Network Working Group
Internet-Draft
Expires: October 20, 2005

A. Newton, Editor
Message Anti-Abuse Working Group

April 21, 2005

Considerations for the use of the Sender Policy Framework
draft-newton-maawg-spf-considerations-00

Status of this Memo

Copyright Notice

Abstract

This document describes issues and considerations when deploying the Sender Policy Framework for the purposes of authenticating sending domains with Internet email.

Table of Contents

1. Introduction

The Sender Policy Framework (SPF) describes a DNS TXT resource record to be used for authenticating the sending domain in the transmission of Internet email (from one domain to another). SPF sits atop the current email standards and is intended to disallow certain abusive use cases of email. To accomplish this task, SPF places new restrictions upon the email transmission from one domain to another.

This document describes issues and considerations regarding the use of SPF that network and mail administrators may have no need to consider otherwise.

These considerations are collected from the membership of the Message Anti-Abuse Working Group. This document is intended to be for informational purposes only.

2. SPF Background

The Sender Policy Framework is an email sender authentication scheme that defines a syntax of a DNS TXT resource record. Through various devices, this DNS TXT resource record lists allowed, disallowed, neutral, and unknown source IP addresses. SPF relies upon the difficulty of spoofing IP addresses with SMTP connections as the basis for its authentication mechanism; senders are matched to the source IP address of an SMTP connection via various identities. Authorization to use the IP address by a sender is given in the DNS TXT resource record.

2.1 SPF Variants

SPF describes itself as a framework, and as such there are many variants of SPF. When deploying SPF, it is crucial to understand which variants of SPF are being advertised as a sender of email and the intentions of the sending SPF domain as a receiver of email. Otherwise, it is likely that non-forged email may not be delivered while forged email may be delivered. The following is a partial list of SPF variants.

2.1.1 Original SPF

The specification in [1] (and its immediate predessor) was the first specification considered to contain the basis of SPF. This specification uses the email address presented by a sending server in the SMTP MAIL FROM command or the fully qualified domain name presented by a sending server in the SMTP HELO or EHLO command if no email address is given in the MAIL FROM command. It should be noted that this specification is no longer being advanced through standardization processes by its authors. However, it is important because many implementations cite it as the specification to which they comply.

2.1.2 SPF Classic

There exists some ambiguity regarding the term "Classic SPF" or "SPF Classic" as it was originally used to distinquish original SPF (Section 2.1.1) from a variant of SPF known as Sender ID (Section 2.1.3). However, the more recent use of the term "SPF Classic" refers to [5]. This specification differs from original SPF (Section 2.1.1) in the following ways:

- DNS lookups that fail to find TXT records will use an algorithm to find DNS zone cuts and requery the DNS further up the tree.
- Checking of the fully qualified domain name given in an SMTP HELO or EHLO command is optional even if an address is given in the SMTP MAIL FROM command.

2.1.3 Sender ID

This variant of SPF is defined by [2] , [3] and [4]. It differs from the above specifications in the following ways:

- The email identity used is determined by the publisher of the SPF record.
- It defines a new email identity that is selected from the headers of the email message using an algorithm called the Purported Responsible Address (PRA).
- The syntax of the record can either be the SPF syntax specified by original SPF (Section 2.1.1) (identified by the "v=spf1" version identifier) or a slightly different SPF syntax (identified by the "spf2.0" identifier).
- The email address selected by the PRA algorithm can be transmitted by the sender to the receiver in the SMTP MAIL FROM command using an ESMTP extension called SUBMITTER.

It should be noted that some proponents of SPF Classic (Section 2.1.2) do not consider Sender ID (Section 2.1.3) to be a legitimate variant of SPF. This may cause confusion when determining the compliance of software.

2.2 SPF Email Identities

All SPF variants use a domain name to lookup an SPF DNS TXT resource record. These domain names are taken from different parts of the SMTP transaction as email is transmitted from one mail server to another.

