



## **Identity Management: On the “Identity = Data + Policies” Model**

Marco Casassa Mont  
Trusted Systems Laboratory  
HP Laboratories Bristol  
HPL-2004-14  
February 5<sup>th</sup> , 2004\*

E-mail: marco\_casassa-mont@hp.com

identity  
management,  
policies, policy  
management,  
identity model,  
enforcement,  
accountability,  
privacy

Digital identities are fundamental to enable digital interactions and transactions on the web. The current digital identity model, based on the “identity = data” paradigm, starts showing its limitations when addressing people's expectations about their identities (in terms of preferences, privacy, trust, etc.) and providing them with degrees of assurance that expectations will be met. An alternative model is introduced, based on the “identity = data + policies” paradigm, along with an underlying policy management framework. Details are given on how this model can address the above issues and how the framework can be implemented. Related technologies and work done by HP Labs Bristol are presented and discussed.

# Identity Management: On the “Identity = Data + Policies” Model

Marco Casassa Mont  
Trusted Systems Laboratory  
Hewlett-Packard Laboratories, UK  
marco\_casassa-mont@hp.com

## Abstract

*Digital identities are fundamental to enable digital interactions and transactions on the web. The current digital identity model, based on the “identity = data” paradigm, starts showing its limitations when addressing people’s expectations about their identities (in terms of preferences, privacy, trust, etc.) and providing them with degrees of assurance that expectations will be met.*

*An alternative model is introduced, based on the “identity = data + policies” paradigm, along with an underlying policy management framework. Details are given on how this model can address the above issues and how the framework can be implemented. Related technologies and work done by HP Labs Bristol are presented and discussed.*

## 1. Introduction

Digital identities are more and more important in everybody’s life. They are collected, stored and managed by service providers, government agencies and enterprises to enable people to access web services and digital information, customise their interactions and transactions; they are also used by organizations to obtain financial insights, knowledge about individuals and competitive advantages.

Digital identities comprise a variety of digital information, including personal data, financial details, profiling information.

Current digital identities are based on the “*identity = data*” paradigm i.e. they are basically modeled as an aggregation of identity attributes. The fact that part of this data might have degrees of certification and be associated to preferences does not substantially invalidate the above statement: preferences expressed by users can be easily ignored or bypassed.

Today’s management of identity information is very pragmatic and reflects organizations’ needs to process large amounts of data in simple and efficient ways. Despite the increase of computing power and progress made in the identity management space, little has been done to take on board people’s expectations (in terms of preferences, privacy, trust, etc.) about their digital

identities and provide them with degree of assurance that these expectations will be satisfied.

This paper analyses in more details some of the limitations of the “identity = data” model and introduces an alternative model, based on the “identity = data + policies” paradigm, where policies describe expectations. It describes how related issues, including policy stickiness to identity information, policy enforceability, accountability and trust can be addressed by the new model.

Work done in this area by HP Labs Bristol is presented and discussed.

## 2. Need for a Paradigm Shift

In the digital world, digital identities are perceived, used and managed differently from what happens for “identities” in the physical world.

In the physical world, people are not only a matter of “attributes and identifiers”, such as a name, an identity card, a passport or credit card details: in other words, the identity of a person cannot be exclusively characterized by these aspects [1].

People have expectations, behaviours, desires, fears, preferred ways of dealing with situations on a contextual basis, etc. People might keep their anonymity and, at the same time, express and convey some of the above aspects. All of them contribute to shape their identities.

To progress in the areas of digital identity and identity management it is necessary to take these aspects into account. In particular it must be possible to capture people’s expectations, desires and policies and provide mechanisms to enforce them in a variety of contexts.

In the current digital world, digital credentials like X.509 certificates [2,3] can be used to represent identity attributes, along with a certification of their validity and trustworthiness: they offer little support to represent and enforce people’s expectations, preferences and policies.

Organizations (such as e-commerce sites, service providers, enterprises, etc.) allow people to specify simple preference policies (for example related to data protection and privacy issues, customization, etc.) but in many cases their associations to identity information is weak and can be easily broken, bypassed or ignored.

The trustworthiness and validity of digital identities is usually checked by data receivers (e.g. service providers, organizations, etc) with thoroughness that depends on the involved risks and costs. Additionally, the management and enforcement of related identity policies are also done by these parties.

People need to trust these receivers that they will act honestly without misusing their information.

