Algorithmically determining Store-and-forward MTA Relays using DomainKeys

Peter Ludemann
Yahoo! Inc
701 First Ave
Sunnyvale, CA 94089

Miles Libbey
Yahoo! Inc
701 First Ave
Sunnyvale, CA 94089

ABSTRACT

Store-and-forward MTA relaying servers have frequently presented problems to various antispam techniques, such as IP-based reputation or email authentication. Algorithms that find email relaying servers can use knowledge about a domain's outbound IP addresses combined with cryptographic domain authentication frameworks such as DomainKeys. This paper presents one such algorithm.

1.Why find relaying servers?

In this paper, the term relaying will be used to describe the situation where messages intended for an email user are systematically and automatically delivered to a non-local address. Many industry terms are used for this action: forwarding, redirection, lifetime email addresses, etc. The feature is prevalent in university's alumni email accounts and Internet access providers. Unmoderated mailing lists also have many of the characteristics of forwarding servers. In this paper, the term is specifically for MTA relays and not used to describe the use of a Forward button or Bounce/redirect option in a MUA.

Email that has traversed a store-and-forward MTA relay is generally indistinguishable from a forgery to a receiving system. Without email authentication technology, the connecting IP address is the only data piece of an email that is not forgeable. In a relaying scenario, the IP address connecting to the final recipient's mail server is not one associated with the message's originator. Instead, the connecting IP address will be present in a Received header, perhaps several below the one describing the final network hop. Generally, this data about the originator is not trustable by the receiving system, as a spammer can pretend to forward an email by adding a faux Received header at the top of their transmission. However, if a receiver had a list of servers that it trusted to properly relay, rules could be developed to parse Received headers to find the originator and then apply IP-based reputation filters, or authenticate the email using path-based models such as SPF or Sender ID that are ineffective in relaying situations.

Relay servers are also likely false positive candidates for sender reputation. With spam email between 65-85% of normal traffic, relaying servers will likely redirect similar percentages of spam. This rate of spam would mark the relaying server's reputation as negative, because the spam rate would be orders of magnitude worse than a best practices sender's spam complaint rate per message. As a result, relays are more likely to be treated as second class (or worse) mail, experiencing deliverability problems such as tagged false positives and degraded performance from greylisting [4] and teergrubing [5], or even message rejection. If a receiver can algorithmically determine a forwarding server, different rules could be applied to avoid this treatment.

This need to reliably determine relay servers creates a transient trust dilemma for the receiving system. If the receiving system blindly trusts Received headers to determine relays, it may enable the spammer to forge email and slip by filters, an unacceptable risk. If it could algorithmically determine the auto-forwarding servers, these risks would be significantly mitigated.

2.The algorithm

The combination of a cryptographic email authentication solution such as DomainKeys [1] and outbound email servers (potentially from published SPF [7] or Sender ID [8] records) allow a receiver to algorithmically find auto-forwarders. If the receiver receives an email that is DomainKey verified and the connecting IP is known not to be an authorized outbound sending IP for that sender, the receiver can reliably determine the email has either been relayed or traversed a mailing list that has not modified the message's content. Further removing those emails that either contain a List-ID: header (assuming the initial email did not contain it) or a bounce address (2821.From) in a different domain, the receiver should be left with a reliable set of relayed emails.

For example, a mail sent from a yahoo.com user to another yahoo.com user should have the last hop information (from the first Received header) that looks something like this:

```
Received: from 209.191.85.211
     (HELO smtp101.mail.mud.yahoo.com)
```

If instead the Received header looks like this:

```
Received: from 142.103.6.59 (EHLO alumni.cs.ubc.ca)
```

and the Return-Path header like:

```
Return-Path: <dairyman88@yahoo.com>
```

then we can conclude that `142.103.6.59 (alumni.cs.ubc.ca)` is a relay server. Similarly, mail that is DomainKey-signed from Gmail and had come directly from Gmail typically has a *proxy.gmail.com IP address on the topmost Received line, so anything else would indicate that the mail had gone through a relaying server.

3.Experiment

From January to March 2006, Yahoo! Mail applied this algorithm to emails that were DomainKey signed and verified by a Yahoo! owned domain (for instance, yahoo.com, yahoo.co.uk, and yahoogroups.com). The data sets were limited to users' spam and not-spam report data for ease of collection, limited scaling needs, and privacy concerns. The processing script was run twice a week, at 01:01 on Mondays and Fridays.

The amount of mail processed between Friday and Monday was generally much smaller than from Monday to Thursday due to user usage patterns.

4.Results

At the beginning of the experiment, the algorithm caught several thousand new relaying servers each week. Within a few weeks the rate became steady at between 300 and 700 new relaying IP addresses, and continues at this rate at publication time. A total of 8,151 relaying IPs have been found. This is considerably more than previous efforts at cataloging relaying IP addresses such as trusted-forwarder.org [6]. The individual server's reverse DNS entries break down to approximately 49% .com, 24% .net, 5% .edu, 3% .org and 15% other (.uk, .de, etc.). The top .net entries are mainly ISPs (comcast.net, earthlink.net, etc.). As for names, about 38% had "mail", "smtp", "mta", "host", "relay", or "server" in a second-level DNS name.

