On The Reliability of Authentication of Identity

Peter Alterman, Ph.D.
Director of Operations, Office of Extramural Research
National Institutes of Health
And
Senior Advisor to the Chair, Federal PKI Steering Committee

Abstract
The goal of authenticating identity and of binding it to an electronic token of one sort or another has been a challenge to implementers of electronic business applications. Commonly-relied upon identity credentials are inadequate to the job of assuring individual identity for binding to electronic credentials. Combining personal information available in multiple databases with common identity credentials does offer reasonable assurance of identity. Mathematical algorithms for assessing the strength of identity assertion based upon this hypothesis are proposed, in order to advance the goal of automating authentication of identity.

Author Keywords: Authentication, Identity Proofing, Identity Binding, Electronic Commerce

All communication regarding this paper should be forwarded to the author. Contact information:

Voice: (USA) 301-496-7998
Fax: (USA) 301-496-0232
Email: peter.alterman@nih.gov
Post: 1 Center Drive, MSC 0155
9000 Rockville Pike
Bethesda, MD 20892
USA

On the Reliability of Authentication of Identity

The goal of authenticating identity and of binding it to an electronic token of one sort or another has been a challenge to implementers of electronic business applications. While there have been no end of imaginative cryptographic developments that can ensure reliable digital signature and/or file encryption, no generally recognized algorithm exists for determining the reliability of the identity that is bound to the products of those imaginative crypto tools.

Specious, but nevertheless long-winded arguments have been advanced claiming that individuals have multiple identities, that is, if I have two email accounts under two names, I have two identities. In fact, what I have in this case is two credentials associated with my persistent identity; the error is in confusing identity credentials with identity; similarly, if I am a father and the budget officer for a large corporation, I do not have two identities, I have two roles: the error is in confusing identity with role.

This paper examines some key issues necessary to assert a specific level of reliability to an identity, and proposes an algorithm to apply to determine the reliability of an identity with a standard metric. The issues and approaches addressed here are different from those related to law enforcement or national security, where many initial assumptions regarding individual rights, applicable law and cost-benefit calculations are quite different.

Within the context of transacting electronic business, most discussions of authentication of identity focus on authenticating a digital credential to a system or an infrastructure; this paper focuses on authenticating an actual identity, not an electronic identity credential. Thus, this discussion is antecedent to any process of binding identity to an electronic credential, addressing the question, “How can you be sure of the identity you are binding to an electronic credential?”

The process of authentication begins with presentation and examination of identity credentials. In the real world, this usually consists of a driver’s license, perhaps with one or two major credit cards. Driver’s licenses are issued by State governments in the U.S.; credit cards are usually issued by banks to account holders who satisfy several common criteria. The most commonly issued U.S. Federal credential used as an assertion of identity is the Social Security Number, even though the SSN is not unique and statute forbids its use for identification purposes.

Foundation identity credentials (that is, credentials such as the driver’s license that are used by other processes to issue identity credentials such as a passport) are not at this time issued by the U.S. Federal Government (the one exception being the military ID card). They are issued by State and local governments and by private corporations. Another salient fact is that no one identity credential is particularly reliable. Anyone living around a major college or university is aware of the unreliability of the driver’s license as an identity credential; driver’s licenses are regularly and easily counterfeited by students wishing to gain access to alcohol. The sorry state of the Virginia Department of Motor Vehicles in issuing fraudulent driver’s licenses to illegal immigrants, and to some of the terrorists responsible for the 9-11 attacks in Washington and New York, was big news in the Fall of 2001.

Interestingly enough, so-called in-person proofing is another form of presenting a portable credential. In the case where an individual comes in to a registration station and presents a credential, what is being proofed is the congruence between a biometric and a token that includes that biometric (usually the face matching the photograph on the credential). This form of in-person proofing is only as good as the credential presented, which as we have seen may be easily counterfeited. Where in-person proofing includes matching a presented biometric and/or credential to an authoritative database, the process is only as reliable as the database. In either case, the actual fact of an individual presenting him- or herself for inspection does not determine the strength of identity assurance. Rather, it is the reliability of the resource that the person's biometric is matched against that determines the reliability of the proofing process.

