Hettinga's Best of the Month

From Contributing Editor Bob Hettinga
Email: rah@shipwright.com
URL: The e$ Home Page

"... however it may deserve respect for its usefulness and antiquity, [predicting the end of the world] has not been found agreeable to experience." -- Edward Gibbon, 'Decline and Fall of the Roman Empire'

At 09:35 AM 5/7/98 +0100, Mainz, Reinhold wrote:

Hello, Carl,

I think, our little (and helpful) discussion in this case will soon be ended, so I have copied it again to E-CARM.

Hallo Reinhold,

several people wrote to me to ask for this discussion to be kept on the list, so I would keep it there anyway.

I've consoladated your two messages into one, here, and changed the subject line.

Of course a person's name is important in real life. The trouble is that the name is meaningless to the software checking a certificate chain and performing some action as a result (e.g., allowing an electronic prescription for drugs to be processed). That software cares only about the key and its attribute. The name just gets in the way (or, more usually, gets ignored).

But I can use a person`s common name in my local application, too. And then it might be good to have a prooved name (see comments in the next mail).

As far as I can tell from this and the other things you wrote, you want the name to be a valid identifier. Sometimes that is in another application. Sometimes that is in the 3-D world.

This is one sign of the major problem with use of names on the scale of the Internet based on common everyday practice. Names work in our small, physical communities. We expect them to work everywhere. They don't.

However, we are tempted to force them into working. One way is to look at the common name in a database and then use other information in the database to eliminate obvious non-matches to the person we're thinking about.

In actual fact, this is a process of human guessing. It makes all other security concerns effectively meaningless because it is so easily mistaken or subject to active attack (compared to the difficulty of breaking a cryptographic key, for example).

Your technical view is O. K., but for many reasons in healthcare system - I think - it would be good to have a certificate for some attributes together with the real name for using the name outside of electronics. It is correct that this is a process of human guessing, but today this a process of human guessing, too.

This is a place where we who do cryptography may part company with the man in the street. There are no certainties in cryptography, just as there aren't in life, but we use near-certainties in cryptography. We work hard to reduce the probability of failure of a cryptographic key from one part in 10^{30} to one part in 10^{40}, for example. The kind

of human operations you are talking about have probable failure rates of one part in 10^{3} to 10^{6}. Once you include such an operation in the process, then all the work we do with cryptography is made meaningless. In other words, you can do without strong cryptography entirely and still have the same system probability of failure if you insist on using a human guess as part of the process.

If you weren't trying to use the results of this strong cryptography that would be OK, but in these situations you are using those results. Therefore, you are implicitly relying on the extremely low probability of failure that strong cryptography gives you -- and then throwing that probability away by introducing a human guess.

It is true that current processes in the 3-D world use these human steps, but their probability of failure was once much lower than the 10^{-3} I estimated above. That is because the community in which the process was applied was much smaller.

For example, I came from a family of 6. If my sister said "Jim just fell off his bicycle", we knew exactly who she meant by the name "Jim". There was no probability of error at all even though the name "Jim" has only a tiny amount of entropy.

As the community gets larger, error is introduced or names have to grow in entropy. At some point, the community gets large enough that the entropy of common names is exceeded and the only possible result is error.

I don't know if this is happening in Germany, but here in the US there is an attack on 3-D world security mechanisms called "True Name Fraud". For example, someone will assemble information about a person (name, address, phone, employer, credit card numbers, social security number, ...) and then forge a picture ID with some of that information but a picture of the attacker rather than the named person. The attacker will then go into a store and open a charge account under the name of the real person, and use that account to charge merchandise. He then takes the merchandise and sells it on the black market, destroys the false ID he had created and pockets the money from that sale. This is a sample of the failure of names as identifiers in the 3-D world, as a result of the community getting larger. When we all lived in small towns, this attack on the security system could not have happened. The person at the store where you were applying for credit would have known the named person and therefore know that the attacker was not that person.

The same could happen in the health care system -- with the result that someone could pose as a doctor or write unauthorized prescriptions for controlled drugs -- buy those drugs and then sell them on the street. Right?

I believe that part of what we are doing with public key signatures and authorization certificates is providing a solution for a problem that already exists in the 3-D world and is merely made much worse in the electronic world. We have the advantage that the problem is so much worse that we know we have to find a solution to it -- but the solution we find can presumably help in the 3-D world where the same problems exist, only on an almost tolerable level.

