

The AGUPI Paradigm

I want to read only mail that is
1) authenticated, and
2) from a reputable sender, or
3) from an accredited sender

When I give someone my email address, that no longer means I want them to mail me. It means I want them to add me to their whitelist, so I can mail them.

The Aspen Framework

Authentication: Unified SPF + Crypto

Mail senders are expected to

- 1) publish Unified SPF records for all domains and MTAs which send mail,
- 2) sign messages using crypto where possible
- 3) implement BATV/SES where possible.

Mail receivers are expected to

- 1) perform Unified SPF checks on all 4 identities
- 2) check crypto where available
- 3) reject at the edge before DATA

ISPs are expected to

- 1) support SMTP AUTH so users can phone home
- 2) rate limit outbound mail to control zombies
- 3) limit users to mail identities they have authority to use by rewriting cross-customer forgeries of the return-path, and adding Sender headers.

Forwarders are expected to

- 1) implement SRS or SUBMITTER
- 2) prepend Resent-From

Reputation: Karma.com

The source of all reputation is the human decision that a message is spam. (Cloudmark)

MUAs have "this is spam" buttons. Those buttons double as unsubscribe/opt-out triggers for legitimate bulk mail, and as killfile triggers for nonbulk mail.

Those buttons submit information about the offender to a local centralized server, following the SpamNet example, except organized by ISP.

Total traffic flow minus spam reports = inferred ham. ISPs aggregate this information and submit to a reputation clearinghouse.

An alternative to hierarchical aggregation is Mark Langston's Gossip Project, which is P2P.

Reputation can also exist on the personal scale.

Accreditation: The Market

Folks who have no reputation and want to jumpstart the process may wish to sign up with an accreditation system.

There are many kinds of accreditation: putting down a bond is only one kind. (Bonded Sender)

Giving certified proof of your identity is another. The honest have nothing to hide, but spammers can't give out an address where they might get sued (or arrested). (Verisign VDL)

What about small new domains who don't want to engage accreditation? We can just greylist the first time we see mail from them, and hope that the reputation system will react fast enough to recognize them the next time they try to send mail.

What about slackers?

Domains that don't publish SPF records get a default record: we assume they are happy with "v=spf1 a/24 mx/24 ptr -all". If they aren't happy with that, they should publish SPF. Even a "v=spf1 ?all". This needs to be coordinated, of course.

At the Personal Level, Two Kinds of Trust

Direct Trust: I have sent mail to Alice.
I will read mail from Alice.

Trust by Introduction: Alice mails me and CCs Bob.
I will read mail from Bob.

See also <http://loaf.cantbedone.org/> and Twingle by Simon Cozens

An Abuse Reporting Standard

"This is spam" buttons trigger a report to a local server. Those reports contain essential information about the alleged violation, possibly including the original, verbatim. The local server passes them on to the responsible sender domain for followup action. (AOL scomps) It also scores the sender negatively in the reputation system.

Karma.com: the Clearinghouse

Patterned on the model of Visa, the reputation clearinghouse is jointly owned by all participants (producers and consumers of reputation). The revenue model sees fees flow from consumers to producers. ISPs are motivated to share data because that lowers their usage fees: they become a producer. Joint ownership reduces friction and seeds future governance.

There are two types of commercial email

1) marketing 2) transactional
Smart senders separate the streams.