Digital Credentials

Dr. Stefan Brands
brands@credentica.com
Credentica Inc.
Version 2, January 2003

1. Introduction

Often at least one of the parties in a transaction needs to know whether the other party is authorized to perform a certain action. Typically, authorization is granted on the basis of a person's privileges, personal characteristics, reputation, identity, group membership, willingness to provide value in exchange, and so on. In all these cases, the verifying party must rely on the inspection of one or more tangible objects issued by trusted third parties. Examples of such *credentials* are coins and banknotes, stamps, medical prescriptions, cinema tickets, voting ballots, membership cards, access tokens, diplomas, passports, and drivers' licenses.

Physical credentials are increasingly prone to counterfeiting, however, and are unsuitable for use over mobile networks, the Internet, and other electronic environments. To overcome these fundamental drawbacks, a transition to a fully digital form of credentials is inevitable. Section 2 of this paper provides an overview of the most secure and flexible technology for constructing digital credentials to have come out of the cryptographic research community thus far. Section 3 examines an alternative approach, based on digital identity certificates, and shows why this naïve approach is fundamentally insecure, does not scale, and invades privacy.

2. Digital Credentials

Digital Credentials are basic cryptographic constructs, much like digital signatures but much more powerful. They are issued to applicants by trusted parties, referred to as Credential Authorities. Each Credential Authority has its own key pair for digitally signing messages. When issuing a Digital Credential to Alice, the issuing Credential Authority through its own digital signature binds one or more *attributes* to a Digital Credential public key, the secret key of which only Alice should know. The whole package that Alice receives is called a Digital Credential; for instance, a government agency could issue to Alice a single credential that specifies her name, number of kids, marital status, and citizenship, all neatly tied to a single public key by means of one digital signature of the government agency. Although the sequences of zeros and ones that make up the Alice's public key and the signature of the Credential Authority are unique for each Digital Credential issued, the Credential Authority cannot learn who obtains which sequences; they are *blinded* by Alice during the issuing process.

Since a Digital Credential is just a cryptographically protected sequence of zeros and ones, it can be transferred over electronic networks and by smartcards, and can be verified with 100 percent accuracy by computers. To show her Digital Credential to Bob, Alice sends her Digital Credential public key and the signature of the Credential Authority. She also digitally signs a nonce, using her secret key. A nonce is a random number, the concatenation of Bob's name and a counter, or any other fresh data provided by Bob. Bob cannot replay Alice's information for his own benefit in another transaction, since in each showing protocol execution a new nonce must be signed; this requires knowledge of Alice's secret key, which never leaves Alice's device. At the same time, Alice can selectively disclose to Bob a property of the attributes in her Digital Credential, while hiding any other information about them. Alice's signature on Bob's nonce doubles up as a proof that the disclosed property indeed holds true.

