An overview of the Sender Policy Framework (SPF) as an anti-phishing mechanism

Stefan Görling

Royal Institute of Technology, Stockholm, Sweden

Autobiographical note

Stefan Görling is a doctoral researcher at the department for Industrial Technology and Management at the Royal Institute of Technology (KTH), Stockholm. He is currently working on a research project studying the use of domain-names in Sweden.

An overview of the Sender Policy Framework (SPF) as an anti-phishing mechanism

Structured abstract

Purpose – The purpose of this paper is to provide an overview of the Sender Policy Framework (SPF) and discuss its merits for adoption as an anti-phishing mechanism.

Design/methodology/approach – All active domains in the .se zone were probed to determine if they have an SPF-policy. This data collection step is combined with a theoretical discussion of the SPF standard and related initiatives.

Findings – This paper finds that the adoption ratio is very low. Few seem to be interested in deploying it, despite it being designed for easy adoption and being consistently implemented in several popular anti-spam solutions.

Practical implications – Despite the low adoption ratio the standard merits implementation by IT/IS managers and software vendors.

Originality/value – This study analyzes the adoption ratio, which is valuable both for software vendors and endorsers of the standard. It also provides an overview of the standard itself as an attempt to avoid common misconceptions about the role of the standard and its relation to other anti-spam initiatives.

Keywords – E-mail, Sender Policy Framework, SPF, Anti-Spam, Phishing

Paper type – Research paper

Introduction

In the last four years, the problem of e-mail spam grew to menacing proportions. Ferris Research estimates that the amount of spam increased five-fold between 2003 and 2005, and that lost productivity due to spam could be as high as $50 billion (Keizer 2005). In 2004, Bill Gates predicted that spam would be eradicated within two years (CBS 2004). The growth of the phenomenon has also inspired increased activities in scholarly research. An academic description of the problem was provided by Pavlov et al. (2005).

As spam filters grew more common, forging senders on e-mail messages gained in popularity. By doing so, the messages gains trustworthiness both in the eyes of the filters and the recipients, thus increasing the likelihood that the message will reach its intended target. Forged sender addresses are not only used in regular commercial spam, but are also commonly used when spreading computer viruses and launching phishing attacks.

Since Internet was conceived as an open network, the current specifications for e-mail messaging have no built-in authentication or authorization. The recognition of this problem is far from new, the first standard draft with modifications to address this issue was published in 1987, long before e-mail became a common communication tool (IETF 1987).

Limitations of prior standards prevented their wider adoption (cf. Gutmann 2002). They are far from transparent, depending on the user to manually exchange cryptographic keys and individually sign their messages. Until a few years ago, the occurrences of e-mail forgery were relatively rare, so few troubled to implement these mechanisms. It was not until the problems of spam and phishing exploded anew, that the new, more transparent initiatives were proposed.

One of these new technologies is Sender Policy Framework (SPF), a system designed to validate sender addresses in e-mail messages.

The purpose of this paper is to present an overview of the existing anti-spam and anti-phishing initiatives and describe what SPF is, to what extent it is adopted, and why it might be a good idea to promote its wider implementation.

An overview of various anti-spam initiatives

The following initiatives aim to solve the problem of unsolicited and fraudulent e-mail. These solutions tend to address well a particular subset of the problem, thus most of them are complementary.

Redesign – Suggestions to deprecate SMTP protocol in favor of something better circulate periodically. Due to the extensive work needed for such a major alteration, few concrete suggestions were put forward so far. Two of the more recognized initiatives were the IM2000 (IM2000.org 2006) and the SMTPng (SMTPng.org 2002) initiatives. Current efforts do not seek to replace the current standards but rather to improve them by adopting additional standards.

Filters – Most corporations and Internet Service Providers (ISPs) deployed spam filters, which attempt to remove spam before it reaches its intended recipients. Some popular methods include, but are not limited to keyword/string matching, white/black-listing (e.g. Hird 2002) and statistical methods such as Bayesian rules (Sahami et al. 1998; Graham 2002). These methods often induce a "Red Queen" effect where spammers and anti-spammers constantly compete.

