

Network Working Group
Internet-Draft
Obsoletes: [4408](#) (if approved)
Intended status: Standards Track
Expires: January 6, 2013

S. Kitterman
Agari
July 5, 2012

**Sender Policy Framework (SPF) for Authorizing Use of Domains in Email,
Version 1
draft-ietf-spfbis-4408bis-03.txt**

Abstract

E-mail on the Internet can be forged in a number of ways. In particular, existing protocols place no restriction on what a sending host can use as the reverse-path of a message or the domain given on the SMTP HELO/EHLO commands. This document describes version 1 of the Sender Policy Framework (SPF) protocol, whereby a domain can explicitly authorize the hosts that are allowed to use its domain name, and a receiving host can check such authorization. This document obsoletes [RFC4408](#).

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 January 6, 2013.

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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	5
1.1.	Protocol Status	5
1.2.	Experimental History	6
1.3.	Terminology	6
1.3.1.	Keywords	6
1.3.2.	Imported Definitions	6
1.3.3.	Mail From Definition	6
1.3.4.	Deprecated	7
2.	Operation	8
2.1.	The HELO Identity	8
2.2.	The MAIL FROM Identity	8
2.3.	Publishing Authorization	8
2.4.	Checking Authorization	9
2.5.	Interpreting the Result	10
2.5.1.	None	10
2.5.2.	Neutral	11
2.5.3.	Pass	11
2.5.4.	Fail	11
2.5.5.	Softfail	11
2.5.6.	TempError	12
2.5.7.	PermError	12
3.	SPF Records	13
3.1.	Publishing	13
3.1.1.	DNS Resource Records	13
3.1.2.	Multiple DNS Records	14
3.1.3.	Multiple Strings in a Single DNS record	14
3.1.4.	Record Size	14
3.1.5.	Wildcard Records	14
4.	The check_host() Function	16
4.1.	Arguments	16
4.2.	Results	16
4.3.	Initial Processing	16
4.4.	Record Lookup	17
4.5.	Selecting Records	17
4.6.	Record Evaluation	17
4.6.1.	Term Evaluation	18
4.6.2.	Mechanisms	18
4.6.3.	Modifiers	19
4.6.4.	Evaluation Limits	19
4.7.	Default Result	19
4.8.	Domain Specification	20
5.	Mechanism Definitions	21
5.1.	"all"	21
5.2.	"include"	22
5.3.	"a"	23
5.4.	"mx"	23

5.5.	"ptr" (deprecated)	24
5.6.	"ip4" and "ip6"	25
5.7.	"exists"	26
6.	Modifier Definitions	27
6.1.	redirect: Redirected Query	27
6.2.	exp: Explanation	28
7.	The Received-SPF Header Field	30
8.	Macros	32
8.1.	Macro Definitions	32
8.2.	Expansion Examples	35
9.	Implications	36
9.1.	Sending Domains	36
9.1.1.	DNS Resource Considerations	36
9.1.2.	Administrator's Considerations	37
9.2.	Mediators	37
9.2.1.	Mailing Lists	37
9.2.2.	Forwarding Services and Aliases	37
9.3.	Mail Services	39
9.4.	MTA Relays	40
10.	Security Considerations	41
10.1.	Processing Limits	41
10.2.	SPF-Authorized Email May Contain Other False Identities	41
10.3.	Spoofed DNS and IP Data	42
10.4.	Cross-User Forgery	42
10.5.	Untrusted Information Sources	42
10.6.	Privacy Exposure	43
11.	Contributors and Acknowledgements	44
12.	IANA Considerations	45
12.1.	The SPF DNS Record Type	45
12.2.	The Received-SPF Mail Header Field	45
13.	References	46
13.1.	Normative References	46
13.2.	Informative References	47
Appendix A.	Collected ABNF	49
Appendix B.	Extended Examples	52
B.1.	Simple Examples	52
B.2.	Multiple Domain Example	53
B.3.	DNSBL Style Example	54
B.4.	Multiple Requirements Example	54
Appendix C.	Change History	55
Author's Address		57

1. Introduction

The current email infrastructure has the property that any host injecting mail into the mail system can identify itself as any domain name it wants. Hosts can do this at a variety of levels: in particular, the session, the envelope, and the mail headers. Although this feature is desirable in some circumstances, it is a major obstacle to reducing Unsolicited Bulk email (UBE, aka spam). Furthermore, many domain name holders are understandably concerned about the ease with which other entities can make use of their domain names, often with malicious intent.

This document defines a protocol by which domain owners can authorize hosts to use their domain name in the "MAIL FROM" or "HELO" identity. Compliant domain holders publish Sender Policy Framework (SPF) records specifying which hosts are permitted to use their names, and compliant mail receivers use the published SPF records to test the authorization of sending Mail Transfer Agents (MTAs) using a given "HELO" or "MAIL FROM" identity during a mail transaction.

An additional benefit to mail receivers is that after the use of an identity is verified, local policy decisions about the mail can be made based on the sender's domain, rather than the host's IP address. This is advantageous because reputation of domain names is likely to be more accurate than reputation of host IP addresses. Furthermore, if a claimed identity fails verification, local policy can take stronger action against such email, such as rejecting it.

1.1. Protocol Status

SPF has been in development since the summer of 2003 and has seen deployment beyond the developers beginning in December 2003. The design of SPF slowly evolved until the spring of 2004 and has since stabilized. There have been quite a number of forms of SPF, some written up as documents, some submitted as Internet Drafts, and many discussed and debated in development forums. The protocol was originally documented in [[RFC4408](#)], which this memo replaces.

The goal of this document is to clearly document the protocol defined by earlier draft specifications of SPF as used in existing implementations. This conception of SPF is sometimes called "SPF Classic". It is understood that particular implementations and deployments will differ from, and build upon, this work. It is hoped that we have nonetheless captured the common understanding of SPF version 1.

