Access Management Based On Digital Credentials Part 1

Dr Stefan Brands

Credentica and
McGill School of Computer Science

Abstract

This paper is the first of a two-part series on how to carry out digital access management in a highly secure yet un-intrusive manner that scales seamlessly across organizational boundaries.

This first installment provides an overview of Digital Credentials, the cryptographic primitive that is at the heart of the new approach to access management. Digital Credentials are cryptographically protected binary strings that provide fine-grained privacy control as well as strong protection against counterfeiting, pooling, discarding, copying, lending, remote extortion, and other frauds.

Introduction

Corporate and government organizations are increasingly making their information resources accessible in electronic form over open networks. They do so to facilitate information sharing with dispersed business units, affiliate and partner organizations, providers of outsourced services, remote workers, suppliers, consultants, and so on. The ultimate goal of the transition is to cut cost and to increase revenue by enhancing productivity, reducing errors, and opening up a range of new opportunities.

As information is shared with more and more parties and across organizational boundaries, the need for security measures to keep information out of the hands of unauthorized parties goes up dramatically. Traditional information security products such as firewalls, anti-virus software, intrusion detection systems, and vulnerability assessment tools are rapidly becoming outdated; they work well when applied to protect information that resides within a single administrative domain, but they are unsuitable to protect information that is intended to be shared across organizational boundaries. With trust domains becoming logical rather than physical, security must be tied directly to the information itself rather than to the perimeter of its repository. This requires the application of cryptographic authentication and encryption techniques.

The move towards stronger security measures in access management is an undeniable trend at present. For an examination of what this trend means in several markets with a strong reliance on access management, see [1].

A Naïve Approach to Security in Access Management

It is widely believed that strong security is best achieved by building access management applications on top of a Public Key Infrastructure (PKI). In this approach, a trusted central authority provides each potential access requestor with a digital identity certificate through an enrollment process. Each digital identity certificate contains a unique secret key and a corresponding public key, to enable its holder to cryptographically demonstrate ownership of the certificate without revealing the certificate's secret key. Access providers rely on the digital identity certificates presented by access requestors in much the same way as government organizations traditionally rely on Social Security Numbers. That is, they use them as (strongly authenticated) pointers to records stored in central databases; this enables access providers to look up any information about access requestors that they deem appropriate to make an authorization decision.

This straightforward approach works well for managing access to information in small-scale single-enterprise domains, where privacy expectations towards certificate authorities, access providers and central parties are virtually nonexistent (after all, they are all part of the same organization), scalability is hardly an issue, and fraud can be contained quickly. However, as an organization's information sharing needs expand to an increasing number of parties, the fundamental shortcomings of the PKI approach are becoming painfully clear. Indeed, a careful analysis shows that building multi-domain access management systems on top of PKI fundamentally provides poor security, does not scale, and violates many of the privacy principles codified in law in most countries; see [1] as well as the second installment of this paper. With regard to privacy, what is at stake is not only the privacy of access requestors but also of access providers, who cannot prevent central parties from being able to learn the identities of their customers, their peak hours, transaction details, and so on.

This mismatch between the growing security needs of access management and PKI is not all that surprising. PKI was invented in 1978, at the dawn of modern cryptography, when open networks were hardly existent, let alone organizations seeking to share their information over such networks. The design goal of PKI was to facilitate message encryption, with entity authentication (through digital identity certificates) serving to prevent man-in-the-middle attacks. However, authentication as used for message encryption is a far cry from managing the access of authorized parties to sensitive resources, with all its own privacy, security, and performance needs.

What This Paper is About

The objective of this two-part paper is to introduce a superior alternative for access management that enables the separation of the concepts of authentication and authorization in a secure, scalable, and privacy-friendly manner.