Transaction costs – Spam will exist only as long as it is profitable. Currently, the transaction cost of sending e-mail is so low, that even given an extremely low response rate between 0.0075% and 5.6% (Mindlin 2006), the profit margins are still substantial. Increasing transaction cost might render spam less profitable. This increase could be achieved either by creating an e-mail tax (cf. Kraut et. al 2003), solving a mathematical puzzle and thus paying with CPU-resources (Back 2002), or by forcing re-transmits of each e-mail (known as greylisting, cf. Harris 2003).

Legislation – Spam-related laws and policies were passed in several countries. They usually take one of the following forms: opt-in, opt-out, or flagging advertisements. They had little effect because senders of unsolicited e-mails tend to be hard to identify and also because most of it originates from foreign countries. A compilation of spam-related legislation is provided by Sorkin (2006).

Block ports – Nowadays, the majority of spam is delivered by Botnets. Botnets consist of large numbers of broadband-connected home computers that were hijacked by computer viruses or other malware. To reduce the amount of abusive traffic, many ISPs block network access on mail-related ports (primarily, TCP port 25), so that all outbound e-mail has to pass through valid mail servers belonging to the ISPs (Seitzer 2005). While providing an easy way to reduce the number of abuse reports, these ISPs violate the principles of net neutrality by limiting users' ability to run certain applications.

Sender authentication – Sender authentication verifies that e-mail originated from its claimed sender. SPF, a type of sender authentication, is the focus of this paper (IETF 2006a). Content verification provides tools to confirm that e-mail's content was not altered during transmission. Built-in sender authentication and content verification are not included in the current SMTP standard. Standard extensions covering these areas include but are not limited to PGP (IETF 1998), S/MIME (IETF 1999), and DomainKeys (dkim.org 2006).They will be discussed in the next section.

Initiatives described above address somewhat orthogonal parts of the spam issue. Applied together, correctly implemented, they should mitigate much of the problem.

Standards for sender authentication

Spam is far from a homogenous phenomenon. It includes a wide range of unsolicited e-mail messages such as corporate newsletters, product and service advertisements (pharmaceuticals, mortgages, porn, etc.), fraud attempts (phishing, 419 messages, a.k.a 'nigerian scams'), and political and religious pamphlets.

Sender authentication primarily addresses problems with fraud such as phishing, where the real sender masquerades as another entity, usually a well-known financial organization. It also increases accountability of real senders, which in the future might alleviate the overall spam problem. Additional important context where such authentication is desirable is when transmitting sensitive information. One recent example of this was when a Swedish convict was wrongly released after the reception of a falsified fax message.

Original standards (OpenPGP and S/MIME) appeared partly with these tasks in mind. They are not transparent: both the sender and the receiver must have compliant mail clients, and manually enter some security information at least during the configuration time. The failure to implement these standards in a user-friendly fashion severely limited the adoption of these protocols (cf. Whitten & Tygar 1999).

Standard/ Mechanism	Transparent to user	Standardized	Authentication	Content verification	Public/ Private key-based	Protects from phishing
OpenPGP	No	Yes, since 1996	Yes, to individual	Yes	Yes	No
S/MIME	No	Yes, since 1998	Yes, to individual	Yes	Yes	No
SPF	Yes	Yes, experimental 2006	Yes, to domain	No	No	Yes
DomainKeys /DKIM	Yes	No, working on draft	Yes, to domain	Yes	Yes	No

SPF was designed with transparency and ease of adoption in mind. As a result, it has limited functionality compared to all the other standards. In particular,

- it can only verify sender's domain, i.e. whether the mail originated at BigBank.com rather than JaneDoe@BigBank.com.
- it gives no guarantees with regard to the actual ownership of BigBank.com by BigBank Inc.
- it does not include any encryption or content verification.

All of these standards provide sender authentication, but only SPF protects against phishing by default. While a message signed with OpenPGP, S/MIME or DomainKeys can be authenticated and even verified as not having been modified, they provide little guidance of what to do with a mail that is unsigned.

With SPF, the policy of the claimed originating domain determines how unsigned messages should be handled. This makes SPF complementary to all other standards listed, as it provides an actual default policy for unsigned messages.

Using SPF, a fraudulent message from BigBank would be removed before reaching the recipient, while the use of other standards would need the user to manually inspect the message and realize that the lack of signature is a sign of fraud.

