

Private Credentials

Zero-Knowledge Systems, Inc.

November 2000[†]

Abstract

The current approach to digital certificates and Public Key Infrastructures ignores the privacy rights of individuals, groups, and organizations. Digital certificates can be followed, traced, and linked instantaneously as they move around. Unless drastic measures are taken, individuals will soon be forced to communicate and transact in what could be the most pervasive electronic surveillance system ever built.

This white paper explains our concerns, and offers a new vision of the future built around Private Credentials. We describe the functionality of Private Credentials, and show how they are beneficial to both individuals and organizations, not only with respect to privacy but also in terms of efficiency and security.

Our exposition presumes that the reader is familiar with the concept of a digital signature, but does not require any technical or mathematical knowledge.

1 Introduction

“The real danger is the gradual erosion of individual liberties through the automation, integration, and interconnection of many small, separate record-keeping systems, each of which alone may seem innocuous, even benevolent, and wholly justifiable.”

— U.S. Privacy Protection Study Commission, 1977

Individuals and organizations often have a legitimate need to verify the identity or other personal information of the individuals they communicate or transact with. With the advancement of computing and network technologies, paper-based certificates and other traditional methods of certifying a person’s credentials are increasingly vulnerable to counterfeiting, unauthorized duplication, theft, and destruction. They also require physical transport or proximity, making them costly or useless in electronic environments.

Digital Certificates

Digital certificates are widely perceived to be the most promising technique for safeguarding electronic environments. Digital certificates are just sequences of zeros and ones (of a special mathematical

[†]An earlier version of this white paper was distributed in September 2000.

structure), and so they can be verified with 100 percent accuracy by computers and can be transferred electronically without human intervention or noticeable loss in time. Owing to their special mathematical structure, it would take millions of years to forge a digital certificate, even if all the world's computing power could be tapped.

Similar to passports, diplomas, drivers' licenses, and other traditional certificates, digital certificates can specify any kind of data. But, not all certificates are created equal.

Privacy Dangers

Under the current approach to digital certificates, individuals are given an *identity certificate* that will form the basis of all their communications and transactions. It is like an electronic version of a passport, but much more secure. However, the corrosive effect on privacy is enormous. Each identity certificate can be instantaneously and automatically followed around and traced to its holder's identity as it moves through the system. This enables organizations and individuals to compile extremely precise personal dossiers, containing detailed information about a person's financial situation, medical history, lifestyle, habits, preferences, whereabouts, and so on. The dossiers can be compiled, linked, and updated in real time without human intervention. Moreover, individuals will be unable to repudiate their actions, because these are digitally signed by them; this results in all sorts of risks, including legislative risks.

These exceptional surveillance powers are enjoyed not only by the organizations that a person directly communicates or transacts with, but also by their unscrupulous employees, by hackers, by intelligence agencies, by private and public organizations who acquire dossiers, and by any organization that issues digital certificates. Typical representatives of the latter group are financial institutions, governments (local, state, and federal), insurance companies, health care providers, post offices, public transport organizations, and consumer reporting bureaus.

Ineffective Privacy Approaches

It is a widespread misconception that the encryption of stored and transmitted data suffices to address privacy concerns. Confidentiality protects data only against wiretappers and other outsiders; it does nothing to prevent the parties to a communication or transaction, or anyone with authorized access to stored data, from tracing, linking, selling, or misusing the data in whatever manner they see fit. Consider, for instance, the flurry of recent reports in the press about Web sites that accidentally revealed credit card data and other personal information of their clientele to the public at large.

Equally ineffective is the recommendation by standards bodies and identity certificate providers to issue identity certificates that specify a pseudonym instead of a real name. Pseudonymous certificates that can only be obtained by certificate applicants who identify themselves do not prevent tracing and linking; they are similar to Social Security numbers, credit card numbers, and health registration numbers, but with better authentication. The alternative of not requiring certificate applicants to identify themselves offers better privacy, but is highly undesirable in most applications. Namely, it is hard or impossible to protect against lending, copying, discarding, and other misuse of certificates. The certificate issuer cannot even contain the damages due to fraud. Furthermore, it is difficult to register without identification, certificate applicants cannot build a long-term relation with the certificate issuer, and all the actions of a certificate holder are still linkable.

Private Credentials

In view of this worrisome situation, Zero-Knowledge Systems has set out to develop and implement technologies that empower individuals, groups, and organizations to communicate and transact privately and securely. Central to our approach is the notion of *Private Credentials*. While identity certificates are similar to (but more invasive than) passports and other paper-based identity documents, Private Credentials are similar to coins, stamps, and public transport tickets, none of which contain information that can be correlated to the identity of their holder.

Private Credentials are not only more secure and efficient than their paper counterparts, but more functional too. For instance, a Private Credential holder can disclose a specific portion of the data that has been encoded into a Private Credential, without revealing any other information; this is much like using a marker to cross out data on a paper-based certificate, but much more powerful. Also, a Private Credential can be presented in such a manner that the verifier is left with no evidence at all of the disclosed property (much like waving a passport when passing customs, a practice customary in several countries) or only with partial evidence (much like presenting a paper-based certificate with crossed-out data fields so that a photocopy can be made).

Benefits to Organizations

Private Credentials are advantageous not only to individuals, but also to organizations that rely on the verification of digital evidence. Amongst others, they prevent issuers and other central parties from competing unfairly; they minimize the scope for law enforcement intrusions on databases; they reduce the scope for discrimination and identity fraud; they foster fair competition with respect to the collection and use of personal data; they are the cheapest and most effective way to comply with as many of the privacy principles of codes of conduct and privacy legislation as possible; they improve transaction finality; and, they cultivate goodwill among customers.

Difference with Infomediaries

Information intermediaries (also known as infomediaries) claim some of the privacy benefits of our Private Credentials. An infomediary, however, is just a broker of personal information. Infomediaries require their customers to funnel all their communications and transactions through their company, so that the data can be organized into profile categories and sold to anyone interested in it. Once an individual surrenders identifiable personal data to an infomediary, there is no sure way to have it destroyed after it has served its original purpose. Any privacy policies an infomediary may voluntarily live up to today may be ignored or bypassed in the case of acquisitions, mergers, bankruptcy, or management changes. Moreover, the databases of an infomediary become an appealing target for hackers, malicious insiders, and overzealous law enforcement agents. And, the ease of getting a subpoena makes it possible for just about anyone to gather information from an infomediary.

In sum, infomediaries do not give individuals control over their own information. In contrast to their “trust-us” approach, Private Credentials ensure that identifiable personal data never gets out in the first place; individuals need not trust anyone else with their personal data.

Applications of Special Interest

Private Credentials can be used to design all sorts of electronic communication and transaction systems, since complete disclosure of personal data is often not necessary at all. Applications of special

interest include, but are not limited to: electronic cash; digital pseudonyms for public forums and virtual communities (such as Internet newsgroups and chat rooms); access control (to Virtual Private Networks, subscription-based services, Web sites, databases, buildings, and so on); digital copyright protection (certificates permitting use of works); electronic voting; electronic patient files; electronic postage; automated data bartering; online auctions; financial securities trading; pay-per-view tickets; public transport ticketing; electronic food stamps; road-toll pricing; national ID cards (but with privacy); permission-based marketing; Web site personalization; multi-agent systems; collaborative filtering (i.e., making recommendations to one person based on the opinions of like-minded persons); gift certificates; loyalty schemes; electronic gambling; and, medical prescriptions.

Our Private Credentials are not complementary to identity certificates, but encompass them as a special case; one always has the option of disclosing an identifier. The major contribution of Private Credentials is that they enable individuals to determine for themselves when, how, and to what extent information about them is revealed to others, and to what extent others can link or trace this information.

Overview of Contents

In Section 2 we go into more detail on the subject of identity certificates, and how they invade privacy. In Section 3 we argue that technological solutions to the privacy problem are inevitable, and describe the basic privacy objectives they should meet. In Section 4 we explain how our Private Credentials protect privacy and how they are beneficial in other ways. Section 5 describes software-only as well as hardware-based techniques to protect against all kinds of fraud, including reuse, copying, lending, and discarding of Private Credentials. Section 6 concludes with an outlook for the future.

2 Identity Certificates

“Subtler and more far reaching means of invading privacy have become available to the Government. [...] the progress of science in furnishing the Government with means of espionage is not likely to stop with wiretapping.”