Usually data receivers dictate their “rules of engagements”. This happens not only for people but also for organizations sharing their customers’ data with other organisations, for example in supply chain contexts, federated identity contexts, etc.

Various legislations and laws are available to address some of the related issues, like data protection and privacy [4], but they are hard to enforce, especially in contexts where identity information flows across organizational and geographical boundaries.

A paradigm shift is necessary in order to:

- Allow people (or third parties acting on their behalf) to explicitly dictate how their identity information must be managed and used, via a strong association of policies (i.e. expectations) to the identity information. Policies become an integral part of identity;
- Provide a framework for policy enforcement and accountability management, driven by these policies.

The remaining part of this paper explores the “identity = data + policies” model and related issues. As a significant case, we analyse its implications for privacy management. Related research and work done at the Trusted Systems Laboratory, HP Labs, Bristol are also discussed.

### 3. The “Identity = Data + Policies” Model

The “identity = data + policies” model is based on the concept that people’s expectations (including requirements and obligations) are explicitly associated to their identity information via policies, as shown in figure 1:

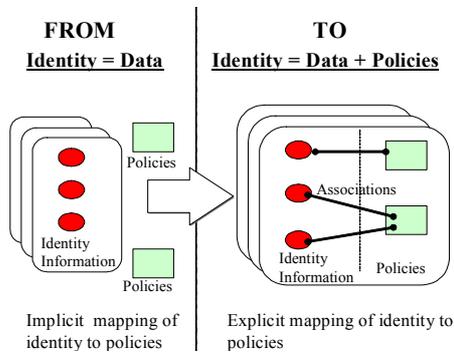


Figure 1. Paradigm Shift for Identity

Ideally, policies stick to the identity information both when stored and transmitted:

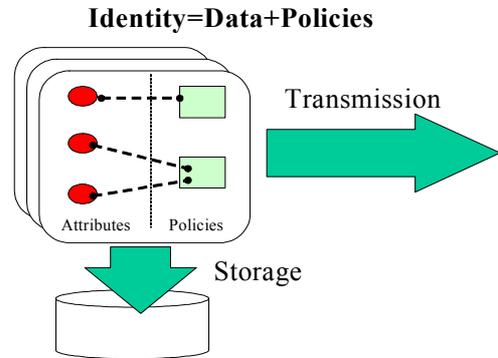


Figure 2. Stickiness of policies during storage and transmissions

The focus is on people and their requirements. Policies must be seen as tools available to people - or trusted third parties acting on their behalf - to dictate: constraints and obligations on the way their identity information is going to be managed; preferences; rules of engagement, etc.

Policies can focus on different aspects, including:

- **privacy:** policies define conditions and constraints on how data must be handled, disclosed to other parties, protected, etc;
- **authorization:** policies dictate who can access what and under which conditions;
- **obligation:** policies define constraints that need to be satisfied and fulfilled, potentially over a long period of time (such as retention policies);
- **preferences:** policies define preferences when multiple choices are available, for example in the way identity information is handled, disclosed or used;
- **trust:** policies dictate trust requirements to be satisfied by the involved parties;
- **control:** policies allow people to be involved in the management and monitoring of their identity data.

A few examples of policies, from a user’s perspective, follow:

- Store my identity information in a private way so that only legitimate people or applications can access it, subordinate to their intent and the information purpose;
- Do not disclose any of my personal details to entities X, Y Z;
- Delete my identity details immediately after my transaction has been accepted;
- Delete my personal information after X years;
- Use my identity details X when dealing with entity Y;

- Notify me via e-mail, every time you use some of my identity information;
- Ask for my authorization (via a predefined communication channel) every time you need to disclose this attribute to a third party;
- Interact with this trusted authority to state your intentions before obtain the current values of these attributes. You will be audited.

Policies might need to be modified over time (by entities who are entitled to), depending on contextual situations, specific circumstances, etc. The next sections analyse a few issues related to the stickiness of policies to identity information and ways to address them.

People must be offered the possibility to directly manage their digital identities and be assured that their personal information is used or managed according to their expectations. In case they do not care, have no time or are not able to directly manage their personal data, they should still have the option to delegate this effort to trusted third parties (that will act on their behalf).

### 3.1. Issues

The act of storing identity information or disclosing identity information to third parties is critical: once confidential information is available to other parties, it can potentially be misused and the control lost.