The Sender Policy Framework

As previously described, the Sender Policy Framework (SPF) is not a general anti-spam measure but a mechanism specifically designed to mitigate problems with fraud and phishing. It is used to validate that the message was sent by the sender domain specified in the "MAIL FROM:" address of the message envelope. For example, when a mail from paypal.com requests revalidation of users' account and password, SPF can check if it is indeed sent from the domain paypal.com or is part of a scam.

This section summarizes how SPF avoids falsified sender addresses in e-mail messages. A more detailed technical overview of this standard is provided by Wong (2004).

Deployment and usage of SPF is quite simple. A domain's policy specifies which mail servers may send e-mail using this domain in the sender address. The policy is then published as a DNS record, either as a text (TXT) or a specific SPF record:

An example policy could look like:

```
bigbank.com. IN TXT "v=spf1 +mx a:office.bigbank.com/28 –all"
bigbank.com. IN SPF "v=spf1 +mx a:office.bigbank.com/28 -all"
```

The policy in this example states that:

- [+mx] mail servers specified in MX records for this domain are authentic
- [a:office.bigbank.com/28] if the claimed originating mail server is in this address range, it also authentic
- [-all] all other mail servers are invalid.

Bigbank.com's policy contains complete information about which mail servers are authorized to send mail claiming @bigbank.com address. The policy always contains a "catch-all" rule specifying what to do with e-mail that does not come from one of these servers. This rule must take one of four forms: blacklist unauthorized messages, whitelist authorized messages while being very suspicious to the rest, whitelist authorized messages while being neutral to the rest, whitelist everything. Ideally,

the last two rules exist only to allow incremental adoption of the standard in order to make sure the policy is working correctly before fully deploying it.

The principle of SPF validation is as follows. When an e-mail is received, an SPF-enabled mail server, spam filter or client (the check could be implemented on one or several of these) will make a DNS query in order to see whether the claimed originating domain has a published policy. If this is the case, the policy is compared with the IP address from which the message was received.

Therefore, if a fraudster is trying to forge an e-mail with a source address of billing@bigbank.com, and BigBank Inc. created and published a policy stating that all e-mail from bigbank.com must be sent from mail.bigbank.com, the receiving end should conclude that the mail is forged and thus prevent the user from being fooled by the attack.

DNS is used as a database because of its convenience: the infrastructure is already deployed and has proven to be stable and scalable. If SPF were widely deployed, it would create additional load on the DNS infrastructure. It should be manageable due to the caching functionality in DNS.

The validation procedure on the receiving end is also fairly efficient, as it does not contain any CPU-intensive algorithms, such as cryptographic key calculations. The policy has to be retrieved from DNS or from a local cache, and then parsed and compared to the IP address from which the e-mail was received. This adds extra load on the receiving end for each e-mail message, but should be comparable and even lower than that of several other anti-spam solutions in operation today.

As it is stored in DNS, each individual policy is globally visible. But since it resides within the scope of the domain name, it can only be altered by the domain owner. Anyone can determine that only servers on the bigbank.com office network are allowed to send @bigbank.com e-mails, but no one can forge a message to appear to be coming from the bigbank.com mail server without gaining access to the server itself. As SPF in itself does not contain any cryptographic measures and relies on the DNS infrastructure, it is sensitive to attacks such as DNS spoofing. In order to be fully secure a wider deployment of DNS-SEC is necessary.

It should be noted that there was a controversy whether SPF is a good solution at all. Most of the criticism brought forward against it relied on misconceptions with regard to what SPF is and what problems it aims to solve. As mentioned above, it is not intended to solve the spam problem in full. Critique of its design targets two areas: the use of TXT records in DNS and the underlying philosophy of the standard, which expects mail to originate from certain well-defined servers (a fair summary is provided by Wikipedia, 2006-12-22).

The use of TXT records is not perfect from a technical standpoint, since it puts structured data into unstructured format. It was chosen in order to improve chances of wider deployment. An alternative SPF record structure is available, but may take many years before all DNS server software supports it.

The other issue is whether or not anyone should be able to submit any e-mail using an arbitrary message transfer agent (MTA). As described, the general idea of SPF is to restrict MTAs, which are allowed to send e-mail from certain domains but not from others.

Previously, anyone could send mail using somename@anywork.com using MTA associated with his or her home ISP provider. If SPF were used, all messages have to be send to one of the authorized corporate servers, instead of simply choosing the closest MTA.

