disclosure polices) can be used to allow a selective disclosure of any aggregation and combination of confidential information; they can be associated in a fine-grained way to any kind of attribute.

They can be composed and extended in a very flexible way. We use an XML-based representation as matter of convenience. Any kind of constraint, obligation and permission can be added, as long as the Trust Authority (TA) and the receivers understand its semantics.

The receiver of the encrypted information (for example an identity provider or an e-commerce site) can programmatically interpret the associated disclosure policies by means of a policy engine.

A further HP Labs technical report will provide more detail about other aspects of our sticky policies, including hierarchies of policies, composition of policies and their mapping at different levels of abstraction (service, application, system and OS).

3.2.2 Policy Enforcement

TCPA integrity checking mechanisms [15] can be used to allow the TA('s) platform to be checked out by the user (to make sure that the TA will operate as expected) and/or the recipient of the data (to help the recipient decide whether the TA can be trusted with the information that the recipient needs to provide to the TA in order for the decryption key to be issued).

An analogous approach may be used with other types of Trusted Platform that use a trusted hardware device as a root of trust, and not necessarily just those compliant with the TCPA specification. For example, the enforcement could be provided by using similar mechanisms within Microsoft's Palladium [12].

Furthermore, Trusted Operating Systems (OSs) can be used to increase security and trust, for example by storage of sensitive information that the receiver needs to disclose to the TA within one or more separate OS compartments. The technology required to implement the above solution is currently available and has been developed by HP Labs, Bristol, UK.

In particular, TCPA integrity checking mechanisms can be used to allow:

(1) the TA(’s) platform to be checked out by the user (to make sure that the TA will operate as expected) and/or
(2) the TA's platform to be checked out by the recipient of the data (to help the recipient decide whether the TA can be trusted with the information that the recipient needs to provide to the TA in order for the decryption key to be issued) and/or
(3) the recipient's platform to be checked out by the TA and/or
(4) the recipient's platform to be checked out by the user
(5) analogous checking for further forwarding, e.g. a further entity's platform to be checked out by the recipient before forwarding on, etc.)

In general, it is the TA(s) that controls the disclosure of data, and not the receiver. However, this is not always the case. Protection can be given against the disclosure policy being contravened, in at least two ways:

- Via the receiver’s own platform, via enforcement mechanisms on that platform that enforce policies defined by data wrappers or tags, or enforce the platform’s policies relating to treatment of data. For example, this enforcement could be carried out at the OS level or by passing control to a TCPA-compliant Trusted Platform’s Trusted Platform Module (TPM) and only allowing data to be disclosed if special software protected by that TPM judged that it was appropriate to do so. Note that the correct operation of such mechanisms should be checked by the TA before release of the disclosure key, via an extension of the TCPA integrity checking process, as described above.
- If the data is disclosed to a third party using the mechanisms described in this document (which of course, it might not be – it could be given to a third party by many different means), the TA could check that this disclosure has been carried out according to the specification of the (original) disclosure policy, and both refuse to release the key to any third party and report the receiver’s behaviour in some appropriate way.

3.2.3 Multiple TAs

To enable an electronic transaction involving user’s confidential data, the receiver might pass the overall encrypted data or any portion of it to another third party (for example another identity provider). It might decide to encrypt portions of this data by using additional policies. This third party has to interact again with a TA as described above.

The receiver may have to use multiple TAs in order to access the data. For example, one TA might be competent with respect to security platforms and other might be competent in privacy, so it would make sense for both to carry out checks before allowing an entity to access data. In this case, the user might encrypt the data using a disclosure policy that specifies that it is necessary to use two (or more) sub keys in order to decrypt the data, and each of the TAs would provide one of these keys. Multiple keys might be needed to decrypt the same piece of data, or different data fields might be encrypted using different keys.

There is another case where multiple TAs might be needed: when data is forwarded from the receiver on to another entity. Here, there are two different types of case:

- Either the receiver uses the same TA, in which case it could just send on the encrypted message it received from the original sender (or, if desired, it could use a different disclosure policy and therefore obtain a different encryption).