It is for these reasons that distinctions between "strong" and "weak" credentials, that is, between a picture ID and a library card, for example, are not really useful. Either is as likely to be counterfeited as the other, and databases of identifiers have been hacked.

Reliability of identity can be built up from a series of credentials and records, however. This is an example of the principle that many bits of somewhat reliable data may aggregate into a bit of quite reliable information. If an individual presents a driver's license, automobile registration and insurance card for the same vehicle, all of which have the same name and address, that is, if they are mutually referential, a much stronger case can be made that the series of credentials reliably defines an identity. Add a mortgage account, a checking account, voter registration records, medical insurance account, and the overall confidence one has in the individual's identity grows even greater. Add to this list access to medical records (undesirable for reasons other than identity proofing, but then we are speaking here in the abstract) and credit history and the confidence in the individual's identity rapidly rises towards certainty, that is, the electronic credential issuer is just bout 100% sure the individual presenting all these credentials – onerous as that surely would be – is who he or she claims to be. A series of identity credentials that are not mutually-referential are weaker than a series of identity credentials that are mutually-referential, because mutual reference provides supportive evidence of the validity of each of the credentials: if they all agree, they are likely all valid, or, as noted above, all authenticate a common identity (real or not). The factors in this model exhibit several key features worth noting: first, there are multiple credentials from unrelated sources; second, the credentials are mutually-supportive; third, they demonstrate duration of the identity over time.

	Credentials	**In-Person Proofing**	**Records**
Strengths	Well-known and accepted; existing process to acquire; biometrics can offer strong identity verification; multiple unrelated credentials can reinforce each other	Positive match between face and credential guarantees a "real person"	Multiple unrelated records provide strong reliability of identity; can be automated
Weaknesses	No way to avoid counterfeiting potential; even biometrics have false positives and false negatives and vulnerabilities in back-end processes; extremely vulnerable to fake identity	Dependent on validity of credential or token back end; actually difficult to match biometric to record; cannot be automated	Privacy issues; requirement for multiple, unrelated records to avoid spoofing

Table 1: Strengths and weaknesses of different kinds of identity credentials.

No model is completely foolproof and a concerted effort to create a seamless record of identity satisfying the three conditions above can succeed, although at a cost arguably as high as, or higher than, the benefits to be reaped by such fraud. Furthermore, if all these credentials depict an identity other than the one the individual was born with, they may still verify the existence of an identity that is functional within society, or that is "doing business as" the identity presented. Consider the case of the movie star who changes her name from Frances Gump to Judy Garland. An interesting epistemological discussion can be based on this point.

The dimension of duration over time reminds us that there are two general kinds of identity credentials: portable credentials designed to assert identity, such as the driver's license, and embedded credentials, such as deeds of trust, which record incidents and relationships in a person's life; that is, the personal record. Most of have an umbilicus of records and identity credentials that stretch from the present right back to our births. One can get at the question of whether the identity being proffered is actually the identity the individual was born with by examining records that have significant duration, that is, that go back in time a substantial number of years.

An example of a business case that uses this model implicitly is credit reporting. In the U.S. a few companies compile personally-identifiable records on business transactions by citizens and sell that information to firms deciding whether to extend credit to individuals. Moody's and Dun and Bradstreet, among others, use a similar approach for evaluating the creditworthiness of businesses for investors and lenders. Clearly, then, this concept of evaluating identity is implicit in the credit and lending markets of the world.

Those whose lives are more marginal, who don not have my of these connections, for whom credentials and records are in fact more difficult to assert, have identities that are more difficult to validate by third parties; this makes them no less real, no less valid,; it simply makes them more difficult to identity proof.