The following is an example of an SMTP transaction. The first two columns are provided for illustrative reasons and are not part of the SMTP transaction. The first column is a line number, and the second column contains an "S:" or "C:" to show who said what in the SMTP conversation. Lines with an "S:" (for server) indicate the mail server receiving the email, and lines with a "C:" (for client) indicate the mail server sending the email.

```
1   S:  220 milo.example.org SuperDuper Mail Server v1.5
2   C:  HELO felix.example.net
3   S:  250 milo.example.org Ok.
4   C:  MAIL FROM: <bob@example.com>
5   S:  250 Ok.
6   C:  RCPT TO: <alice@example.org>
7   S:  250 Ok.
8   C:  DATA
9   S:  354 Ok.
10  C:  Subject: This is an example email
11  C:  From: Bob <bob@example.net>
12  C:  To: Alice <alice@example.com>
13  C:
14  C:  This is the body of the email message.
15  C:  It is two lines long.
16  C:  .
17  S:  250 Ok. 42548455.00000B74
```

Figure 1: An Example SMTP Transaction

The domain name given on line 2 (felix.example.net) of the SMTP transaction is the HELO identity, and it is suppose to be the fully qualified domain name of the mail server sending the email. Original SPF (Section 2.1.1) and SPF Classic (Section 2.1.2) specify that an SPF DNS TXT resource record by that name is to be consulted to determine if the source IP address of the SMTP connection is authorized to act on behalf of the HELO identity.

The domain name of the email address in Line 4 (example.com) is the MAIL FROM identity used by original SPF (Section 2.1.1) and SPF Classic (Section 2.1.2) An SPF DNS TXT resource record by that name is to be consulted to determine if the source IP address of the SMTP connection is authorized to act on behalf of the MAIL FROM identity.

Lines 10 through 12 are the headers of the email message. The PRA algorithm selects an identity from these headers. In this example, the PRA specifies the email address in Line 11 and the use of the domain name in that email address (example.net) as the Purported Responsible Domain (PRD). An SPF DNS TXT resource record by that name is to be consulted to determine if the source IP address of the SMTP connection is authorized to act on behalf of the PRA identity. See [3].

3. Identity Usage

Proper deployment of SPF requires an understanding of the proper usages of identities used in the tranmission of email from one mail server to another. Improper deployment of SPF may result in inappropriately high confidence of protection against certain classes of mail forgery or may result in the loss of certain types of message transfers or both.

3.1 User Agents

Given the example in Figure 1, the following is typical of what will be seen by an end-user with a mail client.

```
Subject: This is an example email
From:    Bob <bob@example.net>
To:      Alice <alice@example.com>
Date:    Wed, 06 Apr 2005 20:52:11 -0400

This is the body of the email message.
It is two lines long.
```

Figure 2: User Agent Example

The identity used with the SMTP MAIL FROM command is bob@example.com but the end-user sees the message as being from bob@example.net (note the difference between example.com and example.net). Therefore, SPF's use of the identity in the SMTP MAIL FROM command will not stop users from seeing a forged identity.

Since the PRA is the only identity verified that is part of the email message, verification of the PRA is the only part of SPF that attempts to insure that end-users see authorized email addresses. However, the PRA algorithm does not always select an identity that is shown to end users by mail clients. Therefore, SPF is not guaranteed to prevent end users from seeing forged identities.

Additionally, the PRA identity only focuses on the address specification part of a header in an email message. It does not validate against the display name part of a header in an email message. Using the example above:

```
   From: Bob <bob@example.net>
```

The display name portion in this header is "Bob" and the address specification portion in this header is "<bob@example.net>. Many email clients show only the display name portion of the header. Therefore, it is possible to have a positive validation against the PRA without having a positive validation against the information given to the user. For example:

From: Bob Smurd <badguy@example.net>
would be seen as:
From: Bob Smurd
yet yield a positive result.

3.2 Forwarding

Forwarding of email from one mail server to another may prevent proper delivery of messages using the MAIL FROM and PRA identities. Consider the following email forwarding scenario.

