A Privacy Paradigm that Tradeoffs Anonymity and Trust

Giampaolo Bella [1,2], Francesco Librizzi[1] and Salvatore Riccobene [1]
[1]Dipartimento di Matematica e Informatica, Universita' di Catania
Viale A.Doria, 6 — I-95125 Catania, ITALY
[2]SAP Research Labs, 805 Avenue du Dr Maurice Donat, 06250 Mougins, France
{*giamp,librizzi,sriccobene*}*@dmi.unict.it*

Abstract: **The protection of a principal's private information over a network is problematic. It is even more so in non-hierarchical settings such as the Web 2.0, where each node both provides and requires data. Two paradigms exist to tackle these issues: one relies on the principal's trust over the network, the other one insists on the principal's anonymity. A new paradigm is advanced as a natural tradeoff between the two: it sees the principal act using her real identity but only circulate statistical information about the resource she requires. This signifies that the privacy requirement is shifted from the principal's identity to her data. If evaluated with respect to the existing ones, the new paradigm appears simpler and more lightweight.**

1 INTRODUCTION

Ensuring privacy of a principal's data while they traverse a network of computers is far from trivial. Most security protocols appeared in the last three decades, from Kerberos to SSL, only aim at transmitting data confidentially but assume that the initiator is willing to share them with the responder. This simple assumption is unacceptable over any network where privacy is a sensitive property.

This is particularly the case with the Web 2.0, for example, a revolution of the Internet architecture towards social networking services such as Wikipedia and blogs. The Web 2.0 aims at bringing down the classical hierarchies of information flows: each principal is active, so that she can share with others any type of information of her choice, such as personal opinions, experiences, pictures, and so on. There are millions of principals who both send and receive data over the Internet.

To illustrate the privacy concerns that are typical of the Web 2.0, let us adopt it as a setting for electronic purchases. A principal contacts a node to obtain the desired resources, and enters her personal information. In processing the request, the node may collaborate with others, sharing (some of) the principal's personal information, such as what she wants to buy. Moreover, finalising the purchase requires the node to interact with a Bank to manage the payment and with a shipment society to deliver the purchased items. The result is that the principal's private data have flown through several nodes, each handling them according to its own privacy policy, when the principal perhaps remains unaware of the flow.

There exist two main paradigms to let a principal safeguard her privacy over the Web 2.0. One rests on the principal's trust over the network, so that she will accept to transmit her identity and required resources, but will be able to suspend or resume trusted nodes' treatment of her personal data. The other one is based on anonymity, so that data are linked with a pseudonym and not with the real principal's identity.

We advance a privacy paradigm that appears to be the natural tradeoff between the two. There are various real-world contexts where the above paradigms are not applicable, because the principal does not sufficiently trust the network to circulate her identity and required resource name or the anonymity is not allowed. Our paradigm applies here because it removes both the need for principal's anonymity and for her trust in the entire network.

The main idea underlying our paradigm is that the principal uses her real identity but only circulates statistical information about the resources she is looking for. The statistics are computed by the first node that receives her request and is interested in doing business. This feature has a double advantage. Not only is the principal, who might be a casual customer, entirely relieved from the burden of doing statistics, but these will be more suitably calculated by a merchant node according to the most appropriate business rules.

Moving the computation of statistics away from the principal arguably raises the risks of privacy breach. But we upgrade and tailor an existing privacy-enforcing protocol to prevent the node that does the statistics from disclosing the raw, private data without being indicted. Statistical data will be orchestrated through the network and each node will send matching offers to the principal by certified email. Finally, the principal will only disclose the very resource name to the chosen node, though via some fair-exchange scheme. The principal's trust is only confined to this end node.

The organization of this paper is simple. The next section shows an outline of the existing paradigms (§2), then our paradigm is explained in its three phases (§3). Some conclusions terminate the presentation (§4).

