DomainKeys Identified Mail (DKIM)

Murray Kucherawy
The Trusted Domain Project
<msk@trusteddomain.org>

PART ONE

Origins

- Phishing was beginning to appear in earnest early in the last decade
  - Spoofing was already rampant, but less painful
- Yahoo! developed an experimental domain-level signing method for email and deployed it in 2003
  - This was *DomainKeys*
- After a successful initial test with a few partners, they sought review and advice from outside experts
  - Cisco proposed a variant they thought was better, which created a lot of confusion and contention
- An industry consortium (ESTG) was formed that lasted for about 18 months to develop and test a merged proposal to send toward standardization through the IETF
  - With some gentle FCC pressure

Internet Engineering Task Force

- Produces standards documents based on a model of “rough consensus and running code”
- Meets three times per year in changing locations
- Divided into several *areas*
  - General, Applications, Security, Routing, Transport, Internet, Real-time Applications, Operations & Management
- Formed a DKIM working group on that basis under the Security area
- Obviously a community interested in doing the work
  - There were already implementations
- Working group was chartered to produce the protocol document and associated informational documents

Toward Standardization

- Email, unlike other Internet application protocols such as HTTP or FTP, has a long history of being accommodating, adaptive, flexible
  - This has big benefits, but also big costs
- This also makes developing something new and getting it widely adopted a real chore
  - The email community is very broad, with a lot of perspectives and varied experiences to be considered
  - Extremely sensitive to disrupting the deployed base
- Took about 18 months to produce the first DKIM standards document

How DKIM Signing Works

- Prepares the message for signing
  - May *canonicalize* it to anticipate some mutations
  - Whitespace addition, header re-folding, extra blank lines, header re-ordering, new header fields, removed header fields
  - May set certain other handling parameters (e.g., signature expiration, body length)
- Digitally signs the message using a key that will be associated with a domain
- Places that signature in a new header field
- Sends the message
- Note: Signing domain is not necessarily related to any other domain in the message, visible or otherwise
  - It may or may not be an *author domain signature*

How DKIM Verifying Works

- Extracts the signature from the message header
- Parts of the signature tell the verifier where the public key can be found
  - It's in the DNS of the signing domain
- Repeat the "prepare for signing" steps with what was received and see if the result verifies against the signature
- If the signature verifies, we say the signing domain took *some* responsibility for the message

Some Important Points

- DKIM signatures are not reliably associated with individual users
- A valid signature imparts no guarantees of valid content or safe intent, so “pass” doesn’t always mean “good”
  - Indeed, spammers were some of the first to adopt
- Similarly, an invalid signature doesn’t automatically mean sinister intent
  - There are lots of valid things that break signatures, like mailing lists
- Worth repeating: The signing domain might not be the same as the From: domain or any other domain in the message

So What Do We Have?

- We have a domain name whose owner is at least partly responsible for this message
- This allows some useful choices that weren't reliable before
  - Policy agreements, like Yahoo!-eBay
    - This led to ADSP, and eventually DMARC
  - Collecting data for a while allows one to observe patterns in domain behaviours:
    - Correlation between short-lived domains and abusive behaviour
    - Domains whose mail is low quality (many complaints) or high quality (few complaints)
    - This is effectively domain *reputation*

A New Approach

- We have spent decades trying to identify bad actors and bad content, and block those
  - But once identified and filtered, they just move
    - New IP address, new domain, or simply a subdomain
- Maybe we should focus on giving preferential treatment to good actors
  - Heavily filter or rate-limit everyone except sites that earn good reputations via DKIM
  - No benefit to moving around; benefit to sitting still and behaving

What Next?

- In addition to conducting open source reputation R&D, TDP is developing the *Affil* program
  - Some authority, such as a government, publishes a list of domains owned by legitimate institutions in a class (banks, charities, etc.)
  - Using DKIM-verified domain names, one can determine whether a piece of email came from such an institution by checking the list
  - Mailbox providers indicate visually if the authority has endorsed the signing domain as legitimate

What Next?

- With the momentum toward IPv6, email will also need to make the transition
  - Some popular defenses on IPv4 aren't portable
    - The address space is just too big to track, even at the 64-bit aggregation level
    - Reverse DNS is not expected to be reliable
  - Email authentication via DKIM and SPF will become mandatory, at least for IPv6 if not universally

Preparing for DKIM Signing

- The reputation of your domain name will be based on the quality of your mail stream
- To begin signing your mail:
  1. Audit your policies
  2. Develop a key rotation policy
  3. Scan your mail on its way out
  4. Audit your systems for “leaks”
  5. Think about DNSSEC
  6. Make nice with your DNS people
  - You will be asking them to put funny things in the DNS, and creating load for them

Preparing for DKIM Verifying

- DNS load will increase a bit varying with your inbound volume
- You have to answer the “so what?” question
  - Try a reputation system?
  - Have an internal favourites list?
  - Just observe and record for a while?
- You have to figure out if unsigned mail should get special treatment
- No right answers yet
  - Ask four experts and you’ll get five opinions
  - We’re still mainly using IPv4 so we haven’t been forced to come up with best practices

PART TWO

DKIM Adoption?

