Submission of paper to ISSE 2009

Title: Visual representation of electronic identities

By Stefan Santesson

Abstract:

Since the EU directive on electronic signatures was published in 1999, national certification authorities have issued millions of certificates and qualified certificate. But have you actually seen one?

It is a paradox, considering that development of standards for electronic identification using Public Key Cryptography has been going on for a bit over 20 years by now, that we still have no generic solution or standard for how to display a certificate based identity to a human being.

This also says something about our approach to this entire problem space as a community. For decades we have brought the sharpest minds of technology and legislation into projects, forums, associations and standards development activities and we have discussed liability, policy and profiles, but very seldom have we actually discussed the importance of a functioning human visual experience.

Now this of course has a reason, but to seek the reason behind the lack of focus on human interface issues we have to look in many directions.
When X.509 was originally designed, it was created as an extension to the X.500 directory framework, which at the time rested on the idea to create one universal naming structure based on well-defined attributes ordered in a hierarchical structure. This initially erased any need to define a generic structure for expressing different types of identities and to worry about how to communicate that identity to a user. X.509 was an instrument for X.500 names.. period.

Once we realized that we had to fit names that had nothing to do with X.500 into the X.500 naming structure in order to express them in X.509 certificates, we had gotten ourselves into a deep hole from which there was no easy way out. The way we created X.509 and implemented it, it was a pain to define new attributes. This lead to an overuse of certain attributes for various data, which in turn made it impossible to know what information that actually was stored in an attribute. Does the country attribute express the country of residence or country of citizenship? Nobody knows just by looking at a certificate.

Many efforts has been done to try to solve this problem but they have been to complex and have not solved all problems that must be solved in order to allow a totally generic certificate viewer to display a certificate to a human in a meaningful manner, regardless of certificate type or issuer.

Some efforts have been done in RFC 3739, the IETF Qualified Certificate X.509 profile, by allowing issuers to define an object identifier (OID) that would define the semantics of attributes. The problem is that all viewers would have to learn all these OIDs by manual configuration, and that is to hard.

RFC 3709 introduced graphical elements such as Community Logo, Issuer Logo and Subject Logo. The interest in using these graphical images has been significant but implementation has been slow. This solution still leaves the UI maker with unsolved issues, such as knowing what attributes to display and how to label them correctly (to write out what type of information a certain attribute really contains).

Ideas have been discussed to make one big certificate image that could contain all relevant information for a human viewer, as this of course would be the simplest solution. The problem has been the lack of a good open and standardized generic format to present a scalable image with textual and graphical elements that would be displayed uniformly over all platforms. Until the Portable Document Format became an international standard in 2008 (ISO32000:1-2008)

This has lead to a new and intensified effort to once and for all solve this long going and yet unsolved problem by writing a new standard in the Internet Engineering Task Force (IETF). The current editorial team includes representatives from VeriSign and Adobe as well as original editors of the past standards efforts.

The basic idea for this new and very simple standard is to allow a certificate issuer (CA) to design a complete visual representation of a certificate in one combined image file. The standard features the use of PDF as format for that visual representation but is extensible to allow any current or future visual representation format. This “Certificate Image” is then bound to the certificate signature by including a hash over the certificate image in the certificate together with a URI, identifying the location of the image.

The visual representation must be designed to represent a complete visual representation through which the typical user (viewer) will understand the major aspects of the certificate:

- Certificate type (the type and purpose of this electronic identity token)
- Certificate Issuer
- Certificate Subject

The issuer is free to use any combination of graphical and textual elements in various colours. It may include branding marks, logotypes and other graphical design elements as well as textual labels expressed in various languages. Each certificate image can be marked with a language code to allow the viewing application to retrieve he image matching the client’s current language configuration.

The logic deployed by the certificate viewing application is extremely simple as it simply involves raw display of a signed image file (in the form of a FDF document or other suitable image format) once the certificate has been adequately validated. This is simple enough to allow virtually any application to display a certificate image when relevant, such as when connecting to a secure web server or validating the signature on signed software.

The response to this standards effort has been extremely positive and promising with serious discussions with the global community of certificate issuers and certificate viewing platforms. One thing seems clear. If we can find a working solution to this problem, acceptable to all major market players, then we would truly bring electronic identification to a new level.