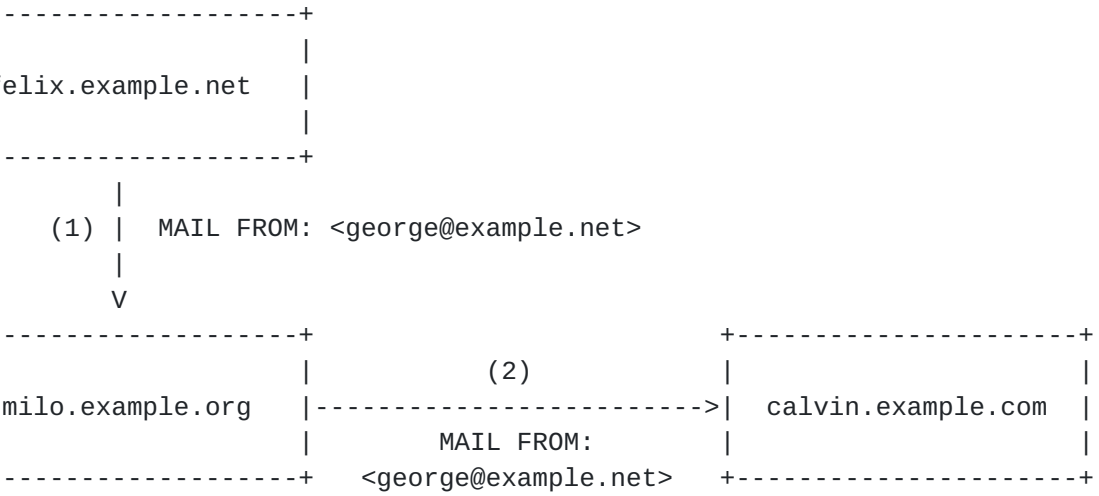

Figure 3: Mail Path Example

In this scenario, MAIL FROM checking from felix.example.net to milo.example.org would produce valid results. However, a MAIL FROM check from milo.example.org to calvin.example.com would produce invalid results as george@example.net would cause the SPF record for example.net to be consulted even though the SMTP connection is originating from example.org.

While the intent of PRA algorithm is to properly detect the last forwarder of an email, it relies upon behaviors not found in all mail servers and programs. Therefore, it is quite possible that PRA checking would also have the same results as MAIL FROM checking.

For mail services with a small and fixed number of known forwarding relationships, this problem may be overcome using the include mechanism. In this case, "include:example.org" placed in the SPF record of example.net would produce proper results. However, there are some drawbacks to using the "include" mechanism (see Section 6.3.3).

However, it is not always possible to know forwarding relationships

or to produce SPF include mechanisms for all known forwarding relationships. To overcome this problem, it is advisable to do the following:

- For MAIL FROM checking, rewriting the MAIL FROM identity to point to the domain of the sending server will cause the proper SPF records to be consulted. In the example above, if the server milo.example.org used a MAIL FROM identity of george@example.org instead of george@example.net, then the records pertaining to milo.example.org would be consulted. It is important to note that all upstream notifications are to be directed to the MAIL FROM identity, therefore MAIL FROM rewriting will require milo.example.org to properly handle any bounces of the message it is sending.
- To use PRA checking, a forwarder should insert the proper PRA compatible header or headers into the message. See [3].

Note that the HELO identity will not cause false positives with forwarding. As noted above, the MAIL FROM identity is used as the address to which bounces should be sent in case of errors, and many forwarding processes legitimately do not rewrite MAIL FROM for this reason. However, such a reason does not exist for the HELO identity and there are no legitimate reasons for a mail server to use the identity of another.

3.3 Publisher's Intent

As noted in Section 3.2, to prevent the consultation of inappropriate SPF records, senders may adopt strategies of rewriting MAIL FROM or inserting PRA compliant headers or both. The strategy picked should be reflected in the SPF record, and therefore receivers of email should follow the intent of the published SPF record regarding which identity is to be checked.

With the given SPF variants (Section 2.1), there are three types of SPF records that give scope to the identifier to be checked.

1. v=spf1
2. spf2.0/mfrom
3. spf2.0/pra

For the purposes of backwards compatibility, Sender ID (Section 2.1.3) interprets the first record to be equivalent to "spf2.0/mfrom,pra", which is to say that both MAIL FROM and PRA checks are to occur. To avoid confusion, senders should explicitly publish spf2.0/mfrom records. If a sender has not taken prepatory steps to accomodate PRA checking, a non-committal SPF record of "spf2.0/mfrom ?all" will signal that all PRA checking against this domain will have unknown results.

4. Multihoming of Mail Servers

Because SPF uses IP addresses as the key to authentication, special care must be given with mail servers that have more than one IP address, especially if they are not all listed in the SPF record. Here is one such common scenario: a separate publicly addressable network interface is given to a server for the sole purpose of remote management. In these cases, changes to the routing fabric of the Internet may cause mail service to switch away from the intended network interface to one not intended to service SMTP traffic.

To avoid this problem, mail servers should be explicitly bound to the network interfaces published in the SPF records.

5. Usage of Headers for Check Status

Original SPF (Section 2.1.1) and SPF Classic (Section 2.1.2) defines the Received-SPF trace header to be added to email messages that have undergone SPF checking. Because these are trace headers, multiple sets of Received-SPF headers may appear in a single email, each set being added by a previous mail hop. Without proper care, simple filtering (such as with an unscoped regular expression) may have unexpected results. And like Recieved headers, Received-SPF headers are an easy target for forgery. Mail filters and clients should not use their contents to determine the disposition of email messages.