— Justice Louis Brandeis in *Olmstead v. United States*, 277 U.S. 438, 473-474 (1928)

Under the predominant approach, individuals are given a digital *identity certificate*. This is a digital signature of a trusted entity, called the *Certificate Authority (CA)*, that binds a *public key* of an individual to his or her name. The name could be the individual’s real name, but it could also be a Social Security number or any other data that can readily be associated with the individual.

When communicating or transacting with another party, the *certificate holder* sends his or her digital certificate and uses the corresponding secret key to authenticate the communication or transaction. The authentication prevents wiretappers and malicious certificate verifiers from replaying messages, and may also serve to provide non-repudiable evidence of the transaction. The certificate *verifier* can verify the name–key binding of the presented certificate without the involvement of the *CA*, by applying a trusted copy of the *CA*’s public key. By using the identity certificate as an authenticated pointer into all sorts of online and offline databases, the verifier can look up any information about the certificate holder it is interested in; this is similar to the way Social Security Numbers are used, but with much better authentication.

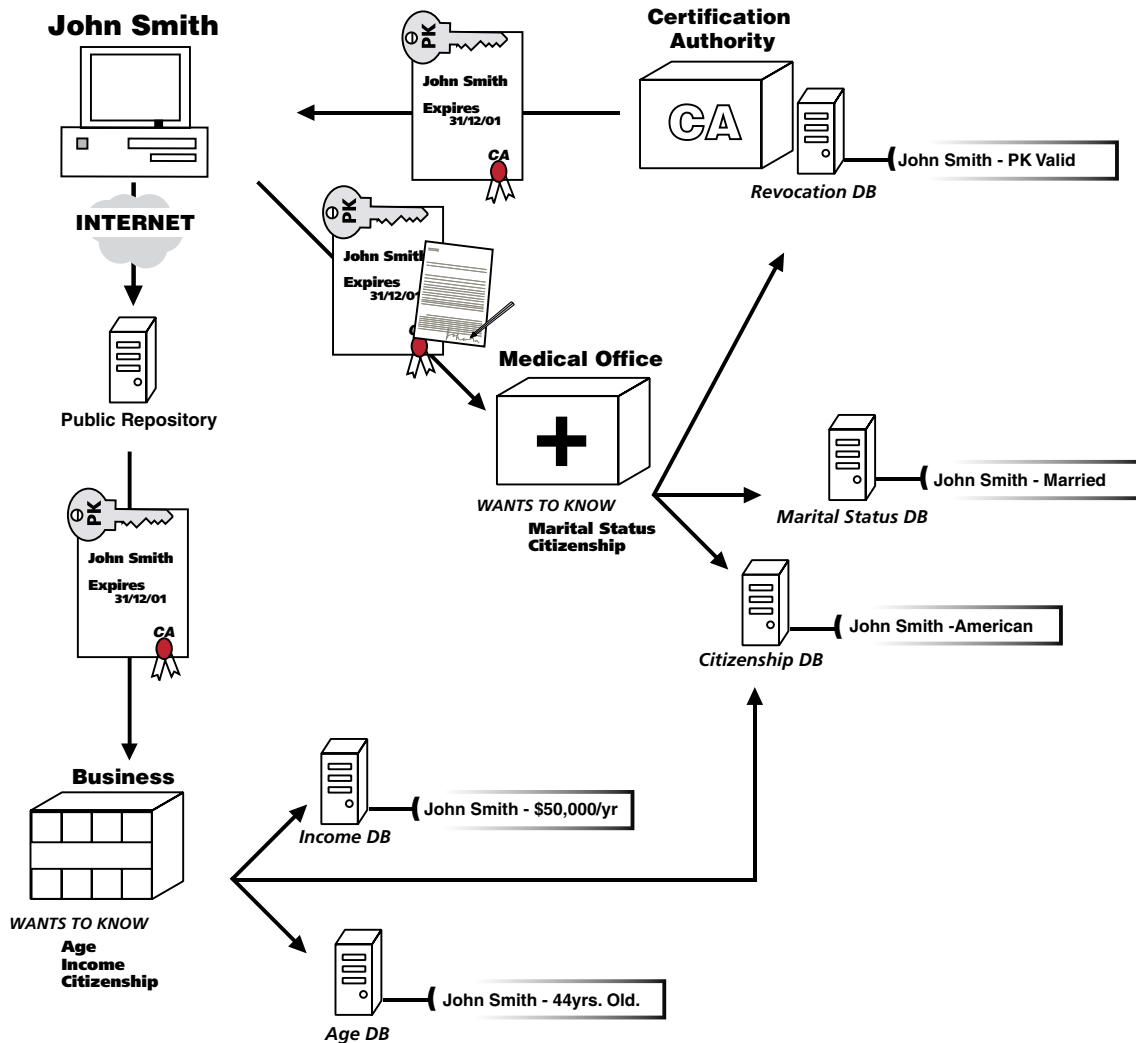


Figure 1: The identity certificate model.

Figure 1 illustrates the identity certificate model. The figure shows a user named John Smith receiving from a CA a certificate that binds his name to his public key. The certificate also specifies an expiry date and possibly other data (not shown here). To engage in an authenticated transaction with a medical office, John Smith sends his certificate together with his own digital signature on a message that contains a unique data field. The signature prevents replay attacks and also gives the medical office cryptographically non-repudiable evidence of the transaction. The medical office uses the certificate to retrieve John Smith's personal data (here, his marital status and citizenship) from various databases. These databases may or may not be proprietary, and they may be online or offline. The medical office also consults a revocation database to make sure that John Smith's certificate has not been revoked. Furthermore, the figure shows John Smith sending his certificate to an online public repository. From here it is retrieved by a business. The business uses John Smith's certificate to look up whatever personal data it is interested in (here, John Smith's age, income, and marital status). Note that the business does not need John Smith's involvement or awareness to do so, since his digital signature is not needed.

A system relying on the distribution and management of digital certificates is called a *Public Key Infrastructure* (PKI).

The Pursuit of Identity

The identity certificate business is worth tens of billions of dollars according to stock market valuations. Identity certificates are already widely used on the Internet, for the purpose of authenticating e-mail, Web servers, and software. The Web browsers of all major software manufacturers have built-in capabilities for storing, sending, and verifying identity certificates. Also, identity certificate standards exist for various industry segments, all based on the X.509 framework [43] of the International Telecommunications Union.¹ X.509 was specifically designed to work with X.500, an online database that contains globally unique identifiers for persons, devices, and anything else that can be assigned a “distinguished name.”

But these developments are only the beginning. Around the world (especially in Europe and Asia), public telecommunication organizations, transport organizations, municipalities, health care providers, ministry departments, financial institutions, and other influential organizations are planning to provide their customers with identity certificates that will be the sole means of participating in their systems. Many industrialized countries are even planning federal PKIs using identity certificates.²

The trends of wireless connectivity and computerization of devices dramatically increase the need for secure authentication. In the near future, identity certificates may be built into any piece of software or device able to communicate with other devices or with individuals, such as mobile phones, watches, televisions, cars, and computerized household appliances.

These developments reflect the widespread belief that identity certificates are essential for secure electronic communications and transactions. This persistent misconception is perpetuated in the press and in popular accounts on the subject of computer security.³

Privacy Problems of Identity Certificates

At Zero-Knowledge Systems we are deeply concerned about the corrosive effect of identity certificates on privacy. Our concerns are manifold:

¹Standards include: X9.55 [3, 5] (an ANSI-adopted standard developed by the American Bankers Association and targeted at the financial services industry), PKIX [1, 19] (a draft Internet standard by the Internet Engineering Task Force), Privacy Enhanced Mail [13, 46] (PEM, an e-mail standard proposal by the Internet Engineering Task Force), Fortezza (the U.S. federal standard for secure e-mail and file encryption in the defense system), Secure/Multipurpose Internet Mail Extensions [22, 23] (S/MIME, an e-mail standard proposed by RSA Security), Secure Socket Layer version 3.0 [34] (SSL, developed by Netscape to support server and client authentication and session encryption), and Secure Electronic Transactions [48] (SET, proposed by MasterCard and Visa for securing card-not-present credit card transactions).

²On the forefront are the United States of America (the NSA, the IRS, the FBI, the Department of Defense, the Social Security Administration, and 19 other federal agencies are running pilot projects, and Access Certificates for Electronic Services will provide for public electronic access to government services and information), the United Kingdom (its CLOUD COVER initiative was started to stimulate the growth of a government-wide PKI), Australia (the Public Key Authentication Framework will result from the Gatekeeper federal infrastructure program and efforts by the Certification Forum of Australia), Canada (in 1995, the Treasury Board endorsed the Government of Canada Public Key Infrastructure), and Hong Kong (in November 1999, the Hong Kong postal service started issuing identity certificates to most of the 6.5 million residents). For a 1999 snap-shot overview of the PKI initiatives in 26 member countries of the Organisation for Economic Co-operation and Development (OECD), see the OECD Working Party on Information Security and Privacy [51].

