Ensuring Privacy for E-Health Services

Yee, George; Korba, Larry; Song, Ronggong

Proceedings of The First International Conference on Availability, Reliability and Security (ARES 2006), 2006

National Research Council Canada
Conseil national de recherches Canada

Institute for Information Technology
Institut de technologie de l'information

NRC · CNRC

*Ensuring Privacy for E-Health Services **

Yee, G., Korba, L., Song, R.
April 2006

* published in the Proceedings of The First International Conference on Availability, Reliability and Security (ARES 2006). April 20-22, 2006. Vienna, Austria. NRC 48462.

Canada

Ensuring Privacy for E-Health Services[1]

George Yee, Larry Korba, and Ronggong Song
Institute for Information Technology
National Research Council Canada
{George.Yee, Larry.Korba, Ronggong.Song}@nrc-cnrc.gc.ca

Abstract

The growth of the Internet has been accompanied by the growth of e-health services (e.g. online medical advice, online pharmacies). This proliferation of services and the increasing regulatory and legal requirements for personal privacy have fueled the need to protect the personal privacy of service users. Existing approaches for privacy protection such as access control are predicated on the e-service provider having possession and control over the user's personal data. In this paper, we propose a new approach to protecting personal privacy for e-health services: keeping possession and control over the user's personally identifiable information in the hands of the user as much as possible. Our approach can also be characterized as distributing personally identifiable information only on a "need to know" basis.

1. Introduction

In order for e-health services to be successful, personal privacy must be protected. As defined by Goldberg et al in 1997 [1], privacy refers to the ability of individuals to *control* the collection, retention, and distribution of information about themselves.

In this work, an e-health service is performed by application software (service processes) that is owned by a provider (usually a company); the service is accessible across the Internet. Further, the provider has a privacy policy that spells out what consumer personal information is needed to perform the service and how the personal information will be handled. The consumer has a personal privacy policy that defines what personal information she (we use "she" and "her" to stand for both sexes) is willing to disclose and how that information is to be handled by the provider. Examples of current e-health services are myDNA.com (health information), WebMD.com (health information and technology solutions provider), e-med.co.uk (online medical consulting) and Walgreens.com (online pharmacy).

Our approach uses selective disclosure of the user's information and a smart card, in conjunction with the user's personal privacy policy, to keep *control* of the user's personally identifiable data in the hands of the user as much as possible, rather than in the hands of the service provider. This approach is motivated by the fact that once consumer personal information is in the hands of a service provider, it can be very difficult to detect that the provider has violated the consumer's privacy preferences. It is a hard problem to guarantee that a provider will not circumvent any kind of privacy protection access control that might be in place.

In the literature, some components of our proposal exist, but not, as far as we can tell, assembled into the approach presented in this work. For example, Clarke [9] wrote about smart cards (he actually was complaining that their use destroys privacy), anonymity, and the use of pseudonyms and trusted third parties. Laudon [10] suggested that individuals could sell their private information in an information market, and thus maintain control over their private information (the maintaining control part is similar to what we propose here but the means for doing so is completely different). However, Laudon's proposal is flawed in that it does not discuss the potential abuse of private information in a market setting (e.g. theft). Smart cards have been around for over 3 decades and have been applied across many domains including e-commerce [2, 3]. Their computational, memory, and security features make them ideal for portable data applications requiring security [3]. As another example, Lategan & Olivier [11] present an approach called PrivGuard for protecting private information by classifying it based on the purpose for collecting the information, i.e. how the information will be used, and then designing methods for protecting each information class. If the information is only required for validation, it is encrypted so that the validation can be performed without decryption. If the information is used for purposes other than validation, a system based on Kerberos (using tickets) in conjunction with trusted third parties is used to give the consumer (the

information owner) as much control over the information as possible. Our approach is similar in that we also strive to give the owner of the information as much control over the information as possible, and our approach also makes use of trusted third parties to actually access the private data. We also look at how the private information will be used but only to separate out personally identifiable information for different treatment. However, our approach differs from that of Lategan & Olivier in at least the following ways: a) our approach is more lightweight with lower overhead (i.e. no encryption of private information, no ticket system for access), b) our approach is meant to apply only to specific types of services whereas PrivGuard applies generally, and c) we use trusted third parties to actually perform part of the service (the part requiring personally identifiable information) whereas PrivGuard only uses them to grant tickets and to access the private information.