In this first installment, we introduce the notion of *Digital Credentials*, the underlying state-of-the-art cryptographic primitive that is at the core of the new approach. Digital Credentials are the culmination of two decades of scientific advances by dozens of professional cryptographers starting in the early eighties, following the invention of PKI. Digital Credentials were specifically designed to overcome fundamental security and functionality problems of paper documents and other tangible objects traditionally used for establishing a person's privileges, characteristics, identity, and so on. In contrast to physical "credentials," which are increasingly prone to counterfeiting and cannot be transferred electronically, Digital Credentials are cryptographically protected against all manner of fraud and can be transferred instantaneously across the globe. They support all the traditional authentication strengths, ranging from software-only protection to military-grade two-factor and three-factor security. What's more, Digital Credentials allow fine-grained control over one's privacy, have security properties that go well beyond what can be achieved for their traditional counterparts (in ways that may at times seem counter-intuitive), and can be implemented efficiently using low-cost smartcards without cryptographic coprocessors.

In the second installment of this series, we will explain how to implement secure access management based on Digital Credentials. Our access management architecture will rely on Digital Credentials in four basic manners:

- To directly implement access privileges, entitlements, and any other attributes that access requestors show to access providers to allow them to make local authorization decisions;
- To implement privacy-enhanced digital identity certificates (usable as digital pseudonyms where identification is not required) that function much like Social Security Numbers but that allow the secure separation of different spheres of activity;
- To authenticate the entries of electronic records stored in central or federated databases; and
- To implement digital audit trails and digital receipts that witness details of access requests.

We will also give examples in the second installment of how to apply the resulting architecture to Electronic Health Record management, national ID chip-cards, e-government, and Digital Rights Management. At the end of the second installment we will compare the new approach with the naïve PKI approach described above, and will show how the latter falls short by far even if one were to replace digital identity certificates by attribute certificates (a PKI-derived primitive that only superficially resembles Digital Credentials).

Before we dive into Digital Credentials, we take some time to examine the meaning of privacy and how it relates to security. This will hopefully lead to a better appreciation of the utility of privacy-enhancing technologies in general and of Digital Credentials in particular.

Privacy Backgrounder

Privacy is generally defined as *the claim of individuals to determine for themselves when, how, and to what extent information about them is communicated to others*. This definition, which is at the basis of almost all information privacy legislation and regulations around the world, emphasizes that the control over the release of identifiable personal information should rest with the data subject. Taking away control from the data subject takes away his privacy. As such, legislation that moves control out of the hands of the data subject does not do anything for privacy, nor do "trusted" third party solutions that are unilaterally imposed upon data subjects. Of course, data subjects may volunteer to delegate some or all of their control to others – as long as they have the choice and ability to keep that control.

The Relation Between Privacy and Security

There is much confusion about the relation between privacy and security. Security is generally defined as the extent to which information can be stored and transmitted in such a manner that access to the data is limited to authorized parties. Many people believe that privacy and security are opposite goals that must be "balanced." Underlying this belief are implicit assumptions that have no basis in reality, such as the assumption that privacy equals anonymity which in turn prevents accountability. In the next section, when we discuss Digital Credentials, we will see several counter-examples that show that privacy and accountability are not necessarily opposites.

In order to better understand the relation between privacy and security, consider the *Fair Information Principles* codified in 1980 by the OECD. These principles form the basis of most of today's public privacy policy and legislation around the world. Specifically, the eight OECD privacy guidelines are *openness, collection limitation, purpose specification, use limitation, data quality, individual participation, security safeguards*, and *accountability*. (Since they do not take technological advancements into account, several countries actually go beyond them. For instance, the Canadian CSA code augments and subtly modifies these eight principles, to include consent, limiting disclosure and retention, data accuracy, individual access, and challenging compliance). Now, note that *security safeguards* is only one of the eight principles; in other words, security safeguards are *necessary* to achieve privacy, but they are not *sufficient*. In fact, technologies that are aimed only at providing security safeguards do nothing to address the most urgent privacy concerns, which come primarily from secondary use by insiders: security safeguards are primarily aimed at protection against access by unauthorized outsiders.

Ironically, many of today's prevalent security technologies have a highly adverse impact on two of the most important privacy principles: collection and use limitation. They are, in fact, *privacy-invasive* technologies. Here are two examples:

- *Surveillance Cameras*. While the security risk of being assaulted in public places may go down by placing these areas under permanent surveillance by high-quality surveillance cameras, very few people would argue that such an intrusive surveillance measure protects their privacy.