Policies can dictate how identity information has to be managed before and after its disclosure. However, there are important issues that need to be addressed:

- **Stickiness of policies:** How can policies be strictly associated to identity information? If it is possible to break this association or ignore these policies or if their entire management is delegated to the data receivers, we are back to a situation of reliance and trust on the receivers. Is it possible to prevent this from happening? If not, how can we provide mechanisms to minimize this and enforce accountability?
- **Trust:** today people have to trust data receivers that they will handle identity information in a law compliant way and expectations will be fulfilled. People must trust receivers' technical competence, the adequacy of their IT infrastructures, storage and processes and their compliance to security standards.  
Can this "trust dependency" be shifted from data receivers to third parties trusted by people? Can compliance checks be done upfront, prior to any data disclosure, to verify that data receivers satisfy a predefined set of legal and technical requirements? Could trusted third parties do this on behalf of people?

- **Enforcement:** How can people be sure that the policies associated to identity information are going to be enforced, both when this information is stored and sent to other parties? How can data receivers be forced to go through the required policy enforcement steps? How can we reduce the risk of fraudulent or unintentional misuses of identity information?
- **Accountability:** Can people be assured that data receivers will fulfill their promises and enforce people's requests and preferences (in terms of privacy, confidentiality, etc.)? How can data receivers be made more accountable for their actions, in particular for the way they handle identity information? How can they be audited and non-repudable evidence collected?
- **Monitoring:** How can people (or third parties acting on their behalf) monitor what happens to their identities, once they have been disclosed? How can people (or third parties acting on their behalf) directly control and manage the destiny of their confidential information?

It is possible to address the above issues by:

- Making policies become integral part of the identity information i.e. leveraging the "identity = data + policies" model;
- Providing an underlying policy management framework to reduce people's reliance on data receivers and provide accountable policy enforcement mechanisms.

Despite the fact that this topic is still under research, systems and solutions can be built right now to provide parts of the required enforceability and accountability functionalities.

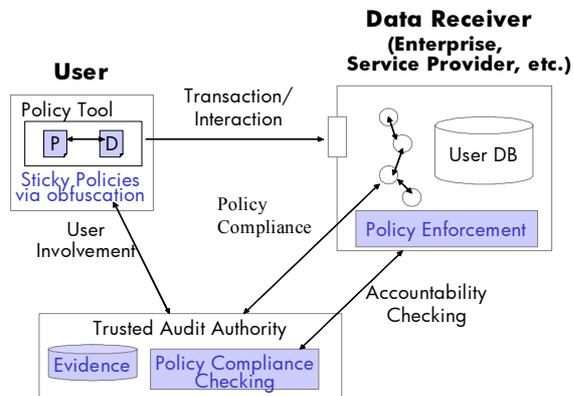
## 4. Policy Management Framework

This section describes a policy management framework underpinning the "identity = data + policies" model, how it can be used to address the issues described above and how HPL Bristol technologies can be used to implement some related aspects.

The framework is based on four core mechanisms:

- Mechanisms to strongly associate policies to identity data;
- Trusted Audit Authorities, i.e. trusted third parties, to enforce (aspects of) policies, audit data receivers and provide users with degrees of control on their information;
- Mechanisms to control the storage and flow of identity (data + policies) and enforce (aspects of) policies at the operating system level;
- Mechanisms to check for the integrity of IT platforms used to manipulate and manage identity information.

Figure 3 contains a high level architecture that shows how these components fit together to deal with an **identity flow** scenario i.e. identity information exchanged among multiple parties as a consequence of electronic transactions and interactions:



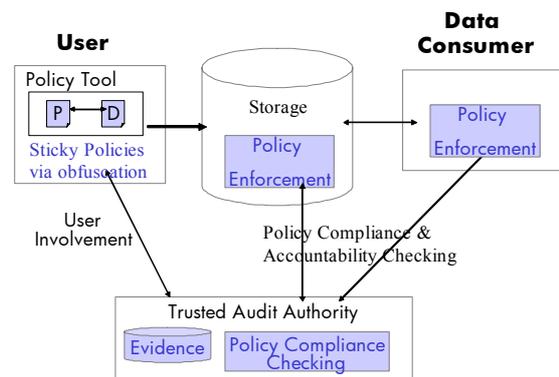
**Figure 3. Identity Flow: Model Components**

Identity information is obfuscated and strongly associated to policies before it is sent to data receivers: this association can be achieved by using traditional cryptographic and enveloping techniques [5] or alternative cryptographic schemas (like IBE [6,7]). Related details can be found in the next section. Only the fulfillment of these policies enables the data receiver to access the identity data. This process is mediated by one or more trusted services via their policy compliance checking mechanisms. These services can be provided by trusted third parties, referred by this paper as Trusted Audit Authorities (TAA). The TAA will disclose the decryption keys (needed to de-obfuscate the data) only if the associated policies are fulfilled. The acceptance of policies and the evidence collected by the TAA makes data receivers more accountable.

