Trusted Agents that Enhance User Privacy by Self-Profiling

Siani Pearson
Trusted E-Services Laboratory
HP Laboratories Bristol
HPL-2002-196
July 15th , 2002*

E-mail: Siani_Pearson@hp.com

privacy, distributed profile, agent, trusted computing platform, integrity checking

In this paper, a method is described for user self-profiling when engaged in e-commerce over the Internet, by which customers can have greater control over profiles relating to their behaviour or preferences and can exploit this information without revealing their identity. This is achieved using trusted agents that exploit Trusted Computing Platform technology (cf. www.trustedpc.org).

* Internal Accession Date Only Approved for External Publication
AAMAS Workshop, "Special track on privacy" Bologna, Italy, 15th July 2002

Trusted Agents that Enhance User Privacy by Self-Profiling

Siani Pearson
Hewlett Packard Laboratories
Stoke Gifford
Bristol, BS34 8QZ. UK.
+44-117-3128438

Siani_Pearson@hp.com

ABSTRACT

In this paper, a method is described for user self-profiling when engaged in e-commerce over the Internet, by which customers can have greater control over profiles relating to their behaviour or preferences and can exploit this information without revealing their identity. This is achieved using trusted agents that exploit Trusted Computing Platform technology (cf. www.trustedpc.org).

General Terms

Security.

Keywords

Privacy, distributed profile, agent, Trusted Computing Platform, integrity checking.

1. INTRODUCTION

Presently customer profile information (e.g. shopping habits) tends to be gathered by companies as customers interact with them (also through questionnaires, special offers, etc.), and this information is sold between companies. The central idea presented in this paper is to provide an alternative to the privacy violations associated with such an approach by a customer developing and securely recording their own profile (of e-shopping habits etc.) that pertains to a single registered identity that may or may not be anonymous. This profile can be made available either free or at low cost to companies in order that they can alert the customer to savings that the customer could have made by buying their products over others, special offers, and so on.

2. MOTIVATION FOR CUSTOMER SELF-PROFILING

In the field of networked computing systems there are many reasons why a business should form a profile of a user. For example, in a commercial context a supplier desires to obtain a profile of each customer including characteristics such as the type, quantity or frequency of product purchases. This customer profile then allows the supplier to offer incentives such as discounts appropriate to a customer's profile.

Typically, these customer profiles are held by the supplier, but give only a partial picture of the customer. Suppliers often desire to learn more about each customer, but a complete profile can only be obtained by combining profiles held by many different suppliers. Information sharing between a large number of suppliers requires a high degree of co-operation, and may impact upon the privacy and personal freedom of the customer (for example, personal details are often bought and sold without customers' knowledge or consent [17]).

This commercial context is just one example, but there are many other situations where user profiling is desirable. For example, personal data can be used to customise the client interface [13]. As a result several initiatives have recently been proposed related to ownership and server-side storage of customer data, such as Microsoft's My Services [11] and the 'Liberty Alliance' proposal for an open standard [9].

This paper provides an alternative to such approaches by using personalisation technologies to help the user. With the mechanism described in this paper, the customer develops a profile him/herself, on his/her own terms, with anonymity if required, and makes that information available to his/her own advantage. A benefit of this approach is the potential for intelligent interaction between the company and the secure customer profile database (e.g. through multivariate analysis [20]), without the database itself being fully divulged. Also, if the profile is anonymous the customer may be more inclined to divulge detailed information about shopping habits, etc. Furthermore, company offers may be strengthened by being highlighted as having been originated through secure self-profiling.

Companies may or may not be able to identify the real person/home address, etc., but they will be able to email offers to the anonymous trusted individual or, less intrusively, leave offers for collection over the web. Companies may not necessarily be able to download an entire profile, but rather interact with it.

Protection of the customer's identity can obviously benefit the customer but may also make life easier for the provider due to

limiting their data protection liability. In addition, making only relevant information available may be advantageous for the provider, particularly because providers are nowadays often bombarded with huge amounts of mostly irrelevant information out of which they have to mine a small relevant data set.

2.1 Example Scenario

An example of a scenario would be a customer who (as an anonymous trusted individual) flies London-New York regularly and buys through a particular e-company. Secure software records this information, and all other commercial interactions. Other companies can access this information (perhaps at small cost paid to the customer) from the customer's platform, knowing it is correct and secure, and can alert the customer of their superior products. From the customer's perspective this would be "reverse junk mail" that would leave them in credit. However, the company may not know other details of the customer, for example which newspaper that person reads.

3. TRUSTED COMPUTING PLATFORM TECHNOLOGY

This new approach for providing user self-profiling makes use of a new computer technology – Trusted Computing Platforms – for provision of trusted pseudonymous identities, hardware protection for secrets and an independent mechanism for verifying the trustworthiness of the agents that carry out the self-profiling. This section introduces the concept of Trusted Computing Platforms before looking at the general model for self-profiling in the next section.

Computer platforms are ubiquitous; they are central to the growing reliance on electronic business and commerce, and the need for information protection is increasing, particularly on client platforms. However, the degree of confidence in software-only security solutions depends on their correct installation and execution, which can be affected by all other software that has been executed on the same platform. Experts conclude that trusted hardware is needed as the basis for security solutions.