³See, for instance, Adams and Lloyd [2], Feghhi, Williams, and Feghhi [32], Ford and Baum [33], Lewis [47], and Zimits and Montano [59].

(Traceability) All of an individual's communications and transactions can be automatically traced on the basis of his or her identity certificates. In this manner, dossiers can be compiled about the individual's habits, behavior, movements, preferences, characteristics, and so on. Reasons why a CA may get to see the certificates that are shown to verifiers include: verifiers may be incited to deposit to the CA a copy of the transcript of each transaction they engage in (e.g., to enable fraud detection); verifiers may resort by default to online certificate validation by the CA (as with credit card transactions); and, the verifiers may belong to the same entity as the CA (this usually is the case in so-called closed systems, such as road-toll pricing).

In case the communications and transactions of certificate holders are not securely encrypted, wiretappers see the same information as verifiers. If the certificate issuing process can be monitored also, wiretappers can learn everything the CA knows.⁴

The CA can easily link each dossier to the identity of the individual to whom it pertains. For verifiers and wiretappers, linking dossiers to identities is typically a simple task as well: either separate entries or the aggregated contents of a dossier reveal the identity, or the match can be made in another way (e.g., on the basis of voice or facial recognition or by tracing the source of an Internet connection).

(Linkability) Each verifier can store all the certificates presented to it, and can link them on the basis of their key holder identifiers, public keys, or CA signatures. Different verifiers can exchange and link their data on the same basis. Furthermore, all the dossiers compiled by linking and tracing the actions of participants in one PKI can be tied to the dossiers compiled in other PKIs.⁵ With the trend or at least the capability of organizations to merge their databases at ever decreasing cost, it is naive to believe that linkable data that is submitted to different locations will never be linked. Motivations for linking include security management, mergers, marketing objectives, and inventory management.

(Non-repudiable evidence) Any digital signatures made by certificate holders can be added to their dossiers; they form self-signed statements that cannot be repudiated. As the European Commission [21] notes, "Digital signatures could even bring significant law enforcement benefits as they allow for example messages to be attributed to a particular reader and/or sender." In a similar manner, anyone who gets to see a digital certificate, by wiretapping a communication or by consulting a certificate revocation list (CRL), has convincing evidence that the identity and any other information signed by the CA belong together. Even outsiders can obtain this information, by retrieving certificates from mail servers and other public repositories. This leads to all sorts of risks, including legislative risks; see, for instance, Bohm, Brown, and Gladman [7].

(Discrimination) Certificate holders cannot control which databases are consulted by verifiers. As a result, certificate holders subject themselves with each interaction to the possibility of being discriminated against on the basis of data that is not relevant to the situation at hand.

⁴Parties that actively monitor the Internet and other telecommunication infrastructures, or that have the capability to do so, include intelligence agencies and Internet service providers. The U.S. Communications Assistance for Law Enforcement Act [26, 27, 28, 49] and similar legislation in other countries require the telecommunications industry to build wiretapping capabilities into their infrastructures.

⁵In systems where each key holder is assigned a globally unique identifier, it is easy to automatically link the actions of key holders across different PKIs. The use of local names makes it more difficult to link transactions across different PKIs, but with today's network resources and linking power the barrier is very low. Moreover, different local names of the same individual can be linked when CAs cooperate.

The discrimination could go about without the certificate holder being aware of the source of the discrimination, the nature of the data used against him or her, or even the mere fact of the discrimination.

(Errors) There is no guarantee that the data that is consulted by verifiers is accurate. When data records do not represent an individual's true situation, eligible individuals may lose their insurances, loans, housing, jobs, good reputations, and so on. Data records may not reflect an individual's true situation for a variety of reasons: the data may be outdated, may have been composed by drawing incorrect inferences from other sources of data, may have been corrupted when transcribing information in oral form or in writing, may have been modified or destroyed by hackers and other outsiders, may have been modified or destroyed by authorized database users or unauthorized insiders, and may be the result of misattributions due to identity theft. Errors spread throughout systems and accumulate as data is disseminated and merged, and so victims may be affected by the same errors over and over again.

(Loss of control) Certificate users have no control over what happens with their data after their transactions come to an end. Typically, disclosed data ends up being used for purposes beyond the original purpose for which it was disclosed. These purposes need not be known at the time the data was collected; the mere existence of data often suffices to tempt organizations to use it in whatever way they see fit to suit their needs and desires. The public and private sectors will inevitably find new uses to improve the efficiency, security, or reach of their operations, to gain or maintain a competitive edge. Law enforcement agencies will seek access to the data believing that it will advance their investigative practices.

(Revocation issues) Many things can happen that require the revocation of a certificate. While revocation is an exceptional event, the task of verifiers to check the revocation status of unexpired certificates is not. Either they must download regularly a digitally signed update of a CRL, or they must resort to an online certificate validation service.

Since CRLs are distributed to all verifiers, and potentially to anyone who requests them, entities can collect data about key holders they have never communicated or transacted with. Furthermore, a CA can falsely add public keys to its CRL, to block the communications and transactions of targeted certificate holders.

Online certificate validation services are even worse. The provider of an online certificate validation service learns in real time who is communicating with whom and can falsely deny access to certificate holders. Also, by consulting validation services anyone can verify not only negative data about a certificate holder (i.e., revoked certificates) but also positive data (i.e., the fact that a certificate is in good standing).

Functional and Organizational Drawbacks of Identity Certificates

Identity certificates not only violate the privacy rights of individuals, they also pose many problems for certificate verifiers:

- The transaction process requires a sufficient delay to identify and correct frauds or other undesirable conditions. This may result in organizations not being able to serve as many customers as they could otherwise, or in customers leaving and going elsewhere (especially when browsing the Web).

- Because certificate holders are not guaranteed that their transactions will be authorized, there is significant uncertainty in the transaction process. Business may be lost on the basis of erroneous or irrelevant data, or simply because the online connection fails (e.g., due to peak load or a natural disaster).
- In case the representatives of an organization are spread out geographically, central database verification may be expensive (due to communication costs or the difficulty of dealing with peak load) or may simply not be an option because of the absence of network connections.
- Requests for central database look-up may be dishonored for many reasons and may be expensive (many large databases are operated by commercial organizations such as consumer reporting bureaus).
- It is difficult for organizations to protect their online databases against misuse by hackers and intrusion by insiders. This exposes organizations to incidents that might incur legal liability or damage their reputation.
- The trend is for governments to require organizations that maintain databases with personal data to adhere to (legal or self-enforced) privacy standards. This adds significant compliance costs.
- The possession of data about the personal preferences and lifestyle of individuals enables organizations to discriminate against their customers in a variety of ways. This increases the likelihood of complaints and encourages legislative action.
- Distribution of CRLs (or their updates) creates a lag between the time a certificate becomes invalid and when it appears on the next CRL update. If validity periods are long, CRLs will grow and additional computing resources are needed for searching and storing them. Online certificate validation avoids the need for verifiers to manage their own versions of a CRL and to deal with certificates they are not interested in, but is costly, hard to scale to large communities, and suffers from all the security problems of the central database paradigm. Either way, certificate revocation seriously reduces the finality of transactions.

Identity certificates are also problematic for *CAs*:

- It is often much harder, more error-prone, and more costly for a *CA* to establish a person's identity than to establish privileges and other non-identity attributes. In PKIs where organizations are interested only in non-identity attributes, not relying on identity can bring substantial savings in cost and time, and can reduce the risk of identity fraud.⁶
- It is inconvenient for millions of individuals to make a physical appearance before a *CA*. In November 1998, market researcher INTECO Corp. found that only 64% of Internet users would be willing to appear in person to have their identity verified for a digital certificate.
- Identity certification may expose a *CA* to potentially unbounded liability; see Bohm, Brown, and Gladman [7], Geer [35], Gladman, Ellison, and Bohm [38], Guida [41], Kaufman Winn [44], and Kaufman Winn and Ellison [45].

⁶For details on identity fraud, see Cavoukian [11], the Federal Trade Commission [30], the General Accounting Office [36], Givens [37], and the U.S. National Fraud Center [58].

