DMARC Overview

Murray Kucherawy
<superuser@gmail.com>

History

- In 2009, a few large industry players started discussions about how to expand on the DKIM and SPF services to improve phishing countermeasures
- A small consortium of companies was formed to develop a specification and some implementations to test it
  - Grew to ~15 members
- Released to public for scrutiny and feedback in January 2012
  - Includes a web site and public mailing list
  - Document has undergone a few revisions since then
  - Ran an interoperability event

Overview

- Phishing is an expensive problem
- There are some protocols that provide authentication layers on top of email, but by themselves they aren't enough
  - They protect invisible things
- We need something that runs on top which:
  1. Uses available, deployed authentication schemes
  2. Increases detection of From: field abuse
  3. Provides strong but "scalable" policy options
  4. Adds comprehensive reporting capabilities

Policy Component

- Attempts to determine if the domain found in the From: field of a message was used by an authorized author
  - SPF and DKIM don't attempt to validate use of the From: field, but that content is virtually always shown to the user
- If the domain validated by DKIM or SPF matches the From: field domain, the message passes the DMARC test
- If not, policy action can be taken by the receiver

Policy Component

- Policy is retrieved from the DNS of the domain found in the From: field
  - Can request that a message be quarantined or rejected if it fails the DMARC test
  - Optional separation of policy in terms of domain vs. subdomains
- Domain owner can also select a percentage of mail to be thus affected, allowing for experiments and gradual roll-out

Why DKIM and SPF?

- SPF determines path authorization
  - Validates use of the MAIL FROM domain
- DKIM confirms association of the content with a domain name (the signer)
  - Validates use of the “d=“ domain
- They have obvious failure modes, but they don’t overlap much
- The union of their “pass” modes appears to be quite sufficient for DMARC’s goals

Reporting Component

- Supports two modes of reporting
  - **Failure**: details about every message that fails the DMARC test, using work done by the MARF working group
  - **Aggregate**: daily summaries of mail that failed the DMARC test and were subjected to policy action
- Has shown to be enormously valuable in finding phishing perpetrators, identifying infrastructure "leaks", and debugging
  - Helpful in identifying email sending partners that aren't configured properly for authenticated email
  - Also useful in with M&A infrastructure monitoring

Subdomains

- An easy way around prior policy work (e.g., ADSP) is to use a subdomain
  - You could protect example.com itself with ADSP, but then attackers can just use security.example.com
  - The DKIM WG had a protracted battle about how to deal with this, and eventually didn't
- DMARC needs a way to plug this hole

Subdomains

- Use the *public suffix* list to decide where to ask for policy if there's not a specific one
  - So for security.example.com, we know to also ask example.com for a policy
- Trent will talk more about this

Implementation

- Open source implementations
  - One complete package, one set of open source extensions to a commercial MTA
- Some patches and modules to commercial MTAs available
- Numerous proprietary implementations
  - All of this has actually been a useful secondary shakedown of SPF and DKIM implementations
  - Also has provided a lot of signal to spam trap operators
- Intermediaries do report processing on behalf of domain owners
- Estimated coverage of 60+% of global user mailboxes

Questions?

- Ask away!