1.2. Experimental History

This document updates and replaces [RFC 4408](#) that was part of a group of simultaneously published Experimental RFCs ([RFC 4405](#), [RFC 4406](#), [RFC 4407](#), and [RFC 4408](#)) in 2006. At that time the IESG requested the community observe the success or failure of the two approaches documented in these RFCs during the two years following publication, in order that a community consensus can be reached in the future.

SPF is widely deployed by large and small email providers alike. There are multiple, interoperable implementations.

For SPF (as documented in [RFC 4408](#)) a careful effort was made to collect and document lessons learned and errata during the two year period. The errata list has been stable (no new submissions) and only minor protocol lessons learned were identified. Resolution of the IESG's experiment is documented in [[draft-ietf-spfbis-experiment](#)].

1.3. Terminology

1.3.1. Keywords

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)].

1.3.2. Imported Definitions

The ABNF tokens "ALPHA", "DIGIT", and "SP" are defined in [[RFC5234](#)].

The token "local-part" is defined in [[RFC5321](#)].

"dot-atom", "quoted-string", "comment", "CFWS", "FWS", and "CRLF" are defined in [[RFC5322](#)]

1.3.3. Mail From Definition

This document is concerned with the portion of a mail message commonly called "envelope sender", "return path", "reverse path", "bounce address", "5321 FROM", or "MAIL FROM". Since these terms are either not well defined or often used casually, this document defines the "MAIL FROM" identity in [Section 2.2](#). Note that other terms that might superficially look like the common terms, such as "reverse-path", are used only with the defined meanings from normative documents.

1.3.4. Deprecated

There are several [[RFC4408](#)] features that are marked "deprecated". In the context of this document, deprecated means that senders SHOULD NOT publish SPF records that make use of the feature in question. Its use is NOT RECOMMENDED and it might be removed entirely in future updates to the protocol. Such features do, however, remain part of the SPF protocol and receiving systems MUST support them unless this specification says otherwise.

2. Operation

2.1. The HELO Identity

The "HELO" identity derives from either the SMTP HELO or EHLO command (see [[RFC5321](#)]). These commands supply the SMTP client (sending host) for the SMTP session. Note that requirements for the domain presented in the EHLO or HELO command are not always clear to the sending party, and SPF clients must be prepared for the "HELO" identity to be malformed or an IP address literal. At the time of this writing, many legitimate emails are delivered with invalid HELO domains.

It is RECOMMENDED that SPF clients not only check the "MAIL FROM" identity, but also separately check the "HELO" identity by applying the `check_host()` function ([Section 4](#)) to the "HELO" identity as the `<sender>`.

2.2. The MAIL FROM Identity

The "MAIL FROM" identity derives from the SMTP MAIL command (see [[RFC5321](#)]). This command supplies the "reverse-path" for a message, which generally consists of the sender mailbox, and is the mailbox to which notification messages are to be sent if there are problems delivering the message.

[RFC5321] allows the reverse-path to be null (see [Section 4.5.5 in RFC 5321](#)). In this case, there is no explicit sender mailbox, and such a message can be assumed to be a notification message from the mail system itself. When the reverse-path is null, this document defines the "MAIL FROM" identity to be the mailbox composed of the localpart "postmaster" and the "HELO" identity (which might or might not have been checked separately before).

SPF clients MUST check the "MAIL FROM" identity if a completed "HELO" check has not reached a definitive policy result by applying the `check_host()` function to the "MAIL FROM" identity as the `<sender>`.

2.3. Publishing Authorization

An SPF-compliant domain MUST publish a valid SPF record as described in [Section 3](#). This record authorizes the use of the domain name in the "HELO" and "MAIL FROM" identities by the MTAs it specifies.

If domain owners choose to publish SPF records, it is RECOMMENDED that they end in "-all", or redirect to other records that do, so that a definitive determination of authorization can be made.

Domain holders can publish SPF records that explicitly authorize no hosts if mail should never originate using that domain.

When changing SPF records, care must be taken to ensure that there is a transition period so that the old policy remains valid until all legitimate email has been checked.

[2.4.](#) Checking Authorization

A mail receiver can perform a set of SPF checks for each mail message it receives. An SPF check tests the authorization of a client host to emit mail with a given identity. Typically, such checks are done by a receiving MTA, but can be performed elsewhere in the mail processing chain so long as the required information is available and reliable. At least the "MAIL FROM" identity **MUST** be checked, but it is **RECOMMENDED** that the "HELO" identity also be checked beforehand.

Without explicit approval of the domain owner, checking other identities against SPF version 1 records is **NOT RECOMMENDED** because there are cases that are known to give incorrect results. For example, almost all mailing lists rewrite the "MAIL FROM" identity (see [Section 9.2.1](#)), but some do not change any other identities in the message. The scenario described in [Section 9.2.2](#), sub-[section 1.2](#), is another example. Documents that define other identities should define the method for explicit approval.

It is possible that mail receivers will use the SPF check as part of a larger set of tests on incoming mail. The results of other tests might influence whether or not a particular SPF check is performed. For example, finding the sending host's IP address on a local white list might cause all other tests to be skipped and all mail from that host to be accepted.

When a mail receiver decides to perform an SPF check, it **MUST** use a correctly-implemented `check_host()` function ([Section 4](#)) evaluated with the correct parameters. Although the test as a whole is optional, once it has been decided to perform a test it must be performed as specified so that the correct semantics are preserved between publisher and receiver.

To make the test, the mail receiver **MUST** evaluate the `check_host()` function with the arguments set as follows:

<ip> - the IP address of the SMTP client that is emitting the mail, either IPv4 or IPv6.

<domain> - the domain portion of the "MAIL FROM" or "HELO" identity.

<sender> - the "MAIL FROM" or "HELO" identity.

Note that the <domain> argument might not be a well-formed domain name. For example, if the reverse-path was null, then the EHLO/HELO domain is used, with its associated problems (see [Section 2.1](#)). In these cases, `check_host()` is defined in [Section 4.3](#) to return a "none" result.