3 Incorporating Privacy

“One of the greatest delusions in the world is the hope that the evils in this world are to be cured by legislation.”

— Thomas B. Reed, speech, 1886

In today’s computerized world, one cannot rely on the goodwill of others to protect one’s own privacy.

Why Technology?

Many countries attempt to protect privacy through the adoption of privacy legislation. While legislation discourages systematic abuse by the private sector, it is insufficient in most other respects:

- Legislation does not deter criminals who commit their crimes using computers: in today’s networked environment they can operate over large distances and across jurisdictional boundaries, making it difficult to enforce laws and to prosecute suspects. Global harmonization of laws is a daunting task that may never be accomplished at all due to cultural differences.
- Attempts to stop organizations from following privacy-invading practices conflicting with privacy legislation can easily take many years.
- It can be very hard for data collectors and auditors to determine whether privacy legislation has been lived up to.
- New technologies develop much faster than law. With each new consumer technology it can take many years to understand its privacy implications and develop adequate policies.
- Laws cannot protect against theft or modification of personal data stored in computer databases by hackers, nor can it protect against misuse by employees and other individuals authorized to access databases.
- Even in democratic societies laws may be amended, changed, exempted, overturned, or simply ignored.
- To check compliance with privacy legislation it is unavoidable that regular audits be performed on the databases of organizations; this increases the accessibility of personal data records, which broadens the scope for abuse.

The approach of privacy through self-regulation, promoted most notably in the United States, is even less effective:

- Marketers and other data miners have major commercial incentives to use personal data in all manners they see fit. Large profits can be made by using and selling consumer profiles; not using collected personal data for new business purposes is perceived as a serious waste of resources.
- Self-regulation requires the participation of the entire industry, but this is an unrealistic goal: new companies may prefer to not comply, companies that agree to guidelines may in fact not comply with them, and there will always be short-term incentives to ignore voluntary privacy measures.

- Organizations can rarely be held liable to compensate consumers for damages caused by the misuse of personal data, not in the least because individuals often have no clue as to the origin of privacy breaches. This makes it very difficult to enforce self-regulation.

Surveys consistently confirm the ineffectiveness of self-regulation.⁷

Privacy legislation does have a place. It should be used to supplement privacy-protecting technologies, though, not as a replacement. Legislation is needed to set the boundaries of what kinds of personal data may be bartered for what purposes, under what circumstances (if any) a verifier may refuse access, on what grounds (if any) a CA may refuse applicants, and so on. Also, legislation should provide a right to judicial remedies and should facilitate the prosecution of fraudsters.

Privacy Goals

To empower individuals to control their own data, a number of privacy goals must be met:

(Anonymity) Anonymity serves as the base case for privacy: without the ability to remain anonymous, individuals cannot control their own privacy. Moreover, user identifiers more often enable than prevent fraud, as the dramatic rise in identity fraud over the past decade has made clear.⁸ In many environments, misuse can only be prevented through anonymity; disciplines such as treatments for medical conditions have long acknowledged this.

(Control) In many situations full anonymity is not beneficial to anyone. Surveys indicate that most individuals are willing to trade personal data depending on a number of factors, including the benefits they will receive in return. More importantly, often at least one of the parties to a transaction has a legitimate need to verify previous contacts, the affiliation of the other party to a group, the authenticity of personal data of the other party, the eligibility of the other party to perform certain actions, and so on.

(Unlinkability) Untraceability of an isolated transaction is not sufficient to prevent linking of different transactions that originate from the same individual. Without unlinkability, all an individual's past and future transactions become traceable as soon as the individual is identified in a single one of these.⁹ Without unlinkability, individuals cannot control how much data they actually disclose, since the aggregate information obtained by linking different transactions typically reveals much more than the information that was willingly disclosed in each separate transaction.

(No self-authenticating records) In case a certificate holder authenticates certificate showings by means of a digital signature, he or she leaves a permanent self-authenticating record that can be verified by anyone. This gives coercive powers to the receiver and anyone else who sees the signed statement. If transactions are untraceable, little harm may come from this, but there is always the possibility that signed data disclosed in one anonymous transaction is linked to later

⁷See, for instance, surveys conducted by or on behalf of the Center for Democracy and Technology [12], the Electronic Privacy Information Center [24, 25], the European Commission [53], the Federal Deposit Insurance Corporation [29], the Federal Trade Commission [31], and OMB Watch [50]. For other critiques, see the American Civil Liberties Union [4], Budnitz [9, 10], Clarke [20], Rotenberg [54, 55], and Varney [57].

⁸Since 1996, calls on identity theft have been the number one topic on the hotline of the U.S. Privacy Rights Clearinghouse.

⁹Identification can be very simple. By way of example, the call history of the holder of an anonymous prepaid phone card can often be determined because the home phone number is among the numbers dialed.

transactions in which an identifier is revealed. Individuals should be able to authenticate their messages and data without leaving a self-authenticating record.

(Smartcard implementations) Implementation in smartcards¹⁰ or other tamper-resistant devices offers many security benefits, and so there is a need to preserve the listed privacy properties in these devices.

Identity certificates do not meet any of these basic privacy requirements.

Individuals should not have to trust anything but their own computer and the software it runs. This trust can be leveraged by compact operating systems, open source code, independent code reviews, digital certification of source code and executables, open market availability of computing devices, and anti-virus software.

4 Private Credentials

“One of the major issues that may need to be tackled in the context of the information infrastructure is surveillance by either the public or private sectors, especially via transactional and life style data.”

— Organisation for Economic Co-operation and Development [52]

Zero-Knowledge Systems is devoted to developing privacy-enhancing products that enable individuals to decide for themselves how their personal data is collected and used, how it is modified, and to which extent it can be linked or traced. These products do not address privacy as an afterthought, but treat it as a primary design objective.

Central to our approach are Private Credentials. Conceptually, Private Credentials are similar to blind signatures, first proposed by Chaum [16, 17]. Blind signatures and the constructs Chaum built from them (see [14]) have been instrumental in providing insight in how electronic privacy may be achieved. However, they have serious drawbacks that render them inadequate to overcome the problems of identity certificates; for details, see the Appendix.

We now explain how our Private Credentials achieve all the privacy properties discussed in Section 3, and more.

Issuing and Showing

Two basic stages can be discerned in the life cycle of a Private Credential: an issuing protocol and a showing protocol. Their characteristics are as follows:

(Issuing protocol) In the issuing protocol, an individual obtains from a *CA* a public key (for which the individual knows a secret key) and a digital signature of the *CA*. The *CA*'s signature binds the individual's public key to one or more *attributes*. An attribute is any type of information specified in a standard format. The whole package is called a Private Credential. A Private

¹⁰A smartcard is a plastic card with the shape and thickness of a conventional credit card, containing one or more embedded integrated circuits including a microprocessor capable of making arithmetic decisions. Smartcards can process data in intelligent manners, by taking actions based on secret data that never needs to leave the card. Memory access and input/output are guarded against unauthorized access, and the card can disable itself after a false PIN has been entered several times. Tampering with a smartcard in order to get to its contents can set off an alarm in the card that blocks it or overwrites the memory contents with all zeros.

Credential can specify any number of attributes. For instance, a “demographic” credential can specify its owner’s age, income, marital status, and residence, all neatly tied to a single public key, by means of one digital signature of the *CA*.

Although the sequences of zeros and ones that make up the individual’s public key and the *CA*’s signature are unique for each Private Credential issued, the *CA* cannot learn who obtains which sequences. At the same time, the individual cannot prevent the *CA* from encoding the attributes into the Private Credential, or more precisely, into his or her secret key. An explanation of how these two properties are achieved involves mathematics beyond the scope of this paper; the interested reader is referred to Brands [8].

(Showing protocol) To show a Private Credential to a verifier, the Private Credential holder sends the public key and the *CA*’s digital signature to the verifier, and uses the secret key to authenticate a message. The message must include a “nonce” (such as a number unique to the verifier concatenated to a random number chosen by the verifier, an indication of time and date, or a sequence number), to prevent the verifier or a wiretapper from “replaying” the Private Credential. By ensuring that the message is unique, knowledge of the Private Credential’s secret key (which the Private Credential holder never reveals) is required whenever the Private Credential is shown to a verifier.

