On granting limited access to private information

Frans A. Lategan[*] Martin S. Olivier[†]

ABSTRACT

We distribute our private information on an ever increasing number of computers daily, and we effectively give target organisations carte blanche to do what they want with our private information once they have collected it. We have only their privacy policy as a possible safeguard against misuse of our private information. In this paper we describe a classification of private information based on the purpose it is acquired for. We also propose a method by which we can grant limited access to our private information, and thus enforce the terms of their privacy policies. Private information is also revealed at the last possible stage, further reducing the possibility of misuse. This safeguards private information in four of the five categories mentioned.

Keywords: Privacy, access control
Length: 5000 words

1. INTRODUCTION

In most cases where a subject is granted access to data, such access is limited. Typical limits are usually an expiry date on such access, restrictions on what can be done with the data or restrictions on where the data can be accessed from. No organisation would give a subject carte blanche access to their data on a vague promise to take good care of it. Unfortunately this is exactly what happens when individuals supply their private information to some target organisation on the internet — that target organisation might have a privacy statement, but effectively has total control of the private information, and might resell, redistribute or modify the data without the owner's knowledge or consent. Furthermore the owner can not "take back" his or her private information, since there is no secure way to force the target organisation to delete the data.

We argue that private information can be categorised according to the purpose the information is required for. In each of such categories, a different approach can be used to safeguard the private information, while still allowing it to be used for the intended purpose.

In this paper we present a protocol to allow an individual to supply his private information to a target organisation in a way that limits the access such an organisation has to the information. To do this, we require several well known tools such as tickets, public key encryption and a trusted third party.

This paper is structured as follows: in Section 2 we shall supply some background information. Our classification of private information is outlined in Section 3. We then present a quick overview of our protocol in Section 4, followed by the actual implementation in Section 5 and a summary in Section 8.

2. BACKGROUND

For effective electronic commerce every individual i is forced to reveal some private information m about herself at some stage to a target organisation o. Even if i trusts o to perform the electronic transaction, it does not imply that i would like o to keep a permanent record of m in a database. Unfortunately i has no control over m as soon as m is disclosed to o, and stored in o's database. These databases are increasingly compromised, misused, sold, or even made freely available to the public over the Internet. See [1, 2, 3, 4] for a sample of some of the concerns over the safeguarding of private information.

One of the ways that private information can be safeguarded, is by using privacy policies (see [3]). A lot of research has gone into automating such policies, as described in the P3P protocol [5], but these policies are still not enforceable. However, P3P allows an individual i to define a set of acceptable usage rules of private data m. Any organisation o with a published privacy policy P that does not violate these rules can request i's browser to automatically supply those parts of m that are required. This saves i from having to read and interpret every site's privacy policy before sending private information to such a site, and from having to manually enter the private information every time. Although P3P can define arbitration authorities for disputes, adherence to published policies is not enforced; safe transmission of data between parties is also not described.

[*] Department of Computer Science, Rand Afrikaans University, PO Box 524, Auckland Park, 2006, South Africa, e-mail: fransl@discoveryhealth.co.za

[†] Department of Computer Science, Rand Afrikaans University, PO Box 524, Auckland Park, 2006, South Africa, e-mail: molivier@rkw.rau.ac.za

WWW10, May 1-5, 2001, Hong Kong.
ACM 1-58113-348-0/01/0005.

Kerberos is an authentication service that issues tickets granting an individual i access to resources on a network, without i having to log on each time such access is required — a valid ticket is all that is required. We use similar concepts in our proposed protocol. More information on Kerberos can be found in [6, 7, 8].

Secure Electronic Transaction (SET) is a protocol developed jointly by Visa and Mastercard to allow the secure exchange of credit card information between a buyer and an online merchant. It does not allow the merchant direct access to the buyer's payment details, thereby protecting the privacy of the payment details. Unfortunately SET does not protect the privacy of any other part of the buyer's personal information. Our proposed protocol remedies this. More information on SET can be found in [9].