As long as that CA is specialized and authorized to make pronouncements about health care attributes, then I can make a solid security decision. If the CA is authorized to make general purpose German identity certificates with health care attributes as just one feature, then that CA has acquired a power it probably doesn't deserve to have. For example, when Ross Anderson studied this problem in England, he found that there were different organizations empowered to certify physicians from those to certify nurses. That system would require two different CAs, one for each of those functions.

What I am trying to say (and maybe saying too much) is that the popular talk about a general purpose CA that just binds (name,key) together is almost certainly a wrong idea.

Probably we need several CAs in Germany, too. We are trying to build only one for the healthcare system. If this doesn`t work, we will build a PCA (same security guidelines, policy, organisation = same level of trust).

There are two different things involved here, but the salesmen for commercial CAs try to focus people on just one of the two...the one that the commercial CAs take care of.

That one is security of the cryptographic and personnel processes. These are important -- and the more certificates a

CA issues, the more important this kind of security is. For example, when I issue certificates to identify my friends, I am not likely to come under attack the way a commercial CA is because my certificates don't constitute that much value to an attack.

The other thing involved here is authorization to make statements. A CA making statements (issuing certificates) about healthcare (e.g., declaring someone to be a physician) should be run by the organization that today makes such statements. The root key for that CA should belong to that organization. This is a different organization from the one that issues national ID cards or driver's licenses, or .... Here in the US, we have the state issuing a license to practice medicine and a different state agency issuing driver's licenses. I don't know the name of the agency certifying doctors, but the Maryland driver's license is issued by the MVA (Motor Vehicle Administration). No one would propose that the MVA start issuing licenses to practice medicine on the theory that the MVA has a lot of experience issuing (driver's) licenses.

Commercial CA salesmen talk about separating CA and RA functions, so that the RA can live in the agency qualified to issue the authorization while the CA is centralized and runs protected cryptography. This is a theoretically workable design, but it has security flaws. Perhaps the biggest is that someone who infiltrates the CA can issue certificates appropriate to any of the RAs that feed that CA. The justification for this split is that cryptography is expensive. That's a false claim. Heavy security is expensive, but the level of security a commercial CA talks about is not necessary for me issuing personal certificates, for example. As the domain of the CA shrinks, so does the level of threat. As it turns out, every agency that keeps paper files already has security in place for protecting those files from attack -- and that physical security adequate for the paper files is the same physical security you need for an electronic version of the same process. In a small operation, the CA does not have to be on-line, so there is no need for elaborate network security (firewalls, network administrators, ...).

In other words, as far as my analysis is concerned, the split between CA and RA is more to attempt to build a business for commercial CAs than to answer a real security need. The proper solution is to have a CA where there would be an RA.

From the logic of what is happening when certificates are issued, there should be CA root keys for each kind of authorization and different certificates for each. There should not be a single certificate giving all the attributes associated with a person. If the latter course is followed, then with just a few people you end up with one issuing key empowered to make pronouncements about every attribute in the society. This gives a huge aggregation of power to the CA that holds that issuing key. That aggregation of power would not be tolerated in the physical world, issuing physical certificates. As I said, no one would seriously suggest that the MVA start issuing licenses to practice medicine, just because the MVA has experience issuing licenses.

It might be argued that a certificate verifier wants only one or two root keys to remember and therefore a central CA with only one root key is superior. However, if I have software checking validity of electronically delivered prescriptions, the only root key I need is the one from the agency issuing certificates that give permission to write prescriptions. I don't need to validate certificates for permission to drive or certificates giving the legal title to my house.

From: "Mainz, Reinhold"
To: "'Carl Ellison'"
Cc: "'e-carm@c3po.kc-inc.net'"
Subject: AW: AW: AW: [E-CARM] Re: business with an unknown party
Date: Thu, 7 May 1998 09:47:43 +0100

However, your public key does not need to be tied to the CA's public key in order for your key to be unique. All public keys are globally unique. If they are not, through some fluke, then the public key cryptosystem has been broken.

But I get a unique identification with name and attributes besides the key.

It is not true that the name and attributes give you a unique identification. I don't have any good examples from the medical profession, but I wouldn't be surprised if there were two doctors with the same name, somplace within

Germany. I know that if the attribute you're talking about is postal address, for example, but my friend and co-worker Donald Eastlake III and his son Donald Eastlake IV live at the same address. Both are adults. You have no way of knowing whether the pair is (II and III), (III and IV) or (IV and V) when you encounter just one of their names, and a certificate that lists the full name plus the mailing address and phone number doesn't let you uniquely identify the person unless you already know the number attached to their name. I knew Donald as "Donald" for ages before I learned he was "III". During that time, I could not have disambiguated the unique names. I also didn't know his address during that time, so I would have been faced with all the Donald Eastlake entries in the country -- not just the two at his address.