These factors, combined with increasing privacy issues and emerging ebusiness opportunities that demand higher levels of confidence, have led the *Trusted Computing Platform Alliance* (TCPA) to design a specification for computing platforms [19] that creates a foundation of trust for software processes, based on a small amount of hardware. The specification is intended for use in the real world of electronic commerce, electronic business, and corporate infrastructure security.

3.1 Trusted Platforms

A *Trusted Platform* (TP) – sometimes also called a Trusted Computing Platform - provides most of the basic features of a secure computer, but does so using the smallest possible changes to standard platform architectures. Essentially, it is a normal open computer platform that has been modified to maintain privacy. 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. (This is the TCPA protected storage mechanism).
2. A mechanism for the platform to prove that it is a Trusted Platform while maintaining anonymity (if required). This is discussed further in the following subsection.
3. A mechanism for a platform to show that it is executing the expected software: the integrity of a Trusted Platform, 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 trust decision itself can only be made by the entity that desires to use the platform, and will change according to the intended use of the platform, even if the platform remains unchanged. As an integral part of this process, the entity needs to rely on statements by trusted individuals or organizations about the proper behaviour of a platform.

For further discussion of these capabilities, see [14].

Physically, a Trusted Platform is distinguished by the use of 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 [22], is physical to prevent forgery, tamper-resistant to prevent counterfeiting, and has cryptographic functions to provide authenticity, integrity, confidentiality, guard against replay attacks, make digital signatures, and use digital certificates as required (further explanation of such terms is given in [16]).

Trusted Platforms are an unfamiliar concept, even to security specialists, but since the release of TCPA specification v1.0 in February 2001 and its backing by major industry players, they are set to become widespread.

3.2 TCPA Pseudonymous Identities

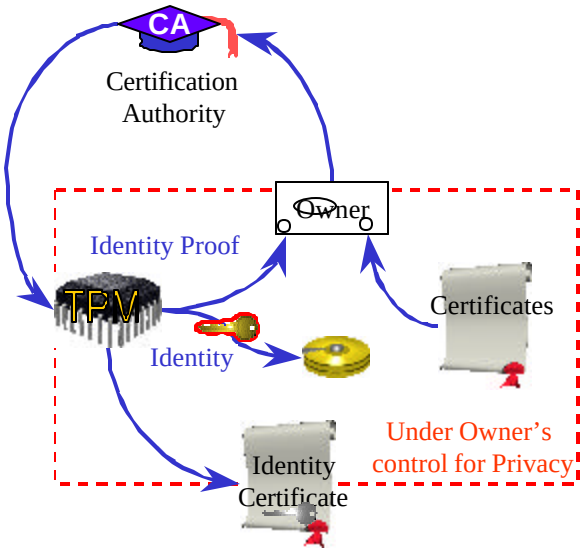

Figure 1. Creation of TCPA pseudonymous identities.

Proof that a platform is a genuine Trusted Platform is provided by cryptographic attestation identities and the process is illustrated in Figure 1. Attestation identities (also called 'pseudonymous identities') prove that they correspond to a Trusted Platform and a specific identity always identifies the same platform.

Each identity is created on the individual Trusted Platform, with attestation from a PKI Certification Authority (CA). Each identity has a randomly generated asymmetric cryptographic key and an arbitrary textual string used as an identifier for the pseudonym (chosen by the owner of the platform). To obtain attestation from a CA, the platform's owner sends the CA information that proves that the identity was created by a genuine Trusted Platform. This process uses signed certificates from the manufacturer of the platform and uses a secret installed in the TPM. That secret is known only to the Trusted Platform and is used only under control of the owner of the platform. In particular, it is not divulged to arbitrary third parties, unlike the cryptographic attestation identities.

The most valuable methods of using a TPM identity are:

Integrity signature. To prove that some particular data existed in a particular platform when the platform was in a particular state, the platform creates a digital signature over those data plus the current state (integrity values) using a platform identity (TPM identity).

Certifying other keys. One use of a TPM identity is to sign a statement about secondary asymmetric keys available to the TPM. These secondary keys may be used for any legitimate purpose by the OS or by applications within the platform.

These properties of TPM identities enable a number of possibilities, including providing confidence in the platform. A platform identity can enable a user to trust the platform where he/she is working. In addition, a TP enables users or organisations to verify not only that a remote platform is the correct platform, but also that it is running the correct software. These methods will be used to strengthen the trustworthiness of the self-profiling system presented in this paper.

3.3 Privacy

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 attestation identities. The origin of a specific identity cannot be tracked further, except by the Certification Authority (CA) that issues a certificate for that attestation identity. So appropriate selection of CAs enables the owner to control traceability from an attestation identity to the certificates that attest to a specific TPM and a specific platform. Identities can only be correlated with other identities by the CA that certifies these identities – and the owner has sole choice of that CA. So the owner can choose a CA whose policy is not to correlate identities, or whose policy is to correlate identities, according to the wishes of the owner. Different identities are used for different purposes and in particular, separate identities would usually be given to different users of the Trusted Platform. This property is exploited within the profiling system described in the following section.
- Each user's data can be kept private and even the platform owner or administrator cannot access that data without the necessary access data. Hence a platform could be owned and used by a single person (which would often happen in the case of consumers or small businesses), or owned by one entity and used by another entity. This would be typical in a corporate environment, where the IT department is the owner and the user is the individual who is issued with the platform. In the following section, this property is used to protect the privacy of the user's profile, even from a 'superuser' (whether administrator or hacker).