- Or it uses a different TA, in which case the third party would have to apply to that TA to get the decryption key, etc., as described above.

3.2.4 Accountability Management

If the receiver discloses data in a way that is not contemplated by the policies he previously agreed, there is an audit trail (at the TA(s) site(s)) showing that he/she actually understood and agreed with those policies.

In case of identity or profile thefts, the audit information can be used to pin down a list of potential "offenders" and carry on forensic analysis. Enforcing the tracing and auditing of disclosures makes the information receivers more accountable.

3.2.5 Running Personal TA services

Owners of identity and profile information can run their own TA services to have first hand understanding of what happens to their information and make ultimate decisions. In this case, the information owners can directly use the TCPA integrity challenge to check that the remote IT environment (of the receiver) has not been corrupted, before proceeding with the data disclosure. Alternatively users can periodically interact with the TA to monitor the disclosure state of their confidential information.

4. Discussion

The idea of using trusted third parties to mediate the access to confidential information is not new. There are well-known related issues, including why a person or an organisation should trust a third party. Multiple approaches have been analysed and described in the literature, including branding, certifications and seals, presence on the market and historical information. This fundamental aspect is not covered in this paper as it is out of its scope. From our perspective companies that are trusted in the real world can be trust authorities.

Multiple trusted third parties can be involved in order to minimise the risk of having to trust or rely only on an entity. In our specific case, multiple Trust Authorities (Tracing Authorities) could be involved in the process of issuing IBE decryption keys. In addition, information owners can run their trust authorities.

We believe that the value we bring in this area is in the mechanisms we provide to associate "tamper resistant" disclosure polices (sticky policies) to confidential data, the interaction model adopted to force requestors to be traced (audited) and the technology used to check the integrity and trustworthiness of remote IT environments.

In term of obfuscation of users' data, traditional RSA cryptography (based on public/private keys), PKCS#7 enveloping techniques and PKI can be used to provide functionalities similar to IBE's. For example the Trust Authority's X.509 identity certificate can be used to encrypt a symmetric key, generated by the user. This symmetric key can be used to encrypt users' confidential information along with a hash value derived from the associated sticky policies. We believe IBE technology simplifies the management of obfuscated data by providing a model that naturally fits with the required interaction model. In addition, in case of multiple trusted third parties (TAs) are used, we believe that IBE technology scales better than using an analogous approach based on RSA and PKI technology.

A Trust Authority (Tracing Authority) is the right place to implement tracing and auditing activities. Requestors do need to interact with the Trust Authority to obtain an IBE

decryption key. They need to provide their contextual credentials, as mandated by the disclosure policies (sticky polices): this information is logged accordingly and can be used to make them accountable. The auditing and tracing effort is effective also to audit users' behaviours, as the Trust Authority is a trusted bridge with users.

It is important to notice that once confidential information has been disclosed to a requestor and it is in clear text (at the requestor site), it can be potentially misused. In our model, in case of leakages and misbehaviours, the tracing and auditing information can be used for forensic analysis to pin down responsibilities.

Current literature, including papers [8] and [9], recommends that enterprises define their own privacy and security polices, in a way that it is compliant with laws and legislation. To programmatically implement these policies they need policy engines integrated with traditional authentication and access control components. The model and technical solution described in this paper are complementary to the above aspects: they leverage IBE technology along with a TA service infrastructure to reduce the involved risks by increasing accountability, keep users in the disclosure loop and avoid unauthorised disclosures of information. In this context TCPA technology is used to do pre-emptive trust and security checks.

An evolution of the proposed model might include emerging tagged-OS (currently under research and development in TSL, at HP Labs, Bristol) to enforce (parts of) sticky polices directly at the OS level.