2 EXISTING PARADIGMS

2.1 Trust: Suspending and Resuming Data

Waidner and Schunter design a suite of protocols (briefly addressed as WS protocol) to allow a principal to manage her private data across a network of trusted nodes [1]. Nodes are trusted in the sense that they will conform to the protocol, as clarified below. Their treatment is in the context of electronic purchase, an "on-line retail scenario" as they address it. But it is meant to generally applicable to a Web 2.0 environment.
Their protocol sees a principal transmit her personal data (name, surname, what she wants to buy and so on) to a relevant node in the trusted network, e.g. a bookseller. The bookseller may collaborate with the other nodes in the trusted network in order to fulfill the principal's request. Waidner and Schunter suggest that each node have a privacy panel that allows the principal to manage her personal information at the various collaborating nodes'. She can view the node's privacy policy, learn to whom the node disclosed her data, and also block or delete her data. The principal bundled her data with various ACLs (Access Control Lists) to specify who can do what on them, and with a DF (Data Flow) matrix to indicate her intended flows. Both of them are digitally signed by the originator, but we argue that each intermediate node might affix its own signature to the plaintext, if only it were not trusted not to do so.
The disclosure can be seen as a delegation. Initially, the principal delegates a node to handle her data, and then the node delegates another one, and so on. The principal can block or unblock the use of her personal data at a node's through dedicated protocols. The block message will be propagate to all nodes that received the principal's data.
The block protocol sees the principal begin by sending an authenticated block request message to the node. If this node ever disclosed the data to others, then it forwards the block request message, othewise it responds to the calling node with a signed block response message as an acknowledgement. Along each delegation path, the delegation response messages are nested.
The unblock protocol is simpler. It sends the unblock request messages through the delegation graph but requires no response messages.

2.2 Anonymity: Using a Pseudonym

The previous paradigm works correctly if the network is trusted, but nothing confirms that the nodes will conform to the protocol. The dual paradigm disposes with such a trust entirely, and provides the principal with anonymity or pseudonymity.
The DAA (Direct Anonymous Attestation) protocol [2, 3], which is adopted in the TPM v1.2 (Trusted Platform Module) specification by TCG (Trusted Computing Group), is the best-known protocol aiming at principal's anonymity. Its key ideas are portrayed in Figure 1. It is useful to spell out the DAA protocol in the same context of the WS protocol, electronic purchases.

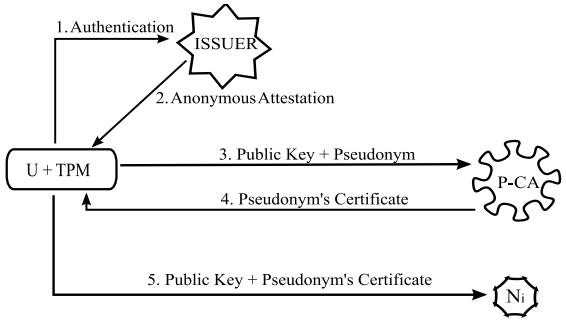

Figure 1: Principal's anonymity

An Issuer authenticates a principal through her pre-existing certificate, and issues an anonymous attestation (message) for her — steps 1 and 2. This message — encrypted with the principal's public key — states that the principal is genuine but does not reveal her identity. In step 3, the principal generates a public key, chooses a pseudonym, and submits them to an entity called P-CA (Privacy Certificate Authority). The P-CA verifies that the principal has valid attestation, and that her pseudonym is computed out of information that is present in the attestation. If this twin check succeeds, then the P-CA will release the certificate for the pseudonym's public key, step 4. Through this valid-attested purchase certificate the user can access to the merchant node — step 5 in Figure 1. Moreover, it protects the principal's identity and the TPM's endorsement public key (that is unique for each TPM by construction),which are only used in the initial phase with the Issuer.
The attestation to be anonymous, an attacker must be unable to link the pseudonym with the principal's identity. Because the Issuer is the only entity that can resolve that link, the DAA protocol protects it by adopting a group of Issuers and a group signature scheme [4] [5] [6]. Moreover, the protocol separates the Issuer from the P-CA, so to make a successful attack would require collusion with both authorities. The protocol can be additionally strengthened by having the P-CA release one-time certificates [3]. Some real-world applications adopting a DAA protocol already exist [7].

3 A TRADEOFF ANONYMITY/TRUST

A coarse simplification of the Web 2.0 is as a network of nodes that collaborate sharing information. Each node exposes its functionalities as web services, such as those provided by mush-up applications. They offer mechanisms (applications, protocols, etc.) to enable access to one or more resources, and can share principals' data [8].
We have seen two existing paradigms that are two dual worlds. In the first, the principals have to rely on the network (ACLs and DF matrix), whereas in the second the principals do not

rely on the network and use the anonymity system. But in some applications the two paradigms cannot be applied. For comparison, we present also our paradigm in the same context without hindering its generality.

Our aim is to conjugate the benefits of both paradigms by removing as much as possible the trust into the network (as the DAA protocol already does) and by shifting the privacy enforcement mechanism from the principal's identity (i.e. shifting from anonymity) to the principal's actual data so as to avoid the necessity of a privacy certification authority. Of course, we cannot aim at removing the principal's trust from every node in her network. At least one node that will eventually provide the required resource or product must be trusted to keep the principal's privacy.