The rest of this paper is organized as follows. Section 2 gives examples of privacy policies. Section 3 presents our approach for using selective disclosure and smart cards to protect privacy for e-health services. Section 4 gives an example of applying our approach. Section 5 presents our conclusions and plans for future research.

2. Privacy policies

Since our approach involves privacy policies, it is useful to have examples of privacy policies to work with. Figure 1 (adapted from [5]) gives an example of a consumer privacy policy along with the corresponding provider privacy policy for an online pharmacy. The provider's privacy policy specifies the provider's requirements for personal information. *Policy Use* in a policy indicates the type of online service for which the policy will be used. Since a privacy policy may change over time, we have a *valid* field to hold the time period during which the policy is valid. The required fields (e.g. collector, what) of these policies are derived from Canadian privacy legislation in the form of privacy principles [5]. These principles are representative of privacy legislation in many countries, including the European Union. These are minimum privacy policies in the sense that the fields *collector*, *what*, *purposes*, *retention time*, and *disclose-to* form the minimum set of fields required to satisfy the legislation for any one information item. Each set of such fields is termed a *privacy rule* describing a particular information item.

Note that all information that the consumer discloses to a provider is considered personal information and described in the consumer's personal privacy policy. Some of this information is personally identifiable information (PII), i.e. the information can identify the consumer. For example, "name", "address", and "telephone number" are PII. There may be other information described in a personal privacy policy that is not personally identifiable information (non-PII), i.e. the information by itself cannot identify the consumer. For example, the selection of Aspirin as a medication at an online pharmacy cannot normally identify the consumer.

Privacy policies need to be machine-readable and may be expressed using a XML-based language such as APPEL [4], which can be used to express privacy preferences for both consumer and provider.

Policy Use: *Pharmacy* Owner: *Alice Consumer* Valid: *unlimited*	Privacy Use: *Pharmacy* Owner: *A-Z Drugs Inc.* Valid: *unlimited*
Collector: A-Z Drugs Inc. *What:* name, address, tel *Purposes:* identification *Retention Time:* unlimited *Disclose-To*: none *Collector:* A-Z Drugs Inc. *What:* drug name *Purposes:* purchase *Retention Time:* 2 years *Disclose-To*: none	*Collector:* Drugs Dept. *What:* name, address, tel *Purposes:* identification *Retention Time:* 1 year *Disclose-To*: none *Collector:* Drugs Dept. *What:* drug name *Purposes:* sale *Retention Time:* 1 year *Disclose-To*: none

Figure 1. Example consumer (left) and provider (right) privacy policies

3. Using selective disclosure and smart cards to protect privacy

Our goal is to protect an e-health service user's privacy according to her personal privacy policy. Our answer to privacy protection is simple: *remove the user's PII from the possession and control of the provider*. We accomplish this by having the user's personal information in a smart card, called a *privacy controller*, owned by the user and in her possession. The personal information in the privacy controller can only be controlled by the user. We propose that an e-health service can be partitioned into a *primary service* (the service that the user has engaged) that does not require PII and *support services* that do require PII but are trusted to maintain the anonymity of the user. The user's privacy controller selectively discloses PII only to the support services that require the user's PII (i.e. on a need-to-know basis). This is done automatically without the user's intervention. Paypal.com [8] is an example of such a support service provider for payment collection. Further, the privacy controller

smart card will process the user's personal information (both PII and non-PII) according to her privacy policy. The user is anonymous to the primary service provider at all times. We further require the services of a trusted authority (e.g. a Certificate Authority with an extended role) to program the smart card to act as a privacy controller, to keep the true identity of the user should there be a need to recover it (e.g. in legal proceedings), and to distribute the smart card. Table 1 summarizes our terminology. Figure 2 illustrates our approach.