More detailed information on Trusted Platform technology can be found in [14].

4. A GENERAL APPROACH FOR PROVIDING SELF-PROFILING

Trusted customer self-profiling may be implemented using TCPA technology by using software associated with one or more TCPA anonymous identities associated with the user on his or her Trusted Computing Platform. This platform shall be referred to in this paper as the 'client platform', to distinguish it from a platform which communicates with the client platform in order to

obtain profile information, which shall be referred to as the 'enquiry platform'. Such platforms need not just be a desktop PC – they could be any type of computing platform, including laptop, server, Personal Digital Assistant (PDA), printer, or mobile phone.

4.1 Self-Profiling

By *profiling* we may understand 'the process of inferring a set of characteristics (typically behavioural) about an individual person or collective entity and then treating that person or entity (or other persons or entities) in the light of these characteristics' [3]. *Self-profiling* is the process of creating such a set of characteristics about oneself.

Correspondingly, in the context of this paper a *profile* refers to a set of user preferences or settings, which is the result of capturing certain user information (such as name, address, purchase preferences, etc.) and transforming them into a usable form.

Figure 2. Example user self profile.

Figure 2 shows an example user self-profile. The user self-profile comprises a user identity combined with one or more profile characteristics. The user identity comprises a certificate signed by a Privacy-CA, the certificate including a text identity label and a public identity key. Each of the profile characteristics may take any suitable form, and a profile characteristic is optionally verifiable with reference to an endorsement.

In use, the user self-profile is preferably supplied within a response signed by the Trusted Platform Module. By providing the user self-profile in a signed response, an enquirer has a high degree of confidence that the user self-profile has been formed in a trusted manner.

4.2 Trusted Self-Profiling Agents

Preferably, the software used for self-profiling will take the form of various agents that are independent of user control and can decide for themselves how to interpret the information they receive and act upon it. In this way the agents can act on behalf of the user, while their independence facilitates trust by other parties, in ways that will be explained below. Agents are used in order to make it easier to build a trusted system for self-profiling, and in particular to enhance user privacy by autonomously manipulating, capturing, requesting, examining and otherwise operating upon profiling records.

As is argued by Negroponte [12], agents can embody user profiles: 'the concept of 'agent' embodied in humans helping humans is often one where expertise is … mixed with knowledge of you'. Within our system, agents on each client machine would build up one or more *profiles* corresponding to users of that machine.

Preferably, third parties will check out the user's computer using TCPA integrity check mechanisms to make sure that the information they receive about the user's profile(s) can be trusted. In this case there would be one profile per TCPA identity (a 'label' chosen by the user).

The profile should preferably be stored in an encrypted form, and further confidentiality could be achieved by using TCPA 'protected storage' mechanisms (see [22]) to ensure that the profile is only released to an enquirer with the consent of the client's TPM. Furthermore, the agents could be protected by the platform's TPM (for example, via integrity checks at boot and even periodically during runtime, with the additional option of being located within the TPM). Further discussion of such protection is given in subsection 4.2.1 below.

A Trusted Platform could maintain a log of platform activity, signed using the platform's identity to establish the authenticity and integrity of the log. This log could be used as trusted input when generating a profile for the user, and either could be protected against unauthorized alteration, by using the TCPA 'protected storage' mechanisms.

Previous solutions involve external companies gathering information about users, which gives users less control over such information and potentially infringes their privacy and also opens them up to receiving multiple unwelcome mailings. This new solution involves the companies gaining useful information without the customers' actual identity being revealed. Customers will gain more control over their own profiles relating to their behaviour, and can profit from it without being exploited or having their behaviour tracked back to them personally.

4.3 Trustworthiness of this System

Trusted Platforms use the following definition of trust: an entity can be trusted if it always operates as expected for the intended purpose [21]. A platform cannot itself decide whether it is trusted because trust depends on the intended use of that platform. Only a user can decide whether the platform is trusted for the purpose intended by that user. So, the platform reports information to the user to enable that decision to be made. For further details, see [22].

Both local and remote entities can trust the mechanism proposed in this paper, because trust is provided via both of the following:

- Special processes in a Trusted Platform that dynamically collect evidence of (the platform and the agents') behaviour and provide evidence of this behaviour. This information provides the means of knowing whether the system (in the sense of 'platform plus agents') *can* be trusted.
- Social trust to provide confidence (a) in the mechanisms that collect and provide evidence of this behaviour as well as (b) that particular values of evidence represent a system that is in a "good" state. This information therefore provides the means of knowing whether a platform and the agents *should* be trusted.