5. Current and Future Work

Two core technologies are used to implement our model: IBE and TCPA. They are currently available at HP Labs and on the market. In particular TSL (HP Labs, Bristol) has implemented an optimised version of the IBE code that provides IBE cryptography functions with a performance comparable to RSA-based code. TCPA chips and PCs are available on the market, for example supplied by IBM.

We currently have simple implementations of most of the components required by our technical solution, including a *Trusted Authority service*, a *user add-in* to author sticky policies and a policy driven (and context aware) authorization engine [3]. Work is in progress to build a non-repudiable logging and auditing system [1].

We are refining our model and learning by building the system. Our aim is to provide a first demonstrator of our technical solution.

A decision still needs to be made about the specific scenario where this solution is going to be deployed. The current list of candidates includes a federated e-commerce scenario and a B2B (supply chain) scenario.

A further HP Labs report will describe our progress in this area along with the lessons we learnt.

6. Conclusion

It is more and more important to defend and preserve people's privacy on the Internet, against unwanted and unauthorised disclosure of their confidential data. Despite laws, legislations and technical attempts to solve this problem, at the moment there are no solutions to address the whole set of involved issues.

In this paper we specifically address two important problems: letting users be more in control of their personal data and making enterprises and organisations be more accountable of their behaviours, whilst dealing with users' confidential data.

We introduced and described a model based on "sticky policies", to strictly associate "tamper resistant" privacy polices to obfuscated data, along with trusted tracing services. We described a technical solution where IBE technology coupled with TCPA are used to solve the above problems.

These core technologies are available at HP Labs Bristol, along with simple implementations of most of the required solution components. Work is in progress to build a full working prototype and make experiments in real-world contexts.

7. References

[1] A. Baldwin, S. Shiu - ACTS: an Accountable Service for Controlled Sharing of Actions. HPL-2002-334 [HP Restricted], 2002

[2] D. Boneh, M. Franklin - Identity-based Encryption from the Weil Pairing. Crypto 2001, 2001

[3] M. Casassa Mont, R. Brown - PASTELS project: Trust Management, Monitoring and Policy-driven Authorization Framework for E-Services in an Internet based B2B environment. HPL-2001-28, 2001

[4] L. Chen, K. Harrison, A. Moss, D. Soldera, N. P. Smart - "Certification of Public Keys within an Identity Based System", LNCS 2433, ed. G. Goos, J. Hartmanis and J. van Leeuwen, Proceedings of Information Security, pp. 332-333, 2002.

[5] C. Cocks - An Identity Based Encryption Scheme based on Quadratic Residues. Communications-Electronics Security Group (CESG), UK. http://www.cesg.gov.uk/technology/id-pkc/media/ciren.pdf , 2001

[6] F. Gallegos, D. P. Manson, S. Allen-Senft - *Information Technology Control and Audit*. Auerbach, 1999

[7] R Housley, W. Ford, W. Polk, D. Solo - RFC2459: Internet X.509 Public Key Infrastructure Certificate and CRL profile, IETF, 1999

[8] G. Karjoth, M. Hunter - A Privacy Policy Model for Enterprises, IBM Research, Zurich - 15th IEEE Computer Foundations Workshop, June 2002

[9] G. Karjoth, M. Schunter, M. Waidner - Platform for Enterprise Privacy Practices: Privacy-enabled Management of Customer Data - 2nd Workshop on Privacy Enhancing Technologies, Lecture Notes in Computer Science, Springer Verlang – 2002

[10] Liberty Alliance Project - http://www.projectliberty.org/, 2002

[11] Microsoft - Microsoft .NET Passport, http://www.microsoft.com/netservices/passport/ , 2002

[12] Microsoft Corporation - White Paper on Palladium, June 2002. Available via http://www.microsoft.com/presspass/features/2002/jul02/0724palladiumwp.asp

[13] S. Pearson (ed.) - Trusted Computing Platforms, Prentice Hall, 2002.

[14] S. Pearson - "A Trusted Mechanism for User Self-Profiling in E-Commerce", Selected Papers from Special Track on Privacy and Protection with Multi-Agent Systems, LNAI journal, Springer, 2003.