Table 1. Summary of terminology

Term	Definition
Primary service	the service that the consumer wants to engage, e.g. online purchase of medication; does not require PII
Primary service provider (PSP), or provider	the provider of the primary service
Primary service processes, or service processes	the software processes of the primary service
Support service	a service that supports the primary service, such as a payment service or a delivery service; a support service requires PII (e.g. credit card number)
Support service provider (SSP)	the provider of a support service, e.g. Paypal.com; also known as a trusted third party
Trusted Authority	This is a role with the following responsibilities: a) program the smart card to act as a privacy controller, b) distribute the smart card, c) guard the true identity of the user, and d) approve and certify support providers. This role may be suitable for an existing Certificate Authority (CA) (as used in Public Key Infrastructure) that we then call an eCA (for extended CA), or a government office such as the office of a privacy commissioner; also known as a trusted third party.

3.1. Privacy controller and service process requirements

As the controller processes the user's privacy policy, for each privacy rule component, the privacy controller must:

a) *Collector*: Confirm that the collector named by the service processes is the collector specified in the user's policy.

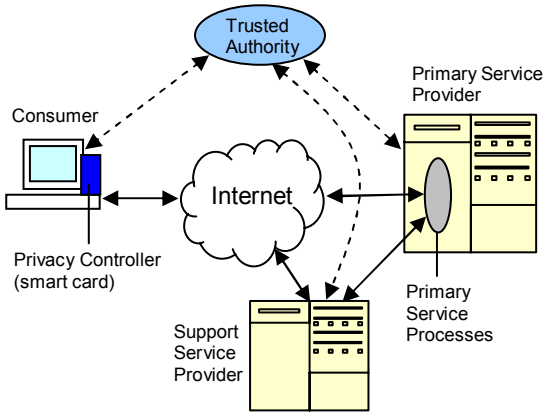

Figure 2. Using selective disclosure and smart card to protect user privacy

b) *What*: Confirm that the information item requested by the service processes is as specified in the user's policy.

c) *Purposes*: Confirm that the purposes for which the information will be used are as specified in the user's policy.

d) *Retention Time*: Destroy the user's personal information at the end of its retention time. This is important for avoiding eventual storage overflow.

e) *Disclose-To*: Confirm that the receiving party in the case of a disclosure request is the party specified in the user's privacy policy.

The service processes must cooperate with the privacy controller where necessary in order to carry out the above requirements (e.g. provide the provider's privacy policy to the privacy controller).

These requirements dictate the functionality of the privacy controller and the service processes. The privacy controller, in acting to ensure compliance with the user's privacy policy, runs in two phases as described below. Phase 2 can only be reached if phase 1 is successful (if phase 1 is unsuccessful, the consumer and provider can enter into negotiation [6, 7] failing which the consumer can try another provider).

Privacy controller processing for user privacy policy compliance

In phase 1,

- Establish a connection with the PSP and download its privacy policy and SSP information.
- Check the user's privacy policy to make sure that the collectors of PII are the SSPs indicated by the PSP. Make corrections if there are mistakes.
- Verify that the privacy rules in the provider's privacy policy matches (comparing privacy

policies for a match is outside the scope of this paper but see [5]) the privacy rules in the user's privacy policy. If this verification fails, inform the user and terminate (or negotiate privacy policies as indicated above). Otherwise, proceed to phase 2.