5.Weaknesses of the algorithm

The experiment analyzed email sent through a free email system. If a spammer knew the experiment was being run, they could have sent themselves an email and auto-forwarded it through a system they controlled. This would result in additional fake relay machines being found. If this algorithm is used in a production environment to find trusted relay machines, care should be taken to avoid this gaming scenario. In non-free email systems, this risk is heavily mitigated as the spammer should not be able to obtain an account to receive messages. The algorithm does not rely on using the free system as the message originator—it was chosen for convenience because the IP addresses were readily accessible by the authors; most of the originated messages from Yahoo! Mail and Groups are DomainKey signed; and the large user base increases the chances that messages will be sent and received by a large number of relaying servers.

The algorithm makes no attempt at discovering relay servers that have merely changed IP addresses. It is possible that the algorithm found the same servers over and over. To mitigate this risk, we should examine the IP and rDNS pairs to find mismatches; however, this work has not been performed yet. This should find servers that have been moved, avoiding relays that operate on transient IP addresses.

If the algorithm occasionally makes a mistake and marks a site as a relay when it isn't, the effect is relatively benign. If an antispam system uses the algorithm's results to reduce the impact of an IP in its decision making process, the impact of incorrectly identified relays should not be as severe as if an antispam system does not know that an IP is a relay and consequently marks it as a spam site. If an IP is identified as a relay, the effect should be to ignore the sender IP when applying antispam rules, instead using other characteristics of the message. Because a piece of spam typically has multiple indicators of being spam, removing one piece of information (the connecting IP) still allows the antispam system to identify spam using other techniques. Users in general prefer false negatives to false positives, because of the tediousness of looking through the spam folder for mistakes; mistakenly identifying a few IPs as relays, and lowering the effectiveness of spam detection for messages sent from them is a reasonable price to pay for reducing highly visible false positives.

6.Summary

Cryptographic email authentication protocols verify the originators of emails. When combined with path information, email administrators and antispam systems can reliably find relay servers that are servicing their users.

The Messaging Anti-Abuse Working Group claims that 80% or more of 2006 Internet e-mail is considered spam [9]. With relaying servers generally forwarding all email (including spam) for their customers, IP based reputation algorithms frequently mistakenly rate a relay's IP address negatively or as a major indicator of a message being spam. This can have a dramatic impact on false positives, mis-tagging or rejecting all messages legitimately destined for a user, regardless of the originator.

With the information provided by this algorithm, last hop IP reputation checks and path-based authentication techniques could be downgraded, to force an antispam system to place heavier weight on other message characteristics when making a spam judgment. More reliable rules could be developed to determine the originating IP by parsing Received: headers for specific relays to correct for the missing information.

It is worth noting that some relay servers are used to both relay and initiate email. For instance, the IEEE sends out newsletters and other member-oriented material from the same servers that provide its relaying service for users. The algorithm presented does not determine such cases, it merely shows servers that are minimally relays. These cases provide an even greater challenge for IP based reputation. On one hand, they originate solicited bulk messages. On the other, they forward untagged phishing emails.

The relay information could also be used to identify the risks of using a technique that does not work in relaying cases for a particular deployment. For instance, if an enterprise email system finds only a few relaying services connecting to it, the enterprise might decide the risk of deploying a relaying-unfriendly technique is worthwhile, while a consumer system might find and conclude the opposite.

The authors believe that if the data set were expanded to the full set of email received by the Yahoo! Mail system, many more store-and-forward relaying servers would be found than with the experimental data. The number of relay servers found by the algorithm is not rapidly decreasing.

7.REFERENCES

[1] Delany, M. *Domain-based Email Authentication Using Public-Keys Advertised in the DNS (DomainKeys)*. http://www.ietf.org/internet-drafts/draft-delany-domainkeys-base-03.txt (September 2005)

[2] Crocker, D. *Internet Mail Architecture*. http://www.ietf.org/internet-drafts/draft-crocker-email-arch-04.txt (March 2005).

[3] Hutzler, C., Crocker, D., Resnick, P., Sanders, R., and Allman. E. *Email Submission: Access and Accountability*. http://www.ietf.org/internet-drafts/draft-hutzler-spamops-05.txt (Oct 2005).

[4] Harris, E. *The Next Step in the Spam Control War: Greylisting*,

http://projects.puremagic.com/greylisting/whitepaper.html

[5] Donnerhacke, L. *Teergrubing FAQ* http://www.iks-jena.de/mitarb/lutz/usenet/teergrube.en.html

[6] *Trusted-Forwarder Whitelist.* http://trusted-forwarder.org/.

[7] Wong, M., Schlitt, W. *Sender Policy Framework (SPF) for Authorizing Use of Domains in E-mail*. (June 2005)

[8] Lyon, J., Wong, M. *Sender-ID: Authenticating E-mail* http://www.ietf.org/internet-drafts/draft-lyon-senderid-core-01.txt. (May 2005)

[9] MAAWG. *Email Metrics Report - March 2006*. http://www.maawg.org/about/FINAL_4Q2005_Metrics_Report.pdf (March 2006)