These are just simple counter-examples of the claim that name+attributes is unique. The general problem remains that a name has too few bits of entropy to disambiguate among a large group of people.

You can extend the name with natural attributes (like address), but still not get uniqueness.

You can extend the name with attributes chosen for uniqueness (e.g., e-mail domain names, or a national ID number, or a public key), but those names are not in common use in the real world. Therefore, if the purpose of this unique name is for a human to identify the named person (in person or in records), then the constructed unique name fails. If the constructed name includes the person's common name then you're introducing a security flaw. That is because the person who encounters this unique name may know it is unique (and therefore to be trusted as an identifier) and may also know only one person with the common name inside that unique name -- and may quite naturally jump to the conclusion that the constructed unique name belongs to the same person. For example, Alice is issued a certificate to the constructed name "Alice Johnson 997-42-3715 32969" where the number is Alice's national ID number concatenated with a globally issued ID number for the issuer of the national ID number. Therefore, this is a globally unique name. Bob knows only one Alice Johnson and encounters this certificate. Bob doesn't know his friend Alice's national ID number, but because he knows only one Alice Johnson, he assumes this certificate belongs to the one Alice Johnson he knows.

In this process, Bob has been tricked by his own continued use of the practices with which he grew up in the 3-D world into reducing the unique name "Alice Johnson 997-42-3715 32969" to the non-unique name "Alice Johnson", without knowing he had done that. The designers of the unique name and certificate don't know he had done that either. They would be horrified if they found out -- but that's still what Bob did.

If the name chosen for uniqueness does not contain the common name as a sub-field (e.g., is a public key or the cryptographic hash of a public key) then this kind of error can not happen. People complain, when faced with a unique name that is the hash of a public key, that this is human-unfriendly and that it forces people to adopt new procedures. They protest, "I am not a number -- I have a name!" They worry that they can't do what they normally do in their daily lives. That is not a flaw, it is a feature. What people normally do in their daily lives is what Bob did in the failure example.

I would suggest that you, as a Gedankenexperiment, consider two different versions of the certificate you're planning to issue: one with the name and no attributes; one with attributes and no name. My claim is that the first will be useless for your needs and the second will be completely useful.

You are right for electronic contacts. But perhaps I want to contact the person by phone and then I want to know the name. If what you want is the postal address, e-mail address or phone number of the keyholder, then those are attributes that should be bound to a key. The certificate that binds those attributes to the key should probably be signed by the keyholder himself, not by a CA. That is because it is the keyholder who is the final authority on how he can be reached. The keyholder is alsothe final authority on the name he chooses to be called. For example, I prefer for people to call me "Carl". I am the authority for that statement.

No, the keyholder isn`t the final authority on his name. In Germany the person`s name is given by his birth certificate. The name can be looked up in the person`s "Personalausweis", an obligatory identity card. In the other way we would have to consider pseudonyms. Pseudonyms may be good for electronic commerce but usually not for health care.

There are different kinds of name here and they are quite separate. That separation is indicated partly by the fact that

certificates for the different kinds of name would be issued by different people.

The name on the birth certificate is chosen by the child's parents. Therefore, they are the authorities on this kind of name. They should issue any certificate binding the full name to a key, but of course by the time the individual has grown old enough to have a cryptographic key, the parents' role is over and there is no reason for them to issue any naming certificates. Still, they are the real authorities here. The birth certificate may include a registration number. The Personalausweis certainly does. Those numbers are a different kind of name and they are chosen by a government agency. Any certificate binding a national ID number or birth certificate number to a key should be issued by the government agency responsible for allocating those numbers. Meanwhile, the child's siblings may give the child a nickname (Hansl. (?)), while the child may become such a huge fan of old rock and roll that he decides to ask his friends to call him Sting. Those names are still names and they function, in certain communities, but they are issued by children and the last one is issued by the keyholder himself.

The fact that a person's full name (chosen by his parents) appears on two unique documents (the birth certificate and the Personalausweis) does not make the full name unique. The full name is still not a valid identifier. It is the ID numbers on those documents that are the unique names.