IETF (2006b) published a draft standard specifying that mail should always be submitted to such an authorized MTA. Some professionals regard this policy as breaking the original principles of how Internet e-mail was built, violating the maxim known as Occam's razor. In reality, there is already a large number of restraints on MTAs. Most of them were put in place in order to disallow relay of arbitrary e-mail precisely due to the general spam problem.

Some minimal organizational will likely be required to deploy this standard. Users who are used to sending corporate mail from the road or from home, might have to change their e-mail settings. Messages bearing corporate domain must be submitted to an authorized mail-server using a VPN connection or the widely supported message submission mechanism as specified by IETF (2006b).

As every policy, this one must be kept current in order to be functional. After publishing a new record to DNS, it must be kept updated as new MTAs are added or removed, in order to avoid false classifications.

The current adoption of Sender Policy Framework

SPF standard is supported by large online movers such as AOL.com, Gmail.com and Internet Engineering Task Force (IETF). Having experienced a number of phishing attacks in Sweden during the last year (EuroSecure, 2005; F-Secure, 2006), we decided to study the adoption rate of this promising standard.

These numbers should be helpful in motivating IT/IS managers and interesting to advocates of sender authentication standards because it tracks effectiveness of voluntary standard deployments.

Our dataset contains all Swedish domains (domains ending in .se), which were active on the 3rd of February 2006. We considered a total of 385,862 domains. Apart from these domains, there exists a number of .se domains which were registered, but are not currently in use (have no DNS records). These domains are not included in this study, since they have no servers associated with them.

Not all .se domains are hosted in Sweden. Anyone from outside Sweden may register an .se domain name, and anyone in Sweden may use domains not ending in .se. However, this data set is somewhat unique as it includes all domains within a certain top-level domain (TLD).

All DNS records of these domains were scraped as a part of a larger project. In particular, all domains were queried for the existence of mail servers (MX records) and Sender Policies (as defined by SPF).

An SPF record may describe valid servers in a number of ways (a, mx, ptr, ip4, ip6, exists, include). Each of these mechanisms can be prefixed with one of four qualifiers (fail '-', softfail '~', pass '+', neutral '?')[1]. The collected records were analyzed in order to see how the actual live policies were constructed by those who took the effort to implement them. For this purpose, the flags in the record were extracted and classified.

[1] This is a simplification, please see RFC 4408 for a more fine-grained explanation.

The analysis showed that the adoption ratio is extremely low. Among the 385,862 domains, only 1.63% (6,286) of all domains have a published SPF policy. Counting only domains having active mail settings (330,163 domains with MX records), SPF usage ratio is 1.9%

Out of the identified SPF records, 54.5% (3,430) of them contained only the minimal catch-all rule ("v=spf1 ?all") specifying that domain has no policy. This policy is basically useless. (12.3%) 775 domains contained the softfail catch-all rule ("~all"), a state between valid and failed where unidentified sender is unlikely to be authorized, but that we are not sure enough to drop the message. A state which may be useful when testing a new policy.

In total 19.6% (1,233) domains have a catch-all specifying that a mail is forged if it does not originate from one of the listed servers. These are the only domains with what we consider to be a full SPF implementation. They constitute 0.37% of all domains.

Studying the use of operators in the policies, we found that 98.5% (6192 records) contained one or more "a" operator, 36.5% (2296 records) contained "mx" operator(s), 25.0% (1573) contained "ip4" operator(s), 11.3% (710) contained "include" operator(s), 5.6% (350) contained "ptr" operator(s), 0.05% (3) contained exist operator(s), none contained any "ip6".

Discussion

SPF is no silver bullet in solving the spam-problem. Sender authentication is just a first step in the quest to reduce spam and fraudulent behavior on the Internet (cf. Watson 2004). Neither is it a fully perfect solution for verifying the sender. Several other initiatives and debates on how to proceed with better standards are underway.

The main argument for adopting SPF is that it is an existing standard and that in many cases are fairly easy to implement within the current infrastructure.

The classification and quantification of the nature of the domains using SPF and not in Sweden is beyond the scope of this paper, nevertheless it is interesting to note that the authors have found no SPF-records while probing major Swedish banks and other randomly selected financial institutions[2].