Although invalid, malformed, or non-existent domains cause SPF checks to return "none" because no SPF record can be found, it has long been the policy of many MTAs to reject email from such domains, especially in the case of invalid "MAIL FROM". Rejecting email will prevent one method of circumventing of SPF records.

Implementations must take care to correctly extract the <domain> from the data given with the SMTP MAIL FROM command as many MTAs will still accept such things as source routes (see [[RFC5321](#)], [Appendix C](#)), the %-hack (see [[RFC1123](#)]), and bang paths (see [[RFC1983](#)]). These archaic features have been maliciously used to bypass security systems.

[2.5. Interpreting the Result](#)

This section describes how software that performs the authorization should interpret the results of the `check_host()` function. The authorization check SHOULD be performed during the processing of the SMTP transaction that sends the mail. This allows errors to be returned directly to the sending MTA by way of SMTP replies.

Performing the authorization after the SMTP transaction has finished can cause problems, such as the following: (1) It might be difficult to accurately extract the required information from potentially deceptive headers; (2) legitimate email might fail because the sender's policy had since changed.

Generating non-delivery notifications to forged identities that have failed the authorization check is generally abusive and against the explicit wishes of the identity owner.

[2.5.1. None](#)

A result of "none" means that no records were published by the domain or that no checkable sender domain could be determined from the given identity. The checking software cannot ascertain whether or not the client host is authorized.

2.5.2. Neutral

The domain owner has explicitly stated that he cannot or does not want to assert whether or not the IP address is authorized. A "neutral" result MUST be treated exactly like the "none" result; the distinction exists only for informational purposes. Treating "neutral" more harshly than "none" would discourage domain owners from testing the use of SPF records (see [Section 9.1](#)).

2.5.3. Pass

A "pass" result means that the client is authorized to inject mail with the given identity. The domain can now, in the sense of reputation, be considered responsible for sending the message. Further policy checks can now proceed with confidence in the legitimate use of the identity.

2.5.4. Fail

A "fail" result is an explicit statement that the client is not authorized to use the domain in the given identity. The checking software can choose to mark the mail based on this or to reject the mail outright.

If the checking software chooses to reject the mail during the SMTP transaction, then it SHOULD use an SMTP reply code of 550 (see [\[RFC5321\]](#)) and, if supported, the 5.7.1 enhanced status code (see [\[RFC3463\]](#)), in addition to an appropriate reply text. The `check_host()` function will return either a default explanation string or one from the domain that published the SPF records (see [Section 6.2](#)). If the information does not originate with the checking software, it should be made clear that the text is provided by the sender's domain. For example:

```
550-5.7.1 SPF MAIL FROM check failed:
550-5.7.1 The domain example.com explains:
550 5.7.1 Please see http://www.example.com/mailpolicy.html
```

2.5.5. Softfail

A "softfail" result should be treated as somewhere between a "fail" and a "neutral". The domain believes the host is not authorized but is not willing to make that strong of a statement. Receiving software SHOULD NOT reject the message based solely on this result, but MAY subject the message to closer scrutiny than normal.

The domain owner wants to discourage the use of this host and thus desires limited feedback when a "softfail" result occurs. For

example, the recipient's Mail User Agent (MUA) could highlight the "softfail" status, or the receiving MTA could give the sender a message using greylisting, [[RFC6647](#)], with a note the first time the message is received, but accept it on a later attempt based on receiver policy.

2.5.6. TempError

A "temperror" result means that the SPF client encountered a transient error while performing the check. Checking software can choose to accept or temporarily reject the message. If the message is rejected during the SMTP transaction for this reason, the software SHOULD use an SMTP reply code of 451 and, if supported, the 4.4.3 enhanced status code.

2.5.7. PermError

A "permerror" result means that the domain's published records could not be correctly interpreted. This signals an error condition that requires manual intervention to be resolved, as opposed to the temperror result. If the message is rejected during the SMTP transaction for this reason, the software SHOULD use an SMTP reply code of 550 and, if supported, the 5.5.2 enhanced status code. Be aware that if the domain owner uses macros ([Section 8](#)), it is possible that this result is due to the checked identities having an unexpected format.

3. SPF Records

An SPF record is a DNS Resource Record (RR) that declares which hosts are, and are not, authorized to use a domain name for the "HELO" and "MAIL FROM" identities. Loosely, the record partitions all hosts into permitted and not-permitted sets (though some hosts might fall into neither category).

The SPF record is a single string of text. An example record is the following:

```
v=spf1 +mx a:colo.example.com/28 -all
```

This record has a version of "spf1" and three directives: "+mx", "a:colo.example.com/28" (the + is implied), and "-all".

3.1. Publishing

Domain owners wishing to be SPF compliant must publish SPF records for the hosts that are used in the "MAIL FROM" and "HELO" identities. The SPF records are placed in the DNS tree at the host name it pertains to, not a subdomain under it, such as is done with SRV records.

The example above in [Section 3](#) might be published via these lines in a domain zone file:

```
example.com.          TXT "v=spf1 +mx a:colo.example.com/28 -all"
smtp-out.example.com. TXT "v=spf1 a -all"
```

When publishing via TXT records, beware of other TXT records published there for other purposes. They might cause problems with size limits (see [Section 3.1.4](#)).

Domains publishing records SHOULD try to keep the number of "include" mechanisms and chained "redirect" modifiers to a minimum. Domains SHOULD also try to minimize the amount of other DNS information needed to evaluate a record. [Section 9.1.1](#) provides some suggestions on how to achieve this.

3.1.1. DNS Resource Records

SPF records MUST be published as type TXT [[RFC1035](#)]. The character content of the record is encoded as [[US-ASCII](#)]. Use of alternate DNS RR types has been dropped. See [Appendix A](#) of [[draft-ietf-spfbis-experiment](#)].

3.1.2. Multiple DNS Records

A domain name MUST NOT have multiple records that would cause an authorization check to select more than one record. See [Section 4.5](#) for the selection rules.