So what one can say from this analysis is that while it is not necessarily possible to authenticate any particular identity credential, it is quite possible to authenticate an identity. The corollary statement is that it is easy to counterfeit a credential, but it is much more difficult to counterfeit a life.

To recap: the reliability which one can have about an identity, then, is related to the following factors: number of identity credentials issued by unrelated sources; interrelationship among credentials, and duration of personally-identifiable credentials over time.

There is, in addition, another factor which directly bears on the reliability of identity authentication, and which is linked directly to individual transactions: indemnification. If the issuer of an electronic credential indemnifies the relying party to an electronic transaction for the full amount of potential loss, then the relying party's requirement for having a reliable identity bound to that electronic credential falls towards zero. Quite simply, if the relying party is fully insured against loss, the relying party is indifferent to the strength of identity binding. This is a rhetorical worst case, of course. Nevertheless, it does point out another key factor in determining the degree to which reliability of identity authentication may be calculated, or quantified. The indemnifier is a third party vouching for the asserted identity and the extent of that voucher may be calculated by the degree to which he is willing to indemnify the dollar (or Euro) value of the transaction.

Thus, the factors involved in quantifying reliability of authentication of identity may be shown in the following relationship:

(1) $Ri = \frac{(n*r*d),}{n}$ where

Ri = the reliability of identity authentication (inverse of risk)

n = number of identity credentials issued by unrelated entities, not based on each other. A birth certificate and a passport are each issued by unrelated entities (satisfying the first requirement) but the passport relies on the birth certificate for key identity data (failing the second requirement).

r = the degree to which key data types in each identity credential correspond to each other, e.g., if "birth date" or "color of eyes" is the same in each identity credential or record.

d = duration of identity credential over time

The metrics for each of these factors have to be refines, especially n. Nevertheless, some beginnings can be made. For example, Ri will most likely be a percentage, as the goal is 100% assurance of identity. d ranges from birth through current life, so it may be represented reasonably as percentage of the individual's lifespan. r may also be seen as a percentage. Asserting r = 100% (or 1.00) for all n would mean that all the presented credentials display the same name, address, age, etc. To arrive at this figure, one would calculate the number of identity fields in each (name, address, Soundex number, insurance policy number, etc.) and determine how many overlap, then determine how many are equivalent.

A clear advantage of using a percentage figure for Ri is that it is amenable to rounding, so the question "how much authentication is good enough" can be reformulated as a requirement for reliability at a specific percentage figure, e.g., 70% assurance as determined by a standard methodology.

A second formula is required to calculated reliability in the presence of indemnification or third party assertion of identity.

(2) $A = \frac{Ri+(Ri*I)}{2}$ where

A = the indemnified reliability probability

Ri = as above

I = the percentage of indemnification of a particular transaction. Since this variable is related to the value of any particular transaction, it can be written as a percentage of the value of the transaction, ranging from 0% to 100%

The reasoning behind this second algorithm is based upon the premise that there will be a certain reliability of identity calculable from the first formula (1), while the indemnification function will further decrease perceived risk (and increase reliability). Multiplying the two factors is a first attempt to estimate the additional reliability conferred by indemnification above the "intrinsic" reliability Ri. The number is divided by two to normalize the result, that is, a maximum of 100% assurance of identity plus 100% indemnification of the transaction.

These are admittedly crude first attempts to quantify the factors which determine reliability of authenticated identity and their relationships to each other. Nevertheless, the models presented here, and their refined successors, are greatly needed in the developing electronic commerce environment. With a reliability algorithm and an indemnified reliability algorithm, determination of reliability of authentication of identity may be automated, with concomitant benefits in citizen service, reduced risk of doing business electronically for all parties, and streamlined issuance of secure digital credentials. In general, the better the calculation of reliability, the better the calculation of risk and the better the calculation of profit and loss, and of whether to move traditional business models to the Internet.