[15] TCPA - Trusted Computing Platform Alliance Main Specification v1.1, www.trustedcomputing.org 2001

[16] W3C - The Platform for Privacy Preferences 1.0 specification (P3P 1.0). http://www.w3.org/tr/p3p - W3C Proposed Recommendation, 2002

Appendix A: IBE Cryptography Schema

The IBE cryptography schema [2, 4, 5] has two core properties:

- **1st Property**: any kind of string can be used as an IBE encryption key (public key). This "string" consists of any sequence of characters or bytes such as **a role**, a text, a name, an e-mail address, a picture, a list of terms and conditions, etc. Information is encrypted by using this string along with a "public detail", uniquely associated to a specific trusted third party, referred in this paper as *trust authority (TA)*. This trust authority is the only entity that can generate the correspondent IBE decryption key. It only relies on a local *secret* that is a critical resource and needs to be properly protected;

- **2nd Property**: the generation of an IBE decryption key (associated to an IBE encryption key, i.e. a string) can be postponed in time. In other words an IBE decryption key can be generated (by a trust authority) a long time after the correspondent IBE encryption key was created.

Figure A.1 shows the basic IBE interaction model:

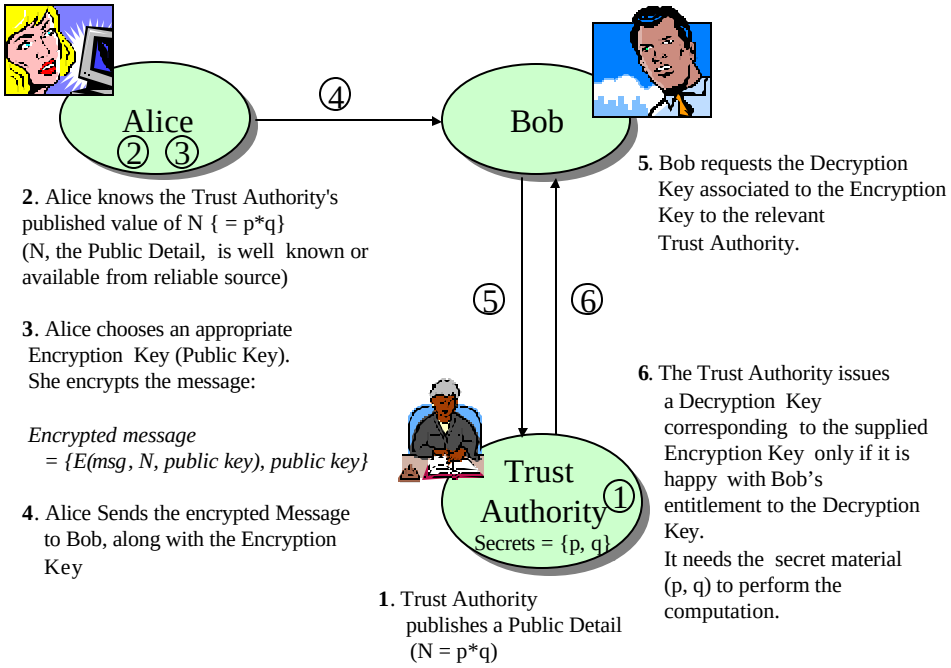

Figure A.1: High-level IBE Interaction Model

Three players are involved in the above interaction model: a sender of an encrypted message (Alice), the receiver of the encrypted message (Bob) and a trust authority in charge of issuing decryption keys.

Alice wants to send an encrypted message to Bob. Alice and Bob trust a third party, the trust authority (TA). The following steps take place:

1. During the TA's initialisation phase, the TA generates a secret (stored and protected at the TA site) and a correspondent "public detail" that is publicly available.