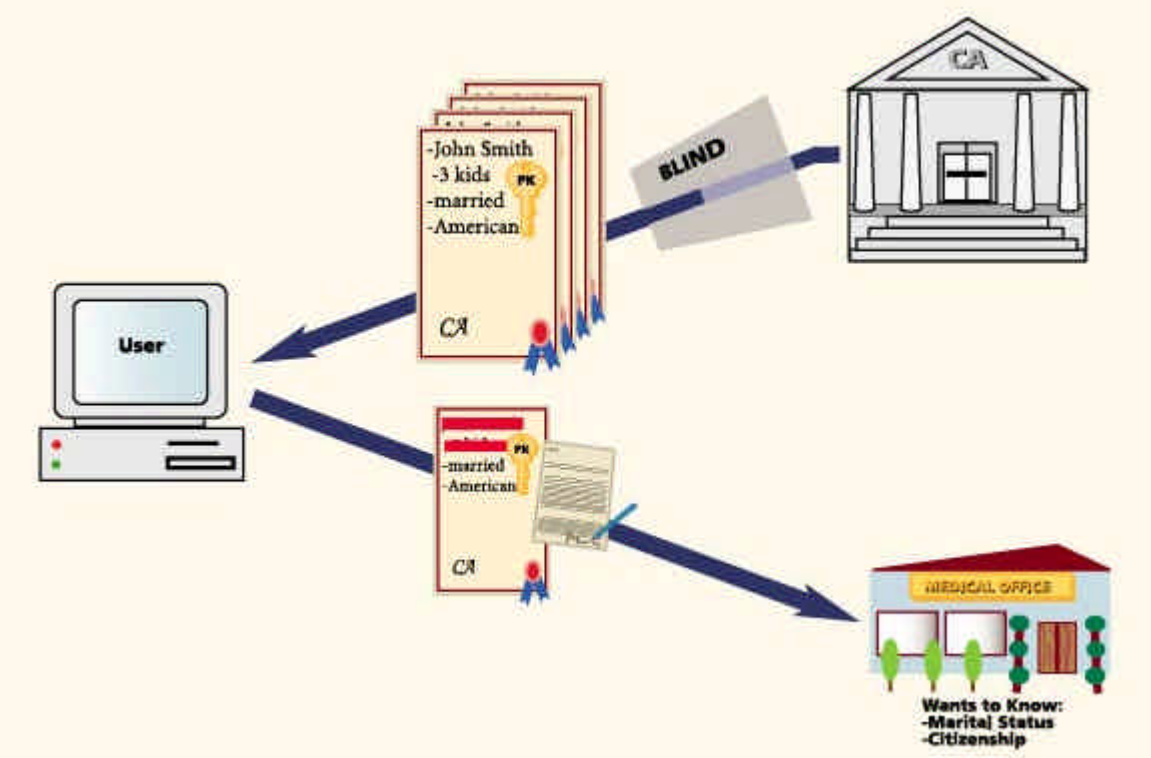

Issuing and showing a Digital Credential

The class of properties that a Credential holder can selectively disclose is much larger than what can be done with a paper-based certificate and a marker. For instance, the holder of a Digital Credential that specifies her exact age can prove eligibility for a discount pass without revealing whether anything more about her age beyond that she is either a minor or a senior.

A detailed description of how these basic properties are achieved in a highly practical manner is outside the scope of this paper. A technical overview can be found in [1], and the full details are described in a book published by MIT Press [2]. As explained in these references, by carefully exploiting the basic properties of Digital Credentials, one can realize all the following features:

- **(Privacy)** Digital Credentials accommodate fully adaptable levels of privacy ranging from user-driven anonymity to government/enterprise-mandated identification. They support automated negotiation of credential information, ensuring that only the minimum credential information needed to meet the authorization requirements of the service provider is disclosed; this minimizes the risk of identity theft, and preserves privacy. The selective disclosure technique can be applied not only to attributes encoded into a single Digital Credential, but also to attributes in different Digital Credentials, possibly certified by different Credential Authorities.[1]
- **(Strong accountability)** Digital Credentials offer audit capability for non-repudiation and to assess compliance with regulatory requirements, through digital audit trails and receipts that facilitate automated dispute resolution. Malicious parties, including Credential Authorities, cannot frame the holder of a Digital Credential by making it look as if he or

[1] Rather than encoding many attributes into a single Digital Credential, it may be preferable to distribute them across multiple Digital Credentials. This helps avoid the aggregation of an individual's attributes by a single Credential Authority, improves efficiency when many attributes need to be encoded independently, and removes the need to update certificates more frequently than otherwise needed.

she participated in a transaction, even if they would have unlimited computing power or special knowledge of trapdoor information. Audit trails can be kept in the form of role-based digital signatures; in the case of abuse, the transaction originator will not be able to disavow having conducted the transaction.