Clearly, both aspects of trust are necessary when designing online systems, quite apart from additional social guarantees of privacy and security (as discussed for example in [19]). Processes in a Trusted Platform provide information about the behavior of a platform, but that information cannot be trusted unless someone vouches for the method of providing the information and for the expected value of the information. In our case, it will be necessary for third parties (such as the software developers of the agents) to vouch as to their trustworthiness and provide metric information that would allow both local and remote entities to judge whether the agents were operating as expected or had been corrupted. Such mechanisms will be discussed further in the following subsection. Further discussion about trust in agents, trust in virtual societies and analysis of the trustworthiness of Trusted Platforms may be found for example in [4], [5] and [2; 14] respectively.

4.3.1 Checking trustworthiness of the agents

Each agent may be integrity checked to ensure that the agent is operating as expected and has not been modified or substituted in an unauthorized manner. This process would involve a trusted third party (usually the vendor of the agent software) publishing or otherwise making available a signed version of the integrity measurements that should correspond to a genuine agent. Upon boot, each agent may be integrity checked with reference to its signed version and the corresponding (trusted third party's) public key certificate, and not be trusted for use if this integrity check fails. If the integrity check fails, it may be arranged that the complete platform integrity fails. The integrity checking is performed in an analogous manner to the platform integrity checking process, namely by measuring integrity metrics and comparing these with certified correct metrics. For example, by reading and hashing the agent code to produce a first hash; reading and decrypting the signed version using the public key certificate to produce a second hash; and comparing the first and second hashes (see [22] for example, for further explanation of the TCPA integrity-checking process, and [16] for further explanation of cryptographic notions).

The TPM can be used to provide protected storage for logs, digests, application-related data, agents, etc. via TCPA protected storage mechanisms so that such data cannot be interpreted by unauthorized entities. However, if these data are not stored within the TPM itself or within other tamper-resistant hardware, they will not be protected against unauthorized modification or deletion — although alteration to such data can be detected (for example by storing a digest within the TPM).

Therefore, the agents can be protected by the TPM both while not being used and while executing, by means of the integrity checking process described above, preferably at least some of them being stored within the TPM, and by the agents themselves running within a protected environment such as the TPM or within a suitably isolated compartment.

Note that both the profile creation and enquiry systems will still be Trusted Platforms, because, apart from the TPM, all other security functions (and ordinary software) can operate as normal processes in a software environment that has been found to be trustworthy enough for some particular purpose. Furthermore, such mechanisms for checking whether the software state of these platforms is trustworthy, and whether the systems are indeed operating in the expected manner have been briefly described.

4.3.2 Trustworthiness of client and enquiry platforms and protection of client side user privacy

The agents, combined with the TCPA technology, provide *trusted monitoring* of the way that the profiling information is produced (which can be using a wide range of techniques, including randomisation of actual measurements) – the TPM then signs this information to show that it has been gathered using trustworthy technology and then sends it out to authorised parties. Note how this is much more trustworthy for the receiver of this information than having profile information just based on something the user produces, as that cannot necessarily be trusted.

It is possible that a single device is able to perform the functions of both the user platform and an enquiry platform, perhaps acting at times as a user platform and at other times as an enquiry platform. In any case, it can be in the interest of the user to demonstrate that their platform is trustworthy to enquiry programs if the user wishes to benefit from targeted or customized services, or potentially increase the value and hence the payment for their profiling information, as already mentioned in Section 2.

However, it is a potential problem of any self-profiling solution that an obvious way for the user to protect their privacy is to lie. The more the profile generation is automated via the agents rather than directly input via the user, the less of a problem this is, so long as the agents are designed not to lie! However, in compensation the system must protect the client side user privacy (with its profile characteristics). This is achieved in the following ways:

1. Protection of the stored profile using encryption and hardware-based storage of the decryption key(s) (preferably, using the TCPA protected storage functionality). Authorization data is needed in order to gain access to data stored via the TPM, and this cannot be overridden even by the platform owner or administrator, so the profile need not be accessible to anyone without the say-so of the user (or, more practically, an agent acting on behalf of the user).
2. Access being granted to the profile only if the client platform's software environment is in the expected state (e.g. has not been hacked); this is possible because TCPA provides (protected storage) functionality for sealing data to a platform and software environment in this way.
3. Integrity checking of the enquiry platform (if this platform is a TP), preferably coupled with business and policy-level checks on the corresponding enquirer, before the client platform releases profile information to the enquiry platform. Here, it would be extremely beneficial to have an agent, or combination of agents, on the client platform that were able to analyse the integrity metrics returned by the enquiry platform, the information about the enquirer, the type of profile information to be released etc. and compare these with a policy associated with the user on the client platform in order to determine whether it would be an appropriate situation in which to release selected policy information.

4.4 Creation and Usage of Self-Profiles

In this section a method is described for allowing self-profiling by a user, as well as a method for allowing such a user self-profile to be accessed by enquirers, such that the user self-profile is trusted by the enquirers to be accurate and reliable.

In the following description it is not assumed that the user is the owner of the user platform: the method described is also applicable to situations where the owner of the user platform allows access by one or more users.