Because SPF is not the only type of email authentication in use, mail servers should use the Authentication-Results header (See [6]). This header has the advantage of working with multiple authentication schemes and is intended to allow a mail server to communicate the status of an authentication check to mail filters and mail clients.

6. DNS Usage

6.1 Mechansim Lookups

Network administrators need to be conscious of the fact that SPF records can create more load on their DNS servers than just that of querying the SPF records.

SPF Classic (Section 2.1.2), which has the most strict upper bounds on DNS lookups, allows for 10 SPF mechanisms that may trigger subsequent DNS lookups. In turn each mechanism may require more than one DNS lookup to fulfill the requirements of the mechanism. The theoretical maximum to conduct an SPF check is 111 DNS lookups. For instance, the appearance of one "mx" mechanism in an SPF record could result in 10 DNS lookups in the process of following the targets of the queried DNS MX records.

To avoid lengthy processing times and excess load on DNS, the use of the "ip4" and "ip6" mechanisms is recommended. Use of the "include", "a", "mx", "ptr", "exists" mechanisms, the "redirect" modifier, and the %{p} macro should carefully consider the total number of DNS lookups incorporated into an SPF record.

Additionally, SPF Classic (Section 2.1.2) places an upper bound of 20 seconds on the duration needed to process the SPF check. SPF checks that take longer than 20 seconds or require more than 101 DNS lookups will result in message delivery rejection. Use of an SPF profiler is recommended to determine the number of DNS lookups and the potential check duration.

These restrictions also need to be taken into consideration when performing SPF checks. While 111 DNS lookups with a 20 second timeout is tolerable for infrequent email reception, receivers and senders may need to reach bilateral transaction agreements that bypass SPF in cases where SPF records cannot be formulated with more tolerable values.

6.2 Zone Cut Issues

SPF Classic (Section 2.1.2) introduces an algorithm to find the DNS zone cut of administrative domains. For example, if an SPF check is conducted against a non-existent domain name of prattle.example.net, the SPF record at example.net will be found. This algorithm is an attempt to overcome undesirable behaviour in DNS wildcards (which are not recommended for use by SPF Classic).

This algorithm has implications with regards to the confidence attached to HELO checking. Should a misconfiguration occur, a HELO

check may inadvertently consult the wrong SPF record. Because SPF records related to HELO can be stricter than SPF records relating to MAIL FROM and PRA, this would result in the loss of security.

Synthesized DNS records, whether provided by the standard DNS wildcard device or other means, are an administrative issue. Other methods can be employed to get the proper DNS behaviour by the publisher of an SPF record.

Therefore, disabling the DNS zone cut algorithm in an SPF processor is recommended.

6.3 Publication Procedures

6.3.1 Initial Publication

Because there are many unknowns in the paths email messages may take through the Internet, many publishers of SPF records make non-committal assertions regarding their message delivery (usually by ending an SPF record with "~all"). To a receiver, a non-committal assertion may not have any affect on judging the disposition of email, and issues with regard to SPF processing may not be noticed until a more aggressive SPF record is published.

SPF publishers should follow these steps on initial publication of an SPF record:

1. Initially publish the SPF record with a neutral assertion (i.e. end it in "?all").
2. Once there is a high degree of confidence that the SPF record accurately reflects message delivery, lower the TTL on the SPF record to a value that allows the record to be quickly propagated within the DNS should it need to be changed rapidly.
3. Change the SPF record so that it makes the more aggressive assertion of softfail (i.e. end it in "~all").
4. If no adverse problems are found after a sufficient period of time, change the SPF record so that it makes the most aggressive assertion (i.e. end it in "-all").
5. Once there is a high degree of confidence that the SPF record is not causing adverse mail delivery problems, increase the TTL on the SPF record to a more reasonable value.

It should be noted that decreasing the TTL on the SPF record will result in higher DNS query load.

Due to unknowable forwarding relationships with the Internet email infrastructure, it may not be possible for all domains to publish SPF records with an agressive assertion. Use of these publication procedures may even lead to the conclusion that email should not be

subjected to SPF checks (see Section 7.2).

6.3.2 Continued Publication

Once an SPF record has been established, a publisher should put in place proper procedures for the maintenance and continued publication of the SPF record.

Over time, it is highly likely that some changes will need to be made to the contents of the SPF record. Because the SPF syntax is seemingly simple, administrators may be tempted to modify the record without full knowledge of the SPF syntax. Failure to correctly modify an SPF record may result in message delivery rejections. To avoid this problem, any SPF record should be run against an SPF syntax checker before the new record is published.

Additionally, it might be useful to store SPF records in a version control system. This allows quick reversion back to a previous record should a problem be discovered. It will also help in the analysis of mail problems by allowing past records to be studied.

