Avenues for a Diploma Thesis on
E-Mail Sender Authentication

Julian Mehnle

Based on my introductory presentation on e-mail sender authentication, I see two major avenues for a diploma thesis on that field:

Coordinating Sender Authentication Methods ("SPFv3")

All e-mail sender authentication methods are not perfect. Each has specific advantages and inherent limitations compared to others. For example, *path authentication* methods like *SPFv1/Sender ID (SPFv2)* have problems with *alias-style forwarding* (where the sender address is not rewritten during forwarding), whereas *payload authentication* methods like *DK/DKIM*, *PGP*, and *S/MIME* have problems with messages being mutilated by forwarders and mailing lists (thus breaking signatures). This challenges the notion of there ever being a "single final solution" to sender address forgery.

A promising prospect for overcoming this problem appears to be the combination of several authentication methods, allowing them to compensate each other's deficiencies. However, some coordination protocol would be required to enable identity (domain, address) owners to specify a conglomerate of authentication methods that fit their specific e-mail sending infrastructure. (For example, a domain owner might prefer one authentication method over others, but want to specify another one as a fallback for the cases where the preferred method has problems.) The motive would be to create a common declaration language, thereby supporting an "evolutionary market" of authentication methods from which domain owners can choose whatever seem to be the best methods, while giving receivers definite policy instructions.

Originally, *SPF* had started out as *"Sender Permitted From"*, matching *SPFv1's* functionality of merely declaring – through its various mechanisms (`ip4`/`6`, `a`, `mx`) – authorized IP addresses for MTAs sending mail on behalf of a domain. However, after the *SPFv1* specification was mostly feature-frozen, it was realized that the potential of *SPF's* fundamental concept was far greater than that: it is conceivable to extend *SPF* to support the declaration of other sender authentication methods in identity owners' *SPF* records, thus yielding a *"Sender Policy Framework"*. In light of this prospect, *SPF* was so renamed in 2004, however it was already too late for such a significant design extension. Since then, the *SPF* project has concentrated on refining the *SPFv1* specification to produce an *IETF* RFC and on improving implementations. Efforts towards an *SPFv3* to provide an umbrella for other authentication methods have not been undertaken so far.

A thesis about *SPFv3* as a sender authentication coordination protocol would specifically cover the tying of several known e-mail sender authentication methods into *SPFv3*, but could also include research on the semantics of the trust relationships supported by the various cryptographic methods (e.g., "all messages from this domain are signed by keys that are in turn signed by a master key ID xxx"), the interfacing with public key distribution for the specific purpose of e-mail sender authentication (e.g. RFC 4398), as well as on appropriate PKI architectures. Finally, an initial implementation of *SPFv3* could also be developed.

A Domain/User-Grained Reputation System

The concept of e-mail sender authentication addresses merely the sender address forgery problem. While nowadays most e-mail abuse is indeed accompanied by sender address forgery, abusers are known to adopt to new security measures quickly, so combating forgery does not equate to combating abuse. *Content-based* classification of e-mail messages has a good (though not excellent) efficacy record, however experience indicates that this approach leads to a never-ending battle between abusers trying to subvert content filters and receivers

being forced to improve them. With current statistics reporting about 85% of all e-mail sent as being abusive, with a growing tendency, and given the recent rise of sender authentication methods, *identity-based* AKA *reputation-based* classification of messages seems to be worth being explored.

There are IP-address-based reputation systems today already, such as *SPEWS*, *SpamHaus*, or *SpamCop*. IP addresses, however, are a cheap commodity nowadays, given abusers' common practice of mass-hijacking end-users' PCs and building huge bot-nets. As a result, the efficacy of those basic reputation services is limited. As evidenced by recent anti-abuse conferences of the e-mail industry, domain-based and e-mail-address-based reputation systems are considered "the next big step" in abuse control by the leading e-mail service providers and anti-abuse vendors.

However, the semantics of domains and e-mail addresses are vastly different from those of IP addresses. Many findings that have already been learned from IP-address-based reputation systems need to be rediscovered for domain-based and e-mail-address-based reputation systems, and some new ones have to be learned. For example, IP addresses are often dynamically allocated to end-users, while e-mail addresses are usually kept for years, often even multiple ones per user. Also, due to the multitude of roles of e-mail sender identities (MTA domain name, envelope sender, the various header fields, *PGP/S/MIME* key IDs), all reputation for an identity cannot be simply lumped together as with IP addresses but must be treated according to the identity's specific roles.

Reputation can be highly subjective – identity-based reputation much more so than IP-address based one, due to the varied nature of e-mail sender identities. Where current IP-address-based reputation systems mostly use objective criteria in making a binary decision for listing an address as good or bad, a domain/user-grained reputation system has to take individual receivers' subjective criteria into account. For example, one receiver might prefer not to receive mail from domains that have been registered less than a month ago, whereas another might not care about domain age at all or be forced to accept mail from young domains.

Last but not least, every sophisticated reputation system needs input, and from several sources. On top of technical input (such as domain age), user input is required. But how to avoid poisonous input from malicious users? To that end, the possibility of meta reputation within the system could be explored.

Again, development of a working implementation of a domain/user-grained reputation system would likely be in order.