2. Alice trusts the TA. She retrieves the public detail from the TA site;

3. Alice wants to send a message to Bob. She defines an appropriate IBE encryption key (public key) to encrypt this message. The IBE encryption key can be any type of string, for example Bob's role or Bob's e-mail address. Alice's message is encrypted by making use of this IBE encryption key and the TA's public detail.

4. Alice sends the encrypted message to Bob, along with the IBE encryption key she used to encrypt the message.

5. Bob needs the decryption key associated to the above IBE encryption key, to decrypt Alice's message. Bob has to interact with the trust authority. He might have to provide additional information (credentials) to prove he is the legitimate receiver of the message.

6. The trust authority generates and issues to Bob the IBE decryption key (associated to the IBE encryption key chosen by Alice) if it is satisfied by Bob's "credentials". The trust authority might decide to generate the IBE decryption key depending on the fulfilment of specific constraints as specified by the correspondent IBE encryption key. For example a trust authority might issue an IBE decryption key to Bob only if he is compliant with a well-defined list of terms and conditions. Please notice that the IBE public key (i.e. a string), used to encrypt the document, would directly specify the list of these terms and conditions.

Appendix B: TCPA

The *Trusted Computing Platform Alliance* (TCPA) is an industry alliance formed in October 1999 that focuses on developing and standardizing Trusted Platform technology. The TCPA specification, released in February 2001, is designed to be independent of the type of platform (PC, server, PDA, printer, mobile phone, and so on), although the technology for creating a Trusted PC has so far been fully specified. The specification is intended for use in the real world of electronic commerce, electronic business, and corporate infrastructure security. The technology addresses such questions as: "How can I trust a remote system that is not under my control?"

B.1 Trusted Platforms

A *Trusted Platform* is a computing platform that has a trusted component, probably in the form of built-in hardware, which it uses to create a foundation of trust for software processes. Trusted Platforms get their name from the fact that they enable either a local user or a remotely communicating user to trust a platform for some particular purpose. A behavioural definition of trust has been adopted by TCPA: *An entity can be trusted if it always behaves in the expected manner for the intended purpose.*

The computing platforms listed in the TCPA specification are one such type of Trusted Platform. A (TCPA) Trusted Platform has improved data protection and identification; it enables users to decide whether it is safe to use the platform for sensitive tasks, and maintains user privacy. It provides most of the basic features of a secure computer, but does so using the smallest possible changes to standard platform architectures.

It does this by providing the following basic functionalities:

1. Protection against theft and misuse of secrets held on the platform. Such secrets are rendered unintelligible unless the correct access information is presented and the correct programs are running.
2. A mechanism for the platform to prove that it is a TP while maintaining anonymity (if required).
3. A mechanism for a platform to show that it is executing the expected software: the integrity of a TP, including the integrity of many components of the platform (such as BIOS, OS loader and so on) can be checked by both local users and remote entities. This mechanism is used to provide the information needed to deduce the level of trust in the platform.

The architecture of a TP has to be fundamentally different from existing computing platforms in that it must include cost-effective security hardware (roughly equivalent to a smart card chip) that acts as the "root of trust" in a platform. This device is called a *Trusted Platform Module* (TPM). The TPM, as described in [15], is physical to prevent forgery, tamper-resistant to prevent counterfeiting, and has cryptographic functionality.

The TCPA architecture is designed to provide immediate, intermediate, and long-term benefits to users. Some features will be available immediately, while other features require further software development (expected shortly). The most advanced features require a public key infrastructure and are designed for use by e-services.

B.2 Privacy using TCPA

Platform privacy is already an issue, because of identification of platforms from MAC and IP addresses, for example. However, TCPA technology is designed with privacy protection in mind, and provides the following features:

- The owner has complete control over activation of the TPM (the manufacturer and users can also turn it off).
- The owner has complete control over generation of TCPA identities.
- Each user's data can be kept private and even the platform owner or administrator cannot access that data without the necessary access data.
- The revelation of secrets can be prevented unless the software state is in an approved state.

For further discussion of TCPA capabilities, see [13].