3.1.3. Multiple Strings in a Single DNS record

As defined in [[RFC1035](#)] sections [3.3.14](#) and [3.3](#), a single text DNS record (either TXT or SPF RR types) can be composed of more than one string. If a published record contains multiple strings, then the record MUST be treated as if those strings are concatenated together without adding spaces. For example:

```
IN TXT "v=spf1 .... first" "second string..."
```

MUST be treated as equivalent to

```
IN TXT "v=spf1 .... firstsecond string..."
```

TXT records containing multiple strings are useful in constructing records that would exceed the 255-byte maximum length of a string within a single TXT record.

3.1.4. Record Size

The published SPF record for a given domain name SHOULD remain small enough that the results of a query for it will fit within 512 octets. This will keep even older DNS implementations from falling over to TCP. Since the answer size is dependent on many things outside the scope of this document, it is only possible to give this guideline: If the combined length of the DNS name and the text of all the records of a given type is under 450 characters, then DNS answers should fit in UDP packets. Note that when computing the sizes for queries of the TXT format, one must take into account any other TXT records published at the domain name. Records that are too long to fit in a single UDP packet MAY be silently ignored by SPF clients.

3.1.5. Wildcard Records

Use of wildcard records for publishing is not recommended. Care must be taken if wildcard records are used. If a domain publishes wildcard MX records, it might want to publish wildcard declarations, subject to the same requirements and problems. In particular, the declaration must be repeated for any host that has any RR records at all, and for subdomains thereof. For example, the example given in [[RFC1034](#)], [Section 4.3.3](#), could be extended with the following:

X.COM.	MX	10	A.X.COM
X.COM.	TXT	"v=spf1 a:A.X.COM -all"	
*.X.COM.	MX	10	A.X.COM
*.X.COM.	TXT	"v=spf1 a:A.X.COM -all"	
A.X.COM.	A	1.2.3.4	
A.X.COM.	MX	10	A.X.COM
A.X.COM.	TXT	"v=spf1 a:A.X.COM -all"	
*.A.X.COM.	MX	10	A.X.COM
*.A.X.COM.	TXT	"v=spf1 a:A.X.COM -all"	

Notice that SPF records must be repeated twice for every name within the domain: once for the name, and once with a wildcard to cover the tree under the name.

Use of wildcards is discouraged in general as they cause every name under the domain to exist and queries against arbitrary names will never return RCODE 3 (Name Error).

4. The `check_host()` Function

The `check_host()` function fetches SPF records, parses them, and interprets them to determine whether a particular host is or is not permitted to send mail with a given identity. Mail receivers that perform this check **MUST** correctly evaluate the `check_host()` function as described here.

Implementations **MAY** use a different algorithm than the canonical algorithm defined here, so long as the results are the same in all cases.

4.1. Arguments

The `check_host()` function takes these arguments:

- `<ip>` - the IP address of the SMTP client that is emitting the mail, either IPv4 or IPv6.
- `<domain>` - the domain that provides the sought-after authorization information; initially, the domain portion of the "MAIL FROM" or "HELO" identity.
- `<sender>` - the "MAIL FROM" or "HELO" identity.

The domain portion of `<sender>` will usually be the same as the `<domain>` argument when `check_host()` is initially evaluated. However, this will generally not be true for recursive evaluations (see [Section 5.2](#) below).

Actual implementations of the `check_host()` function might need additional arguments.

4.2. Results

The function `check_host()` can return one of several results described in [Section 2.5](#). Based on the result, the action to be taken is determined by the local policies of the receiver.

4.3. Initial Processing

If the `<domain>` is malformed (label longer than 63 characters, zero-length label not at the end, etc.) or is not a fully qualified domain name, or if the DNS lookup returns "domain does not exist" (RCODE 3), `check_host()` immediately returns the result "none".

If the `<sender>` has no localpart, substitute the string "postmaster" for the localpart.

4.4. Record Lookup

In accordance with how the records are published (see [Section 3.1](#) above), a DNS query needs to be made for the <domain> name, querying for type TXT only.

If all DNS lookups that are made return a server failure (RCODE 2), or other error (RCODE other than 0 or 3), or time out, then `check_host()` exits immediately with the result "temperror". Alternatively, for a server failure (RCODE 2) result, `check_host()` MAY track failures and treat multiple failures within 24 hours for the same domain as "permerror".

This alternative is intended to shorten the queue time of messages that cannot be accepted, by returning a permanent negative completion reply code to the client, instead of a transient one. Saving queries is accomplished according to [\[RFC2308\]](#).

4.5. Selecting Records

Records begin with a version section:

```
record          = version terms *SP
version         = "v=spf1"
```

Starting with the set of records that were returned by the lookup, discard records that do not begin with a version section of exactly "v=spf1". Note that the version section is terminated either by an SP character or the end of the record. A record with a version section of "v=spf10" does not match and must be discarded.

There should be exactly one record remaining and evaluation can proceed. If there are two or more records remaining, then `check_host()` exits immediately with the result of "permerror".

If no matching records are returned, an SPF client MUST assume that the domain makes no SPF declarations. SPF processing MUST stop and return "none".

4.6. Record Evaluation

After one SPF record has been selected, the `check_host()` function parses and interprets it to find a result for the current test. If there are any syntax errors, `check_host()` returns immediately with the result "permerror".

Implementations MAY choose to parse the entire record first and return "permerror" if the record is not syntactically well formed.

However, in all cases, any syntax errors anywhere in the record MUST be detected.

[4.6.1.](#) Term Evaluation

There are two types of terms: mechanisms and modifiers. A record contains an ordered list of these as specified in the following Augmented Backus-Naur Form (ABNF).

```
terms                = *( 1*SP ( directive / modifier ) )

directive            = [ qualifier ] mechanism
qualifier            = "+" / "-" / "?" / "~"
mechanism            = ( all / include
                        / A / MX / PTR / IP4 / IP6 / exists )
modifier            = redirect / explanation / unknown-modifier
unknown-modifier    = name "=" macro-string
                        ; where name is not any known modifier

name                = ALPHA *( ALPHA / DIGIT / "-" / "_" / "." )
```

Most mechanisms allow a ":" or "/" character after the name.