4.4.1 Creating a Profile

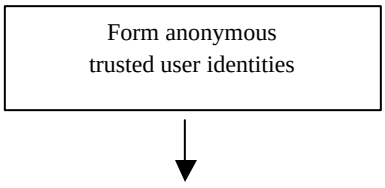

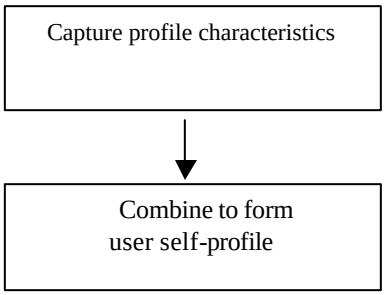

Figure 3. Method for obtaining a user self profile.

Figure 3 shows a preferred way of obtaining a user self-profile. This step is performed in response to a request from an enquirer.

The user self profile is obtained via the steps of:

1. Forming a user identity
2. Capturing at least one profile characteristic
3. Combining the user identity and the captured profile characteristic to form a user self-profile

These will be dealt with in turn.

4.4.1.1 Forming a user identity

Preferably, the first step of forming a user identity actually comprises forming a *trusted* user identity. Such a trusted user identity would be a cryptographic identity, preferably formed using an asymmetric encryption algorithm. As one example, a RSA algorithm (of the type designed by Rivest, Adi-Shamir and Adleman – see for example [16]) is used to form a private identity key and public identity key pair. The public identity key is associated with a text label, and a certificate formed signed by a trusted third party. Ideally, the trusted user identity is formed under a TCPA protocol defined by the Trusted Computing Platform Alliance, in which case the trusted third party is termed a Privacy-CA. The trusted user identity allows an enquirer to trust the accuracy and reliability of the user identity.

The user identity may relate to the user's real identity; for example, the text label contains the user's real name. On the other hand, the user identity may also be anonymous so that it does not reveal the user's real identity. An association between real and anonymous user identity is known, for example, only by a trusted third party such as a Privacy-CA. Preferably, the user identity is an anonymous trusted user identity, which allows an enquirer to trust that the user provides accurate and reliably identity information, without revealing the user's real identity.

The use of 'temporary identities' or 'digital pseudonyms' as a privacy protection mechanism is well known in the context of online electronic commerce transactions. [23] for example, teaches the use of such pseudonyms on a per-transaction as well as on a per-merchant basis. However, the use of TCPA pseudonymous identities in this model conveys the advantages of being

(statistically) unique, difficult to forge or counterfeit, and verifiable to either a local or remote entity. Furthermore:

- A TP identity guarantees that certain properties hold for the platform associated with it, and this is useful information for entities communicating with that platform, even on a one-off basis.
- A TP identity allows linkage of behavior to previous usage of a platform using that same identity. Amongst other things, this allows a business relationship to be built up over time between the TP user and external entities.
- Only the certification authority that issued identity certificates for a TP can correlate a TP identity with other TP identities.

It is possible to form a plurality of user identities, such that a different identity can be used in different contexts, or different identities used at different times in the same context. This allows the user to retain greater control over their user self-profile, by reducing the ability of enquirers to share information about the user.

4.4.1.2 Capturing profile characteristics

The profile characteristics are captured in any suitable form, and the profile characteristics themselves are widely variable depending upon the context in which the user profile is to be employed. For example, profile characteristics can be captured from user inputs, such as user responses to questions concerning the user's interests or preferences. Alternatively, profile characteristics can be captured by recording user behaviour. For example, characteristics are based upon a history of activity on a user's platform, such as by logging relevant events. Yet again, profile characteristics may be supplied from a separate computing platform and then be captured at the user platform. Here, a profile characteristic is formed such as by a commercial supplier and supplied to the user platform to form part of the user self-profile. For example, the profile characteristic is formed as a cookie.

These and other methods for capturing profile characteristics can be employed alone, or in any combination. Preferably, a plurality of profile characteristics are captured, ideally pertaining to many different aspects of the user. The set of profile characteristics preferably represent a complete profile of the user, containing all characteristics of interest to each of a relevant group of enquirers.

Optionally, any one or more of the profile characteristics is verifiable. Verification allows an enquirer to place a relatively high degree of trust in the accuracy of the profile characteristic. For example, a profile characteristic is verified by a profile certifying authority. The profile certifying authority, if satisfied with the accuracy of the profile characteristic, provides an endorsement which is associated with a profile characteristic value to form a verified profile characteristic. The endorsement is suitably generated cryptographically, such as from a private key known only to the profile certifying authority and is verifiable using a public key made widely available by the profile certifying authority. Alternatively, the TPM could directly certify (part of) a profile.

4.4.1.3 Forming a user self-profile

A user self-profile is formed by combining a selected user identity with selected profile characteristics. Preferably, the user self-profile is tailored to the needs to each enquirer, by selecting only a subset of the available profile characteristics which are of interest to the enquirer. The user does not release all of their profile characteristics to any one enquirer, and so maintains control of the complete user self-profile. By selecting amongst plural user identities, the user can maintain a high degree of privacy whilst releasing relevant profile characteristics of interest to enquirers.