- **(Pooling protection)** Different people can be prevented from pooling together multiple Digital Credentials in order to jointly obtain access rights that they do not enjoy separately. Hereto the service provider requires the access requestor to demonstrate that any Digital Credentials that he or she provides all contain the same built-in identifier. Owing to the selective disclosure property, an honest credential holder can demonstrate this without disclosing the identifier.
- **(Lending protection)** Lending of credential information can be discouraged by wrapping the information into a Digital Credential and encoding confidential data of the legitimate owner into it. The legitimate owner can hide this data (again owing to the selective disclosure property), but the Digital Credential cannot be used without actually knowing the confidential data. (Note that this measure does not require credential holders to use tamper-resistant hardware.)
- **(Discarding protection)** Digital Credentials can be used to prevent the discarding of authenticated information that a party would rather not show. A mark for drunk driving, for instance, can be tied into a driver's license Digital Credential that specifies that the holder is authorized to drive. Once again owing to the selective disclosure property, the owner can hide the mark whenever it need not be disclosed.
- **(Dossier-resistance)** A Digital Credential can be presented to an organization in such a manner that the organization is left with no evidence at all of the transaction (much like showing a passport without letting the other party make a photocopy) or such that the verifier is left with self-authenticating evidence of only a part of the disclosed property. Also, the self-authenticating evidence can be limited to designated parties. In case of a dispute, the disclosed property can always be revealed in full (possibly only with the cooperation of all parties).
- **(Secondary use control)** Similarly, credential verifiers can ensure that they are left with digital evidence that proves only what they want it to prove (which may be much less than what the Digital Credential holders themselves selectively disclosed to them). This property enables verifiers to submit, to a third party (for the purpose of fraud detection, status validation, or statistical data gathering), non-repudiable proofs related to Digital Credentials they has verified, while hiding any competitive or privacy-sensitive information that they learned from their clients.
- **(Limited-show credentials)** A limited-use Digital Credential can contain a built-in identifier, value token, or self-signed fraud confession, that will be exposed if (and only if) the Digital Credential is shown more than a predetermined number of times.[2] These *limited-show* Digital Credentials (which can be used to design stamps, coins, tickets, and so on) have no obvious paper-based analogue. The limited-show property holds even when Digital Credential holders are free at each occasion to choose the attribute property that they demonstrate. Limited-show Digital Credentials are highly practical: to be able to compute a built-in identifier in case of fraud, a footprint of a mere 60 bytes must be stored for each Digital Credential shown, regardless of the complexity of the property disclosed and the number of encoded attributes.

[2] Alternatively, copying and reuse can be prevented by resorting to online Digital Credential validation by a central party, but this may pose a serious performance bottleneck.

- **(Reverse authentication)** This property allows the holder of a Digital Credential to demonstrate that he or she is *not* someone listed on a blacklist, without enabling identification. More generally, the Digital Credential holder can demonstrate that the data in the Digital Credential does *not* meet certain conditions, without revealing more.
- **(Recertification and updating)** In many cases one's right to access a service comes from a pre-existing relationship in which identity has already been established. An individual can present a certified public key for recertification or for updating to a Credential Authority, without enabling it to learn the current values of the attributes in the Digital Credential. One particular use of this property is to enable multiple Credential Authorities to certify attributes within the same Digital Credential without knowing all the attributes.
- **(Information can reside anywhere)** Digital Credentials can be held both locally (on a device of the user) or remotely, and can be managed using roaming. In the extreme, organizations can do away entirely with central databases containing sensitive personal information, by securely distributing each database entry to the individuals to whom it pertains; the unique properties of Digital Credentials ensure that unauthorized users cannot modify, discard, pool, lend, or prevent updates of their own credential information.
- **(Smartcard Implementation)** Digital Credentials can be issued to, or embedded in, smartcards and other tamper-resistant devices; this provides a second layer of protection (on top of the cryptographic protections) against loss, theft, extortion, lending, pooling, copying, and discarding of Digital Credentials, and can prevent other kinds of unauthorized behavior. The storage and computational burden for the tamper-resistant device can be off-loaded almost entirely to another user device that need not be tamper-resistant (such as a handheld device with display and keypad, a laptop, or another chip on the same smartcard that need not be trusted by the system provider), while preserving all the smartcard's security benefits; literally billions of Digital Credentials can be securely managed in this manner using a single 8-bit smartcard chip.
- **(Secure multi-application smartcards)** Smartcards can be used as multi-application devices, without introducing any of the privacy and security problems caused by other technologies. Specifically, different application providers can all share the same secret key stored in a user's smartcard to derive the security benefits of that smartcard. The certificates will have uncorrelated secret keys which cannot be determined by anyone including the smartcard supplier, and all Digital Credentials can be revoked separately. The application software on the user's trusted computer ensures that smartcards attacks are impossible, and that different applications using the same smartcard remain fire-walled.
- **(Managed services)** With an increasing number of incompatible authentication mechanisms available, and network identities becoming federated instead of centrally stored, applications that need to make authority decisions will increasingly ask trusted authorities to issue and/or verify the credential information presented by their clients. With Digital Credentials, Credential Authorities can certify sensitive information on behalf or organizations without being able to learn that data, and Revocation Authorities can validate certificates (using OCSP or other standards) without being able to learn the identities of the clients of organizations (even when these expressly identify themselves to the organizations they transact with through the certificates themselves). In this manner, organizations can outsource core tasks related to digital authentication and authorization, without having to provide their managed services providers with competitive data or customer information for which they could incur legal liabilities. Even the role of the