When I last discussed the Personalausweis with a friend in Germany, I was told that there are laws preventing people from using the national ID number except for official purposes. I did not read the laws involved or spend much time on it. However, such a unique identification number runs a strong risk of being the source of a privacy violation because it can be used to gather a dossier about the citizen, if it were to be widely used. In view of the general European and particular German sensitivity about data privacy, I would be surprised if anyone would allow the national ID number to be in a public certificate that was to be used for multiple purposes. Am I wrong there? I know that here in the US I would never allow my Social Security Number to even be known by a commercial CA much less placed in a public certificate.

So, to come back to your response, there are different names.

The keyholder is the final authority on the name he wants to be called when people address him. For example, I prefer "Carl" and don't like "Mr. Ellison" [which sounds to me like someone addressing my father] but do enjoy the few times people used "Herrn Ellison". These are matters of personal preference and I am the authority. Therefore, I should sign any certificate binding such a name to my key.

In an example I use in the NIST paper, if someone wants to know my name and address for the purpose of sending me money, then I am clearly the best issuer of a certificate binding those to my key. On the other hand, if someone wants to know my name and address so they can send out a team of people to repossess my car or to deliver a subpoena or to arrest me, I'm probably not the person you should trust to sign that certificate. In these examples, the two names look identical but they are different in the way they are used -- and especially in the issuer that can be trusted to provide them accurately.

This is because a public key is yet another unique name for an individual (the keyholder). It isn't human-friendly. Someone would need to keep a data record mapping from the public key to human-friendly information (name, address, phone number, ...) -- but that information isn't needed on-line. The programs that evaluate a security policy based on public key certificates will look at the attributes and ignore the names. Any need to keep an audit trail can record the public key -- and then later, in an off-line process, map that to a person's name (if necessary).

But I want to know that the name is correct, too!

OK. Why? What do you derive from the name that has to be correct? Or is this just a case of still thinking in the local 3-D world where names really work as identifiers? I am not being argumentative. I want to know what value you derive from a name so that I can tell what authority would be qualified to sign the certificate attaching that name (or derived quantity) to a key.

In health care system we can`t work with pseudonyms! I must be sure, that the person which I contact by electronics is the same as I contact in real life.

In my view of the world, you will do this by having a unique number issued for health care purposes only, to identify the person uniquely. [I prefer the hash of the public key as that unique number, but in general there is just a unique number.] When you want to make sure that the flesh-and-blood person is the same as the one you met electronically, you then demand that the flesh-and-blood person prove the right to use that unique number. If this is a sequence number issued by your agency, then the flesh-and-blood person can produce a physical certificate giving that number and his picture and his thumbprint and his signature. If this is the hash of a public key, then the flesh-and-blood person can demonstrate while you watch that he can sign a random challenge with the private key that corresponds to the public key you know has his unique number -- and he's proved who he is. If you want to find him (get his address so you can meet and go through this procedure) then you can ask him on-line for the directions to get to his office.

As I was saying earlier, this shows that what you want is for the common name to function as an identifier. What I was saying is that that fails because there are too many people.

I was in Konstanz for a week last year. It's a nice town, small on the scale of some German cities, but maybe too big already for the common name to uniquely identify a resident. This is especially true when you consider all the tourists who come through.

My belief is that the physical world needs to improve its method of identifying people. A national ID number is one attempt at this, but it runs the risk of being abused (e.g., by those who would build dossiers) and so its use should be restricted. Credit card numbers are another attempt (unique numbers, issued for a single purpose). My doctor here in the US has an ID number (from his license to prescribe medicines, I believe) that he puts on his prescriptions when he writes them. To the pharmacist, that number is my doctor's real name. My doctor's common name is printed on the form also and the pharmacist might use that name also -- for example, when he calls the doctor's office and has to know who to ask for -- but that use of the common name is a social one, not a security one. If these names were provided by certificate, my doctor himself should sign (issue) the certificate giving his common name (since that's just what he prefers to be called) while some state licensing agency should issue the certificate giving the doctor's ID number.

My contribution to this attempt is to observe that a public key is unique and a byte string and therefore is yet another specialized name. It has the advantage that it can be generated without coordination with other generators of unique ID numbers. It is not a sequence number, so you don't need a bottleneck at some agency to allocate the number. Of course, a number by itself carries no validity. That validity comes from the physical document (or electronic certificate) that gives it power and that certificate needs to come from someone authorized to make statements on such a subject.

Your turn.... :)

Carl