- *PKI Used for Multi-Domain Access Management*. PKI technology contributes to the *security safeguards* category of the Fair Information Principles, by providing the following basic security safeguards: confidentiality of data in transit (through encryption), user authentication, data integrity, and non-repudiation. As we explained in the Introduction, however, in the overall equation PKI technology turns out to have devastating consequences for privacy when used for multi-domain access control.

Privacy-Enhancing Technologies

Many people mistakenly believe that privacy, in contrast to security, cannot be addressed through the use of technology; only legislation and regulation can protect privacy. In an electronic world, however, there are no grey areas between privacy and inescapable systemic identification. If at the technical level everything is systemically identifiable, privacy legislation becomes virtually meaningless; how can one force participants not to collect identifiable information when they cannot prevent it from being delivered to them? The only way to tackle this fundamental problem is to resort to technology itself.

The terminology *Privacy-Enhancing Technology* (PET) has become popular in privacy, legislative, business, and technology circles alike. PETs provide individuals with technical means for controlling themselves the extent to which information about them that they disclose to others can be associated with their unique identity. (As with the definition of information privacy, this definition can be defined with respect to information about groups of individuals and to institutions as well.) Until remote mind-control technologies turns from fiction into fact, this implies that PETs at the very least must allow individuals to "de-identify" their own personal information before it is disclosed.

This does not mean that PETs *force* individuals to be anonymous or pseudonymous: they merely give them the option to be so. Whether this option is acceptable in a particular situation depends on legislation, contractual agreements, social conventions, etc. Since one can always send along additional information, any PET offers individuals free choice over the entire privacy spectrum between the two extremes of systemic inescapable identification and the maximum degree of privacy attainable through the use of the PET. This spectrum is not one-dimensional but multi-dimensional, depending on the number and nature of parties that the individual considers in his privacy considerations. Indeed, individuals typically prefer different degrees of privacy towards different parties, and want to be able from time to time to delegate partial control over the secondary use of their information to parties they trust. Since it is easy for technology to take away privacy, but not the opposite, most PETs are designed directly for the purpose of enabling one to be fully anonymous or pseudonymous: the privacy slider principle implies that the entire spectrum between inescapable identifiability and full control can be achieved.

Like with the notion of privacy, there is much confusion about the meaning of PETs. False assumptions abound. Among the most popular misconceptions about PETs are the following:

- If the actions of honest individuals are not identifiable, then fraudsters cannot be identified either;
- It is impossible to revoke anonymous/pseudonymous access rights;
- Privacy can only be achieved by relying on Trusted Third Parties that act in everyone's best interests;
- Encryption is cryptography's solution to the privacy problem;
- Smart cards (and other user-held devices) provide privacy when one's personal data is stored inside the device;
- Privacy at the application level is useless if one can be traced at the data transport level; and
- A prerequisite to information security is physical control over the information itself.

Adding to the confusion is that the definition of PETs is currently being watered down. Driven by the desire to gain customer goodwill and to discourage outside intervention, an increasing number of organizations claim to have incorporated PETs into their applications on the grounds of having added a chunk of code that provides some utility that can be interpreted as being in line with one of the Fair Information Principles. Example measures include informed consent pop-up windows, automated notification of privacy breaches, automated data deletion, automatic preference negotiation mechanisms (such as P3P), and so on. However, while these simple tools may certainly help address privacy concerns, they are not PETs, since they do nothing to prevent systemic identification of the actions of honest parties. Indeed, the vast majority of the current private sector experimentations with privacy tools concerns what may more accurately be referred to as Privacy-Unctuous Technologies; they are characterized by exaggerated or insincere earnestness and have the characteristics of snake oil.

Even more worrisome, privacy-invasive technologies (such as biometrics, PKI for large-scale access management, and surveillance cameras) are increasingly referred to as PETs, an unfortunate trend that can be attributed to two factors: (1) a widespread unawareness of what information technology can and cannot do, and (2), the unscrupulous marketing tactics of some vendors of privacy-invasive security technologies.