Digital signatures and private key encryption are also used in this protocol.

3. CLASSIFICATION

Private information can be categorised according to the purpose the information is required for. Various methods to protect the information can then be tailored to effectively address both the privacy and use of the information. In this paper we do not specifically address the protection of information in the first or sixth category of our classification, and information in the fifth category is not protected from the party requiring actual access to it. The protection of private information in the other categories of the classification is discussed in this paper.

The first category that we can divide private information into, is when the private information is used to verify the result of a calculation, such as detailed earnings to verify taxable income, and the exact content of a share portfolio to verify the valuation. In these cases, access to the private information is not required if the calculated result could be verified by the target organisation without it. Such cases can be protected as discussed in [12] and are outside the scope of this paper. In related cases where an aggregate or average of numerical values for several users are required, our protocol could be used — the trusted third party could be requested to calculate such values if allowed by the privacy policy.

The second class in which we can categorise private information is where private information is required by an organisation in order to pass it on to a third party for a purpose directly linked to the transaction being performed. The organisation does not actually require the private information; the third party to which the information is passed will require it. Examples include credit card numbers (to pass on to a bank for a transfer of funds) and a shipping address (to pass on to a shipping company for actual delivery of a package).

In yet another class of private information, an organisation sometimes does not require private information for the current transaction, but might need it for future use. In such cases it might not be feasible to request the private information at the time required; the availability of such information in future has to be guaranteed in order to complete the current transaction. Examples include a credit card number to confirm a reservation when no debit will be performed until checkout (the account might even be paid in cash), and an e-mail address to notify a purchaser of similar items in stock, or of possible improvements or recalls.

In our fourth classification we consider the case where private information is used by organisations to uniquely identify customers or to allow customers to login to an account with such an organisation. Once again the actual private information does not concern the target organisation — it is only collected to be used in a challenge-response system. This can be a social security number, mother's maiden name, birth date, etc.

In our fifth class of private information usage we examine cases where the actual private information is immediately required by the requesting organisation in order to complete a transaction. To actually deliver a package, or actually transfer money organisations sometimes do require the actual contents of the private information. This is the case where it is the hardest to protect such information, as very few workarounds will do.

The final class will consist of all cases where private information is required for a purpose not covered in any of the previous categories. Such cases might include private data collected by an organisation for future marketing purposes or to resell it later to a third party. As such information is obviously not directly linked to the completion of the current transaction, it could be argued that this is exactly the type of misuse an individual would want to prevent, and as such, would not willingly supply it anyway.

We can ask whether the six classes described above are exhaustive. Since the sixth class contains all cases where private information usage is not covered by the first five categories, it follows intuitively that such a classification is indeed complete.

4. OVERVIEW

To effectively present this method, we shall start by giving an example step by step application of the proposed method, and then we shall specify it in more detail.

4.1 Example

Let us use the case where Alice buys books from Bob, to be shipped using FastShipping. Alice's private details m consist of her name m_{name}, her payment details $m_{payment}$, her e-mail address m_{e-mail} and her shipping address $m_{address}$. Bob's privacy policy states that he requires a customer name for identification, payment details for a once-off payment for the order, as well as a shipping address for a once-off shipment of the order. He would also like Alice's e-mail address, to notify her of specials, and would like to distribute it to book clubs and other customers for reference purposes. Alice's privacy policy allows all of the above, except that she does not want her e-mail address distributed to others, but would like to be notified of specials. Alice now gives Bob a TGT T_1 allowing him access to $m_{name}, m_{payment}$ and $m_{address}$ for 7 days, limited to one transaction, and a TGT T_2 allowing him access to m_{e-mail} for 3 months. T_1