Modifiers always contain an equals ('=') character immediately after the name, and before any ":" or "/" characters that might be part of the macro-string.

Terms that do not contain any of "=", ":", or "/" are mechanisms, as defined in [Section 5](#).

As per the definition of the ABNF notation in [[RFC5234](#)], mechanism and modifier names are case-insensitive.

[4.6.2.](#) Mechanisms

Each mechanism is considered in turn from left to right. If there are no more mechanisms, the result is specified in [Section 4.7](#).

When a mechanism is evaluated, one of three things can happen: it can match, not match, or throw an exception.

If it matches, processing ends and the qualifier value is returned as the result of that record. If it does not match, processing continues with the next mechanism. If it throws an exception, mechanism processing ends and the exception value is returned.

The possible qualifiers, and the results they return are as follows:


```
"+" pass
"-" fail
"~" softfail
"?" neutral
```

The qualifier is optional and defaults to "+".

When a mechanism matches and the qualifier is "-", then a "fail" result is returned and the explanation string is computed as described in [Section 6.2](#).

The specific mechanisms are described in [Section 5](#).

[4.6.3](#). Modifiers

Modifiers are not mechanisms: they do not return match or not-match. Instead they provide additional information. Although modifiers do not directly affect the evaluation of the record, the "redirect" modifier has an effect after all the mechanisms have been evaluated.

[4.6.4](#). Evaluation Limits

SPF implementations MUST limit the number of mechanisms and modifiers ("terms") that do DNS lookups to at most 10 per during SPF record evaluation. Specifically, the "include", "a", "mx", "ptr", and "exists" mechanisms as well as the "redirect" modifier count against this limit. The "all", "ip4", and "ip6" mechanisms do not count against this limit. If this number is exceeded during a check, a permerror MUST be returned. The "exp" modifier does not count against this limit because the DNS lookup to fetch the explanation string occurs after the SPF record evaluation has been completed.

When evaluating the "mx" and "ptr" mechanisms, or the `%{p}` macro, there MUST be a limit of no more than 10 MX or PTR RRs looked up and checked. If more than 10 "mx" or "ptr" records are returned for this further lookup, a permerror MUST be returned. This limit is per mechanism or macro in the record and in addition to the lookup limits above.

[4.7](#). Default Result

If none of the mechanisms match and there is no "redirect" modifier, then the `check_host()` returns a result of "neutral", just as if "?all" were specified as the last directive. If there is a "redirect" modifier, `check_host()` proceeds as defined in [Section 6.1](#).

Note that records SHOULD always use either a "redirect" modifier or an "all" mechanism to explicitly terminate processing.

For example:

```
v=spf1 +mx -all  
or  
v=spf1 +mx redirect=_spf.example.com
```

4.8. Domain Specification

Several of these mechanisms and modifiers have a domain-spec section. The domain-spec string is macro expanded (see [Section 8](#)). The resulting string is the common presentation form of a fully-qualified DNS name: a series of labels separated by periods. This domain is called the <target-name> in the rest of this document.

Note: The result of the macro expansion is not subject to any further escaping. Hence, this facility cannot produce all characters that are legal in a DNS label (e.g., the control characters). However, this facility is powerful enough to express legal host names and common utility labels (such as "_spf") that are used in DNS.

For several mechanisms, the <domain-spec> is optional. If it is not provided, the <domain> is used as the <target-name>.

5. Mechanism Definitions

This section defines two types of mechanisms.

Basic mechanisms contribute to the language framework. They do not specify a particular type of authorization scheme.

```
all
include
```

Designated sender mechanisms are used to designate a set of <ip> addresses as being permitted or not permitted to use the <domain> for sending mail.

```
a
mx
ptr (deprecated)
ip4
ip6
exists
```

The following conventions apply to all mechanisms that perform a comparison between <ip> and an IP address at any point:

If no CIDR-length is given in the directive, then <ip> and the IP address are compared for equality. (Here, CIDR is Classless Inter-Domain Routing.)

If a CIDR-length is specified, then only the specified number of high-order bits of <ip> and the IP address are compared for equality.

When any mechanism fetches host addresses to compare with <ip>, when <ip> is an IPv4 address, A records are fetched, when <ip> is an IPv6 address, AAAA records are fetched. Even if the SMTP connection is via IPv6, an IPv4-mapped IPv6 IP address (see [[RFC4291](#)], [Section 2.5.5](#)) MUST still be considered an IPv4 address and MUST be evaluated using IPv4 mechanisms (i.e. "ip4" and "a").

Several mechanisms rely on information fetched from DNS. For these DNS queries, except where noted, if the DNS server returns an error (RCODE other than 0 or 3) or the query times out, the mechanism throws the exception "temperror". If the server returns "domain does not exist" (RCODE 3), then evaluation of the mechanism continues as if the server returned no error (RCODE 0) and zero answer records.

5.1. "all"

```
all                = "all"
```


The "all" mechanism is a test that always matches. It is used as the rightmost mechanism in a record to provide an explicit default.

For example:

```
v=spf1 a mx -all
```

Mechanisms after "all" will never be tested. Any "redirect" modifier ([Section 6.1](#)) has no effect when there is an "all" mechanism.

5.2. "include"

```
include          = "include" ":" domain-spec
```

The "include" mechanism triggers a recursive evaluation of `check_host()`. The domain-spec is expanded as per [Section 8](#). Then `check_host()` is evaluated with the resulting string as the <domain>. The <ip> and <sender> arguments remain the same as in the current evaluation of `check_host()`.

In hindsight, the name "include" was poorly chosen. Only the evaluated result of the referenced SPF record is used, rather than acting as if the referenced SPF record was literally included in the first. For example, evaluating a "-all" directive in the referenced record does not terminate the overall processing and does not necessarily result in an overall "fail". (Better names for this mechanism would have been "if-pass", "on-pass", etc.)