tamper-resistant smartcard can be outsourced, thereby removing the logistical problem of securely distributing tamper-resistant devices to card holders.[3]

Some of these features are counter-intuitive, since they have no physical-world analogue. Depending on the application in which the Digital Credentials are used, one might decide to go with just a few of the listed features. Note that Digital Credentials support all the traditional authentication strengths, ranging from weak software-only protection to military-grade two-factor and three-factor security.

Since a few years, the Digital Credentials technology is being taught at, amongst others, MIT, Harvard law school, Carnegie-Mellon University, University of San Diego, John Hopkins university, Ecole Normale Superiéure (ENS Paris), Swiss Federal Institute of Technology (ETH in Zürich), Helsinki University of Technology (Finland), Leuven University (Belgium), Aarhus University (Denmark), and most of the leading technical universities in Germany. The practicality of Digital Credentials has been well-established as well. Notably, in 2001, Zero-Knowledge Systems in Montreal developed a wireless prototype for RIM's Blackberry as well several demo applications for Personal Computers. Also, from 1993 until 1999, CAFE and OPERA, two major European consortiums co-funded by the European ESPRIT program, implemented and extensively tested a smartcard-based payment system based on the Digital Credentials technology.

3. A Naïve Approach Based on X.509-Style Certificates & PKI

A digital identity certificate binds an individual's public key to his or her true name, Social Security number, or any other data that the Certificate Authority can readily associate with the individual. Numerous digital identity certificate standards exist based on the X.509 framework of the International Telecommunications Union. An infrastructure that revolves around the distribution and management of public keys and digital identity certificates is called a Public Key Infrastructure (PKI). Digital identity certificates were invented in 1978, at the dawn of modern cryptography, to enable the sender of a message to encrypt that message under the public key of the intended recipient, by binding that public key to the recipient's identity.

Similar to the way organizations today use Social Security Numbers, digital identity certificates can be used to manage Digital Credentials, by using them as authenticated pointers into central databases that contain all the credential information. However, this approach fundamentally provides poor security and performance, and has devastating consequences for privacy:

- **(Access right cloning and lending)** X.509-style PKI does not provide software-only protection to discourage certificate holders from transferring (copies of) their access rights and entitlements to other parties: a user's secret key is simply a random number, and so revealing it to someone else has no direct negative consequences for that certificate holder. This defeats the entire purpose of PKI in the context of access management. Even when secret keys are stored in smartcards, the break of a single smartcard suffices to bypass the security of the system. Hackers around the world are already cloning pay-TV smartcards, high-value phone cards, and debit cards, causing hundreds of millions of dollars in damages.
- **(Non-scalable)** The approach of using an X.509 certificate as an authenticated pointer does not scale beyond pre-established administrative domains, since the actual authorization decisions are left to the access provider. When information is shared amongst different trust domains, the parties that have to make authorization decisions may

[3] Although each and every transaction of a Digital Credential holder will now require the real-time involvement of a third party that guarantees protection of the user's secret key, that third party cannot learn any details that could lead to a privacy compromise (other than knowing the transaction times of pseudonymous users).

not have, or be able to access, all the data they need for making an authorization decision: the missing data may reside in databases outside of their control or be otherwise unavailable, and it is difficult to guarantee its completeness and correctness. With increasing numbers of individuals and organizations seeking to share resources, it rapidly becomes infeasible to guarantee the availability and correctness of the data needed to make authorization decisions, even when switching to low-grained role-based access control.