also limits the use of $m_{payment}$ to Bob as beneficiary, and $m_{address}$ to FastShipping. Bob now presents T_1 to our trusted third party S, and requests a ticket t_p for payment using Bob's bank, BBank, as well as a ticket t_s for shipping using FastShipping. S verifies the validity and policies of T_1, and then issues t_p and t_s. Bob now sends t_p to BBank, requesting payment. BBank presents t_p to S, who supplies BBank with Alice's credit card information for the transaction. BBank notifies Bob of the successful transfer. Notice that Bob never knew Alice's credit card details. Bob now sends the package and t_s to FastShipping, who presents t_s to S to get Alice's shipping address, and deliver her books. Note that Bob also did not know Alice's shipping address. For the next three months, Bob can send mail to Alice by requesting a ticket from S with T_2. This ticket is then sent with the message to a trusted messenger service (which might also be S), who will then forward it to Alice. When T_2 expires, S will no longer issue tickets to Alice's e-mail address. The same will happen if T_2 is presented to S by anyone other than Bob, or if Bob requests a ticket for another recipient. Any attempt reuse T_1 during its 7 day validity for another transaction will also be rejected by S. So, Bob can not reuse or redistribute the payment information, unless BBank conspires with him (unlikely — banks are all about trust. If they can not be trusted, public scorn will soon force them to close). Bob also can not reuse the shipping address for similar reasons, unless FastShipping conspires with him (a possibility, especially if Bob does a lot of business with FastShipping, and in reality there are not that many shipping companies with the ability to ship world wide at reasonable rates. This could be prevented by sending the package to S, who can forward it to Alice for a nominal fee, or to use a series of shipping companies, each knowing only the next step in finally delivering the package).

5. IMPLEMENTATION

The general case is stated, as well as certain properties of the tickets to implement the protocol correctly. A TGT will be issued by S to o for each transaction between i and o. Such a TGT will grant o access to such private information m about i stored at S described by the intersection of o and i's privacy policies.

5.1 General Case

When an individual i wants to send some private data m to a target organisation o, o's privacy policy P is checked using P3P. This policy is then used by i to give o a ticket granting ticket $T(o, i, P)$ granting access to m as described in the intersection between i's own privacy policy and P for a limited time. The actual private information is stored at a trusted third party S, for round the clock availability. When o needs part of m, o presents T as well as d, a description of the required subset of m and optionally o_2, another organisation who actually requires the access to S. S then verifies that P has not changed, and sends o a ticket t_{d,o_2}, where $o_2 = o$ if this was not supplied, and if allowed by P. A ticket can not be reused, and can only be used by the party it has been issued to. If o has to send private information to another party o_2, o has to request another ticket for o_2 from S if allowed by P. If o has to reuse information, a new ticket can be requested with T, unless T has expired.

5.2 More on Tickets

Tickets are used to access the actual private information. A ticket granting ticket (TGT) describes the types of access allowed, and is used to request tickets from a trusted third party S that can be used to access the actual data. We describe the use of tickets in more detail, by applying our categorisation as defined in Section 3

5.2.1 *Validation and calculation*

Although not directly addressed by this paper, this method can be used to calculate aggregates and averages on a set of different users' private information. If allowed to do so by the privacy policy, o can request S to calculate such aggregates. Another method to protect the privacy of information for such cases is discussed in [12].

5.2.2 *Third party requirement*

When a target organisation o requires private information m from an individual i to pass on to a third party, o can send a ticket to such a party allowing it to get the data directly from S.

5.2.3 *Future use*

If o would like to store m for some future use, o can just keep the TGT, and request a ticket from S when such access becomes needed. A further benefit is that m remains current, since all updates at S will filter through when o needs to access m. The availability of m is linked to the expiry date on the TGT.

5.2.4 *Identification*

As i's public key is part of the TGT, o can just encrypt a random message with it, and request i to decrypt it using i's private key. No actual knowledge of or access to private information is required by o in such a case.