Different actions involving the same Private Credential can be linked, on the basis of the unique bit sequences that make up the public key or the *CA*’s signature, but they cannot be traced to the Private Credential holder. Linkability is fine in applications where pseudonymity or the ability to build reputations are important or not considered a threat to privacy. To avoid linkability, an individual must use different “copies” of the same Private Credential. That is, instead of the *CA* issuing a single Private Credential to the individual, it should issue multiple Private Credentials that all specify the same attributes. With today’s computers and electronic networks, retrieving several copies of the same Private Credential is hardly less efficient than getting a single one, and one can always reconnect with the *CA* to get a new batch.

Attributes may be distributed across multiple Private Credentials, possibly issued by different *CAs*. This helps avoid the aggregation of an individual’s attributes by a single *CA*, improves efficiency, and removes the need to update Private Credentials more frequently than otherwise needed. As an example of the latter, if a Private Credential specifies the age of its holder, its expiry date can be no later than the person’s next birthday. In this particular case, the problem can be removed by encoding the date of birth instead of the age, but this is not always possible.

Verifiers need not consult databases containing possibly incorrect personal data, and (assuming Private Credentials are shortlived) there is no need to check the validity status of Private Credentials.

The tools that Zero-Knowledge Systems has set out to develop enable anyone to be an issuer of Private Credentials, but of course it is up to verifiers to decide whether to trust an issuer.

Private Credentials by themselves do not protect against wiretapping and traffic analysis. Wiretapping and traffic analysis in the physical world may be inherently difficult. On networks such as the Internet, one can transmit from a computer that is part of a network located behind a firewall, deploy anonymous remailers, pseudonymous remailers, or Mixmaster remailers (see Goldberg, Wagner, and Brewer [40] for an overview), or use pseudonymous services such as Freedom [39].

Below, we describe other Private Credential properties and give examples of the applications they enable. Again, we refer the reader to Brands [8] for a discussion of how these properties are achieved mathematically.

Selective Disclosure

Private Credentials are not only much more secure and efficient than their physical counterparts, but also much more powerful. Namely, the holder of a Private Credential can show any part of its attribute data without revealing any other information. For instance, the holder of a demographic Private Credential can reveal his or her residency when requesting access to a county database over the Internet, while hiding all other attribute information. This is analogous to having a paper certificate specifying demographic data, and using a marker to cross out any data one does not want to disclose.

Our Private Credentials are much more powerful than paper-based certificates. For instance, after having crossed out data fields on the paper-based certificate those fields can no longer be disclosed when using the certificate another time; Private Credentials do not have this limitation. Also, the class of properties one can selectively disclose is much larger than what can be done with a paper-based certificate and a marker. By way of example, Private Credentials allow an individual to buy over the Internet a discounted pass for minors and senior citizens without revealing whether he or she is a minor or a senior. More generally, Private Credential holders can demonstrate attribute formulas that combine linear relations by zero or more of the logical operators “AND,” “OR,” and “NOT;” any other information about the attributes remains unconditionally hidden. Consider, for instance, a Private Credential containing four attributes, each equal to either zero or one to indicate the absence or presence of a certain quality. To demonstrate meeting either all qualities or none of them, the Private Credential holder proves that the four attribute values sum up to four *or* to zero.

Also, Private Credential holders can prove that an attribute’s value lies in a range, without revealing any other information about the value.

The ability to selectively disclose information applies not only to the attributes contained in a single Private Credential, but also to attributes in different Private Credentials (even if they have been issued by different CAs). For instance, different Private Credential holders can jointly demonstrate that their combined attributes meet certain properties without having to pool or lend their Private Credentials.

Reissuance

In many cases one’s right to access a service comes from a pre-existing relationship in which identity has already been established. Our techniques enable the CA to refresh a previously issued Private Credential without knowing the attributes it contains. The attributes can even be updated before the Private Credential is recertified. One application of this is to prevent the CA from learning all of an individual’s attributes. Different CAs can even certify different attributes for the same Private Credential.

By way of example, a doctor could issue a prescription to a patient for 20 doses of a penicillin cure. Each time the patient visits a drugstore to collect some of the doses, the drugstore can verify that the patient is still eligible and can decrement the number of remaining penicillin doses. On the other hand, no drugstore can determine the total number of doses prescribed or the number remaining at the time of a visit, nor can different visits by the same patient be linked. (Using our Private Credentials, the patient could even pay for each dose in untraceable electronic cash, and receive a digital receipt that can be used to get reimbursed by his or her health insurance company.)

Dossier-Resistance

A Private Credential can be presented to an organization in such a manner that the organization is left with no mathematical evidence at all of the transaction. This is like waving a passport when passing

customs (a practice customary in several countries).

Alternatively, a Private Credential can be shown in such a manner that the verifier is left with self-authenticating evidence of a message or a part of the disclosed property. This is like presenting a paper-based certificate so that a photocopy can be made, but only after having crossed out certain data fields between the moment of showing the certificate and the moment of making the copy. In applications where organizations submit all the Private Credentials they receive to a central authority, to enable the latter to compute statistics or to combat fraud, this property prevents the central authority from learning which information an organization's customers have disclosed. At the same time, organizations cannot provide false information to the central authority. In case of a dispute, the disclosed property can always be revealed in full (possibly only with the cooperation of all parties).

Furthermore, the self-authenticating evidence can be limited to designated parties. This has no obvious paper-based analogue.

5 Fraud Prevention

Private Credentials do more than protect privacy. They greatly reduce the risk of identity fraud, simply because identity is not the basis for conducting transactions and communications. Also, the verifier does not need to consult central databases to look up the data it needs in order to decide whether to provide a service, because all the relevant data can be specified by the Private Credentials themselves; this eliminates the security risks of central databases.

The security provided by Private Credentials goes much further, though. This section describes techniques to discourage copying, reuse, lending, and discarding of Private Credentials. These techniques do not rely on tamper-resistant devices or online Private Credential validation. Tamper-resistant devices can offer an additional line of defense, as we will see, but if they are compromised the security of the software-only methods still holds.

Software-Only Techniques for Discouraging Copying and Reuse

Private Credentials can be constructed in such a way that a central authority can compute all the attributes encoded into a Private Credential once that Private Credential is shown more than a predetermined number of times. (Alternatively, copying and reuse can be prevented by resorting to online Private Credential validation by a central party, but this poses a serious performance bottleneck.) These so-called limited-show Private Credentials have no obvious paper-based analogue. The limited-show property holds even when Private Credential holders are free at each occasion to choose the attribute property that they demonstrate; in particular, a built-in identifier can be computed even if a fraudulent Private Credential holder in separate showings of the same Private Credential never discloses any attribute information. The threshold can be set arbitrarily, and the fraud evidence can be obtained in the form of a self-signed confession from the individual to whom the Private Credential was issued. This confession can be made unconditionally convincing, so that not even parties with unlimited computing resources can frame a Private Credential holder.

Figure 2 illustrates the process. It shows a malicious user, named John Smith, using the same one-show Private Credential at two different organizations. In this example, John Smith at the medical office hides all but his marital status, and at the business organization hides all but his citizenship. Each of the two organizations on its own cannot find out any additional information. The organizations submit their transcripts to a central double-show database, either at the time of the transaction or later

on, depending on their own policies and those of the database maintainer. The database maintainer can detect the fact of reuse once it receives both transaction transcripts, and from the two digital signatures the user provided in the transactions it can infer all the attributes in the Private Credential, in particular John Smith’s name.¹¹