However, once data is disclosed it can be misused by the receivers: in addition, the association between identity information and policies can be broken. Depending on the importance of the identity information and the involved risks, the TAA policy compliance checking mechanisms can do further checks, prior to data disclosure, about the trustworthiness and integrity of the data receiver's IT environment and its policy enforcement mechanisms.

Similar issues apply when confidential data is stored and needs to be protected. It is not only a matter of access control as traditional *access control lists* and security mechanisms have a limited expressiveness of the kinds of constraints (and policies) they can dictate and enforce.

Figure 4 contains a high level architecture showing how our components fit together in an **identity storage** scenario:



**Figure 4. Identity Storage: Model Components**

Similarly to what has been explained for figure 3, identity information is stored in an encrypted way along with related policies. When access to this identity information is required, the TAA checks for policy compliance. This might include checking for the policy enforcement mechanisms and contextual properties of the entity asking to access the data.

#### 4.1 Policy Stickiness

The stickiness of policies to identity information is obtained by obfuscating the identity information in a way that its de-obfuscation is a "function of" the associated policies. Any tampering with these policies must prevent the de-obfuscation of data.

This can be achieved by using traditional public key cryptography along with enveloping mechanisms or by alternative cryptographic mechanisms:

- In the former case, the TAA publishes its (certified) public key (and, of course, its correspondent private key is kept secret). A symmetric key is generated by the identity owner (user) and used to encrypt the identity data. The symmetric key and a hash value of the associated policies are encrypted in a package [5] by using the TAA public key. The overall information (encrypted data, clear text policies and package) can now be sent to any data receiver. The TAA is the only entity that can decrypt the above package, check for the integrity of the associated policies, check for their compliance and eventually disclose the symmetric key.
- An alternative approach is based on the IBE technology [6,7,8]. Any kind of strings (including texts, pictures, terms and conditions, etc.) can be used as IBE encryption key. Policies can be used for this purpose. The correspondent IBE decryption key can only be generated by at least a Trust Authority (TA) i.e., in this case, the TAA. The TAA will check the compliance of any

requestor with the policies. The generation of IBE decryption keys (by one or more TAAs) can be postponed in time i.e. until they are actually necessary for decryption purposes. Any tampering with the IBE encryption key will make impossible for the TAA to generate the correct decryption key.

In both cases the disclosure of identity data happens when the decryption key is disclosed by the TAA. In terms of implementing the “stickiness of policies” to identity information, the above mechanisms are conceptually equivalent.

Policies might need to be modified overtime, despite their stickiness to data. Different mechanisms and methodologies can be used to achieve this:

- Setting an expiration date for policies: in this case the obfuscated data will not be accessible anymore after the expiration date. New identity information along with policies must be re-obtained from the user;
- Usage of indirection mechanisms: policies consist of meaningful “labels” containing references to “sites” where their complete definitions can be retrieved. For examples these references could contain URLs pointing to policies stored by TAAs. Users (or third parties acting on their behalf) will be able to modify these policies (but not the labels). The TAA will interpret policies according to their latest versions even if policies changes will be monitored and kept as evidence;
- Usage of on-the-fly refinement of high level-policies by TAAs. Only high level policies are associated to identity information; policy details are refined by TAAs. The user (identity owner or third party acting on their behalf) can modify how policies are interpreted and refined by the TAAs.

Hybrid combinations of the above mechanisms can be used depending on the nature of the identity data and their policies.

## 4.2 Accountability

Data receivers can be made more accountable thanks to the mediation of Trusted Audit Authorities (TAAs).

Data receivers must interact with one or more TAAs in order to get the decryption keys necessary to access the encrypted identity data. This must happen at least the first time identity information is disclosed or accessed.

TAAs check for the integrity of the associated policies and verify that the requestors are compliant with the conditions and constraints dictated by these policies. Policies might require the requestors to fully authenticate and provide signed statements (assertions) about their intention.