The Role of Cryptography in PETs

Two decades of academic research in cryptography have demonstrated that security and privacy are not trade-offs, but that they are mutually reinforcing when implemented properly. A fundamental premise of modern cryptography is that the need to rely for privacy on Trusted Third Parties can be eliminated. Advanced cryptographic PET building blocks have come out of the cryptographic research community that can be implemented in entirely practical manners.

Here is a partial list of PET building blocks, many of which are almost completely unknown outside of the professional cryptography community:

- Restrictive/ordinary/limited-show blind signatures;
- Selective disclosure proofs for strongly authenticated information;
- Strongly authenticated digital pseudonyms;
- Homomorphic encryption;
- Secret sharing;
- Searching on encrypted data;
- Privacy-preserving data mining;
- Private information retrieval;
- Self-revocable unlinkability and untraceability;
- Untraceable broadcasting; and
- Anonymous remailers

In fact, through the use of cryptography it is possible to achieve all kinds of counter-intuitive results that have no analogue in the physical world. For instance, role-based signing can be implemented through the use of PETs, in such a manner that digital signatures cannot be traced to an identifiable person but only to the role they assumed when signing; at the same time, through the magic of cryptography, in case of a dispute, the signer cannot repudiate his action, while all others can prove that they were not involved. Therefore, in case of a serious dispute a judge could electronically request a selected user population to transmit cryptographic proofs that they did not originate the contended action.

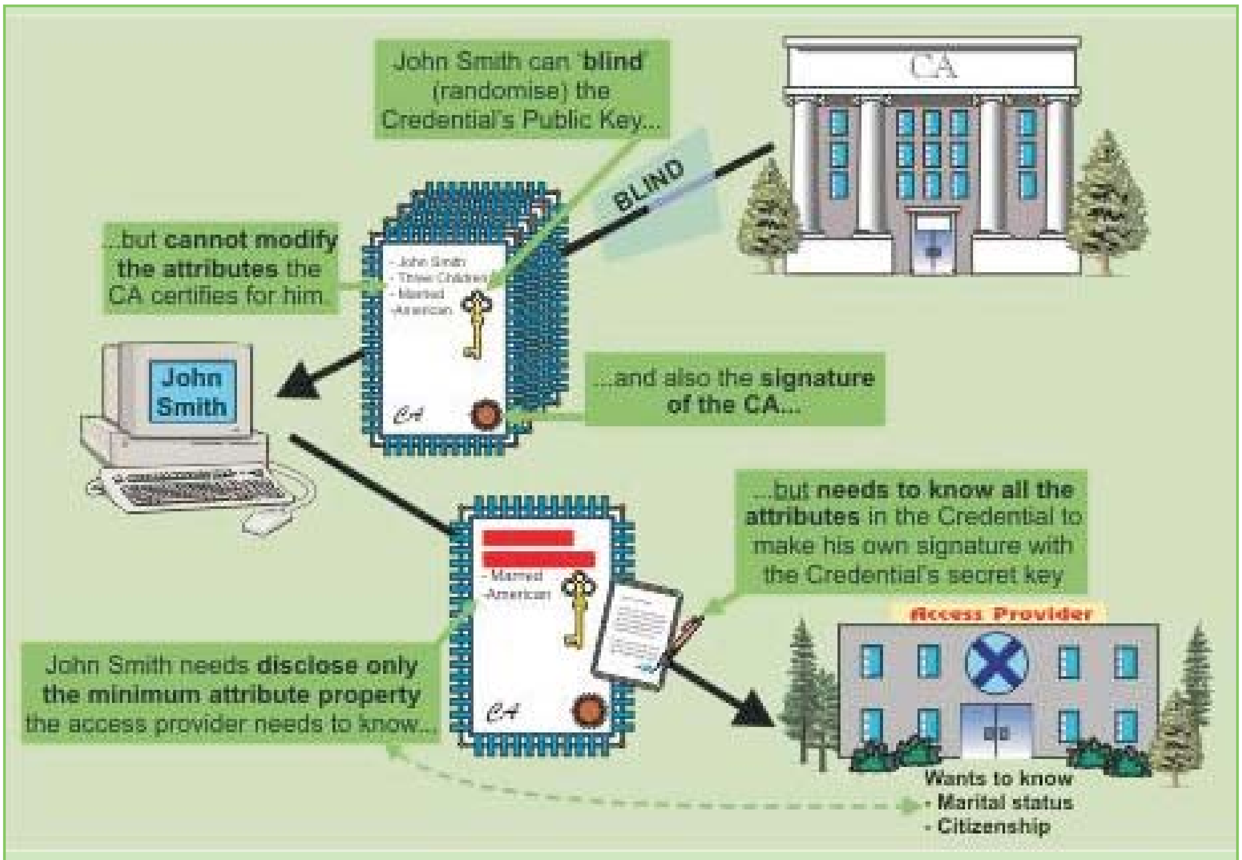
John Smith can 'blind' (randomise) the Credential's Public Key...
CA
BLIND
...but cannot modify the attributes the CA certifies for him
- John Smith
- Three Children
- Married
-American
CA
John Smith
...and also the signature of the CA...
...but needs to know all the attributes in the Credential to make his own signature with the Credential's secret key
- Married
-American
CA
Access Provider
John Smith needs disclose only the minimum attribute property the access provider needs to know...
Wants to know
- Marital status
- Citizenship