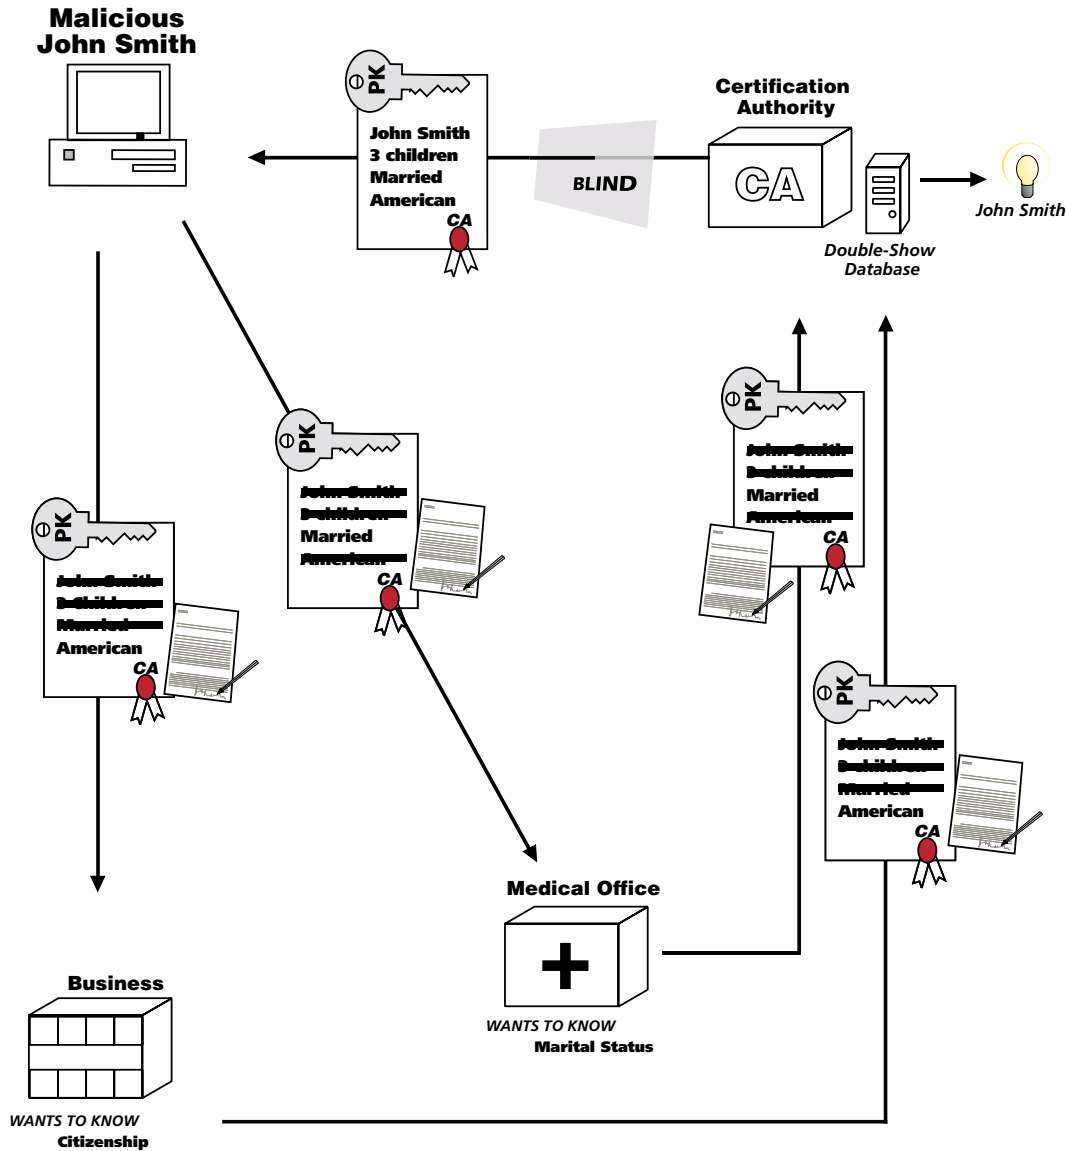


Figure 2: Tracing a fraudster.

Limited-show Private Credentials are highly practical: to be able to compute a built-in identifier in case of fraud, the central authority must store a “footprint” of a mere 60 bytes for each Private Credential shown, regardless of the complexity of the property disclosed and the number of encoded attributes.

¹¹Technically, this is accomplished by encoding the information as the slope of a line in the plane. While the Private Credential itself hides this line, each digital signature that the user makes reveals a random point on the line. One digital signature reveals no information about the slope of the line, but any two points uniquely specify the line.

An obvious use of limited-show Private Credentials is to permit a pre-specified number of uses for tickets, public transport tokens, and so on. Later on we will see that limited-show Private Credentials offer additional security benefits.

Software-Only Techniques for Discouraging Lending

A powerful measure to discourage lending of Private Credentials is to encode into each Private Credential an attribute that the holder wants to remain confidential. An issuer of gender credentials, for instance, could encode into each Private Credential not only a bit indicating the gender of the designated receiver, but also the credit card number or other secret information of the receiver. (The *CA* does not need to know the receiver's secret to be able to encode it.) While the Private Credential holder can hide the secret when showing the Private Credential, it is not possible to show the Private Credential without knowing the secret. Namely, the attributes are a part of the Private Credential holder's secret key, and the secret key is needed to authenticate the verifier's message in the showing protocol. Consequently, lending requires the legitimate Private Credential holder to reveal the entire secret key, which includes all the attributes. Note that X.509 and other identity certificates do not offer software-only protection against lending.

In addition, the *CA* could issue Private Credentials in the form of limited-show Private Credentials. To limit linkability, as we have seen, Private Credential holders must use their Private Credentials only a limited number of times anyway, and the *CA* might as well exploit this by issuing limited-show Private Credentials with built-in identifiers. This subjects a lender to the risk that the borrower uses his or her Private Credential more times than allowed, which would result in the lender being traced; the *CA* can then deny the lender further access to the system, or take other measures. It also reduces the number of times the legitimate holder can continue to use the Private Credential him or herself. (An access control mechanism based on a password, PIN, or biometric can protect against theft of Private Credentials.)

This leaves open the possibility of remote lending: the holder of a Private Credential can provide the "borrower" with a certified public key, and assist in showing it by providing the required responses over a radio link or a network such as the Internet. In communications or transactions that are not face-to-face, remote lending cannot be prevented, regardless of whether Private Credentials or privacy-invading identity certificates are used. Indeed, the "lender" might as well perform the showing and simply relay the provided service or goods to the "borrower." In many PKIs this is not a security problem, especially not if lending occurs only incidentally. In PKIs where large-scale lending is to be discouraged, the *CA* could issue limited-show Private Credentials and establish one account per Private Credential applicant. Large-scale lenders will then be exposed because of their abnormal demand for Private Credentials. Another measure is to charge a fee per Private Credential issued, so that a large-scale lender faces the problem of being compensated for the cost of his or her service without risking exposure. A further measure to discourage large-scale lending is to program smartcards in such a way that they act only when a built-in biometric access control mechanism detects the live presence of the cardholder.

Software-Only Techniques for Discouraging Discarding

To discourage Private Credential holders from discarding Private Credentials that encode unfavorable attributes, the *CA* can encode favorable attributes into the same Private Credentials. Information on

late payments could be encoded into a membership Private Credential for a health club or the like, and marks for drunk driving into a driver's license Private Credential.

This measure does not work when the attributes encoded into the Private Credentials of a Private Credential applicant change over time and become less favorable to the applicant. To limit the ability of Private Credential holders to reuse old Private Credentials stored on a backup, the CA should encode short validity periods. In addition, the CA could issue only limited-show Private Credentials.

The strongest possible software-only protection is to issue only one-show Private Credentials, and to ensure that Private Credential holders cannot show more than one Private Credential at any time. This can be achieved by having the CA recertify (and update) one-show Private Credentials only at the time they are shown. This requires all verifiers to have an online connection with the CA.

Another approach is to rely on tamper-resistant devices such as smartcards. These can be programmed to show Private Credentials in the order in which they were retrieved, and can enter into a suspension or erasure mode in case of tampering.

Incorporating Smartcards

All our Private Credential techniques can be implemented using smartcards and other tamper-resistant devices. These devices can offer strong protection against loss, theft, extortion, lending, copying, and discarding of Private Credentials, and can restrain their holder from other undesired behavior.

Incorporating smartcards requires great care, since their very goal is to shield their internal operations from their holder. As a result, it is virtually impossible to verify that a card does not leak personal data, its card identifier, its access control code, its communication and transaction history, data from other applications running on the same device, and so on. Leakage may take place by emitting electromagnetic radiation that can be remotely captured and viewed, by sending out or receiving radio signals, by sending along additional data when engaging in a protocol, by encoding information in message fields or random numbers specified in the protocol, by timing the delay before transmitting a message, or by halting at a specific step of a protocol.

In recognition of this privacy threat, it is best to not use smartcards as stand-alone devices but in conjunction with user-controlled software-only computing devices. This is the most natural setting in many communication and transaction settings. For example, a smartcard can be used over the Internet only if it is connected to a desktop computer, a notebook, a handheld, a mobile phone, or some other device. By routing all the communications from and to the smartcard through a computer trusted by its user, the user's computer can prevent the smartcard from covertly sending out data, and can sift data before passing it on or halt a transmission in case data fields do not comply with protocol specifications. In addition, any data leakage by or to the smartcard can be blocked. Furthermore, it can be assured that the smartcard cannot learn the Private Credentials of its holder, the information encoded into the Private Credentials, or the properties disclosed when showing a Private Credential. The cardholder can even prevent the smartcard from developing random numbers and other information that would enable a CA to trace the cardholder's transactions once it gains access to the card's contents. These properties hold in the strongest possible sense, namely in the presence of CAs that have access to cryptographic backdoors and conspire with certificate verifiers.