Therefore, the main idea underlying our paradigm is to conceal a principal's private data by means of statistics, and to transmit only the (results of the) statistics over the network. The more a statistic conceals the principal's data, the less trust is required of peer nodes — for example, the principal may not want to share her interest in war books with the network because she does not trust its nodes sufficiently; conversely, she might accept to share a broader statistics such as the fact that she buys books (rather than other goods) because her trust in the network is proportionate to this general datum.

The expected price to pay when trust tends to nothing is a less precise proposal finding, that is a coarser *orchestration*. Arguably, it will be harder to deliver the principal her required book "Fascism (Cambridge Perspectives in History)" upon the sheer indication that she is interested in books.

Our general paradigm of principal's privacy for the Web 2.0 is a tradeoff between the principal's anonymity and her trust over the network. It comprises the following three phases.

1. **Data concealment** is the first and main phase as it operates the main shift of privacy enforcement from the principal's identity to her data. Concealment is done by applying appropriate statistics (§3.1).

2. **Orchestration** then informs the principal of the network nodes providing the resources or the goods that best match her data. Obviously, a deeper data concealment causes a less precise orchestration (§3.2).

3. **Completion** sees the principal choose a node on the basis of its product offer. The principal finally initiates an appropriate security protocol (depending on the application domain) with the chosen node (§3.3).

3.1 Data Concealment

The problem of concealing principals' data was lately tackled in the context of electronic polls by Golle et al. [9]. That scenario sees principals submit their preferences to an electronic pollster. Because the principals' preferences are private data, the pollster is expected to treat them accordingly. However, this cannot be trivially guaranteed, as a dishonest pollster might choose to sell the private data it collects.

Electronic polls usually serve to build statistics, and hence their participating principals are willing to accept publication of the statistics but not of their raw preferences. Golle et al. design (various versions of) a security protocol that we call SEP as it aims at self-enforcing privacy. Due to space limitations, we argue it briefly.

The main idea underlying SEP is that the principals insert inside their preferences some extra cyphertexts that serve as baits. Therefore, if the pollster decides to violate the principals' privacy by publishing their raw preferences, it would be forced to publish the whole collection, including the baits. What frames the pollster as the actual publisher is exactly those baits, which provide the principals with sufficient evidence to accuse the pollster of privacy breach.

This protocol also is fair, as it protects the pollster from unfounded accusations. Precisely, the principals cannot forge preferences with inner baits by themselves, and then make fake accusations to the pollster.

However, the pollster should be able to publish the statistical results without any risk of indictment. To allow this, the function that was used to produce the statistical results must conform to the definition of *differential privacy* [10]:

Definition 1 A randomized function f over data sets gives ϵ-*differential privacy* if for any two data sets X_1 and X_2, which differ in at most one point, and $S \subseteq Range(f)$,

$$\Pr[f(X_1) \in S] \leq \exp(\epsilon) \times \Pr[f(X_2) \in S]$$

The intuition behind this definition is simple. For example, X_1 and X_2 can be thought of as two databases that differ in only one record. A function f satisfies the definition of differential privacy if there is a similar probability (for small values of the ϵ parameter) that the respective applications of f to the two databases yield the same "feature" S. This means that a statistic obtained by using such a function is not significantly influenced by the noise represented by the baits, and therefore its results remain significant.

Now, we explain briefly how SEP works. Its setup phase sees the pollster publish two hash function (h and g), and its public key. The h hash function is used to send bit of real data, whereas the second one g is used to achieve the baits, its image is the set of cyphertext encrypted with the pollster public key. Therefore, only the pollster can decrypt the baits.

When a principal wants to send her preferences to the pollster, she can choose to send a real bit of data or a bait. In the first case, she applies the h hash function and encrypts its output with the pollster public key, finally sends the encrypted value together with her identity. In the second case, she chooses a random value s, applies $g(s)$ and sends it together with her

identity. The pollster cannot distinguish if the received values are baits or real data, because both are encrypted using its public key. Therefore, it can only decrypt and store them for applying further the statistical function.

To accuse the pollster (indictment phase), a principal must exhibit a "sufficient number" (which we purposely do not detail here) of valid exhibits. An exhibit is represented by the principal identity R, the s value and the bit b under indictment. The exhibit is valid if b is obtained by $g(s)$. We remark that $g(s)$ is a cyphertext and the principal cannot decrypt it, therefore she must collaborate with the pollster to achieve $D(g(s))$.