In order to form a user self-profile, our model includes a capture agent for capturing one or more profile characteristics; and a profile agent for combining the user identity and at least one of the profile characteristics, as a user self-profile. Of course, a single agent could be used to carry out both of these functions.

4.4.2 Using a Profile

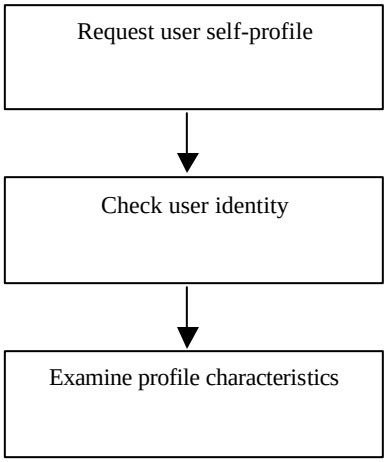

Figure 4. Enquiring about a user self-profile.

A method for enquiring about a user self-profile is also provided, which comprises the following steps:

1. receiving a user self-profile
2. checking a user identity of the user self-profile
3. examining one or more profile characteristics of the user self-profile.

These stages are carried out on the enquiry platform. Figure 4 shows a method for enquiring such a user self-profile.

The user self-profile is preferably received in response to a request sent from the enquiry platform to a user platform. Preferably, the request identifies the enquirer, as well as one or more profile characteristics of interest to the enquirer (either by explicitly naming the profile characteristics of interest, or by providing information which allows suitable profile characteristics to be determined).

The enquirer performs a cryptographic check of the user identity. Where the user identity is a trusted user identity, the enquirer checks a signature of a trusted third party. This check can simply be that the signature is present and in the expected format, or can involve more detailed investigation such as obtaining a signature checking key from the trusted third party. The enquirer may check the public identity key associated with the user identity label, such as by using this key to encrypt a message which can then only be read by a user possessing the corresponding private identity key. Hence, the enquirer may trust the identity of the user with a high degree of confidence.

The enquirer examines the profile characteristics according to the nature of those characteristics. Where the profile characteristics are verifiable, preferably the enquirer verifies those profile characteristics by checking an endorsement. The endorsement is checked using a public checking key made available by a profile certifying authority.

4.5 An Example

As an example of a practical scenario, let us consider the example context where the user's platform allows the user to purchase goods and services over the Internet from a supplier who runs one of the enquiry platforms. The supplier desires to obtain a profile of the user so that the supplier can offer the user incentives, such as discounts, tailored to the interests and preferences of the user. Hence, the user platform creates a user self-profile which can be made available to the enquiry platform of the supplier. This profile can be trusted by the enquirer to be accurate and reliable.

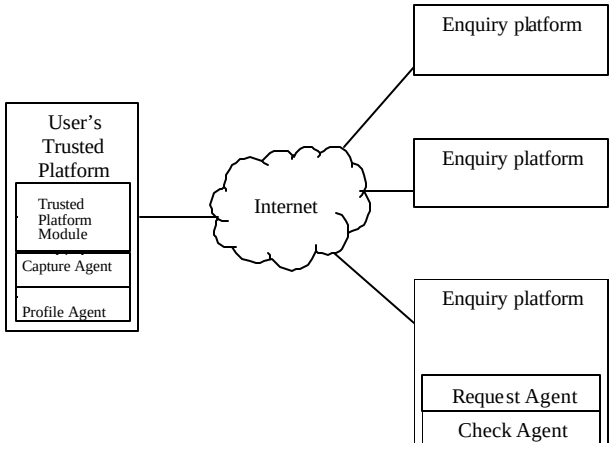

Figure 5. Main components of the system.

An example of such a computing system is shown in Figure 5. The computing system comprises a user platform coupled to an enquirer's platform over the Internet to form a networked computing system. Since the computing network is open, it is particularly advantageous that the enquirer is able to trust the accuracy and reliability of a user self-profile formed at one of the user's machines.

The user's platform is a TP in the form of a palmtop computer. The enquirer's platform is a relatively large and non-portable server operated by a commercial supplier who offers goods through an online store to customers such as the user of the user platform in Figure 5. The server preferably performs many other functions, additional to the enquiry function described here. At least in the initial stages of the transaction it is desired to allow customers to browse the store, although it is also desired to tailor the online store for a particular customer, such as by offering links to products that might be of interest, or by offering discounts or other incentives. The enquiry platform enquires for a user self-profile supplied from the user platform, and in response to the user self-profile the enquiry platform is able to establish a profile of the user. The profile can be used by the enquiry platform for purposes such as to improve the online store for this customer, whilst avoiding the need to hold large quantities of data about customers on the enquiry platform or related equipment run by the commercial supplier. For example, the user profile supplied to the enquiry platform is deleted at the end of a customer visit to the online store, because the profile will be available again from the user platform in a subsequent visit.

The process for forming one or more trusted user identities, according to the TCPA specification, has already been described in the previous section. As part of this process, the Privacy-CA checks the real identity of the user, such as by checking a passport, driving licence, or other paper or electronic identity documents. Note that only the Privacy-CA can collate the credentials, or trace them back to the user. A user may therefore choose a Privacy-CA whose polices meet the user's privacy requirements. The user can himself act as a Privacy-CA if the user has sufficient credibility.