This approach ensures that smartcards cannot be misused for the purpose of surveillance. It has many other advantages over the smartcard-only approach as well. Since smartcard holders can enter their password, PIN, or biometric using the keyboard of their user-controlled computer, and can read messages on their computer's display, fake terminal attacks can easily be prevented. Also, our techniques overcome the computation, communication, and storage burden that are one of the main

reasons why smartcard implementations are stalling worldwide; transactions with our Private Credentials can be completed within as little as 1/20-th of a second by a standard 8-bit smartcard processor, so that demanding applications are entirely feasible. Different PKIs can make use of the same smartcard without being able to interchange personal data (unless the cardholder consents).

Furthermore, the user's computer can keep transaction logs, and allow the user to review and print them. Special software could provide Private Credential management functions, such as helping the user track personal transactions and categorize them. The transaction data could automatically flow into a software program, so that the user need never key in any data.

For improved confidentiality, all data stored on the user's computer can be encrypted using a password-derived secret key.

Verifiers do not need tamper-resistant terminals issued by a CA, owing to the use of public-key cryptography. This makes it much easier to become an acceptor of Private Credentials. Moreover, anyone can make or provide their own enhancements for the user's computer and distribute them on the open market.

6 Outlook

"It is seldom that liberty of any kind is lost all at once."

— Hume

Privacy protection requires that each individual has the power to decide how his or her personal data is collected and used, how it is modified, and to which extent it can be linked; only in this way can individuals remain in control over their personal data. These goals can be achieved through the use of privacy-enhancing technologies that are entirely feasible and secure. When using our Private Credentials, organizations (including CAs and verifiers) cannot learn more about an (honest) Private Credential holder than what he or she voluntarily and knowingly discloses, even if they conspire and have access to unlimited computing resources. Moreover, different actions by the same Private Credential holder cannot be linked, unless Private Credentials are reused. Individuals can prevent secondary use of their data at all times and can ensure the validity, timeliness, and relevance of their data.

Private Credentials are beneficial in any authentication-based environment in which there is no strict need to identify individuals at each and every occasion. As we have seen, Private Credentials do more than protect privacy: they minimize the risk of identity fraud, and overcome many of the efficiency and security shortcomings of identity certificates. The only acceptable role for identity certificates is to facilitate registration with a CA in situations in which identification of certificate applicants is mandatory; this is similar to the way in which drivers' licenses and passports are often used to acquire permits and other kinds of authentication proofs. Even for this purpose Private Credentials can be used, though.

More generally, Private Credentials are not complementary to identity certificates, but encompass them as a special case. To encapsulate an X.509v3 certificate, for instance, the CA can issue a Private Credential that encodes two attributes: the first attribute is the certificate holder's X.500 name, and the second attribute is (a collision-intractable hash of) the concatenation of all the fields except the subject's X.500 name (i.e., the format version, serial number, the CA's certification algorithm identifier, the CA's X.500 name, the validity period, and the certificate holder's public key and the algorithm with which it is to be used). Additional attributes could be encoded to represent X.509v3 extensions or other data. When showing the Private Credential, the holder must disclose (the preimage of) the

second attribute, and may disclose information about the other attributes. To ensure that the second attribute does not serve as a unique identifier, the entropy of at least the X.509 validity period and any extension fields must be restricted, and the serial number should be set to zero or to a hash of the public key (verifiers can compute the hash themselves).

In other words, Private Credential systems can subsume systems based on identity certificates. While we do not suggest to replace all identity certificate systems right here and now, it is important to bring the unbridled spread of identity certificates and other privacy-invading approaches to a halt. Unfortunately, the benefits of protecting privacy by means of technological measures are not widely acknowledged. Several countries have drafted, or are in the process of drafting, legislation that requires public keys to be bound to true names or traceable pseudonyms. As Baker and Yeo [6] point out, “The effect of these provisions will be to make it more difficult, if not impossible, to establish the legal validity of non-identity certificates and to enforce transactions that are authenticated by non-identity certificates.”

The widespread adoption of automated communication and transaction systems that lack elementary privacy provisions is a very dangerous trend. As Swire [56] points out, “The systems in place in one period can have a powerful effect on what systems will develop in subsequent periods. [...] Once the costs of the database and infrastructure are already incurred for initial purposes, then additional uses may be cost-justified that would not otherwise have been.” Holmes [42] notes that “the danger to democratic countries is not that they will openly embrace totalitarianism. It is [...] that they will unwittingly, almost imperceptibly, and with the best of intentions, allow themselves to drift so far in that direction that the final step will then be but a small one.”

Today, the foundations for the communication and transaction technologies of this century are being laid. To avert the scenario of a global village founded wholly on inescapable identification technologies, it is imperative that those involved with the architecture of the information infrastructure rethink their preconceived ideas about security and identity—and build in privacy before the point of no return has been reached.

References

- [1] C. Adams and S. Farrell. Internet X.509 public key infrastructure certificate management protocols. Internet Draft of the PKIX Working Group, May 1998.
- [2] Carlisle Adams and Steve Lloyd. *Understanding the Public-Key Infrastructure*. New Riders Publishing (Macmillan Technology Series), November 1999. ISBN 157870166X.
- [3] American Bankers Association. X9.45-199x: Enhanced management controls using digital signatures and attribute certificates. Working draft, June 1997.
- [4] American Civil Liberties Union. Elements of effective self regulation for the protection of privacy and questions related to online privacy. Letter to Ms. Jane Coffin, Office of International Affairs, National Telecommunications and Information Administration, July 1998.
- [5] American National Standards Institute. American National Standards Committee X.9.55-1995: Public key cryptography for the financial services industry, 1995.
- [6] Stewart Baker and Matthew Yeo. Survey of international electronic and digital signature initiatives. Steptoe & Johnson LLP, Internet Law and Policy Forum, version of April 14, 1999.

- [7] Nicholas Bohm, Ian Brown, and Brian Gladman. Electronic commerce: Who carries the risk of fraud? Foundation for Information Policy Research, July 2000.
- [8] Stefan A. Brands. *Rethinking Public Key Infrastructures and Digital Certificates*. MIT Press, Cambridge, Massachusetts, August 2000. ISBN 0-262-02491-8.
- [9] Mark E. Budnitz. Industry self-regulation of internet privacy: The sound of one hand clapping. Computers, Freedom & Privacy 1999, April 6–8, Washington DC.
- [10] Mark E. Budnitz. Privacy protection for consumer transactions in electronic commerce: Why self-regulation is inadequate. 49 S. Caro. L. Rev. 847, 1998.
- [11] Ann Cavoukian. Identity theft: Who’s using your name? Information and Privacy Commissioner of Ontario, Canada, June 1997.
- [12] Center for Democracy and Technology. Policy vs. practice; a progress report on federal government privacy notice on the world wide web, April 1999.
- [13] B. Chalks. Privacy enhancement for Internet electronic mail – part IV: Key certification and related services. RFC 1424-C, February 1993.
- [14] D. Chaum. Achieving electronic privacy. *Scientific American*, 267(2):96–101, August 1992.
- [15] D. Chaum, A. Fiat, and M. Naor. Untraceable electronic cash. In S. Goldwasser, editor, *Advances in Cryptology–CRYPTO ’88*, volume 403 of *Lecture Notes in Computer Science*, pages 319–327. Springer-Verlag, 1988.
- [16] David Chaum. Blind signature system. In D. Chaum, editor, *Advances in Cryptology–CRYPTO ’83*, page 153, New York, 1984. Plenum Press.
- [17] David Chaum. Security without identification: Transaction systems to make Big Brother obsolete. *Communications of the ACM*, 28(10):1030–1044, October 1985.
- [18] David Chaum. One-show blind signature systems. U.S. Patent ser. no. 4,987,593, January 1991. Filed April 1990. Continuation of abandoned application Ser. No. 07/168,802, filed March 1988.
- [19] S. Chokhani and W. Ford. Internet public key infrastructure certificate policy and certification practices framework. Internet Draft of the PKIX Working Group, work in progress, September 1997.
- [20] Roger Clarke. Internet privacy concerns confirm the case for intervention. *Communications of the ACM*, 42(2), February 1999. Version of 14 October 1998.
- [21] Directorate-General XIII of the European Commission. Ensuring security and trust in electronic communication; towards a European framework for digital signatures and encryption. Communication to the European Parliament, the Council, the Economic and Social Committee and the Committee of the Regions. COM (97) 503, October 1997.
- [22] S. Dusse, P. Hoffman, B. Ramsdell, and J. Weinstein. S/MIME version 2 certificate handling. Network Working Group, Request for Comments no. 2312, March 1998.

