Fighting Spam

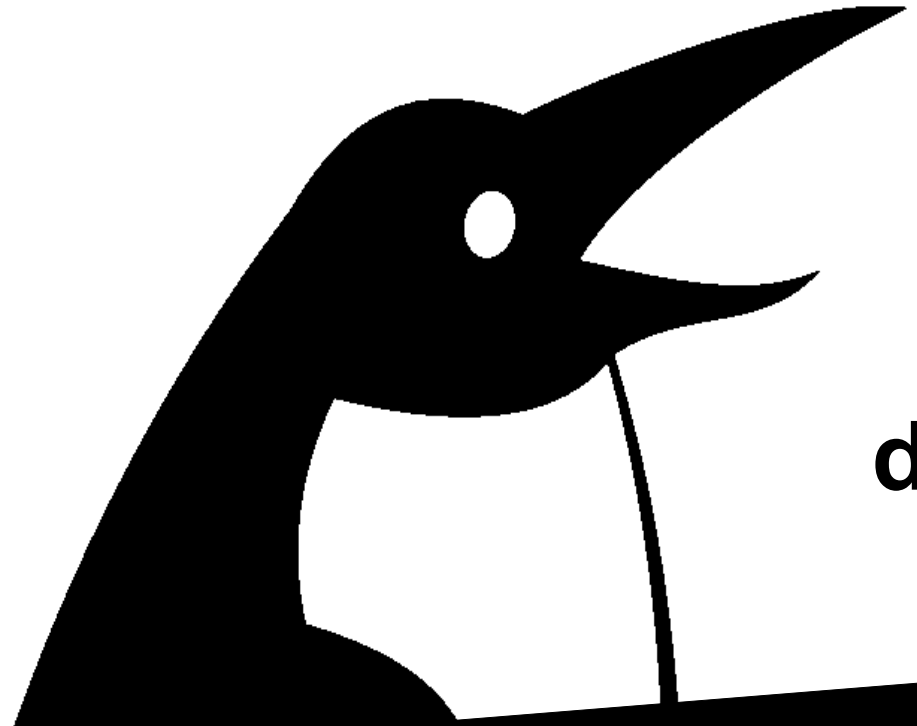

2017-10-18

Dianne Skoll

Roaring Penguin Software Inc.

dfs@roaringpenguin.com

Approaches to Fighting Spam

- Reputation-Based (IP, Domain)
- Authentication (SPF, DKIM, DMARC)
- Behavior-Based (greylisting, botnet detection)
- Content-Based
- Defense in Depth

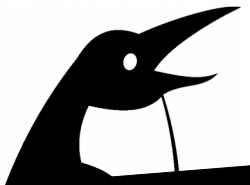

IP Reputation

- Typically implemented by DNSBLs.
- Reactive – IPs are listed only after they spam.
- Some DNSBLs are high-quality. Most are not.
- Few are transparent as to listing and delisting criteria.
- Few have good IPv6 coverage.
- Useful as a first pass to cut down on spam passing to the rest of the filtering stages.

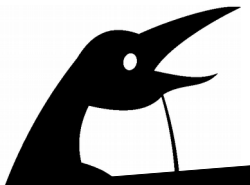

Domain Reputation

- Also typically implemented by DNSBLs.
- Reactive.
- Low to moderate hit rate.
- May be applied to sending domain and/or to domains of URLs in the message body.

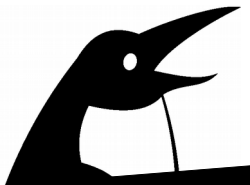

Domain Reputation - 2

- Spammers often register throwaway domains as sending domains.
- Idea: Penalize messages from “newly-seen” domains.
- CanIt 10.1.7 tracks domains seen across all CanIt installations and permits you to (mildly) penalize mail from only-recently-seen domains.

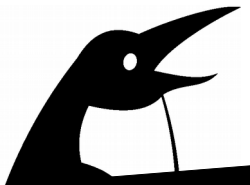

Authentication: SPF

- SPF (Sender Policy Framework) is a mechanism whereby domain owners can declare which machines may send email on their domains' behalf.
- For arbitrary domains, an SPF "pass" is a mild *spam* indicator!
- Spammers are better at setting up SPF than many legitimate administrators.