In this case, the identity-label of the user's trusted identity is an arbitrary text character string which does not reveal the real identity of the user. Such an anonymous trusted user identity

allows the user a greater degree of privacy and increases willingness of the user to provide a detailed self-profile revealing characteristics of interest to an enquirer. However, since the enquirer, being a commercial supplier, is mainly interested in profile characteristics, the real identity of the user is not at this stage particularly important. The anonymous trusted user identity functions simply as a convenient label. The anonymous trusted user identity is of particular benefit at initial stages of a commercial transaction, such as when the user browses the supplier's online store.

The Trusted Platform Module may support several trusted user identities and pseudonymous trusted user identities. One of these identities is selected when in an appropriate context. Here, the user is able to select one of many available identities each of which can be trusted by relevant enquirers. The user can retain a high degree of anonymity, and it is difficult for different enquirers to combine information about the user. Optionally, the selection amongst available identities is automatically rotated in a predetermined pattern, or picked randomly or pseudo-randomly.

4.5.1 Capture agent

As shown in Figure 5, a *capture agent* is on the client platform for the purpose of capturing profile characteristics. The capture agent is preferably part of the Trusted Platform Module. That is, the Trusted Platform Module protects and checks the integrity of the capture agent, and the capture agent could even be stored within the TPM if there is sufficient space. Alternatively, the function of the capture agent may be performed by another part of the platform such as a central computing unit in co-operation with a storage such as a disk storage unit, but this is less desirable as the agent will necessarily be less trustworthy.

The profile characteristics can take any suitable form and can be captured in any suitable manner. The profile characteristics are preferably captured from user inputs, such as by asking the user to fill out a questionnaire on screen. The questionnaire represents, for example, the user's preferences in fields such as sports, leisure, hobbies, financial matters or otherwise. Optionally, profile characteristics are captured by recording user behaviour at the user platform, such as by logging a history of websites visited or any other relevant event. Here, it is preferred for the user to actively control when such logging activities take place. As a third option, profile characteristics are captured at the user platform by downloading from a remote source. In the example context, the supplier creates a cookie which is downloaded to the user platform and is captured as one of the profile characteristics.

4.5.2 Profile agent

As shown in Figure 5, the user platform also contains a *profile agent* for forming a user self-profile based upon a user identity as established by the Trusted Platform Module and one or more profile characteristics captured by the capture agent. As with the capture agent, the profile agent is also preferably protected, checked and/or stored within the Trusted Platform Module. The profile agent can form a user self-profile from a single identity and using all of the available profile characteristics, or else form a user self-profile according to a particular context. Each user self-profile can be stored and maintained on the user platform, or can be formed dynamically such as in response to an enquiry.

Optionally, the user self-profile is signed by the Trusted Platform Module, so that an enquirer is able to establish that the user self-profile has come from a secure source. Here, there is a chain of trust in that the enquirer trusts the trusted user identity because there is trust in the certifying authority, and trusts that the user self-profile has not been subverted because there is trust in the Trusted Platform Module.

4.5.3 Request, check and examination agents

On the enquiry platform there are the following: a request agent for requesting a user self-profile from a user platform; a check agent for checking a user identity of the user self-profile; and an examination agent for examining profile characteristics of the user self-profile. Again, these may be combined into a smaller number of agents.

The check agent checks a user identity supplied as part of the user self-profile. The examination agent then examines the profile characteristics supplied as part of the user self-profile. The profile preferences characteristics show the user's product interests, screen layout and shopping habits, either generally or specific to this supplier or a group of suppliers.

5. EXTENSIONS OF THIS MODEL

This approach is an alternative to other methods for protecting privacy while revealing data, including [7], which deals with a novel approach to surveys which uses some cryptographic techniques and [18], which describes a measure of anonymity and an associated method for protecting the identity of message senders in systems where all network communication is observable by the attacker.

This paper has focused on types of system where profile information is distributed and centralized at the client machines rather than at servers. This gives the advantages of giving the user more control, and being able to provide more information to selected targets, thus giving a better job for personalization of interests (related advantages of peer to peer architectures are considered in [1]). Potentially, the profile information could be shared across multiple devices (see [15], which describes how to migrate context-aware data).

The model described in the previous section includes agents that capture information and form user profiles at the client machines, as well as agents at the enquirers' machines for requesting, checking and examining the profile information they are sent. As an alternative, such agents could be situated at a proxy intermediate between clients and a firewall, and act on behalf of one or more users. Such a model could enhance analogous systems such as that described in [7], which applies to the specific case in which (paid) infomediaries are the custodians, agents and brokers

of customer personal information exchanged via the Internet, while at the same time protecting its privacy.

Additional agents could be added into either model, such as additional privacy or trust-checking agents. For example, agents that preserve privacy for one or more users (by means of combining the method described in this paper with other privacy-enhancing techniques, such as those described in [6]) and which understand, implement and report breaches of user privacy preferences, expressed for example via policies such as P3P [24]. Another example would be agents that check whether proposed services and remote platforms are trustworthy (via TCPA integrity checking combined with, for example, techniques for agents to establish trust amongst themselves and update this trust, as described in [10]). Such checking agents would typically instigate TCPA integrity checks, interpret the results and accordingly convey this information to another agent or human, or otherwise act upon it.