However, once decryption keys have been disclosed to the requestors, data can potentially be misused. TAAs

log all the relevant information exchanged during the interactions with requestors. The collected evidence can be used later on for forensic analysis and to pin down responsibilities.

Additional enforcement mechanisms can be used to mitigate the involved risks.

## 4.3 Policy Enforcement

The enforcement of policies is a key aspect of our model. It can happen:

- **Prior to the disclosure of data:** policy conditions and constraints are checked by the TAA prior to the disclosure of decryption keys. Further assurance on the quality of this enforcement can be obtained if data receivers’ IT systems are checked by the TAA, prior to the disclosure of the data.
- **After the disclosure of data:** policy conditions and constraints are enforced by data receivers.

Both enforcement mechanisms require:

- Collaboration among TAAs and data receivers to assess the receivers’ IT systems and configure aspects of their systems;
- Auditing by competent authorities (it could be a role played by TAAs);
- The availability of trusted technologies that enable the enforcement of policies in IT environments that are not directly under the control of TAAs.

Three technologies and mechanisms can be leveraged to achieve (part of) the above types of policy enforcement:

- **TAA policy checking mechanism:** the TAA uses their policy checking mechanisms to check for the fulfilment of policies, upfront the disclosure of decryption keys. These mechanisms are based on a policy engine driven by policies and contextual information (requestors’ credentials, profiles, IT measures of trust, etc.). This policy enforcement mechanism is a “soft” mechanism: it can be bypassed by misbehaving receivers, once data has been disclosed.
- **Trusted platforms:** data receivers might be asked (via policies’ constraints) to use trusted platforms when dealing with identity information. Emerging technologies such TCG (a.k.a. TCGA [9]) trusted modules - TPMs - can be used by TAAs to obtain measures of SW and HW integrity of remote platforms and compare them against policies. This type of policy enforcement is done upfront to data disclosure: it aims at reducing the risks of disclosing confidential data to IT infrastructure which are badly run or have been compromised.

- **Data tagging at the OS level:** a tagged OS [10] allows tags to be directly associated to data (bytes in memory), instead of data containers (i.e. files, documents, etc.). Tags can be linked to policies which dictate how data must be managed, i.e. if it can be copied or merged with other data, where (and if) it can be transmitted to, if it has to be encrypted, etc. It is the OS that enforces these policies rather than the applications. The OS needs to be a trusted OS. By combining TCG/TPM and tagged OS it is possible to increase the level of assurance on the remote systems that they will act as dictated by the policies and they will enforce them. The data receiver might be asked to make use of tagged OSs, the integrity of which is checked by the TAA (prior to the disclosure of decryption keys): policies could be set by the TAA on the remote receiver's systems. After the disclosure of decryption keys, the enforcement of policies is still done by the receiver's platforms, but this time relying on checked and verified trusted platforms.

#### 4.4 Trust

The trust model of the proposed model is centered on the concept of having one or more trusted third parties, TAAs, mediating identity disclosures.

Users can get degrees of assurance that their policies will be satisfied by relying and trusting a TAA (a known party) instead of having to trust the data receivers.

This trust model can be extended by having multiple TAAs, acting in a collaborative way when dealing with policy checking and enforcement. There is no fundamental reason why users should be prevented from running their own TAAs.

The usage of multiple TAAs mitigates the risk of having to trust and relying on only one party.

#### 4.5 Monitoring

The TAA plays a key role also in providing raw information that can be processed and used to monitor disclosures of identity information.

Monitoring activities can be performed by the TAA to spot anomalous situations and trends and prevent misuses: this can be achieved by analysing and correlating evidence collected during disclosures.

Tools can also be provided to users by the TAA to monitor the disclosures of their personal data. These tools provide simple reports based on information collected by TAAs during disclosures.

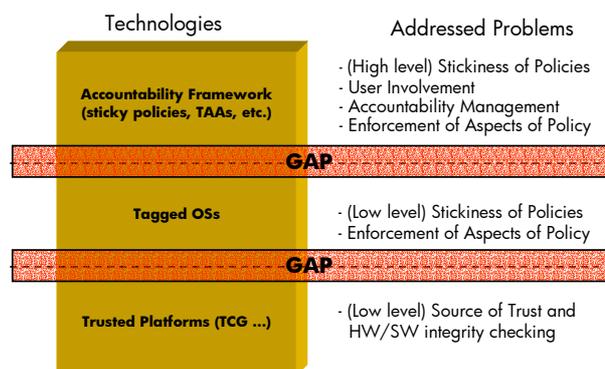
### 5. Discussion

The “identity = data + policies” model is based on the concept of associating policies to identity information to explicitly express them and deal with their enforcement.