- Difficult to measure actively because there's nothing to query from outside to confirm participation or volume
- Passive measurements and anecdotes are the best we have

OpenDKIM

- Open source project including:
  - A DKIM library written in C for use in applications
    - In use at AOL, Yahoo!, Facebook
  - A C application to add DKIM service to MTAs that use Sendmail's milter (postfix, sendmail, Oracle)
  - Several deployment tools
  - Experimental collaborative reputation systems
- Forked in 2009 from Sendmail's *dkim-milter* open source project

Implementation Report

- September 2010-March 2011, 21 sites collected and reported DKIM data to TDP
  - 11.4M messages, 28.3% signed
- Last six months, 10 sites still reporting to TDP
  - 8.4M messages, 48.1% signed, ~90% pass
  - 38607 domains sent passing signatures

DKIM Deployments of Interest

- US Government
  - usps.gov, mail.house.gov, senate.gov, ftc.gov, fdic.gov, dot.gov, va.gov, gsa.gov, fairfaxcounty.gov, cityofboston.gov, eop.gov, hq.doe.gov, ssa.gov, us-cert.gov, us.army.mil
- Financials
  - Bank of America, Wells Fargo, American Express, Citibank, Scotiabank (Canada), Fidelity
- Mailbox Providers
  - AOL, Yahoo, Gmail, Hotmail
- Social Media
  - Facebook, LinkedIn, Yelp, Foursquare
- Mass Mail
  - Constant Contact
- Other Industries
  - United Airlines, Virgin America, Delta Airlines, FTD, Apple (iTunes), Orbitz, Globe & Mail, Aria Las Vegas, ZipCar, Barnes & Noble, Hertz, TiVo, Hewlett Packard, Petro Canada

Author Domain Signing Practices (ADSP)

- A way to make domain-level policy statements, also via the DNS
- Attempts to bind DKIM to the From: domain
- Makes a request to receivers for handling mail without a valid signature
- Highly controversial
  - Interferes with some normal mail operations
    - Mailing lists
  - Can’t protect subdomains
  - Display name attacks are prevalent

DNS Security Extensions (DNSSEC)

- Digital signing of DNS responses
- Allows detection of falsified DNS replies
  - For DKIM, allows detection of fake DKIM keys and ADSP assertions
- Not widely deployed yet, unfortunately
  - Maybe we need something to go very wrong before people will deploy it

Authorized Third Party Signatures (ATPS)

- Allows a domain to indicate that a third-party domain signing its mail should be considered authorized
- Yet another DNS mechanism
- Experimental
  - Much demand for it in the DKIM working group from a vocal minority, but nearly zero adoption so far

DMARC

- Domain-based Message Authentication, Reporting, and Compliance
- Another policy mechanism that also includes reporting
  - "Did this message pass either SPF or DKIM?"
    - "If not, reject it"
  - "Also, please send me daily reports"
- Supports gradual rollout
- Reporting includes aggregate and forensic reports

DMARC

- Like DKIM, developed privately at first by a consortium of interested parties
  - PayPal, Google, Cloudmark, Return Path, Microsoft, Facebook, AOL, Yahoo, Brandenburg InternetWorking, TDP, LinkedIn, Fidelity, Bank Of America, Agari
- Now making its way toward the IETF for development and standardization
- Already protecting over 60% of the world's mailboxes

Authentication-Results

- Header field defined in RFC5451 used to record the results of any of the above authentication schemes for later use by filters or users
- Example:

```
Authentication-Results: medusa.blackops.org;
 dkim=pass (1024-bit key; insecure)
   header.d=email.bhphoto.com
   header.i=@email.bhphoto.com
   header.b=EtU9jhhj;
 dkim-adsp=pass;
 dkim-atps=neutral
```

Why Should You Adopt DKIM?

- With IPv6 coming, the messaging community is planning to apply much more pressure to have mail authenticated or be heavily filtered
  - Maybe even blocked outright
- Operators experimenting with domain reputation can't distinguish one piece of unsigned mail from another
  - Do you really want to be in the same group as all the other non-signing entities out there?
- If you're a good actor, sign your mail!
- If you have a brand to protect, programs like *Affil* can't help you if you're not signing