- **(Central point of attack)** By relying on on-line central databases that all service providers can refer, the door is pushed wide open to devastating abuses of security holes. It is difficult to protect online databases against misuse by hackers and insiders; a summer 2002 survey by Evans Data Corporation, for instance, revealed that of 700 database specialists surveyed, 20% have experienced a direct breach in their database security. According to a 2000 CSI/FBI computer crime survey, 71% of unauthorized break-ins are by corporate insiders. Furthermore, recorded data may be outdated and may be the result of misattributions due to identity theft.

- **(Identity theft)** Systems that systemically rely on user identification give rise to a fraud known as identity theft, whereby fraudsters assume the identities of their victims. According to the U.S. Federal Trade Commission (FTC), identity theft is the fastest growing crime in America, affecting approximately 900,000 new victims in 2001. The FTC expects that its cost will reach USD 8 billion by the year 2005. A recent study by the U.K. Department of Trade and Industry warns that in the not–too–distant future criminals will be as interested in stealing victims' identities as they are in stealing possessions. Notwithstanding the fact that X.509-style PKI provides for message encryption, it seriously increases the risk of identity theft, since its fundamental premise is that of inescapable system-wide identification. Criminals who manage to steal digital identity certificates or to assume the identities of unwitting people will be able to misuse certificates on a global scale, while their victims take the blame. The problem is compounded by the strong reliance on central databases, with all their vulnerabilities. This exposes organizations to potentially unlimited legal liability. (Legislation such as the U.S. E-Sign Law, passed in 2000, and the EU Digital Signature Law, passed in 2001, recognize digital signatures as legally binding.)

- **(Poor performance on low-cost devices)** In principle, both physical and logical security could be integrated by implementing PKI in smartcards or other portable devices (including Web-enabled cell phones and handheld computers). This would allow companies to migrate multiple disparate security systems into one integrated PKI system. However, in the words of the Aberdeen Group in a January 2001 white paper, CPU drain prevents X.509-style PKI from being "a solo building block." Indeed, the computational requirements of processing an X.509-style certificate are well beyond today's popular smartcards and devices. Addressing the problem by adding advanced circuitry (such as cryptographic co-processors) seriously increases the price of these devices. According to the PKI Forum in an April 2002 report, "Price competitiveness is the overriding driving factor for the card industry and will continue to commoditize the cards and components." Moreover, the addition of sophisticated circuitry can easily lead to new weaknesses in the internal defense mechanisms, and adversely affects reliability. These problems worsen in the case of multi-application smartcards, which are desirable to prevent system providers from needing separate card platforms for each individual application and to decrease the number of lost cards. Further compounding the problem is the unsuitability of X.509-style PKI for multi-purpose and multi-application certificates.

- **(Privacy violations)** According to an April 2002 report by Gartner, individuals distrust online authentication systems, their skepticism resting in great part on privacy concerns. In many contexts, to gain access to services and other resources the requestor would prefer to present just enough credentials to be granted access, as has traditionally been the case for the vast majority of transaction mechanisms in non-electronic environments. With X.509-style PKI, however, the real name of the requestor is systematically exposed. Numerous studies show that individuals are increasingly concerned about who has access to their personal information and how it might be used. Failure to protect privacy can damage an organization's reputation, brand image, and valuation. It can also lead to litigation, fines, criminal sanctions, and civil liability. Moreover, the lack of privacy has been shown to be the leading cause of losses in sales opportunities. The fundamental problem with X.509-style PKI is that public keys are globally unique identification numbers that they travel along with every action taken by system participants. They can be automatically linked to the identity of their owners by a myriad of parties, even if the owner's name is not explicitly stated in the certificate. In an attempt to hide the huge privacy problems created by X.509-style PKI, the PKI vendors misleadingly define privacy as "communications are safe from eavesdropping." This works in the context of wiring a message to an intended recipient, but in the context of access management, encryption has very little to do with privacy. Indeed, according to a January 2001 executive white paper by the Aberdeen Group, "as currently sold and implemented, PKI is incompatible with the coming privacy era [and] eliminates even the pretense of ensuring user privacy."
- **(Managed services are intrusive)** A survey released April 2002 by the McAfee security division of Network Associates showed that firms are holding back from outsourcing security primarily due to a strong reluctance to trust a third party. Indeed, there are serious reasons not to trust today's managed PKI services: the providers of online certificate validation services learn in real time the identities of their clients' customers, their peak hours, and other competitive information. Furthermore, Certificate Authorities must know the identity and any other attributes that go into the digital certificates they issue.
- **(Violation of data protection laws)** In response to the growing security and privacy concerns, many countries have enacted data protection legislation that place stringent requirements on use, retention and disclosure of information. Most European countries have adopted national legislation implementing the 1995 European Data Protection Directive, and the United States (at the federal and state levels) has adopted an abundance of sectoral regulations to protect the privacy of personal information. Many other countries have also adopted or are in the process of adopting stringent privacy legislation, based on the "Fair Information Practice" principles of the Organization for Economic and Cooperative Development (OECD) in 1980. Companies that fail to comply may run into serious problems with government, ranging from fines to operational suspension. Data protection laws pertain to "personally identifiable" information, which is any information that can be linked (directly or indirectly) to an individual. Deleting an individual's name from his record does not imply that the record is no longer personally identifiable, since identification may take place indirectly on the basis of social security numbers, health insurance numbers, and so on. In fact, it is quite possible that the unbridled use of PKI will be found unconstitutional when challenged in court.[4]

