

SPFBIS Working Group
Internet-Draft
Intended status: Informational
Expires: November 11, 2012

M. Kucherawy
Cloudmark
May 10, 2012

**Resolution of The SPF and Sender ID Experiments
draft-ietf-spfbis-experiment-08**

Abstract

In 2006 the IETF published a suite of protocol documents comprising SPF and Sender ID, two proposed email authentication protocols. Both of these protocols enable one to publish via the Domain Name System a policy declaring which mail servers were authorized to send email on behalf of the domain name being queried. There was concern that the two would conflict in some significant operational situations, interfering with message delivery.

The IESG required the publication of all of these documents ([RFC4405](#), [RFC4406](#), [RFC4407](#), and [RFC4408](#)) with Experimental status, and requested that the community observe deployment and operation of the protocols over a period of two years from the date of publication to determine a reasonable path forward.

After six years, sufficient experience and evidence have been collected that the experiments thus created can be considered concluded. This document presents those findings.

Status of this Memo

This Internet-Draft is submitted in full conformance with the provisions of [BCP 78](#) and [BCP 79](#).

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at <http://datatracker.ietf.org/drafts/current/>.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on November 11, 2012.

Copyright Notice

Copyright (c) 2012 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

1.	Introduction	3
2.	Definitions	3
3.	Evidence of Deployment	4
3.1.	DNS Resource Record Types	4
3.2.	Implementations	5
3.3.	The SUBMITTER SMTP Extension	5
4.	Evidence of Differences	7
5.	Analysis	7
6.	Conclusions	8
7.	Security Considerations	9
8.	IANA Considerations	9
9.	References	9
9.1.	Normative References	9
9.2.	Informative References	9
Appendix A.	Background on the RRTYPE Issue	10
Appendix B.	Acknowledgments	11
	Author's Address	12

1. Introduction

In April 2006, the IETF published the [SPF] and Sender ID email authentication protocols, the latter consisting of three documents ([SUBMITTER], [SENDER-ID], and [PRA]). Both of these protocols enable one to publish via the Domain Name System a policy declaring which mail servers are authorized to send email on behalf of the selected domain name. There was concern that the two would conflict in some significant operational situations, interfering with message delivery.

Consensus did not clearly support one protocol over the other, and there was significant concern that the two would conflict in some significant operational situations, interfering with message delivery. The IESG required the publication of all of these documents as Experimental, and requested that the community observe deployment and operation of the protocols over a period of two years from the date of publication in order to determine a reasonable path forward.

In line with the IESG's request to evaluate after a period of time, this document concludes the experiments by presenting evidence regarding both deployment and comparative effect of the two protocols. At the end it presents conclusions based on the data collected.

It is important to note that this document makes no direct technical comparison of the two protocols in terms of correctness, weaknesses, or use case coverage. The email community at large has already done that through its deployment choices. Rather, the analysis presented here merely observes what has been deployed and supported in the time since the protocols were published, and draws conclusions based on those observations.

The data collected and presented here are presumed to be a reasonable representative view of the global deployment data, which could never itself be fully surveyed within a reasonable period of time.

2. Definitions

The term "RRTYPE" is used to refer to a Domain Name System ([DNS]) Resource Record (RR) type. These are always expressed internally in software as numbers, assigned by IANA under Expert Review provisions. Assigned RRTYPES also have names. The two of interest in this work are the TXT RRTYPE (16) and the SPF RRTYPE (99).

3. Evidence of Deployment

This section presents the collected research done to determine what parts of the two protocol suites are in general use, as well as related issues like [\[DNS\]](#) support.

3.1. DNS Resource Record Types

Three large-scale DNS surveys were run that looked for the two supported kinds of RRTYPES that can contain SPF policy statements. These surveys selected substantial sets of distinct domain names and queried their nameservers for both RRTYPES that can be used for SPF and/or Sender ID. In the tables below, only syntactically valid replies are counted.

DNS Survey #1 (Cisco)

+-----+	+-----+	+-----+
Domains queried	1,000,000	-
TXT replies	397,511	39.8%
SPF replies	6,627	<1.0%
SPF+TXT replies	6,603	<1.0%
spf2.0/* replies	5,291	<1.0%
+-----+	+-----+	+-----+

The "spf2.0/*" replies are those replies whose payload started with the string "spf2.0/", which are specific requests for Sender ID processing. Domains were selected as the top million domains as reported by Alexa, which monitors browser activity.