Figure 1 - **Issuing and Showing a Digital Credential**

With this background in mind, we are now well-prepared to explain the notion of Digital Credentials in more detail.

Digital Credentials

Digital Credentials are basic cryptographic constructs, much like digital signatures but with much greater functionality. They enable their holders 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. Since Digital Credentials are just sequences of zeros and ones, they can be transferred electronically and can be verified with 100 percent accuracy by computers. At the same time, as we already mentioned, they offer security benefits that go far beyond those of their traditional counterparts.

Issuing and Showing Digital Credentials

Technically, Digital Credentials are issued and shown as follows:

- *Issuing Protocol*. Digital Credentials 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 knows. The whole package that Alice receives is called a Digital Credential. Although the sequences of zeros and ones that make up 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. At the same time, Alice cannot modify the attributes that the Credential Authority encodes into her Digital Credential.

- *Showing Protocol.* 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. To convince Bob that the claimed property is true, Alice's signature on Bob's nonce doubles up as a proof of correctness.

The class of properties that a Credential holder can selectively disclose in the showing protocol 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 age can prove eligibility for a discount pass without revealing anything more about her age beyond that she is either a minor or a senior. At the same time, it is infeasible to demonstrate any property without knowing all the attributes encoded into the Digital Credential (including those that are not disclosed); we will see below how this property can be exploited to counter certain types of fraud.

Figure 1 illustrates the issuing and showing process.

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 of Digital Credentials can be found in [2], and the full technical details appear in [3].

Privacy, Security, and Performance Features

As explained in [2] and [3], by carefully exploiting the basic properties of Digital Credentials, one can realize all the following features:

- *Privacy of Credential Holders.* 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 an access 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. (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). There is no need for Trusted Third Parties to protect one's privacy: even if all the parties that rely on Digital Credentials actively conspire with all Digital Credential issuers and have unlimited computing resources, they cannot learn more than what can be inferred from the assertions that Digital Credential holders willingly and knowingly disclose. This seemingly very strong privacy guarantee is not a new concept: today, when you spend a coin, cast a vote, or use a cinema ticket, you do not automatically and inescapably disclose your identity.

- *Privacy of Credential Verifiers.* In many situations, verifiers may want or need to pass on Digital Credential evidence to central parties, for instance for online revocation status checking, to enable fraud detection on behalf of a population of access providers, to allow statistical data gathering, or to serve as transaction "receipts." Digital Credential verifiers can selectively cover up any or all of the information that Digital Credential holders selectively disclosed to them, before passing on these Digital Credentials. In other words, verifiers can ensure that they are left with non-repudiable digital evidence that proves to third parties no more than exactly what they want the evidence to prove. This may be much less than what the Digital Credential holder selectively disclosed to the verifier. By way of example, consider a patient-physician interaction or a consumer-merchant transaction; while the customer may have no problem identifying himself to his doctor or to the merchant, the latter parties may be worried about disclosing their customer's identity to central parties.