6. CONCLUSIONS

In conclusion, this paper has described how trusted agents can be used to give users control over their profile information (potentially including user preferences and context) such that this information is divulged in a trustworthy manner, under the user's control and in such a manner as to prevent profile building (e.g. retail history associated with the user) by a third party.

7. ACKNOWLEDGEMENTS

Thanks to David Plaquin for drafting Figure 1.

8. REFERENCES

[1] O. Andy. Peer-to-Peer, Harnessing the Power of Disruptive Technology, O'Reilly, 2001.

[2] B. Balacheff, D. Chan, L. Chen, S. Pearson, and G. Proudler, "How can you trust a Computing Platform?" Proceedings of Information Security Solutions Europe Conference (ISSE 2000), Barcelona, Spain, 27-29 September 2000.

[3] L. Bygrave, "Electronic Agents and Privacy: A Cyberspace Odyssey 2001", International Journal of Law and Information Technology, vol 9, no 3, p. 280, Oxford University Press, 2001.

[4] Applied Artificial Intelligence, Special Issue on "Trust in Agents", C. Castelfranchi, R. Falcone, B. Firozabadi and Y. Tan (editors), Taylor and Francis, vol 14, no 8, p.765, 2000.

[5] Castelfranchi, C. and Y.-H. Tan, eds., Trust and Deception in Virtual Societies, Kluwer Academic Publishers, 2001.

[6] H. Chi Wong and K. Sycara, "Adding Security and Trust to Multiagent Systems", Applied Artificial Intelligence, Special Issue on "Trust in Agents", C. Castelfranchi, R. Falcone, B. Firozabadi and Y. Tan (editors), Taylor and Francis, vol 14, no 9, p. 927-941, 2000.

[7] D. Grritzalis and N. Kyrloglou, "Consumer Online-Privacy and Anonymity Protection using Infomediary Schemes", Proceedings of SCCC 2001, p.115-23, IEEE Comput. Soc., 2001.

[8] B. Huberman and T. Hogg, "Protecting Privacy while Revealing Data", Nature Biotech, vol 20, p. 332, 2002.

[9] Liberty Alliance, Liberty Alliance Project, http://www.projectliberty.org

[10] Y. Mass and O. Shehory, "Distributed Trust in Open Multi-agent Systems", Trust in Cyber-societies: Integrating the Human and Artificial Perspectives, R. Falcone, M. Singh and Y. Tan (eds.), LNAI 2246, p. 159-173, Springer, 2001.

[11] Microsoft Corporation, Building User-Centric Experiences: An Introduction to Microsoft .NET My Services (formerly named "Hailstorm"), September 2001. Available via http://www.microsoft.com/net/myservices.asp

[12] N. Negroponte. Being Digital, p. 155, Hodder & Stoughton, London, 1995.

[13] J. Neilson, "Noncommand User Interfaces", Communications of the ACM, vol 36, no 4, p. 83-99, April 1993.

[14] S. Pearson (ed.), Trusted Computing Platforms: TCPA Technology in Context, Prentice Hall, 2002.

[15] S. Riche, G. Brebner and M. Gittler, "Client-Side Profile Storage", Proceedings of International Workshop on Web Engineering, Pisa, Italy, May 25th 2002.

[16] B. Schneier, Applied Cryptography. New York: John Wiley & Sons, 2nd edition, 1996.

[17] K. Scribben. "Privacy@net - An International comparative study of consumer privacy on the Internet", Consumers International, January 2001. (ISBN: 19023 913168).

[18] A. Serjantov and G. Danezis, "Towards an Information Theoretic Metric for Anonymity ", Workshop on Privacy Enhancing Technologies (http://www.pet2002.org/), San Francisco, CA, USA, 14 - 15 April, 2002. Also available via http://www.cl.cam.ac.uk/~aas23/set.ps.

[19] B. Shneiderman, "Designing Trust into Online Experiences", ACM, Special Issue on "Trusting Technology", Communications of the ACM, vol 43, no 12, p..57-59, December 2000.

[20] A. Tabachnick and L. Fidell, Using Multivariate Statistics, 4th ed, pub. Allyn & Bacon, 2000. (ISBN: 0321056779).

[21] Trusted Computing Platform Alliance, Building A Foundation of Trust in the PC, White Paper, January 2000. Available via www.trustedcomputing.org.

[22] Trusted Computing Platform Alliance, TCPA Main Specification, Version 1.1, 2001. Available via www.trustedcomputing.org.

[23] J. D. Tygar “Atomicity in electronic commerce”, Proceedings of the Fifteenth Annual ACM Symposium on Principles of Distributed Computing, Philadelphia, PA, USA, 23-26 May 1996, p. 8-26, New York, 1996. (ISBN: 0897918002).

[24] World Wide Web Consortium, Platform for Privacy Preferences Specification 1.0. http://www.w3.org/TR/P3P/