In phase 2,

- Prompt user for each information item (II) and accept only II of the types specified in the user's privacy policy.
- Store user's II in its personal information store.
- Destroy the user's II if the retention time is up (it has full control over the II in its store).
- Disclose only non-PII to the PSP as described above.
- Accept requests from the PSP to disclose the user's II (PII and non-PII) to SSPs as allowed by the user's privacy policy, passing along the II's retention time. SSPs are not allowed to further disclose the user's PII. Note: the typical public user would normally not be receiving disclosures; in this work, only PSPs or SSPs receive disclosures, e.g. a trusted shipping company receiving an address disclosure for shipping purposes.

Primary service processing

The service processes execute during the controller's phase 2 processing, as follows:

- Perform primary service processing. This includes requesting needed non-PII from the privacy controller.
- If needed, request the controller to disclose information to SSPs.

Figure 3 illustrates the phases of controller and primary service processing using state machines.

3.2. Role of the trusted authority and additional operational details

We assume that the trusted authority is an eCA (could just as well be a privacy commissioner). Prior to the commencement of any e-service, the eCA works to familiarize providers and users with its services. Providers that subscribe to the eCA must arrange their service processes to work with the privacy controller smart card (e.g. conform to smart card interfacing requirements). The smart card is remotely programmed by the eCA to be used as the privacy controller and to work with the e-health primary service providers that have subscribed to the eCA (e.g. download provider's privacy policy, upload user's information). This programming automatically allows the smart card to be

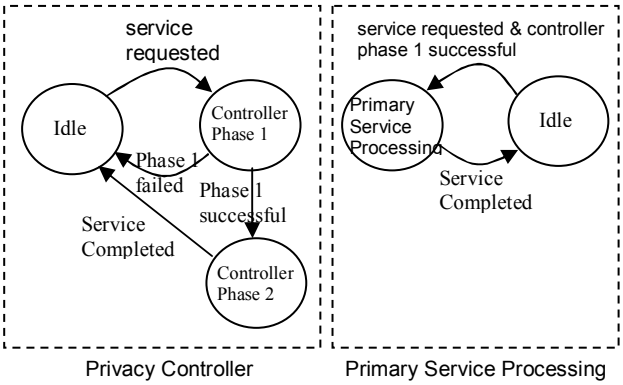

Figure 3. Behaviour of privacy controller and primary service

used with new providers that may subscribe to the eCA later. The eCA distributes these smart cards to service users through local electronics outlets (e.g. Best Buy). When purchased at a local electronics outlet, the smart card only has the ability to automatically connect to the eCA (in addition to normal smart card functions). The eCA also selects and confirms a number of SSPs as trusted parties for business services such as shipping and payment. The eCA can do this by issuing a call for tenders and then doing a thorough background check on the applicants. Further, the eCA issues digital certificates to all primary and support service providers for use in authenticating themselves.

A service user who wants to use PSPs that subscribe to the eCA registers with the authority's web site through a secure channel. After appropriate user credit checks and payment arrangements have been made, the eCA assigns a number of pseudonyms (a different pseudonym for each group of PSP and associated SSPs) to the user and issues the user a corresponding number of digital certificates (for authentication purposes), using a different pseudonym to identify the user in each certificate. The use of a different pseudonym with each PSP-SSP group discourages the PSPs from colluding together to discover information about the user by linking her pseudonym.

To use an e-health service, the user connects the smart card to a USB port on her computer. The user is automatically connected to the eCA's website after mutual authentication (using digital certificates) through a secure channel. The eCA then remotely programs the smart card for use as a privacy controller, instructing the controller to use the user's pseudonyms for identification purposes with all primary and support service providers (a different pseudonym is used for each new group of PSP and associated SSPs) (note: this is done only if the smart card has not been programmed previously). The user is then allowed to