The "include" mechanism makes it possible for one domain to designate multiple administratively-independent domains. For example, a vanity domain "example.net" might send mail using the servers of administratively-independent domains example.com and example.org.

Example.net could say

```
IN TXT "v=spf1 include:example.com include:example.org -all"
```

This would direct `check_host()` to, in effect, check the records of example.com and example.org for a "pass" result. Only if the host were not permitted for either of those domains would the result be "fail".

Whether this mechanism matches, does not match, or throws an exception depends on the result of the recursive evaluation of `check_host()`:

A recursive check_host() result of:	Causes the "include" mechanism to:
pass	match
fail	not match
softfail	not match
neutral	not match
temperror	throw temperror
permerror	throw permerror
none	throw permerror

The "include" mechanism is intended for crossing administrative boundaries. Although it is possible to use includes to consolidate multiple domains that share the same set of designated hosts, domains are encouraged to use redirects where possible, and to minimize the number of includes within a single administrative domain. For example, if example.com and example.org were managed by the same entity, and if the permitted set of hosts for both domains was "mx:example.com", it would be possible for example.org to specify "include:example.com", but it would be preferable to specify "redirect=example.com" or even "mx:example.com".

5.3. "a"

This mechanism matches if <ip> is one of the <target-name>'s IP addresses.

```
A = "a" [ ":" domain-spec ] [ dual-cidr-length ]
```

An address lookup is done on the <target-name>. The <ip> is compared to the returned address(es). If any address matches, the mechanism matches.

5.4. "mx"

This mechanism matches if <ip> is one of the MX hosts for a domain name.

```
MX = "mx" [ ":" domain-spec ] [ dual-cidr-length ]
```


check_host() first performs an MX lookup on the <target-name>. Then it performs an address lookup on each MX name returned. The <ip> is compared to each returned IP address. To prevent Denial of Service (DoS) attacks, more than 10 MX names MUST NOT be looked up during the evaluation of an "mx" mechanism (see [Section 10](#)). If any address matches, the mechanism matches.

Note regarding implicit MXs: If the <target-name> has no MX records, check_host() MUST NOT pretend the target is its single MX, and MUST NOT default to an A or AAAA lookup on the <target-name> directly. This behavior breaks with the legacy "implicit MX" rule. See [\[RFC5321\]](#), [Section 5](#). If such behavior is desired, the publisher should specify an "a" directive.

5.5. "ptr" (deprecated)

This mechanism tests whether the DNS reverse-mapping for <ip> exists and correctly points to a domain name within a particular domain. This mechanism is deprecated and SHOULD NOT be used.

```
PTR                = "ptr"    [ ":" domain-spec ]
```

First, the <ip>'s name is looked up using this procedure: perform a DNS reverse-mapping for <ip>, looking up the corresponding PTR record in "in-addr.arpa." if the address is an IPv4 one and in "ip6.arpa." if it is an IPv6 address. For each record returned, validate the domain name by looking up its IP address. To prevent DoS attacks, more than 10 PTR names MUST NOT be looked up during the evaluation of a "ptr" mechanism (see [Section 10](#)). If <ip> is among the returned IP addresses, then that domain name is validated. In pseudocode:

```
sending-domain_names := ptr_lookup(sending-host_IP);
if more than 10 sending-domain_names are found, use at most 10.
for each name in (sending-domain_names) {
    IP_addresses := a_lookup(name);
    if the sending-domain_IP is one of the IP_addresses {
        validated-sending-domain_names += name;
    }
}
```

Check all validated domain names to see if they end in the <target-name> domain. If any do, this mechanism matches. If no validated domain name can be found, or if none of the validated domain names end in the <target-name>, this mechanism fails to match. If a DNS error occurs while doing the PTR RR lookup, then this mechanism fails to match. If a DNS error occurs while doing an A RR lookup, then that domain name is skipped and the search continues.

Pseudocode:

```
for each name in (validated-sending-domain_names) {
  if name ends in <domain-spec>, return match.
  if name is <domain-spec>, return match.
}
return no-match.
```

This mechanism matches if the <target-name> is either an ancestor of a validated domain name or if the <target-name> and a validated domain name are the same. For example: "mail.example.com" is within the domain "example.com", but "mail.bad-example.com" is not.

Note: This mechanism has been deprecated because it is slow, it is not as reliable as other mechanisms in cases of DNS errors, and it places a large burden on the arpa name servers. If used, proper PTR records must be in place for the domain's hosts and the "ptr" mechanism should be one of the last mechanisms checked. After many yaers of SPF deployment experience it has been concluded it is unnecessary and more reliable alternatives used instead.

5.6. "ip4" and "ip6"

These mechanisms test whether <ip> is contained within a given IP network.

```
IP4          = "ip4"      ":" ip4-network  [ ip4-cidr-length ]
IP6          = "ip6"      ":" ip6-network  [ ip6-cidr-length ]
```

```
ip4-cidr-length = "/" 1*DIGIT
ip6-cidr-length = "/" 1*DIGIT
dual-cidr-length = [ ip4-cidr-length ] [ "/" ip6-cidr-length ]
```

```
ip4-network    = qnum "." qnum "." qnum "." qnum
qnum           = DIGIT                ; 0-9
                / %x31-39 DIGIT        ; 10-99
                / "1" 2DIGIT           ; 100-199
                / "2" %x30-34 DIGIT     ; 200-249
                / "25" %x30-35          ; 250-255
                ; as per conventional dotted quad notation. e.g., 192.0.2.0
ip6-network    = <as per \[RFC 4291\], section 2.2>
                ; e.g., 2001:DB8::CD30
```

The <ip> is compared to the given network. If CIDR-length high-order bits match, the mechanism matches.

If ip4-cidr-length is omitted, it is taken to be "/32". If ip6-cidr-length is omitted, it is taken to be "/128". It is not

permitted to omit parts of the IP address instead of using CIDR notations. That is, use 192.0.2.0/24 instead of 192.0.2.

