Meeting Notes 2008 Nov 16

DDX face to face meeting, IETF, Minneapolis
2008-11-16, 1:30 - 5:30 PM

Attendees: Dave Crocker, Lucy Lynch, Leif Johansson, RL "Bob" Morgan, Chris Bongaarts, Steve Siirila, Ken Klingenstein, Jess Thompson, Jeff Eaton, Fredrik Petta, Paul Russell, Josh Tewell, Mark McCahill, Tony Hayes, Patrik Wallström, Jeff Hodges

Action items:

- RL "Bob" will ask US government contacts about interest in DKIM.
- Everyone will post "user stories" to wiki and mailing list.
- RL "Bob" and Internet2 folks will organize biweekly conference calls.

Bob introduced the activity by observing that it is attempt to bring together many different communities:

- the Internet Society and its Trust and Identity Initiative, promoting a more trustworthy Internet;
- the Internet2 Middleware Initiative promoting more effective and consistent identity and access management and security practices among US research and higher education institutions;
- the InCommon Federation building trust infrastructures for a variety of applications for US R&HE and its partners;
- the global R&HE federation community, including federations in Sweden and many other countries;
- R&HE email administrators via the hied-emailadmin mailing list;
- DKIM standards contributors, primarily those participating in the IETF DKIM WG;
- DKIM adopters and promoters including many large commercial email providers.

Dave Crocker provided an introduction to DKIM and some thoughts on deployment strategy. Highlights:

- Basing trust assessments on incoming messages on domain names rather than IP addresses has several benefits but also some costs, such as needing to receive and process message headers.
- Much anti-abuse practice is based on mistrust: looking for bad behavior and quarantining it. Working with trust is quite different: success creates a trusted channel, and problems are errors, not necessarily abuse.
- DKIM is a way to create a separate channel of more-trusted email that can be handled differently.
- An important feature of DKIM is that third parties (non-authors) can be signers.
- It is important to base signer identity on the base domain name, not the full key selector name.
- Successful DKIM interop has already shown among several implementations and many deployments.
- DKIM-provided domain-name-based authentication is a necessary first step, but reputation service is needed for relying parties to find it useful.
- Pilot project could be based on a "published member list" as minimal source of reputation. This could be useful in many industry sectors. It wouldn't try to compete with existing "email reputation industry".

Discussion amid the talk:

Jesse Thompson noted that many schools are seeing spammers/phishers using stolen university accounts. Will DKIM adoption make unversity accounts even more attractive targets since our mail will be more trusted?

Ken Klingenstein observed that marketing materials for DKIM adoption need to target CIOs separately from email admins in terms of benefits and risks.

Chris Bongaarts said that many sites already put .edu IP addresses on white lists, so this proposal is doing the same thing in a different, probably better, way.

There was lots of discussion about the use of distinct domain names as signing identities, to create different mail streams with different characteristics. A key question is whether a large site would sign all messages it forwards. If the point of signing is to indicate a trusted or high-quality mail stream, does that imply the MTA is guaranteeing that by scanning for spam/phish/virus/etc? If we're not doing that for all messages should we sign? Or is it better only to sign a subset with known characteristics, and indicate its nature by a special domain name? The signing domain is independent of author/sender domain, so this is possible. There is no necessary relationship between domains, eg foo.com and bar.foo.com don't necessarily share policy or anything. ADSP may support declaring the semantics of a signing domain.

In SAML there is a similar distinction between the entityID of an Identity Provider (like the DKIM signing domain) and user attributes (eg userid, email address) that may be sent in an identity assertion.

It is legal and useful to have more than one signature, so different forwarding parties may sign independently. A broken (ie failing to verify) signature from one party doesn't invalidate a good signature from another; probably just means the message was altered in transit.

Authenticating messages can help message receivers even without a reputation service or trust community since they could analyze received messages from source X to establish a pattern hence a confidence level that could be input to filtering.

Is there utility in having conventions for the use of signing subdomains? Subdomain could indicate org hierarchy such as department. Or it could indicate that submitters were authenticated.

Minnesota is deploying a service to support indication of "official" communication from campus authorities. The officialness indicator is in a message header and is interpreted by the campus webmail system which highlights the message. This could be implemented more securely and generally via DKIM. More generally MUAs could re-enable image loading (eg logos) for messages from trusted senders. A logo or other image could be put in or linked from the DNS record to be displayed with verified messages (similarly to logos in X.509 certs). Would large providers like Google/Yahoo/Live buy into a scheme like this? Maybe. The could also be inter-institutional uses.

ADSP lets sending orgs make statements about their signing. In universities (and many other orgs, it's said) there are many MTAs sending mail beyond the borders, so it would be hard to say "we sign everything". ADSP may still be useful in other ways; or controlled subdomains could say "we sign everything".

Graylisting is used by some sites to delay mail delivery from suspect sources (includes other methods too?). Signed mail from trusted sources could bypass graylisting, so improve delivery times for trusted partners.

How much variation is there in the methods different sites use in email filtering? If there is a lot of variation, does that make using common approaches to introducing DKIM more difficult?

Faculty sometimes report delivery delays or failures for messages to/from government funding agencies, causing unhappiness around grant submission deadlines etc. Improving delivery success for these trusted partners would be a real motivator. We should check with USG folks about status and interest.

Pretty much all sites now do initial filtering of messages based on IP addresses as listed in various blocklists. These messages are rejected before their headers are ever seen, thus before verification could happen. Does this make signing moot? No, since better filtering on remaining messages is still valuable. Presumably trusted sources aren't usually in blocklists, or you're already rejecting much useful traffic.

How much CPU does DKIM take? Hard to measure, but Yahoo may have seen 10% CPU increase for signing. Some sites think they have no CPU to spare, some sites say CPU is fine, it's I/O that's the problem.

Is it useful to do less spam/virus/etc processing on signed/trusted messages? Probably not. USC knows they generate lots of spam (stolen accounts etc), so filtering still necessary. Savings on processing time is likely to be low.

Does DKIM utility depend on having a Campus Email Policy regarding things like who can send/receive, what things can be sent, etc? Such a policy could help when using features like subdomain signing to indicate special mail streams, but having a policy isn't necessary to get started.

Patrik Wallström of iis.se gave a presentation regarding DKIM and DNSSEC. DNSSEC is fully supported in .se but still sees little significant use. Various end-user applications might be motivators for DNSSEC, but deployment issues such as broken DNS behavior in SOHO NATs/routers is a big problem. DKIM is appealing because it is a server-based technology. DNSSEC removes a substantial security hole in typical DKIM key-fetching. DKIM-Milter 2.8.0 beta out now - http://opensource.is.se/ .

Is it feasible to include DNSSEC support in the DDX activity? Maybe, but DNSSEC in .edu is not imminent. .gov is being signed, so that might be part of working with USG sites.

In conclusion there was consensus that DKIM technology is worth pursuing by the institutions represented and that a community-based trial has promise. Work needs to be done to describe the benefits, understand the tools available and the deployment decisions to be made.