select which PSP to use. After the user selects the PSP, the website prompts the user to enter a privacy policy to be used with the e-health service if this is the user's first use of the service. The website also displays the SSPs that are used by the selected PSP, along with their require PII, so that the user can enter needed data disclosure rules in her privacy policy. The entered policy or a previously entered policy (they are stored on the eCA's website) is then automatically downloaded to the smart card. At this point, the controller automatically begins phase 1 processing. A pop-up window appears indicating an anonymous connection to the service with successful 2-way authentication through a secure channel and with the provider's privacy policy downloaded (controller phase 1 processing). The privacy controller then checks the user's privacy policy to ensure that PII, if any, are going to the correct SSPs, and compares the user's e-service privacy policy (previously entered) with the provider's e-service privacy policy for compatibility. Compatibility means that the provider's privacy rules do not violate the user's privacy rules (see [5] for more explanation of policy matching). For example, for corresponding privacy rules between the two policies, all fields should match exactly except possibly for retention time. For retention time, as long as the provider's retention time is less than or equal to the consumer's retention time, the two retention times will be compatible. If this comparison of policies for compatibility is successful, the privacy controller initiates phase 2 processing. Otherwise, the privacy controller initiates a privacy policy negotiation session with the provider that takes place via the privacy controller. If this negotiation is successful, the privacy controller can begin phase 2. If neither the original phase 1 nor the negotiation is successful, the user must try a different provider. Once the controller starts phase 2, the provider's service processes are initiated. The latter then requests non-PII from the controller and requests it to send information disclosures (possibly sending PII to SSPs (e.g. address for shipping)) as the service requires. Service output is sent back to the user via the controller-service processes channel.

It follows from the above that the eCA can link the user's pseudonym with the user's real name. This is allowed on purpose, so that when necessary the PSP can request the true identity of the user. For example, this may be necessary in a medical emergency where an e-pharmacy primary service provider needs to contact the user.

3.3. Applicability to e-health services

For our approach to apply to a particular e-health service, it must be possible to partition the service into a primary service and support services. In fact, our approach applies to the following types of e-health services: information services, certain consulting services that do not require PII other than for payment or for payment plus physical delivery, and online pharmacies.

An information service is essentially a medical database that can be queried for information with no PII required, PII required for payment only, or PII required for payment and physical delivery.

A consulting service takes a medical condition as input and returns advice on how to treat the condition. For such a service, if the condition does not identify the user, PII would only be required for payment or for payment and physical delivery. How can one know whether or not the medical condition can identify the user? The condition can be tested by querying a database of medical conditions to obtain estimated numbers of people with the condition. If the number returned is greater than a threshold of 100, for example, the condition may be accepted as being non-identifiable with the user. Higher degrees of not being identifiable can be obtained by using higher thresholds.

An online pharmacy requires PII only for payment and physical delivery.

In conclusion, our approach can be applied to a large number of today's e-health services that are already partitioned or can be partitioned into primary services not requiring PII and support services (e.g. for payment and physical delivery) that do require PII.

3.4. Usability discussion

An important question that comes up with the introduction of any new technology or approach is "How usable is it?" or "Is it likely to be used by people?". We believe our approach is very usable because it is similar to some existing commercial processes and is straight-forward to use. First of all, the process of registering with the eCA is similar to the current way of registering with websites for a service or membership. Second, there is no delay in getting the smart card – the user only needs to pick one up at a local electronics store. Further, the user only has to get the smart card once and can use it with all existing and new providers that subscribe to the eCA. Third, the user only has to plug the smart card in a USB port on her computer to begin the process of connecting to a service.

It might also be argued that people will tend to stay away because the use of a smart card is unfamiliar. However, smart card use has been growing at a high rate, in part because the way they are used is similar to how millions of people use magstripe cards and PINs to access their bank accounts.