DNS Survey #2 (The Trusted Domain Project)

+-----+	+-----+	+-----+
Domains queried	270,351	-
TXT replies	151,310	56.0%
SPF replies	2,813	1.0%
SPF+TXT replies	2,632	<1.0%
spf2.0/* replies	7,285	2.7%
+-----+	+-----+	+-----+

This survey selected its domains from data about observed email traffic and previous SPF and Sender ID evaluations, collected from 23 reporting hosts across a handful of unrelated operators.

During this second survey, some domains were observed to provide immediate answers for RRTYPE 16 queries, but would time out waiting for replies to RRTYPE 99 queries. For example, it was observed that 4,360 (over 1.6%) distinct domains in the survey returned a result of

some kind (a record or an error) for the TXT query in time N, while the SPF query ultimately failed after at least time 4N.

DNS Survey #3 (Hotmail)

+-----+			
Domains queried	100,000	-	
TXT replies	46,221	46.2%	
SPF replies	954	<1.0%	
SPF+TXT replies	1,383	1.4%	
+-----+			

Hotmail's domain set was selected from live email traffic at the time the sample was extracted.

A separate survey was done of queries for RRTYPE 16 and RRTYPE 99 records by observing nameserver traffic records. Only a few queries were ever received for RRTYPE 99 records, and those almost exclusively came from one large email service provider that queried for both RRTYPES. The vast majority of other querying agents only ever requested RRTYPE 16.

3.2. Implementations

It is likely impossible to determine from a survey which Mail Transfer Agents (MTAs) have SPF and/or Sender ID checking enabled at message ingress since it does not appear, for example, in the reply to the EHLO command from extended [\[SMTP\]](#). We therefore rely on evidence found via web searches, and observed the following:

- o A web site [\[SID-IMPL\]](#) dedicated to highlighting Sender ID implementations last updated in late 2007 listed 13 commercial implementations, which we assume means they implement the PRA checks. At least one of them is known no longer to be supported by its vendor. There were no free open source implementations listed.
- o The [\[OPENSPE\]](#) web site maintains a list of implementations of SPF. At the time of this document's writing it listed six libraries, 22 MTAs with built-in SPF implementations, and numerous patches for MTAs and mail clients. The set included a mix of commercial and free open source implementations.

3.3. The SUBMITTER SMTP Extension

The PRA is the output of a heuristic that seeks to scan a message header and extract from it the email address most likely to be the one responsible for injection of that message into the mail stream.

The SUBMITTER extension to SMTP is a mechanism to provide an early hint (i.e., as part of the MAIL command in an SMTP session) to the receiving MTA of what the PRA would be on full receipt of the message.

In a review of numerous MTAs in current or recent use, two (Santronics WinServer and McAfee MxLogic) were found to contain implementations of the SMTP SUBMITTER extension as part of the MTA service, which could act as an enabler to Sender ID.

An unknown number of SMTP clients implement SUBMITTER. Although there is substantial activity showing its use in MTA logs, it is not possible to determine whether they are multiple instances of the same client, or separate client implementations.

An active survey was done of running and publicly reachable MTAs. The MTAs selected were found by querying for MX and A records of a subset of all domains observed by The Trusted Domain Project's data collection system in the preceding 20 months. The results were as follows:

SUBMITTER Survey (The Trusted Domain Project)

+-----+	+-----+	+-----+
MTAs selected	484,980	-
MTAs responding	371,779	76.7%
SUBMITTER enabled	17,425	4.7%
MXLogic banner	16,914	4.6%
+-----+	+-----+	+-----+

Note: The bottom two rows indicate the percentage of responding MTAs with the stated property, not the percentage of selected MTAs.

Based on the SMTP banner presented upon connection, the entire set of SUBMITTER-enabled MTAs consisted of the two found during the review (above) and a third whose identity could not be positively determined.

Of those few responding MTAs advertising the SUBMITTER SMTP extension, 97% were different instances of one MTA. The service operating that MTA (MXLogic, a division of McAfee) reported that about 11% of all observed SMTP sessions involved SMTP clients which make use of the SUBMITTER extension. Note that this represents about 11% of the clients of 4.6% of the responding MTAs in the survey.