It is somewhat surprising that the pollster must decrypt the baits. Would a dishonest pollster help its own accusation in practice? This form of collaboration may be unrealistic in many real-world scenarios. The protocol should be strengthened so as to facilitate valid accusations even when the pollster fails to collaborate.

It seems perfectly realistic to strengthen SEP with the adoption of a PKI (Public Key Infrastructure). We address our modified protocol as SEP+ [11]. In particular, SEP+ wants the pollster to be equipped with a signature key pair, and corresponding signature creation algorithm S and signature verification algorithm V. Upon reception of a pair $\langle R, c\rangle$, where R is the principal's identity and c is a cyphertext representing a bit or a bait, SEP+ requires the pollster to reply with the signed acknowledgement $S(\langle R, c\rangle, t)$, where t is the current timestamp. The principal will verify such a signature using V before she continues the session with the pollster. Should this ack message fail to arrive or arrive late as proven by its timestamp, the principal might decide to quit with the pollster. Therefore, a rational pollster will have interest in computing and sending its signature correctly upon each arrival of a bit/bait.

SEP+ inherits the indictment phase from SEP, but the principal has significant evidence even without the pollster's collaboration. Should the pollster fail to participate by providing correct decryptions, the principal might exhibit (a "sufficient number" of) the pollster's signed acks. They confirm that the pollster accepted a number of ciphertexts thereby accepting to comply with the principal's privacy policy. The integrity of the acks ensured by the digital signatures is crucial here. The PKI also helps pinpoint the pollster because it is securely registered with some certification authority. By contrast, SEP provided no track of the pollster's operations to anyone.

The next phase is data concealment (Figure 2). A principal who wishes to interact with a network to obtain a specific resource begins by finding an "initial" node in that network. A deeper discussion about this search is beyond our interests here. For example, in the context of electronic purchases, the initial node might be accessed through the web site of a price-finding engine; in the context of MANETs, the initial node has some "proximity" relation with the principal.

The principal runs SEP+ with the initial node in order to transmit a large list of interrelated resource names. The interrelation

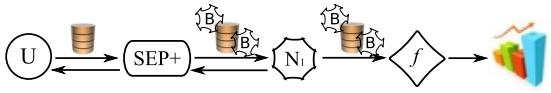

Figure 2: Phase 1: data concealment

may either be syntactic or semantic. Since the list will be input to a differential-privacy (see definition 1) preserving function performing either syntactic or semantic statistics, it must not necessarily include the very resource name the principal requires. For example, if the principal wishes to find a specific file name on Fascism from a file-sharing network, she may submit a list of titles about World War II. Recall that SEP+ prescribes the principal to interpose baits to the list so that the initial node must first apply a differential-privacy preserving statistical function to the list, and only handle its output. Otherwise, it would be indicted as explained above.

It is clear that a decentralized and delocalized application such as the Web 2.0 would make the use of the original SEP inappropriate. It cannot be assumed that any node anywhere in the world would collaborate to its own putative indictment without having signed anything.

3.2 Orchestration

This phase begins when the initial node has prepared the statistical data about the principal request (Figure 3).The initial node begins to transmit the statistical results to a number of participating nodes in the network. Transmission is recursive in the sense that whichever node receives the data forwards them to other nodes depending on its computational resources, anti-DoS heuristics and, above all, on a network reputation system. It is important to remark that the principal's privacy is not affected because only statistical data are treated.

If a node that gets the statistical data feels that it can make a significant offer, then it emails it to the principal using a certified email protocol. Clearly, the level of appropriateness of the offers is bound to be balanced with the type of statistical data that were transmitted. The more privacy-preserving the statistics, the less focused the offers. The orchestration terminates with the principal's choice of the best offer, which contains the very resource that the principal is seeking.

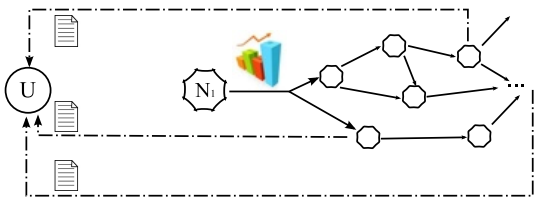

Figure 3: Phase 2: orchestration

Notice that in practice the initial node will begin the orches-

tration regardless of whether it has its own relevant offers to make — the node will have to balance its sole rights of sale with its network reputation. It might decide not to start the orchestration in order to make its own, sole, offer. Because offers are to be made by certified email, the node would be unable to forge more expensive offers by other nodes. Not only would it loose network reputation, but also the principal would realise its selfish behaviour.