- [23] S. Dusse, P. Hoffman, R. Ramsdell, L. Lundblade, and L. Repka. S/MIME version 2 message specification. Network Working Group, Request for Comments no. 2311, March 1998.
- [24] Electronic Privacy Information Center. Surfer beware: Personal privacy and the Internet, June 1997.
- [25] Electronic Privacy Information Center. Surfer beware II: Notice is not enough, June 1998.
- [26] Electronic Surveillance Task Force. Communications privacy in the digital age. Interim Report of the Digital Privacy and Security Working Group, June 1997.
- [27] Federal Bureau of Investigation. The Digital Telephony and Privacy Improvement Act, March 1994.
- [28] Federal Bureau of Investigation. The Digital Telephony and Privacy Improvement Act (update), June 1994.
- [29] Federal Deposit Insurance Corporation. Online privacy of consumer personal information, August 1998.
- [30] Federal Trade Commission. Consumer Identity Fraud meeting. Official Transcript Proceedings before the Federal Trade Commission, August 1996. Washington, D.C.
- [31] Federal Trade Commission. Privacy online: A report to Congress, June 1998.
- [32] Jalal Feghhi, Peter Williams, and Jalil Feghhi. *Digital Certificates : Applied Internet Security*. Addison-Wesley, October 1998. ISBN 0201309807.
- [33] Warwick Ford and Michael Baum. *Secure Electronic Commerce : Building the Infrastructure for Digital Signatures and Encryption*. Prentice Hall, April 1997. ISBN: 0134763424.
- [34] Alan O. Freier, Philip Karlton, and Paul C. Kocher. The SSL protocol, version 3.0. Internet draft, Netscape Communications, November 1996.
- [35] Dan Geer. Risk management is where the money is. *Forum on Risks to the Public in Computers and Related Systems, ACM Committee on Computers and Public Policy*, 20(6), November 1998.
- [36] General Accounting Office. Identity fraud: Information on prevalence, cost, and Internet impact is limited. Briefing Report, May 1998. GAO/GGD-98-100BR.
- [37] Beth Givens. Identity theft – how it happens, its impact on victims, and legislative solutions. Presentation of the Privacy Rights Clearinghouse, May 1997.
- [38] Brian Gladman, Carl Ellison, and Nicholas Bohm. Digital signatures, certificates & electronic commerce. Version 1.1, June 1999.
- [39] Ian Goldberg and Adam Shostack. Freedom network 1.0 architecture. Zero-Knowledge Systems, Inc. white paper, November 1999.
- [40] Ian Goldberg, David Wagner, and Eric A. Brewer. Privacy-enhancing technologies for the Internet. In *COMPCON '97*. IEEE, February 1997.

- [41] Richard A. Guida. Truth about PKI isn't always common knowledge. GCN Spotlight, May 3, 1999.
- [42] Robert Holmes. Privacy: Philosophical foundations and moral dilemmas. In *Proceedings of the 16th International Conference on Data Protection–Facing Dilemmas*, September 1994.
- [43] International Telecommunication Union. ITU-T recommendation X.509, information technology – open systems interconnection – the directory: Authentication framework, June 1997.
- [44] Jane Kaufman Winn. Couriers without luggage: Negotiable instruments and digital signatures. *South Carolina Law Review*, 49(4), 1998.
- [45] Jane Kaufman Winn and Carl Ellison. Regulating the use of electronic authentication procedures by US consumers in the global electronic marketplace. Comment P994312 to the Federal Trade Commission, March 1999.
- [46] S. Kent. Privacy enhancement for Internet electronic mail – part II: Certificate-based key management. RFC 1422, February 1993.
- [47] Jamie Lewis. Public Key Infrastructure architecture. The Burton Group, Network Strategy Report, July 1997.
- [48] MasterCard & Visa. SET secure electronic transaction specification, version 1.0. Book 1: Business Description, Book 2: Programmer's Guide, Book 3: Formal Protocol Definition, May 1997.
- [49] Office of Technology Assessment. Electronic surveillance in a digital age, July 1995. OTA-BP-ITC-149 (Washington, DC: U.S. Government Printing Office).
- [50] OMB Watch. A delicate balance: The privacy and access practices of federal government World Wide Web sites, June 1998. ISBN: 1882583-12-4.
- [51] Organisation for Economic Co-operation and Development. Inventory of approaches to authentication and certification in a global networked society. Document of the Working Party on Information Security and Privacy, October 1999.
- [52] Organization for Economic Co-Operation and Development. Report of the ad hoc meeting of experts on information infrastructures issues related to security of information systems and protection of personal data and privacy, 1996. Paris, General Distribution OCDE/GD(96).
- [53] Charles D. Raab, Colin J. Bennett, Robert M. Gellman, and Nigel Waters. Application of a methodology designed to assess the adequacy of the level of protection of individuals with regard to processing personal data: test of the method of several categories of transfer - final report. Study carried out for the European Commission. Tender No. XV/97/18/D, September 1998.
- [54] Marc Rotenberg. Communications privacy. Prepared Statement Before the Subcommittee on Courts and Intellectual Property of the House Committee on the Judiciary, United States of House of Representatives, Washington, D.C., March 26, 1998.
- [55] Marc Rotenberg. On the European Union Data Directive and Privacy. Testimony and Statement for the Record Before the Committee on International Relations, U.S. House of Representatives, May 7, 1998.

- [56] Peter P. Swire. Financial privacy and the theory of high-tech government surveillance. Draft of September 24, 1998.
- [57] Christine Varney. You call this self-regulation? Wired News, June 9, 1998.
- [58] Norman A. Willox. Identity theft: Authentication as a solution. National Fraud Center, Inc., March 2000.
- [59] Eric C. Zimits and Christopher Montano. Public Key Infrastructure: Unlocking the Internet's economic potential. *IStory*, 3(2), April 1998. The Hambrecht and Quist Internet Research Group.

Appendix: Blind Signatures

In a blind signature scheme (see Chaum [16, 17]), a receiver obtains a signed message from a signature issuer in such a manner that the issuer gains no information about which signed message (from the space of all possibilities) the receiver obtained. Blind signatures cannot provide the functionality and security of Private Credentials, though:

- The blinding prevents the CA from encoding attributes into a certificate. It is possible to represent each combination of attribute values by a different signing key of the CA, but this seriously limits the number of attributes that can be encoded and their value ranges. For instance, the CA cannot encode expiry dates into certificates; at best it can refresh its own signing key, and declare in advance when the entire current issuance of certificates will become invalid.
- A blind signature does not enable one to negotiate a degree of privacy, since verifiers need to know which public key of the CA to apply to verify a certificate. It is not possible, for instance, to encode identifiers that must be disclosed in certain circumstances but may remain hidden in others.
- Certificates that contain unfavorable attributes (i.e., attributes that the holder would rather like to hide, such as a criminal record) can be discarded by their holder. In most applications it is impractical to require all certificate holders to obtain and show attributes that indicate the absence of unfavorable attributes.
- An extortioner can force a certificate applicant to withdraw blind signatures for which the extortioner provides the blinding factors. Because proximity of the extortioner and his or her victim is not required at any time, the extortioner can remain anonymous throughout. The extortioner can use the credential at any convenient time later on, while remaining untraceable.
- Blind signatures cannot prevent or discourage unauthorized lending of certificates.
- Chaum's smartcard techniques require smartcards with large memory and cryptographic coprocessors, to guarantee that the required operations can be performed in reasonable time. More seriously, since attributes are encoded by the smartcard rather than by the CA, physical compromise of a smartcard enables its holder to forge attributes, and to lend out, give away, or distribute copies of new certificates. The CA cannot trace fraud, and containment can only be accomplished by suspending the entire system.

A related technique proposed by Chaum [15, 18], called one-show blind signatures, is not effective either. It makes high computation and communication demands on both the issuer and the receiver, does not extend to limited-show certificates, does not admit zero-knowledge proofs, and cannot be migrated to a setting with smartcards without further degrading efficiency. Also, it does not give certificate holders the ability to selectively disclose information about attributes.