The policy management framework described in this paper show how an implementation of this model can be achieved by leveraging technologies and mechanisms already available today (even if they are at different maturity stages).

However, there are aspects and issues that need to be fully investigated in order to draw conclusions about the feasibility of the proposed model. This is part of our ongoing research. A few important aspects are discussed in the remaining part of this section.

From a technology perspective, we illustrated the usage of three kinds of technologies: cryptographic mechanisms, TCG/TPM (trusted platforms) and Tagged OS. Their relevance and applicability is summarized in Figure 5:



**Figure 5. Technologies and addressed Problems**

The involved technologies are currently available at different maturity stages: whilst TCG/TPMs chips are already available on the market, Tagged Os and solutions underpinning TAAs are still in a research stage. Additional mechanisms are currently under research and development at the trusted platform level such as Microsoft NGSCB [11]. Prototypes and trials need to be done to fully understand the implications – in terms of integration, usability, flexibility, scalability, etc. - of using such technologies.

Further research needs to be done on policies and how to describe and integrate policy aspects (constraints, obligations, conditions, etc.) in a smooth and simple way, at different levels of abstractions. In particular, the integration of sticky policies with stored data is complex and hard to achieve in a way that performance and flexibility are not compromised. Work is in progress in this space.

Transparency and simplicity of use are key requirements, both for end-users and administrators. In particular, end-users require simple tools to set their policies either based on default assumptions/templates or via the mediation of trusted third parties.

It is important to understand what a reasonable trade-off is between increasing the assurance offered to users and the involved IT costs. Even if multiple policy enforcement mechanisms are potentially available, their usage must be subordinated to the involved risks, costs and importance of identity data.

Further analysis needs to be done on the applicability of the proposed model to existing approaches to federated identity management. The work done in the Liberty Alliance Project [12] and similar identity and web service federation initiatives must be leveraged. Specifically it is important to understand if and how TAA functionalities can be implemented by Identity Providers and how to translate these functionalities into standards.

Additional research is done at HP Labs to explore the suitability of this model for distributed Agent Frameworks and how Semantic Web can be leveraged to describe and manipulate policies.

## 6. Applicability of the Model for Privacy Management

Research has been done by HP Labs Bristol to understand the implications of the “identity = data + policies” model for privacy management.

Privacy management is a promising area. Little work has been done so far to provide technological-based solutions to enforce privacy policies and, at the same time, increase people’s assurance by making the involved parties more accountable.

In this context, policies can be seen as “privacy policies”, strictly associated to identity information and enforced by the mechanisms described in section 4.

Figure 6 provides details about the usage of the identity model for privacy management, specifically in an “identity flow” scenario, where identity information is exchanged during multi-party interactions and transactions. Figure 7 describes the architecture of a system implementing the policy management framework. Papers [13,14] provide the details about the addressed scenario and the architectural components.

Privacy policies and their enforcement are a key aspect that is currently under investigation. Current standards, like W3C P3P [15], and proposed specifications, like W3C EPAL [16], focus mainly on policy languages. They do not fully address other important issues, like policy stickiness, enforceability and accountability, trust and monitoring.

We have prototypes of all the core components and we are working to integrate them in a comprehensive solution. The applicability of the model has still to be demonstrated. Part of this work could be done in the context of the EU PRIME project [17], an international project funded by the European Union and focused on research on privacy for identity management in Europe.