5.2.5 *Actual contents*

To actually deliver a package, or actually transfer money (used by a bank) or actually require direct access to m. In this case, the real data is required, and is retrieved from S with a valid, unused ticket. The information is protected in the sense that no intermediary will have access to it.

The first four uses can be achieved without actually revealing m to o. The ticket granting ticket is all that is required. Only in the last case is the actual private information required, but privacy is still protected in a way, since it is only revealed at the last possible stage of any transaction. (And then it is minimal information such as: ship package number 1342 to this address, or transfer $23.54 from account number 352 to account number 2435).

5.3 Properties of Tickets

In order to function as described above, tickets need certain properties, similar to the Kerberos implementation.

id	A unique identifier of the TGT.
key(S)	Public key of the trusted third party S.
i	A unique identifier of the individual owner of the private data.
key(i)	Public key of the individual i.
o	A unique identifier of the target organization o being given access to m.
key(o)	Public key of the target organization o being given access to m.
m^*	An unambiguous description of the private data m being accessed by this TGT. P3P notation can be used.
pol(o)	A copy (or hash) of o's privacy policy at the time of issue.
pol(i)	A copy (or hash) of i's privacy policy at the time of issue.
expDate	Expiry date of the TGT.
reuseCount	Number of times this TGT can be reused.
limit(m_1)	
⋮	
limit(m_n)	Specific constraints pertinent to some part of m, such as shipping company, bank and amount. P3P field identifiers could be used, with a set of valid values for each.

Figure 1: Format of the Ticket Granting Ticket

5.3.1 Security

The TGT and tickets must be tamper proof. The contents should not be modifiable by any party other than S. The necessity for this requirement should be obvious. Such security could be achieved by having S sign them with S's private key. Any party could verify the signature using S's public key. Information in the ticket and TGT that are only meant for certain o could be encrypted by S with their public keys, keeping it private.

5.3.2 Time limit on the ticket granting ticket

Some of the access granted on the TGT might be for a limited time only. This time limit should be visible to a target organisation o, but should not be modifiable by o. This is required so that o can ensure adequate time is allowed for shipping, billing, etc. and also in the case where i subscribes to a service requiring private information, o can ensure that the TGT does not expire before the subscription does. This requirement can be achieved by signing the time limit with i's private key in cases where i set the limit, or S's private key otherwise.

5.3.3 Format of the tickets

The ticket granting ticket should at least contain the fields depicted in Figure 1, while the ticket should look like Figure 2. Tickets can not be reused.

id	A unique identifier of the ticket.
key(S)	Public key of the trusted third party S.
i	A unique identifier of the individual owner of the private data.
key(i)	Public key of the individual i.
o	A unique identifier of the target organization o being given access to m.
key(o)	Public key of the target organization o being given access to m.
m^*	An unambiguous description of the private data m being accessed by this ticket. P3P notation can be used.
expDate	Expiry date of the ticket.

Figure 2: Format of the Ticket

6. DISCUSSION

Some further questions spring to mind, and are discussed here.

The first and probably most important, is: What if S can not be trusted? This could possibly be solved by having more than one trusted third party, and storing either a subset of an individual's private information at each trusted third party. Another solution using more than one trusted third party is to split each information item among several trusted third parties, so that no single S can compromise any information.

Another important question to consider is that of collusion between some of the parties, as already mentioned in Section 4 — say between a big online retailer and its preferred shipping company. Such collusion might them possibly be curbed by using several shipping companies, or by policing by the trusted third parties (they could refuse to give TGT's to such guilty parties for any individuals registered with them, which might make the possible repercussions of such collusion too severe for any o to risk).