3.5. Security measures

Based on the above operating scenarios, the vulnerability areas include: a) storage of personal data, b) distribution of the smart card through local electronic outlets, c) sending data disclosures, d) communication between the privacy controller and the service processes, and between the user and the eCA's web site, e) disclosure of non-PII to the service processes, i.e. although the data is non-PII, could their combinations collected over time compromise the anonymity of the user? f) traceable communications over the Internet, g) dishonest parties masquerading as trusted parties, h) Trojan horse and virus programs in the user's computer, i) Trojan horse and virus programs in the smartcard, and j) the user loses her smart card, either by accident or theft.

We discuss our security measures for each vulnerability area in turn as follows:

a) Storage of personal data: the data is secured on the smart card (processor-enabled) using symmetric encryption (e.g. 3DES). The key for the encryption algorithm can be generated (e.g. using a SHA-2 hash function) by the smart card from the user's password for accessing the card. Further, the smart card incorporates a locking mechanism that locks out any attacker who tries to access the card by trying to guess the password – the locking mechanism can lock the user out, for example, after 5 tries (if a legitimate user is accidentally locked out, the password may be remotely reset by the eCA through a secure connection to the smart card). Thus, the attacker first of all cannot access the card because she does not know the password. Even if the attacker uses some special technology to get at the data, she cannot read it since it is encrypted. Finally, the attacker cannot decrypt the data because she again does not know the password, used to generate the encryption key.

b) Distribution of the smart card through local electronic outlets: the risk is that an attacker could modify the card before it is sold, to i) connect to a fake website controlled by the attacker, or ii) introduce malware into the card that would later play havoc with any programming; possibility i) is defeated by required mutual authentication between the user and the eCA; possibility ii) can be defeated using built-in card self sanity checks together with malware detection software run on the card by the eCA prior to remote programming.

c) Sending / receiving data disclosures: the privacy controller establishes a secure channel (SSL or secure VPN) to the receiving party for use in data conveyance; the sending controller authenticates the receiving party using the receiving party's digital certificate before any data is sent. Receiving parties are pre-screened by the eCA, who issues them digital certificates for authentication purposes.

d) Communication between privacy controller and service processes: the controller establishes a secure channel (SSL or secure VPN) to the service processes to be used for communication purposes. The controller authenticates the service processes using their digital certificates issued to them by the eCA. Similarly, the service processes authenticate the controller using the digital certificate issued to the controller by the eCA. This same secure procedure is used for communication between the user and the eCA's website.

e) Disclosure of non-PII leads to compromising anonymity: we believe that this risk is minimal for most types of Internet e-health services; identity discovery depends on the size of the candidate population, the method of selective disclosure, and the amount of personal data in circulation pertaining to the individual. This risk can be minimized if the candidate population is the whole Internet community. However, some services operate only regionally so this may not apply. The risk may be further minimized by researching and employing more effective methods for selective disclosure.

f) Traceable communications over the Internet: the controller not only establishes a secure channel for communication with the service processes but establishes it using a mix network (e.g. JAP at http://anon.inf.tu-dresden.de/desc/desc_anon_en.html). By so doing, the provider would find it very difficult to trace the identity of the user by way of the user's Internet connection.

g) Dishonest parties masquerading as trusted parties: first, the reputation of the eCA is established (as for a regular CA); for example, the eCA could be subjected to inspection audits and other forms of testing to ensure that processes and responsibilities carried out are trustworthy. After the eCA is established to be trustworthy, it can make sure that all trusted support service providers are indeed trustworthy, using a similar series of inspections and testing as was done for it. The above inspections and testing would be done continuously at various intervals to ensure that these parties remain trustworthy.

h) Trojan horse and virus programs running in the user's computer could modify the user's keyboard entries for a privacy policy or PII and/or redirect them to the attacker. The only defense we can suggest here is for the user to regularly run

malware detection software that identifies and deletes the offending programs.

i) Trojan horse and virus programs running in the user's smart card could modify the user's data communications and/or redirect them to the attacker. A Trojan horse program could also steal the user's smart card password. Again, we suggest regularly running malware detection software that identifies and deletes the offending programs on the smart card. Keeping the smart card stored data and stored password encrypted should help.

j) If the user loses her smart card either by accident or theft, the person who finds the smart card or the person who stole it could masquerade as the original owner and incur services at that owner's expense or could somehow gain access to the original owner's PII. To reduce the risk of this happening, as mentioned in a), the smart card requires a password for access and has a locking mechanism that locks out the attacker after a fixed number of attempts (e.g. 5) to try and guess the password. If the legitimate user were to forget this password, the eCA could reset it through a secure connection between the eCA's website and the smart card.