[4] Some precedents: the Hungarian Constitutional Court in 1991 decided that multi-use personal identification numbers violate the constitutional right of privacy, the Portuguese Constitution states that "Citizens shall not be given an all-purpose national identity number", and SSN legislation in many countries prohibits the use of SSNs beyond very specific purposes.

- **(Unconditional trust required in smartcards)** If the secret key of a digital identity certificate is generated and stored on a personal computer or the like, it is virtually impossible to prevent its compromise, loss, disclosure, modification, and unauthorized use. However, when using X.509-style PKI in combination with smartcards, it is virtually impossible to verify that the cards do not leak their secret keys, card identifiers, access control codes, data from other applications running on the same device, and so on. Moreover, a variety of fake-terminal attacks become possible due to the lack of a trusted display and keyboard on the user's side, and it cannot be guaranteed that the smartcard supplier cannot simply reconstruct all the secret keys. As a result, application providers must have unconditional trust in the honesty of their smartcard suppliers. National defense networks and other critical information infrastructures, as well as enterprises that are interesting subjects for industrial espionage, cannot reasonably place such trust in outsiders.

Attribute certificates, as proposed in the X.509 standard, allow one to circumvent the bottleneck of on-line central databases (much like Digital Credentials). However, they do nothing to address any of the other problems of identity certificates. In fact, attribute certificates introduce new problems that are more serious than the bottleneck problem they solve. Namely, with X.509-style attribute certificates, organizations cannot securely give individuals control over their own information, since this would allow unauthorized individuals to copy, lend, pool, discard, and prevent updating of their own information. Furthermore, all the attributes within an X.509 attribute certificate are systematically revealed when showing the certificate, and so the attributes pertaining to a party must be distributed across many digital certificates. To facilitate fine-grained user-control over which attributes are released to whom, each user must carry an enormous number of certificates that all have to be managed separately; this creates serious scalability problems as well. Moreover, managing just a single attribute certificate using a low-cost smartcard is even more problematic than managing an identity certificate on a smartcard. Finally, the privacy implications of using a myriad of fully identifiable and linkable attribute certificates are even worse than those of digital identity certificates.

As Dr. Peter Swire, then Chief Privacy Counselor to the Clinton Administration, stated in April 2001 (see http://www.law.ohio-state.edu/swire1/EBLCRAprilSwireInterview.doc): "*One of the great mysteries is why digital certificates have not spread yet. [...] we need to have both strong authentication and consumer confidence in privacy for electronic transactions. Until that's solved, I'm not sure we're going to see digital certificates spread.*"

4. References

[1] "A Semi-Technical Overview of Digital Credentials," by S. Brands, International Journal on Information Security, 2003 (to appear). http://www.credentica.com/technology/overview.pdf

[2] "Rethinking Public Key Infrastructures and Digital Certificates; Building in Privacy," by S. Brands, August 2000, MIT Press, ISBN 0-262-02491-8. With a foreword by prof. Ronald L. Rivest. See http://www.credentica.com/technology/book.html for excerpts and endorsements.