6.3.3 Use of the 'include' Mechanism

As stated in Section 2.1, many implementations of SPF adhere to original SPF (Section 2.1.1) even though that current branch of SPF is described by SPF Classic (Section 2.1.2). Unfortunately, both use the same "v=spf1" record identifier, so there exists no easy method to differentiante the two programmatically.

Both versions differ substantially in their error case for the "include" mechanism. Under original SPF (Section 2.1.1), if an "include" mechanism references a non-existent SPF record, SPF processing against all email for the domain making the reference would cease (essentially resulting in a state equivalent to having no SPF record for the domain making the reference).

Under SPF Classic (Section 2.1.2), if an "include" mechansim references a non-existent SPF record, SPF processing against all email for the domain making the reference would result in a PermError state and consequent permanent SMTP rejection of the email.

For SPF publishers defering to the SPF records of other domains (a common scenario for commercial enterprises that out-source their transaction email operations), the existance of the target SPF record should be verified before the use of an associated "include" mechanism. Additionally, SPF publishers should seek assurance of continued SPF publication from the SPF publishers to which they make a reference.

7. SPF Results

7.1 Administrative Utility

SPF recognizes five states with which a publisher can declare their intentional use of SPF: none, pass, fail, neutral, and softfail (PermError and TempError as defined by SPF Classic (Section 2.1.2) are error states and not intential usage states). The differences between none, neutral, and softfail may not be programattically meaningful ([5] specifies neutral to be programmatically equivalent to none).

However, these differences can have meaning to administrators attempting to resolve problems manually. The result of neutral differs from the result of none in that it does indicate that the SPF publisher is aware of SPF checking. The result of none indicates that the SPF publisher is not aware of SPF checking.

7.2 Use of Other Authentication Schemes

There do exist scenarios where mail administrators do not wish to subject their email practices to SPF checks but do wish to offer an affirmative acknowledgment of the practice of using SPF. Such a scenario would be email sending domains that wish to rely on other authentication schemes, such as cryptographic-based signature schemes.

This is easily accomplished with the exclusive use of the "all" mechanism using the pass result. Such as:

```
   v=spf1 +all
or
   spf2.0/mfrom,pra +all
```

8. Security Considerations

SPF depends on the integrity of various parts of Internet infrastructure and has other security considerations that should be understand before the deployment of SPF. Sender ID (Section 2.1.3) clearly enumerates these isses which have implications for all variants of SPF.

Briefly, these issues are:

1. SPF is only as secure as DNS. Should the integrity of DNS be compromised, then SPF becomes much less effective.
2. SPF relies on the difficult nature of IP address spoofing within TCP (the transport used by SMTP).
3. SPF does not always prevent users from seeing forged sender information even when SPF checks return positive results (see Section 3.1).
4. SPF relies upon the linkage between an SPF publisher and their assocation to IP address space. Attacks against the routing fabric of the Internet can break this linkage rendering SPF ineffective.
5. SPF may be used to perpetrate "bounce attacks". While original SPF (Section 2.1.1) and SPF Classic (Section 2.1.2) are less likely to be used for such an attack, they are not immune to it given their described usage of SoftFail.

9 Informative References

[1] Lentczner, M., "Sender Policy Framework (SPF) A Convention to Describe Hosts Authorized to Send SMTP Traffic", draft-mengwong-spf-01 (work in progress), May 2004.

[2] Lyon, J., "Sender ID: Authenticating E-Mail", draft-lyon-senderid-core-00 (work in progress), November 2004.

[3] Lyon, J., "Purported Responsible Address in E-Mail Messages", draft-lyon-senderid-pra-00 (work in progress), November 2004.

[4] Allman, E., "SMTP Service Extension for Indicating the Responsible Submitter of an E-mail Message", draft-katz-submitter-00 (work in progress), November 2004.

[5] Wong, M., "Sender Policy Framework: Authorizing Use of Domains in E-MAIL", draft-schlitt-spf-classic-00 (work in progress), January 2005.

[6] Kucherawy, M., "Message Header for Indicating Sender Authentication Status", draft-kucherawy-sender-auth-header-01 (work in progress), March 2005.

Author's Address

Andrew L. Newton
Message Anti-Abuse Working Group

EMail: anewton@verisignlabs.com; andy@hxr.us
URI: http://www.maawg.org/

Intellectual Property Statement

Disclaimer of Validity

Copyright Statement

Acknowledgment

Funding for the RFC Editor function is currently provided by the Internet Society.