For organizations threatened by phishing-attacks SPF is a way to limit the risk. There are few reasons not to implement this kind of solution while waiting for a better one to become standardized. This study has shown that despite this, the adoption of this standard in Sweden has not happened.

One conclusion of this paper must be that for users of the. se-domains, implementation has not been a priority so far. It is surprising that so little effort has been taken by organizations in order to reduce the risks of spam, computer viruses and phishing-attacks being seemingly sent from their corporate addresses.

There may be several causes why this deployment has not yet taken off. One might be limited knowledge about the existence and principles of this standard, which this paper seeks to address. Another factor might be the lack of visual indicators in software such as e-mail clients that the policy is absent, thus not raising awareness that such a policy might be possible to implement.

[2] To check if a domain has published a policy or not you may use a number of SPF Testers available on the internet, cf. http://www.dnsstuff.com/pages/spf.htm

Limited adoption may also be caused by a failure to recognize the problem of un-validated sender addresses, or not recognizing the SPF-standard to be robust enough to implement.

Organizations affected by phishing attacks, such as banks often respond by publicly stating that they never send e-mails to their customers, thus trying to avoid the problem rather than addressing it. A policy not to send e-mails to your customers might be an easy short-term solution but will most likely proves more complicated as e-banking services advances. Failing to solve the problem with phishing may decrease the adoption speed of such services and lessen overall trust in e-services.

SPF has been backed by a large number of vendors and websites; it is therefore likely that this standard will increase in popularity, despite this slow start. As with many other new technologies, there exists a two-folded problem in implementing it. In order for it to be useful, it must be implemented both by the domain owner as well as at the receiving end (either e-mail clients, mail servers or spam-filters).

Changing Internet standards in order to reduce the openness of the network is not without debate and discussions. It might be so that the search for what is commonly referred to as "the Final Ultimate Solution to the Spam Problem (FUSSP)" leads to decreased interest in adopting partial improvements until this solution has been found. SPF is no final solution to the spam problem; it is not even a general anti-spam method but specifically designed to identify whether the sender of an e-mail message is valid in order to mitigate the problem with fraud and phishing.

There is currently no single standard that will solve the problems with unsolicited e-mail and fraudulent behavior, it must be recognized that several standards must co-exist and work together. It has been the aim of this paper to present an overview highlighting the role of SPF in the overall process of eliminating spam and other fraudulent behavior on the Internet. Therefore an implementation of SPF should be considered even though there might be other, better, standards being developed in the future.

This framework has been designed for simplicity in adoption in mind. The lack of implementation of this standard is therefore worrying as most other standards are more difficult to implement and deploy, thus lowering the likelihood that these will be adopted to any further extent.

Despite not being a full solution, it improves the overall scenario on several areas such as phishing if widely adopted. IT/IS managers currently evaluating this standard should be aware this standard is not currently widely implemented on all levels.

However, as long as they recognize the limits and properties of this standard, the creation of a SPF-policy is desirable. Each new published policy helps the standard to become more widely used and raises security at least for some users. SPF is available, it is ubiquitous and it is easy to adopt. There already exist several major anti-spam software vendors with implemented support for this standard in their products. Simply by publishing a SPF policy, the impact of phishing attempts in your name will decrease.

Limitations

This article is not without limitations. The data collection was carried out over a period of months due to limited computer resources. Changes might have happened during the time of collection, which might affect the results to a minor extent.

The study measures the number of domains having adopted this framework. One might argue that other metrics such as e-mail volumes or similar would be a better value meter.

Contribution to Research

Forged e-mail sender addresses pose a real problem to organizations today. SPF is an endorsed standard, which is easy to implement addressing these concerns. Despite this, the adoption has been found to be very limited. Apparently, these are not reasons enough for corporations to implement SPF. The aim of this paper is to build a broader understanding of the solutions available today.

Further research should address the reasons for the lack of implementation, approaching the question from a more qualitative perspective and try to determine the underlying reasons for this unwillingness to adoption.

Acknowledgements

The author is greatly in debt to IIS for sharing the list of all active domains in the .se-zone. The author also would like to thank Dr David G. Schwartz, Galina Shubina as well as two anonymous reviewers for helpful comments on earlier drafts of this paper. Finally, gratitude goes to the authors of Open Source software used in this project, such as Perl, Net::DNS, MySQL and BIND.

