TrustNetNotes

Notes About A Possible "TrustNet" PlugIn

This page is my braindump for a bunch of PlugIns which I summarize under the name TrustNet. Their intention is to WhiteList people one exchanges mail with based on the WebOfTrust between the the sender and the recipient. Comments welcome. (MalteStretz)

Current ideas:

PGP

- Use PGP to try to find a trustpath.
- Checking the direct trust should be relatively simple if one has access to the user's keyring. But what about server side filtering?
- Is it feasible to query servers for an indirect trust path, too? What's the overhead?
- JustinMason: imo this is definitely a good idea. I'm worried about the CPU overhead of checking GPG sigs, but caching recently-seen "good" sigs in a cache keyed on From-address and first untrusted IP address from the Relays header would help that. Also, gaining access to GPG from perl isn't easy; the CPAN modules are not great. imo the cleanest option may be running GPG directly from a plugin.

This would be a very good idea. The bigist problem is the resorcese and time it would take to check the signatures of a message and then find the chain between the servers key and the senders key.

I think it would be a good idear to give each instalation of spam assasin its own gpg key and use that key to sign the keys of the users of the server and the keys of any other servers that are used at the same site or are cominly comunicated with. This way you are only finding the key chain from one key to the sender and the cache database would be easer to implement.

FOAF

- How can we incorporate FOAF? Querying the website each time has quite some overhead, some caching is needed.
- How to access? XML-RPC or some DNSDB gateway? (Have you noticed that DNS gets abused for quite some things?)
- JustinMason: in thinking about this in the past, I considered that possibly the best way would be to have a crawler run from cron which generate a local cache of the remote data. however, one issue is that FOAF does not specify relays, just email address hashes; so this means that it's vulnerable to spammers faking the From addr. See 'Using From For Whitelisting Problems' below.

Web-O-Trust

- The Web-O-Trust project is relatively dead, maybe we can revitalize it.
- It should be possible to implement the Web-O-Trust syntax in XML and put it into FOAF files.
- JustinMason: I have always argued that Web-O-Trust needs a way to specify various degrees of trust, as well; ie. "this server will never originate or relay spam", "this server is trusted not to be subverted by spammer code, but may relay spam originated elsewhere", etc.

LOAF

- I dont't like the idea at all, but LOAF might be worth looking into, too.
- JustinMason: big problem in my opinion is that the LOAF files are attached to each mail sent. bulky and messy!
- MalteStretz: ACK, thats what I dont like about it, too.

Geo info

- This posting about LOAF made me think that it might be possible to use a website's published Geo information (how near am I geographically to the sender).
- JustinMason: several spammers live near me!
- MalteStretz: but they probably won't publish Geo records 🙂 and if they start to do (probably not targeted ones but ones from high density urban areas), this rule won't work for you but maybe for people living at uncommon places

Querying Addressbooks

- I already implemented a quick hack for to query my KAddressbook from KMail for whitelisting. What about querying LDAP servers?
- JustinMason: see also 'Using From For Whitelisting Problems' below

Social Networks

- I guess quite some of them have some API available so it should be possible to write specific plugins for the services. I'd prefer if they just published FOAF profiles though.
- JustinMason: http://www.tribe.net/ publishes FOAF.

Six/Four

- A friend pointed me to Six/Four, no clue how that could fit in, just noting it here.

Using From For Whitelisting Problems

One common problem that appears when using just email addresses for whitelisting, is that spammers routinely fake the From address to appear to be

- the recipient's email address
- another email address at the recipient's domain
- another email address from the address list in the recip's domain

if we use just the From address with an address-based whitelisting scheme, it will be vulnerable.

The solution is to either:

- use IP address info from the Received headers or the last untrusted relay, and combine that with the address to come up with a combined email-and-ip address, similar to how the AutoWhitelist does it.
- require that any whitelisted address be on a domain that publishes SPF records.
- MalteStretz: cross-checking against SPF for entries for which no routing information is avaliable (ie. addressbook entries etc) is a good idea. another possibility to avoid some FPs could be to exclude all domains which are equal to the recipient's one
- JustinMason: we have theorized that spammers could scrape addresses that occur in conjunction; e.g. together on a mailing list archive page, or a "contact us" page. I don't know if this has happened "in the wild" yet. viruses certainly already do it though.