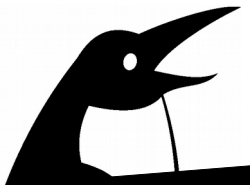

Authentication: SPF - 2

- SPF is useful for *trusted* domains (banks, PayPal, eBay, etc.)
- Adding points on SPF “fail” or “softfail” is useful.
- Subtracting points on SPF “pass” for arbitrary domains is dangerous.
- Subtracting points on SPF “pass” for trusted domains is useful. But you have to know who to trust.
- CanIt supports SPF and flexible SPF rule.

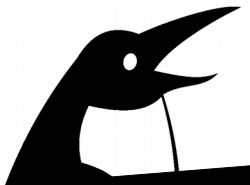

Authentication: SPF - 3

- Because SPF authenticates the sending relay and not the message content, it cannot detect message modification.
- It also makes forwarding of mail annoying; a hack called SRS (Sender Rewriting Scheme) was invented to make forwarding work.
- CanIt supports SRS.

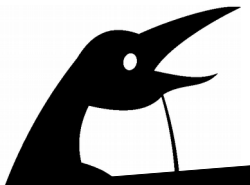

Authentication: DKIM

- DKIM (DomainKeys Identified Mail) is a mechanism that authenticates message headers and bodies.
- More powerful than SPF because it prevents message alteration. Does not break fowarding.
- However, DKIM signatures may be broken inadvertently by sub-par email implementations.

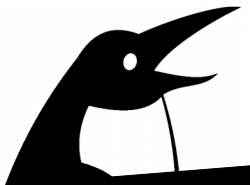

Authentication: DKIM - 2

- As with SPF, a DKIM “pass” for an arbitrary domain does not indicate good mail.
- But a DKIM “pass” for trusted domains is a good indicator of ham. Once again, you need to decide which domains to trust.
- CanIt supports both checking DKIM signatures and DKIM-signing outbound mail.

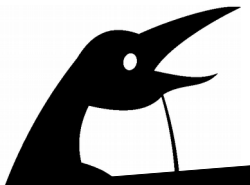

Authentication: DKIM - 3

- Since DKIM keys are typically published using DNS, you should use DNSSec for maximum security.
- This comment applies to SPF also, since SPF records are published using DNS.

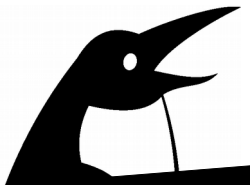

DMARC

- DMARC (Domain-Based Message Authentication, Reporting and Conformance) was designed to patch up DKIM and SPF.
- The problem: Neither DKIM nor SPF specify policy. They merely specify that a mail was or was not relayed via an approved relay, or was/was not correctly signed.
- Recipients have to figure out on their own what to do with this information.

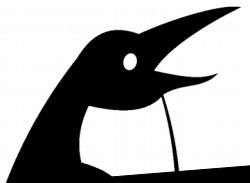

DMARC - 2

- DMARC lets domain owners publish (again via DNS) a policy that says what recipients *should* do if SPF and/or DKIM fails.
- DMARC also includes a reporting mechanism to provide feedback to domain owners that their domain is being spoofed.
- CanIt supports DMARC checking, but not yet reporting.

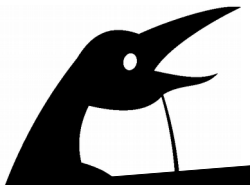

DMARC - 3

- CanIt also lets you override DMARC policy, converting a “reject” to a “quarantine”.
- Why?????
- Because alas, some administrators get the setup wrong and end users complain about rejected email.

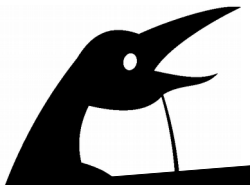

Authentication Summary

- SPF/DKIM/DMARC are useful for letting legitimate mail from trusted domains through.
- Somewhat useful for blocking some forms of spoofed mail.
- Useful for you to use on your outbound mail so you can repudiate spoofed messages, preserving your domain's reputation.
- Not terribly useful to fight spam. :(

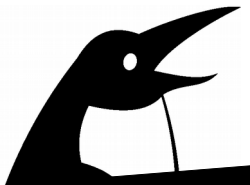

Behavior: Greylisting

- Greylisting forces a newly-seen mail client to spool mail it wants to send for a short time.
- CanIt does this by hashing (part of) the sending IP, the sending email address, the recipient email address and the subject line. The first time it sees this hash, it replies with a 4xx “Try Again Later” response in the SMTP conversation.
- Legitimate clients will retry and get through.