Of course, the better any individual's private information is protected and anonymised, the bigger the risk that such a system could be used for nefarious purposes (imagine someone ordering marijuana from the Netherlands, where it is legal, and having it anonymously delivered to the United States, where it is not. Interception of the package by customs might prevent delivery, but the recipient might not be identified). In cases where a single S is used, a court order might possibly be used to reveal private information about i in cases where criminal intent can be proven. In such a case the expired TGT might be presented to S with the court order or warrant requiring disclosure. S can then supply the original information. This could be even more effective than just storing the individual's private information in the traditional way, as the information stored at S might be more up to date. However, if S is physically located in an area outside the jurisdiction of authorities requiring such disclosure, S might not divulge the private information. This can then be countered by o requiring non-expiring third-party access to the private information m on the TGT for such authorities. The individual's private information could then only be accessed by the relevant authorities, and not by o.

7. RELATED WORK

Namesafe.com [10] provides good protection of identities, addresses and credit card details, but uses services such as Mail Boxes Etc. which are only accessible to users in the United States. The method proposed in this paper does not have such a limitation.

iPrivacy.com [11] encrypts part of the private information, but leaves other parts, such as the city, state and zip code open. The information that is encrypted, such as the address, can the be decrypted by the delivery company. This means that the information could still be out of date if the individual moves, as the encrypted version of an address is stored by the target organisation, instead of a TGT granting right of access to the information. iPrivacy.com's software is obtained from an individual's credit card company, which of course makes it inaccessible to people without credit cards.

8. SUMMARY

We have presented a classification of private information based on the purpose for which it is acquired, and created a protocol to protect private information in several of these classifications.

With this protocol we have effectively prevented unauthorised reuse and redistribution of private information in all cases where the target organisation o did not require direct, unprotected access to an individual i's private data. We have also managed to protect private information by disclosing it at the last possible stage in any transaction, virtually preventing all intermediaries from accessing our private information. However, note that this protocol does not prevent the last stage organisation to store and misuse the actual private information. The impact of this is lessened by the fact that this last stage usually have very little information, which might not be useful per se and the fact that in e-commerce applications these last stage organisations will typically be very large, such as banks and shipping companies, with a lot to lose should they misuse private information.

Further study could possibly integrate this approach with P3P, in order to automate this process and make it totally transparent to the end user.

9. REFERENCES

[1] H. Wang, M. K. O. Lee and C. Wang, "Consumer privacy concerns about Internet marketing.", *Communications of the ACM*, March 1998, 63–70.

[2] B. Thuraisingham, S. Jajodia, P. Samarati, J. Dobson and M. Olivier, "Security and Privacy Issues for the World Wide Web: Panel Discussion", *Database Security XII: Status and Prospects*, Kluwer, To Appear.

[3] L. F. Cranor, "Internet Privacy", *Communications of the ACM*, February 1999, 29–31.

[4] R. Clarke, "Internet Privacy Concerns Confirm the Case for Intervention.", *Communications of the ACM*, February 1999, 60–67.

[5] L. Cranor, M. Langheinrich, M. Marchiori, M. Presler-Marshall and J. Reagle, "The Platform for Privacy Preferences 1.0", *Draft*, http://www.w3.org/TR/P3P.

[6] B. C. Neuman and T. Ts'o, "Kerberos: An Authentication Service for Computer Networks" *IEEE Communications Magazine*, Volume 32, Number 9, September 1994, 33–38

[7] B. C. Neuman and J. Kohl, "The Kerberos Network Authentication Service (V5)" *Web Page*, http://www.faqs.org/rfcs/rfc1510.html

[8] USC/ISI GOST Group, "The Kerberos Network Authentication Service", *Web Page*, http://www.isi.edu/gost/gost-group/products/kerberos/

[9] "SET Home Page", *Web Page*, http://www.setco.org

[10] "NameSafe Home Page", *Web Page*, http://www.namesafe.com

[11] "iPrivacy Home Page", *Web Page*, http://www.iprivacy.com

[12] F. Lategan and M. Olivier, "Enforcing Privacy by Withholding Private Information", — in S. Qing and J. H. P. Eloff (eds), *Information Security for Global Information Infrastructures*, Kluwer, August 2000, 421–430.