We acknowledge that the above measures do not ensure complete security or complete privacy. Such goals are impossible to attain in practice for any sizable useful system. Our aim is only to obtain a reasonable reduction in the risk of an attack succeeding.

4. Application example

Consider an online pharmacy, E-Drugs, Inc. (fictitious name), where prescriptions can be filled online and delivered to the patient the next day. Suppose that E-Drugs, Inc. has subscribed to use the privacy protection services of Privacy Watch, Inc. (fictitious name), the eCA that has implemented our approach. Suppose also that Privacy Watch (PW) has made sure that E-Drugs's service processes can interface to the privacy controller and has added E-Drugs to its web site as a subscriber. In addition, suppose that PW has provided E-Drugs with a digital certificate to be used for authentication purposes. The application of our approach would involve the following steps:

1. Alice, wishing to anonymously fill an electronic prescription, discovers by browsing PW's website that E-Drugs is available as a PW-subscribed seller.
2. (Omit this step if Alice has purchased from a PW seller before.) Alice registers with PW and is assigned a number of pseudonyms (e.g. "Patient21") to be used as identification with different groups of PSPs and associated SSPs, i.e. a PSP may only know Alice as Patient21. She also receives a number of digital certificates from PW, with each certificate incorporating one of the pseudonyms, to be used for authentication purposes. Alice purchases a PW-issued smartcard from a local electronics outlet.
3. Alice connects her smart card to the USB port on her computer. After successful mutual authentication, she is connected to PW's web site through a secure channel.
4. (Omit this step if Alice has purchased from a PW seller before.) PW remotely programs Alice's smart card to be used as her privacy controller.
5. PW requests Alice to select a seller. After she selects E-Drugs, and enters her personal privacy policy on PW's web site (only if not previously entered for this seller), the privacy controller downloads Alice's privacy policy to the smart card. The controller is then connected to the service processes at E-Drugs automatically and anonymously through a secure channel and mix network. After successful mutual authentication, the controller downloads E-Drugs' privacy policy. After making sure that Alice's policy specifies the correct SSP collectors for PII and verifying that her privacy policy is compatible with E-Drugs' privacy policy, the privacy controller requests Alice's electronic prescription, shipping address, and credit card number.
6. Alice enters the requested information (disk location for the prescription) on her computer with the privacy controller making sure that the information corresponds with her privacy policy. The information is securely stored in the privacy controller. Upon request from E-Drugs' service processes, and after checking again with Alice's privacy policy, the controller discloses to the service processes details about the prescription (including the digital signature of the prescribing physician) but withholds Alice's name, address, and credit card number. Upon request from E-Drug's service processes, the controller sets up a secure channel to a trusted payment center (SSP) and authenticates the payment center before disclosing to the center Alice's credit card number. The trusted payment center maintains the patient's anonymity to the outside world by keeping the pseudonym-patient link secret (as do all trusted SSPs). The trusted payment center was designated as trusted by PW beforehand and issued a digital certificate for authentication purposes. Similarly, the controller discloses Alice's name and address

to a trusted shipping center (SSP) that also keeps the pseudonym-patient link secret. Both the trusted payment center and the trusted shipping center use the pseudonym-patient link to link the order to the patient. If the patient tried a re-use attack to fill the prescription more than once, this would be detected by both these support service providers through the pseudonym-patient link.