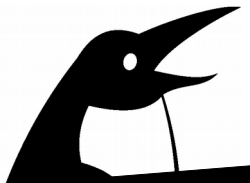

Behavior: Greylisting - 2

- Greylisting is *very* effective at stopping spam software that either doesn't retry or mutates the message with each retry.
- It also has a synergistic effect with IP-based DNSBLs: By forcing a spammer to stay pinned to an IP address, you increase the chances of a DNSBL "catching up".
- Downside: It can introduce delays from new senders. CanIt mitigates by exempting an IP from greylisting for 40 days if it passes.

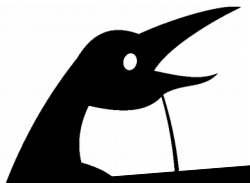

Behavior: Botnet Detection

- The idea here is to detect “similar” messages that come from a large number of different IP addresses.
- This typically indicates a botnet.
- We have experimental code to trawl the mail logs looking for botnets, but the signal-to-noise ratio is not yet good enough to completely automate this.
- Good research topic, however.

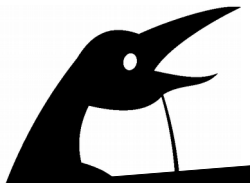

Content-Based Detection

- This is the bread-and-butter of spam filters.
- Heuristics: Hand-written (typically) rules that check messages for spam attributes.
- Learning: Statistical methods (typically Bayes) that learn what is spam vs. ham based on human feedback.
- Malware detection: Virus-scanning, Microsoft Office macro-detection, .exe files, etc.

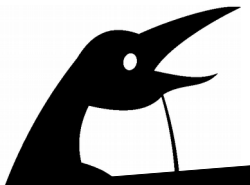

Content-Based Detection - 2

- Content-based tests are the last defense to stop delivery of spam.
- Heuristics require updating and are tricky to create; users often mess up custom rules.
- Bayes is very effective, but can take a while to “lock on” to a new spam trend.
- Targeted attacks are *very* difficult to detect, especially from skilled attackers who mimic style and grammar of legitimate senders.

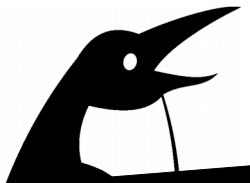

Content-Based Detection - 3

- CanIt has extensive content-based tests:
  - SpamAssassin Heuristics
  - Roaring Penguin SpamAssassin Rules
  - Bayes, including shared Bayes
  - Office Macro detection
  - Filename rules
  - Virus scanner integration
  - Custom Rules
  - Extraction of URLs from messages and PDFs

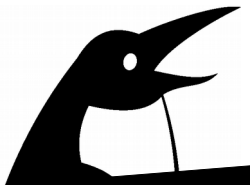

Defense in Depth

- Inevitably, some unwanted mail will get through.
- Often targeted attacks, which are particularly devastating.
- Many attacks rely on a user visiting a web site and either providing her credentials to a malicious actor or downloading a malicious payload.

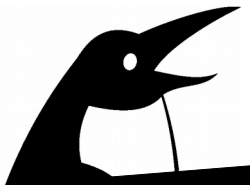

Defense in Depth - 2

- CanIt's URL Proxy feature rewrites URLs to point to CanIt.
- When a user clicks on a link, he is taken to a CanIt landing page that shows:
  - All redirections
  - The name of the server in the URL
  - The location (geolocated from IP address)
  - A configurable warning message not to provide sensitive information

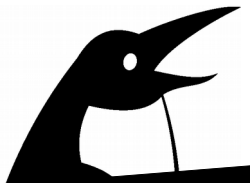

Defense in Depth - 3

- If an email with a malicious URL gets through and then by the time the user clicks on it, CanIt knows that the URL is malicious... the URL proxy *completely prohibits* the user from visiting the original URL.
- Otherwise, the warning serves to alert users to potential scams or malware.

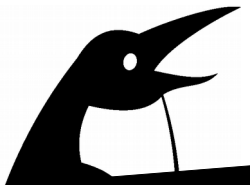

Q & A