5.7. "exists"

This mechanism is used to construct an arbitrary domain name that is used for a DNS A record query. It allows for complicated schemes involving arbitrary parts of the mail envelope to determine what is permitted.

`exists` = "exists" ":" domain-spec

The domain-spec is expanded as per [Section 8](#). The resulting domain name is used for a DNS A RR lookup. If any A record is returned, this mechanism matches. The lookup type is A even when the connection type is IPv6.

Domains can use this mechanism to specify arbitrarily complex queries. For example, suppose example.com publishes the record:

```
v=spf1 exists:%{ir}%.%{l1r+-}._spf.%{d} -all
```

The <target-name> might expand to "1.2.0.192.someuser._spf.example.com". This makes fine-grained decisions possible at the level of the user and client IP address.

This mechanism enables queries that mimic the style of tests that existing anti-spam DNS blacklists (DNSBL) use.

6. Modifier Definitions

Modifiers are name/value pairs that provide additional information. Modifiers always have an "=" separating the name and the value.

The modifiers defined in this document ("redirect" and "exp") MAY appear anywhere in the record, but SHOULD appear at the end, after all mechanisms. Ordering of these two modifiers does not matter. These two modifiers MUST NOT appear in a record more than once each. If they do, then `check_host()` exits with a result of "permerror".

Unrecognized modifiers MUST be ignored no matter where in a record, or how often. This allows implementations of this document to gracefully handle records with modifiers that are defined in other specifications.

6.1. redirect: Redirected Query

If all mechanisms fail to match, and a "redirect" modifier is present, then processing proceeds as follows:

```
redirect          = "redirect" "=" domain-spec
```

The domain-spec portion of the redirect section is expanded as per the macro rules in [Section 8](#). Then `check_host()` is evaluated with the resulting string as the <domain>. The <ip> and <sender> arguments remain the same as in the current evaluation of `check_host()`.

The result of this new evaluation of `check_host()` is then considered the result of the current evaluation with the exception that if no SPF record is found, or if the target-name is malformed, the result is a "permerror" rather than "none".

Note that the newly-queried domain can itself specify redirect processing.

This facility is intended for use by organizations that wish to apply the same record to multiple domains. For example:

```
la.example.com. TXT "v=spf1 redirect=_spf.example.com"
ny.example.com. TXT "v=spf1 redirect=_spf.example.com"
sf.example.com. TXT "v=spf1 redirect=_spf.example.com"
_spf.example.com. TXT "v=spf1 mx:example.com -all"
```

In this example, mail from any of the three domains is described by the same record. This can be an administrative advantage.

Note: In general, the domain "A" cannot reliably use a redirect to another domain "B" not under the same administrative control. Since the <sender> stays the same, there is no guarantee that the record at domain "B" will correctly work for mailboxes in domain "A", especially if domain "B" uses mechanisms involving localparts. An "include" directive is generally be more appropriate.

For clarity, it is RECOMMENDED that any "redirect" modifier appear as the very last term in a record.

6.2. exp: Explanation

explanation = "exp" "=" domain-spec

If check_host() results in a "fail" due to a mechanism match (such as "-all"), and the "exp" modifier is present, then the explanation string returned is computed as described below. If no "exp" modifier is present, then either a default explanation string or an empty explanation string MUST be returned.

The domain-spec is macro expanded (see [Section 8](#)) and becomes the <target-name>. The DNS TXT record for the <target-name> is fetched.

If there are any DNS processing errors (any RCODE other than 0), or if no records are returned, or if more than one record is returned, or if there are syntax errors in the explanation string, then proceed as if no exp modifier was given.

The fetched TXT record's strings are concatenated with no spaces, and then treated as an explain-string, which is macro-expanded. This final result is the explanation string. Implementations MAY limit the length of the resulting explanation string to allow for other protocol constraints and/or reasonable processing limits. Since the explanation string is intended for an SMTP response and [\[RFC5321\]](#) [Section 2.4](#) says that responses are in [\[US-ASCII\]](#), the explanation string is also limited to US-ASCII.

Software evaluating check_host() can use this string to communicate information from the publishing domain in the form of a short message or URL. Software SHOULD make it clear that the explanation string comes from a third party. For example, it can prepend the macro string "%{o} explains: " to the explanation, such as shown in [Section 2.5.4](#).

Suppose example.com has this record:


```
v=spf1 mx -all exp=explain._spf.{d}
```

Here are some examples of possible explanation TXT records at `explain._spf.example.com`:

```
"Mail from example.com should only be sent by its own servers."
```

```
-- a simple, constant message
```

```
"{i} is not one of {d}'s designated mail servers."
```

```
-- a message with a little more information, including the IP  
address that failed the check
```

```
"See http://{d}/why.html?s={S}&i={I}"
```

```
-- a complicated example that constructs a URL with the  
arguments to check_host() so that a web page can be  
generated with detailed, custom instructions
```

Note: During recursion into an "include" mechanism, an `exp=` modifier from the `<target-name>` MUST NOT be used. In contrast, when executing a "redirect" modifier, an `exp=` modifier from the original domain MUST NOT be used.

7. The Received-SPF Header Field

It is RECOMMENDED that SMTP receivers record the result of SPF processing in the message header. If an SMTP receiver chooses to do so, it SHOULD use the "Received-SPF" header field defined here for each identity that was checked. This information is intended for the recipient. (Information intended for the sender is described in [Section 6.2](#), Explanation.)