The best business choice for a node throughout the orchestration, whether to send over the statistical data or stop them, is not obvious. Putative subsequent offers by other nodes might be higher as well as lower. Either way, the typical Web 2.0 setting will force it to balance its choice with its reputation.

Our protocol does not raise the risks of DoS attacks to the principal. She may decide to process the offers only for a limited time window, and to discard them afterwards. Alternatively, and even more simply, she may only process lightweight emails containing the URL with the dedicated offer and discard the others.

3.3 Completion

This final phase begins when the principal has already chosen the node from which to obtain her required resource, and lets the principal come to a formal agreement with that end node (Figure 4). The principal executes a security protocol with the end node in order to settle a secure access to the resource. Obviously, the principal must reveal to the end node the required resource, but the security protocol shall protect its name.

Security here strongly depends on the application domain. It may generically evaluate to mutual authentication and confidentiality. For electronic commerce for example, this phase would require a suitable protocol such as SSL or SET [12]. The latter in particular also settles payment. Other contexts may require a fair-exchange protocol [13] to protect the peers from each other's potential false claims.

Figure 4: Phase 3: completion

As remarked above, the principal must put some trust in the end node. Because the principal is not protected by anonymity, such a trust cannot be removed. The end node will one way or another realise what resource to grant the principal or which good to ship to her. However, this limitation is somewhat shared also with the anonymity paradigm.

The completion phase attempts to counter the risks deriving from the principal's trust on the end node. For example, a fair exchange protocol can require the end node to formally agree with the principal's privacy policy. Should the end node abuse the principal's data, the principal would be able to sue it.

4 CONCLUSIONS

We have advanced a novel paradigm to safeguard a principal's private data on the Web 2.0. It is the natural tradeoff between trusting the network, as the WS protocol does, and distrusting it entirely, as the DAA protocol does. The use of statistical data safeguards data privacy so that the network needs not be trusted any more and the need for anonymity is removed.

Acknowledments This work was supported by the FP7-ICT-2007-1 Project no. 216471, "AVANTSSAR: Automated Validation of Trust and Security of Service-oriented Architectures" (http://www.avantssar.eu).

References

[1] M. Schunter and M. Waidner, "Simplified privacy controls for aggregated sevices - suspend and resume of personal data," in *Privacy Enhancing Technologies*, 2007.

[2] E. Brickell, J. Camenisch, and L. Chen, "Direct anonymous attestation," in *Proceedings of the 11th ACM conference on Computer and communications security*, 2004, pp. 132 – 145.

[3] J. Camenisch, "Better privacy for trusted computing platforms," in *Computer Security – ESORICS 2004*, 2004, pp. 73 – 88.

[4] G. Ateniese, J. Camenisch, M. Joye, and G.Tsudick, "A practical and provably secure coalition-resistant group signature scheme." ser. LNCS 1880. Springer, 2000, pp. 255 – 270.

[5] J. Camenisch and A. Lysyanskaya, "Efficient non-transferable anonymous multi- show credential system with optional anonymity revocation," ser. LNCS 2045/2001. Springer, 2001, pp. 93 – 118.

[6] A. Lysyanskaya and J. Camenisch, "A signature scheme with efficient protocols," ser. SCN 2002, vol. 2576. Springer, 2002, pp. 268 – 289.

[7] E. Gallery and C. J. Mitchell, "Trusted mobile platforms," in *Foundations of Security Analysis and Design IV*, 2007, pp. 282 – 323.

[8] "Reference model for service oriented architecture 1.0," OASIS's web site, http://www.oasis-open.org, 2006.

[9] P. Golle, F. McSherry, and I. Mironov, "Data collection with self-enforcing privacy," in *Proceedings of the 13th ACM conference on Computer and communications security*, 2006, pp. 69 – 78.

[10] C. Dwork, "Differential privacy," in *ICALP 2006*, ser. LNCS 4052. Springer, 2006, pp. 1–12.

[11] G. Bella, F. Librizzi, and S. Riccobene, "Realistic threats to self-enforcing privacy," in *Proceedings of the Fourth International Conference on Information Assurance and Security, IAS'08*. IEEE Press, 2008.

[12] G. Bella, F. Massacci, and L. C. Paulson, "Verifying the

SET purchase protocols," *Journal of Automated Reasoning*, vol. 36, no. 1-2, pp. 5–37, 2006.

[13] J. Zhou and D. Gollmann, "A fair non-repudiation protocol," in *Proceedings of the IEEE Symposium on Research in Security and Privacy*, 1996, pp. 55 – 61.