4. Evidence of Differences

Separate surveys from Hotmail and The Trusted Domain Project compared the cases where the PRA (used by Sender ID) and the [RFC5321](#).MailFrom address (used by SPF) differed. The results of these tests showed that at least 50% of the time the two addresses were the same, but beyond that the percentage varied substantially from one sampling location to the next due to the nature of the mail streams they each receive.

Further, The Trusted Domain Project analyzed approximately 150,000 messages and found that in more than 95% of those cases, Sender ID and SPF reach the same conclusion about a message, meaning either both protocols return a "pass" result or both return a "fail" result. Note that this does not include an evaluation of whether "fail" meant spam or other abusive mail was thus detected or that "pass" mail is good mail; it is merely a measure of how often the two protocols concurred. The data set yielding this response could not further characterize the cases in which the answers differed.

A second analysis of the same nature by Hotmail found that the two protocols yielded the same result approximately 80% of the time when evaluated across billions of messages.

Anecdotally, the differences in conclusions have not been noted as causing significant operational problems by the email-receiving community.

5. Analysis

Given the six years that have passed since the publication of the Experimental RFCs, and the evidence reported in the earlier sections of this document, the following analysis appears to be supported:

1. There has not been substantial adoption of the RRTYPE 99 (SPF) DNS resource record. In all large-scale surveys performed for this work, less than 2% of responding domains published RRTYPE 99 records, and almost no clients requested them.
2. Of the records retrieved, fewer than 3% included specific requests for processing of messages using the PRA algorithm, which is an essential part of Sender ID.
3. Although the two mechanisms often used different email address fields as the subject being evaluated, no data collected showed any substantial operational benefit, in terms of improved accuracy, to using one mechanism over the other.

4. A review of known implementations shows significant support for both protocols, though there were more implementations in support of SPF than of Sender ID. Further, the SPF implementations showed better upkeep and current interest than the Sender ID implementations.
5. A survey of running MTAs shows fewer than 5% of them advertised the SUBMITTER extension, which is a Sender ID enabler. Only three implementations of it were found.
6. Ample options are available in terms of software to implement either of the two protocols.
7. There remain obstacles to deployment of protocols that use DNS RRTYPEs other than the most common ones, including firewalls and DNS servers that block or discard requests for unknown RRTYPEs. Further, few if any web-based DNS configuration tools offer support for RRTYPE 99 records.

6. Conclusions

In light of the analysis in the previous section, the following conclusions are supported:

1. The experiments comprising the series of RFCs defining the SUBMITTER SMTP extension ([RFC4405](#)), the Sender ID mechanism ([RFC4406](#)), the Purported Responsible address algorithm ([RFC4407](#)), and SPF ([RFC4408](#)), should be considered concluded.
2. The absence of significant adoption of the RRTYPE 99 DNS Resource Record suggests that it has not attracted enough support to be useful.
3. Unavailability of software implementing the protocols was not a gating factor in terms of the selection of which to use.
4. The absence of significant adoption of the [[SUBMITTER](#)] extension, [[SENDER-ID](#)], and [[PRA](#)], indicates that there is not a strong community deploying and using these protocols.
5. [[SPF](#)] has widespread implementation and deployment, comparable to that of many standards track protocols.

[Appendix A](#) is offered as a cautionary review of problems that affected the process of developing SPF and Sender ID in terms of their use of the DNS.

7. Security Considerations

This document contains information for the community, akin to an implementation report, and does not introduce any new security concerns.

8. IANA Considerations

This document presents no actions for IANA. [RFC Editor: Please remove this section prior to publication.]

9. References

9.1. Normative References

- [DNS] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), November 1987.
- [PRA] Lyon, J., "Purported Responsible Address in E-Mail Messages", [RFC 4407](#), April 2006.
- [SENDER-ID] Lyon, J. and M. Wong, "Sender ID: Authenticating E-Mail", [RFC 4406](#), April 2006.
- [SPF] Wong, M. and W. Schlitt, "Sender Policy Framework (SPF) for Authorizing Use of Domains in E-Mail, Version 1", [RFC 4408](#), April 2006.
- [SUBMITTER] Allman, E. and H. Katz, "SMTP Service Extension for Indicating the Responsible Submitter of an E-Mail Message", [RFC 4405](#), April 2006.