References

- Back, A. (2002), Hashcash – A Denial of Service Counter-Measure, *http://www.cypherspace.org/hashcash/hashcash.pdf*
- IM2000.org (2006-12-30), IM2000 – electronic mail transport of the future, *http://www.im2000.org*
- CBS, (2004-01-24), Gates: Spam To Be Canned by 2006, *http://www.cbsnews.com/stories/2004/01/24/tech/main595595.shtml*
- dkim.org (2006-12-20), Domain Keys Identified Mail (DKIM), *http://www.dkim.org*
- EuroSecure (2005-10-04), Phishing-attack against Swedish bank Nordea, *http://www.nod32norway.com/news.asp?id=843f549431c747fc*
- F-Secure (2006-04-04), Nordea Phishing is Back, *http://www.f-secure.com/weblog/archives/archive-042006.html#00000849*
- Graham, Paul (2002), A Plan for Spam, *http://www.paulgraham.com/spam.html*
- Gutmann, Peter (2002), PKI: it's not dead, just resting, *Computer vol. 35, Issue 8, p. 41-49*
- Harris, Evan (2003), The Next Step in the Spam Control War: Greylisting, *http://projects.puremagic.com/greylisting/whitepaper.html*
- Hird, Shane (2002), Technical Solutions for Controlling Spam, *Proceedings of AUUG2002, Melbourne, Australia.*
- Keizer, Gregg (2005-02-23), Spam Costs Business Worldwide $50 Billion, *http://www.informationweek.com/story/showArticle.jhtml?articleID=60403016*
- Kraut et al. (2003), Pricing Electronic Mail to Solve the problem of Spam, *http://www.econ.upf.edu/docs/seminars/sunder.pdf*
- IETF (1987), Privacy Enhancement for Internet Electronic Mail: Part I: Message Encipherment and Authentication Procedures, *http://www.ietf.org/rfc/rfc989.txt*
- IETF (1998), OpenPGP Message Format, *http://www.ietf.org/rfc/rfc2440.txt*
- IETF (1999), S/Mime Version 3 Message Specification, *http://www.ietf.org/rfc/rfc22633.txt*
- IETF (2006a), Sender Policy Framework (SPF) for Authorizing Use of Domains in E-Mail, Version 1, *http://www.ietf.org/rfc/rfc4408.txt*
- IETF (2006b), Message Submission for Mail, *http://www.ietf.org/rfc/rfc4409.txt*

- Mindlin, Alex (2006-07-03), Seems Somebody Is Clicking on That Spam, *http://www.nytimes.com/2006/07/03/technology/03drill.html*
- OpenSPF.org (2006-07-01), SPF: A Sender Policy Framework to Prevent Email Forgery, *http://www.openspf.org*
- Pavlov, O., N. Melville, R. Plice (2005), Mitigating the Tragedy of the Digital Commons: The Problem of Unsolicited Commercial E-Mail, *Communications of the Association for Information Systems. 2005. v.16: 73-90*
- Seitzer, Larry (2005-04-08), Shutting Down The Highway To Internet Hell, *http://www.eweek.com/article2/0,1895,1784276,00.asp*
- Sorkin, David E. (2006-12-22), Spam Laws, *http://www.spamlaws.com*
- SMTPng (2002-09-23), SMTP "next generation" Project, *Website archive available at http://web.archive.org/web/20020923132829/http://smtpng.org/*
- Sahami, M. Dumais, S. Heckerman, D. Horvitz, E. (1998), A Bayesian approach to filtering junk e-mail, *AAAI'98 Workshop on Learning for Text Categorization*
- Watson (2004), Beyond Identity: Addressing Problems that Persist in an Electronic Mail System with Reliable Sender Identification, *presented at CEAS 2004*.
- Wikiedia (2006-12-22), Sender Policy Framework Controversy, *http://en.wikipedia.org/wiki/Sender_Policy_Framework#Controversy*
- Whitten A., Tygar JD. (1999), Why Johnny Can't Encrypt – A Usability Evaluation of PGP 5.0, *Proceedings of the 8th USENIX Security Symposium. 1999 ; 169-181.*
- Wong, M. (2004), SPF Overview, Linux Journal, *http://www.linuxjournal.com/article/7327*