- *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 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 access 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 Digital Credential holder can demonstrate this without disclosing the identifier.

- *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.

- *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.)

- *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.

- *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. (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). These limited-show Digital Credentials (which can be used to design the digital equivalent of 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, and even if they conspire with verifiers (who, as we saw, are able to cover any information disclosed to them before passing on the transaction data). 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.

- *Negative 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 security 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 behaviour. 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 without a cryptographic co-processor.

- *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. 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).

Using Digital Credentials to Build Applications

As a result of their much greater security than non-digital credentials, Digital Credentials can safely be used in all kinds of applications where the use of paper or plastic tokens would be insecure. That is, they can be used to implement not only gift certificates, railway tickets, and so on, but also diplomas, work permits, birth certificates, and other objects that traditionally identify their holder in order to deter certain kinds of fraud.

Digital Credentials encompass identity certificates as a special case: an identifier is just one of infinitely many attributes that can be encoded into a Digital Credential, and the Digital Credential holder can disclose it whenever desired. To encapsulate an X.509v3 certificate, for instance, the Credential Authority can issue a Digital 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 Digital Credential, the holder must disclose (the pre-image 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 could be set to a hash of the public key or to zero (verifiers can compute a hash themselves).

Depending on the application in which Digital Credentials are used, one might decide to go with just a few of the listed features. Using a cook-book approach, one can build all manner of secure applications on top of Digital Credentials.

Designing an electronic coin system, for instance, is easy. In this particular application, lending, discarding, and pooling of electronic coins are not attacks that need to be guarded against, while spending a coin multiple times obviously is a highly unacceptable fraud. Consequently, a good choice is to construct each electronic coin as a one-show Digital Credential that encodes two attributes: the denomination and currency of the coin, and its holder's identity or bank account number to facilitate tracing of double-spenders. The first attribute must be disclosed when spending the coin, while the second attribute is never disclosed when spending a coin in order to prevent linking and tracing of payments. In addition, consumers can be given a tamper-resistant smartcard that provides a second layer of protection against double-spending.

Other applications of special interest include digital pseudonyms for public forums and virtual communities; electronic voting; electronic postage; online auctions; financial securities trading; pay per-view tickets; public transport ticketing; electronic food stamps; road-toll pricing; Web site personalization; collaborative filtering; medical prescriptions; on-line age proofs; gift certificates; loyalty schemes; and, electronic gambling. Perhaps the most interesting application of Digital Credentials is an umbrella application: access management across multiple administrative domains. This will be the topic of the second installment in this series.

References

[1] *Secure Access Management; Trends, Drivers, and Solutions*, by S. Brands, Elsevier Information Security Technical Report, Volume 7, Issue 3, pages 81-94, Chez Ciechanowicz (editor), September 2002.

[2] *A Technical Overview of Digital Credentials*, by S. Brands, International Journal on Information Security, 2004 (to appear). www.credentica.com/technology/overview.pdf

[3] *Rethinking Public Key Infrastructures and Digital Certificates; Building in Privacy*, by S. Brands, 315 pages, August 2000, MIT Press, ISBN 0-262-02491-8. www.credentica.com/technology/book.html

About the Author

Dr. Stefan Brands is an expert on the subjects of electronic authentication, digital identity management, payment systems, and information privacy. He is an adjunct professor in cryptography at McGill University in Montreal, and is affiliated with Credentica, an IT security company providing access management software. Previously, he was a principal researcher at DigiCash (1996-1998) and at Zeroknowledge Systems (2000-2001). Dr. Brands conducted his PhD research from 1992 until early 1996 at the Center for Mathematics and Computer Science (CWI) in Amsterdam. His dissertation was approved by professors Adi Shamir, Ron Rivest, and Claus Schnorr, and was published in 2000 by MIT Press with a foreword by Rivest. Dr. Brands is a member of the CSIS Working Group on Authentication in Washington, provides consultancy from time to time, and is a frequently invited speaker at leading industry and academic forums.

He can be reached at brands@cs.mcgill.ca or at brands@credentica.com.