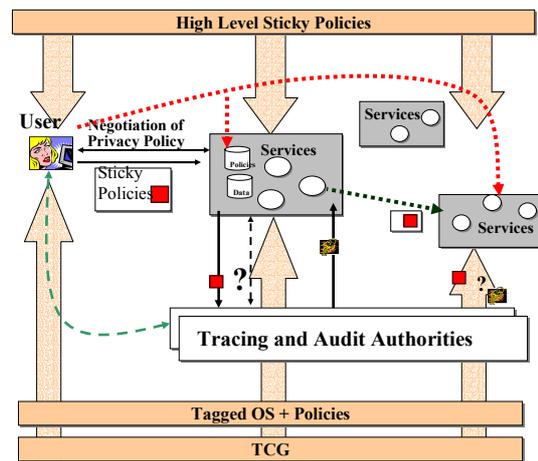


Figure 6. Privacy for Identity Flow

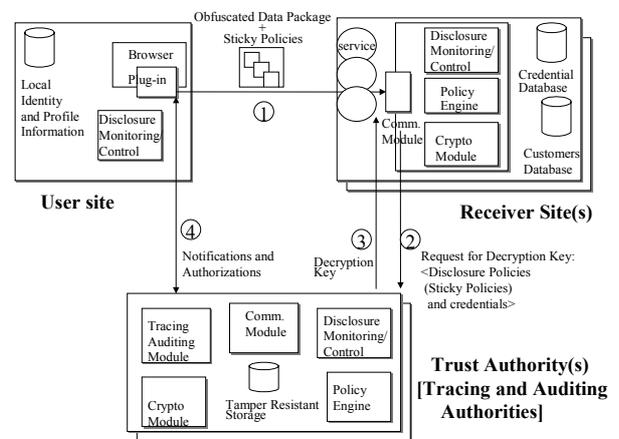


Figure 7. Architectural Details

## 7. Conclusion

The current “identity = data” model shows its limitations when dealing with users’ expectations: little has been done to address issues like enforceability of users’ policies, accountability and trust.

An alternative model able to address to above issues, based on the “identity = data + policies” model, has been introduced and discussed. An underlying policy management framework is currently under research and

development by HP Labs, Bristol along with related technologies. Prototypes of most of the required components are available but a comprehensive solution has still to be implemented and tested in a real-life scenario.

Privacy management is an interesting area to explore the applicability of our work. Further exploration of this space could be done in the context of the EU PRIME project.

## 8. References

- [1] M. Casassa Mont, P. Bramhall, J. Pato, On Adaptive Identity Management, HP Labs, HPL-2003-149, 2003
- [2] R Housley, W. Ford, W. Polk, D. Solo, RFC2459: Internet X.509 Public Key Infrastructure Certificate and CRL profile, IETF, 1999
- [3] D. W. Chadwick, S. Legg, Internet X.509 Public Key Infrastructure Additional LDAP Schema for PKIs and PMIs, IETF, 8 September 2000
- [4] C. Laurant, Privacy International, Privacy and Human Rights 2003, EPIC and Privacy International, <http://www.privacyinternational.org/survey/phr2003/>, 2003
- [5] RSA, PKCS#7, Cryptographic Message Syntax Standard, <http://www.rsasecurity.com/rsalabs/pkcs/pkcs-7/>, 1997
- [6] D. Boneh, M. Franklin, Identity-based Encryption from the Weil Pairing. Crypto 2001, 2001
- [7] C. Cocks, An Identity Based Encryption Scheme based on Quadratic Residues. Communications-Electronics Security Group (CESG), UK, 2001
- [8] M. Casassa Mont, P. Bramhall. IBE Applied to Privacy and Identity Management, HP Labs, HPL-2003-101, 2003
- [9] T CPA, Trusted Computing Platform Alliance Main Specification v1.1, <http://www.trustedcomputing.org>, 2001
- [10] Y. Beres, C. I. Dalton: Dynamic Label Binding at Runtime. In: Proceeding of New Security Paradigms Workshop, August 2003. (2003)
- [11] Microsoft Corporation, Next Generation Secure Computing Base (NGSCB), <http://www.microsoft.com/resources/ngscb/>, 2004
- [12] Liberty Alliance, <http://www.projectliberty.org/>, 2004
- [13] M. Casassa Mont, S. Pearson, P. Bramhall, Towards Accountable Management of Privacy and Identity Management, ESORICS 2003, 2003
- [14] M. Casassa Mont, S. Pearson, P. Bramhall, Towards Accountable Management of Identity and Privacy: Sticky Policies and Enforceable Tracing Services, TrustBus 2003 Workshop, DEXA 2003, 2003
- [15] W3C, The Platform for Privacy Preferences 1.0 specification (P3P 1.0). <http://www.w3.org/tr/p3p> - W3C, 2002
- [16] P. Ashley, S. Hada, G. Karjoth, C. Powers, M. Schunter, Enterprise Privacy Authorization Language (EPAL), IBM, 2003.
- [17] EU PRIME, Privacy and Identity Management for Europe, <http://www.prime-project.eu.org/>, 2004