9.2. Informative References

- [DNS-EXPAND] Faltstrom, P., Austein, R., and P. Koch, "Design Choices When Expanding the DNS", [RFC 5507](#), April 2009.
- [OPENSF] "Sender Policy Framework: Project Overview", <<http://www.openspf.net>>.
- [SID-IMPL] "Sender ID Framework Industry Support and Solutions", October 2007, <<http://www.microsoft.com/mscorp/safety/>>.

technologies/senderid/support.aspx>.

[SMTP] Klensin, J., "Simple Mail Transfer Protocol", [RFC 5321](#), October 2008.

[Appendix A](#). Background on the RRTYPE Issue

SPF was originally developed by a community of interested developers outside the IETF, with the intent of bringing it to the IETF for standardization after it had become relatively mature and ready for the IETF standards track process.

At the time of SPF's initial development, the prospect of getting an RRTYPE allocated for SPF was not seriously considered, partly because doing so had high barriers to entry. As a result, at the time it was brought to the IETF for development and publication, there was already a substantial and growing installed base that had SPF running using TXT RRs. Eventually the application was made for the new RRTYPE as a result of pressure from the DNS experts in the community, who insisted upon doing so as the preferred path toward using the DNS for storing such things as policy data.

Later, after RRTYPE 99 was assigned (long after IESG approval of the document, in fact), a plan was put into place to effect a gradual transition to using RRTYPE 99 instead of using RRTYPE 16. This plan failed to take effect for four primary reasons:

1. there was hesitation to make the transition because existing nameservers (and, in fact, DNS-aware firewalls) would drop or reject requests for unknown RRTYPEs (see [Section 3](#) for evidence of this), which means successful rollout of a new RRtype is contingent upon widespread adoption of updated nameservers and resolver functions;
2. many DNS provisioning tools (e.g., web interfaces to controlling DNS zone data) were, and still are, typically lethargic about adding support for new RRTYPEs;
3. the substantial deployed base was already using RRTYPE 16, and it was working just fine, leading to inertia;
4. [\[SPF\]](#) itself included a faulty transition plan, likely because of the late addition of a requirement to develop one: It said a server SHOULD publish both RRTYPEs and MUST publish at least one, while a client can query either or both, which means both can claim to be fully compliant while failing utterly to interoperate. Publication occurred without proper IETF review,

so this was not detected prior to publication.

It is likely that this will happen again if the bar to creating new RRTYPEs even for experimental development purposes is not lowered, and handling of unknown RRTYPEs in software becomes generally more graceful. Also important in this regard is encouragement of support for new RRTYPEs in DNS record provisioning tools.

Fortunately in the intervening time, the requirements for new RRTYPE assignments was changed to Expert Review, and the posture has become more relaxed. Also, the publication of [[DNS-EXPAND](#)] has provided some useful guidance in this regard. However, there is still a common perception that adding new types of data to the DNS will face resistance due to the lack of appropriate software support.

There are DNS experts within the community that will undoubtedly point to DNS servers and firewalls that mistreat queries for unknown RRTYPEs, and to overly simplistic provisioning tools, and claim they are broken as a way of answering these concerns. This is undoubtedly correct, but the reality is that they are among us and likely will be for some time, and this needs to be considered as new protocols and IETF procedures are developed.

[Appendix B](#). Acknowledgments

The following provided operational data that contributed to the evidence presented above:

Cisco: contributed data about observed Sender ID and SPF records in the DNS for a large number of domains (DNS survey #1)

Hotmail: contributed data about the difference between [RFC5321](#).MailFrom and [RFC5322](#).From domains across large mail volumes, and a survey of DNS replies observed in response to incoming mail traffic (DNS survey #3)

John Levine: conducted a survey of DNS server logs to evaluate SPF-related query traffic

McAfee: provided details about their SUBMITTER implementation and usage statistics

Santronics: contributed data about the use of the SUBMITTER extension in aggregate SMTP client traffic

The Trusted Domain Project: contributed data about the difference between Sender ID and SPF results, conducted one of the detailed TXT/SPF RRTYPE surveys including collecting timing data (DNS survey #2), and conducted the MTA SUBMITTER survey

The author would also like to thank the following for their contributions to the development of the text in this document: Dave Crocker, Scott Kitterman, Barry Leiba, John Leslie, John Levine, Hector Santos, and Alessandro Vesely.

Author's Address

Murray S. Kucherawy
Cloudmark
128 King St., 2nd Floor
San Francisco, CA 94107
USA

Phone: +1 415 946 3800
Email: msk@cloudmark.com