7. Alice receives her order the next day from the trusted shipping center.

5. Conclusions and future research

We have presented a novel approach to protect the privacy of e-health service users based on keeping control of the PII in the hands of the user, trusted SSPs, and an eCA acting as a trusted authority. Our approach involves the use of a smart card acting as a privacy controller in the possession of the user. We chose to use a smart card for its portability, secure storage capability, and the fact that it may be connected to the Internet only for the duration of a service, making it more difficult for attackers to penetrate than if it were connected over long periods of time, as in the case of a regular personal computer. We described a number of security measures to secure our approach.

Some advantages of our approach is that it is very usable, straightforward, employs existing technology, and would be fairly easy to set up. Another advantage is that the privacy controller automatically discloses private information according to the user's privacy policy. The role of the eCA represents a very profitable opportunity for an existing CA. Further, the use of an eCA represents not more of a weakness to the security of the system than the use of a CA for PKI. A possible disadvantage is that the use of selective disclosure may be susceptible to some pattern being discovered among combinations of non-PII that would reveal the user's identity. However, this risk is reduced if the community of service users is Internet-wide.

It can be argued that although our SSPs are "trusted", there is nothing really to stop them from violating the trust placed on them. While this is true, such violations are not easy because they are subjected to monitoring, inspections, and penalties if they are caught. Trusted parties such as Paypal.com and Certificate Authorities exist in reality and do play their trusted roles well. In addition, it is impossible to have perfect security or perfect privacy for any sizable practical system. The best we can do is to minimize the risk of failure to maintain security or privacy. Judged in this light, we believe that this work does achieve a reasonable minimization of such risk.

As future research, we would like to develop improved algorithms for selective disclosure to reduce the risk of patterns in disclosed non-PII that can identify the user. We would also like to build a prototype to show feasibility and to experiment with usability, performance, and scalability.

References

[1] I. Goldberg, D. Wagner, and E. Brewer, "Privacy-Enhancing Technologies for the Internet", *IEEE COMPCON'97,* pp. 103-109, 1997.

[2] K.M. Shelfer, J.D. Procaccino, "Smart Card Evolution", Communications of the ACM, Vol. 45, No. 7, p. 84, July 2002.

[3] M.R. Carr, "Smart card technology with case studies", Proceedings, 36th Annual 2002 International Carnahan Conference on Security Technology, pp. 158-159, Oct. 20-24, 2002.

[4] W3C, "A P3P Preference Exchange Language 1.0 (APPEL 1.0)", retrieved April 22, 2004 at: http://www.w3.org/TR/P3P-preferences/

[5] G. Yee, L. Korba, "Comparing and Matching Privacy Policies Using Community Consensus", Proceedings, 16th IRMA International Conference, San Diego, California, May 15-18, 2005.

[6] G. Yee, L. Korba, "Bilateral E-services Negotiation Under Uncertainty", Proceedings, The 2003 International Symposium on Applications and the Internet (SAINT2003), Orlando, Florida, Jan. 27-31, 2003.

[7] G. Yee, L. Korba, "The Negotiation of Privacy Policies in Distance Education", Proceedings, 14th IRMA International Conference, Philadelphia, Pennsylvania, May 18-21, 2003.

[8] Paypal.com, accessed June 20, 2005 at: https://www.paypal.com/

[9] R. Clarke, "Identification, Anonymity and Pseudonymity in Consumer Transactions: A Vital Systems Design and Public Policy Issue", available as of October 3, 2005 at: http://www.anu.edu.au/people/Roger.Clarke/DV/AnonPsPol.html

[10] K.C. Laudon, "Markets and Privacy", Communications of the ACM, Vol. 39, No. 9, September 1996.

[11] F. Lategan, M. Olivier, "PrivGuard: A Model to Protect Private Information Based On Its Usage", available as of Dec. 14, 2005 from: http://mo.co.za/open/privgrd.pdf

[1] NRC Paper Number: NRC 48462