The Received-SPF header field is a trace field (see [[RFC5322](#)] [Section 3.6.7](#)) and SHOULD be prepended to the existing header, above the Received: field that is generated by the SMTP receiver. It MUST appear above all other Received-SPF fields in the message. The header field has the following format:

```
header-field      = "Received-SPF:" [CFWS] result FWS [comment FWS]
                   [ key-value-list ] CRLF

result            = "pass" / "fail" / "softfail" / "neutral" /
                   "none" / "temperror" / "permerror"

key-value-list    = key-value-pair *( ";" [CFWS] key-value-pair )
                   [";"]

key-value-pair    = key [CFWS] "=" ( dot-atom / quoted-string )

key               = "client-ip" / "envelope-from" / "helo" /
                   "problem" / "receiver" / "identity" /
                   mechanism / name

identity          = "mailfrom"      ; for the "MAIL FROM" identity
                   / "helo"         ; for the "HELO" identity
                   / name           ; other identities

dot-atom          = <unquoted word as per [RFC5322]>
quoted-string     = <quoted string as per [RFC5322]>
comment           = <comment string as per [RFC5322]>
CFWS              = <comment or folding white space as per [RFC5322]>
FWS               = <folding white space as per [RFC5322]>
CRLF              = <standard end-of-line token as per [RFC2532]>
```

The header field SHOULD include a "(...)" style comment after the result, conveying supporting information for the result, such as <ip>, <sender>, and <domain>.

The following key-value pairs are designed for later machine parsing. SPF clients SHOULD give enough information so that the SPF results can be verified. That is, at least "client-ip", "helo", and, if the

"MAIL FROM" identity was checked, "envelope-from".

client-ip	the IP address of the SMTP client
envelope-from	the envelope sender mailbox
helo	the host name given in the HELO or EHLO command
mechanism	the mechanism that matched (if no mechanisms matched, substitute the word "default")
problem	if an error was returned, details about the error
receiver	the host name of the SPF client
identity	the identity that was checked; see the <identity> ABNF rule

Other keys MAY be defined by SPF clients.

SPF clients MUST make sure that the Received-SPF header field does not contain invalid characters, is not excessively long, and does not contain malicious data that has been provided by the sender.

Examples of various header styles that could be generated are the following:

```
Received-SPF: pass (mybox.example.org: domain of
  myname@example.com designates 192.0.2.1 as permitted sender)
  receiver=mybox.example.org; client-ip=192.0.2.1;
  envelope-from="myname@example.com"; helo=foo.example.com;
```

```
Received-SPF: fail (mybox.example.org: domain of
  myname@example.com does not designate
  192.0.2.1 as permitted sender)
  identity=mailfrom; client-ip=192.0.2.1;
  envelope-from="myname@example.com";
```


8. Macros

8.1. Macro Definitions

Many mechanisms and modifiers perform macro expansion on the term.

```

domain-spec      = macro-string domain-end
domain-end       = ( "." toplabel [ "." ] ) / macro-expand

toplabel         = ( *alphanum ALPHA *alphanum ) /
                  ( 1*alphanum "-" *( alphanum / "-" ) alphanum )
                  ; LDH rule plus additional TLD restrictions
                  ; (see \[RFC3696\], Section 2 for background)
alphanum         = ALPHA / DIGIT

explain-string   = *( macro-string / SP )

macro-string     = *( macro-expand / macro-literal )
macro-expand     = ( "%{" macro-letter transformers *delimiter "}" )
                  / "%%" / "%_" / "%-"
macro-literal    = %x21-24 / %x26-7E
                  ; visible characters except "%"
macro-letter     = "s" / "l" / "o" / "d" / "i" / "p" / "h" /
                  "c" / "r" / "t" / "v"
transformers     = *DIGIT [ "r" ]
delimiter        = "." / "-" / "+" / "," / "/" / "_" / "="

```

A literal "%" is expressed by "%%".

"%_" expands to a single " " space.

"%-" expands to a URL-encoded space, viz., "%20".

The following macro letters are expanded in term arguments:

```

s = <sender>
l = local-part of <sender>
o = domain of <sender>
d = <domain>
i = <ip>
p = the validated domain name of <ip> (deprecated)
v = the string "in-addr" if <ip> is ipv4, or "ip6" if <ip> is ipv6
h = HELO/EHLO domain

```

The following macro letters are allowed only in "exp" text:

```

c = SMTP client IP (easily readable format)
r = domain name of host performing the check
t = current timestamp

```


A '%' character not followed by a '{', '%', '-', or '_' character is a syntax error. So

```
-exists:%(ir).sbl.spamhaus.example.org
```

is incorrect and will cause `check_host()` to return a "permerror".

Instead, say

```
-exists:%{ir}.sbl.spamhaus.example.org
```

Optional transformers are the following:

*DIGIT = zero or more digits

'r' = reverse value, splitting on dots by default

If transformers or delimiters are provided, the replacement value for a macro letter is split into parts. After performing any reversal operation and/or removal of left-hand parts, the parts are rejoined using "." and not the original splitting characters.

By default, strings are split on "." (dots). Note that no special treatment is given to leading, trailing, or consecutive delimiters, and so the list of parts might contain empty strings. Older implementations of SPF prohibit trailing dots in domain names, so trailing dots should not be published by domain owners, although they must be accepted by implementations conforming to this document. Macros MAY specify delimiter characters that are used instead of ".".

The 'r' transformer indicates a reversal operation: if the client IP address were 192.0.2.1, the macro %{i} would expand to "192.0.2.1" and the macro %{ir} would expand to "1.2.0.192".

The DIGIT transformer indicates the number of right-hand parts to use, after optional reversal. If a DIGIT is specified, the value MUST be nonzero. If no DIGITs are specified, or if the value specifies more parts than are available, all the available parts are used. If the DIGIT was 5, and only 3 parts were available, the macro interpreter would pretend the DIGIT was 3. Implementations MUST support at least a value of 128, as that is the maximum number of labels in a domain name.

The "s" macro expands to the <sender> argument. It is an email address with a localpart, an "@" character, and a domain. The "l" macro expands to just the localpart. The "o" macro expands to just the domain part. Note that these values remain the same during recursive and chained evaluations due to "include" and/or "redirect". Note also that if the original <sender> had no localpart, the localpart was set to "postmaster" in initial processing (see [Section 4.3](#)).

For IPv4 addresses, both the "i" and "c" macros expand to the

standard dotted-quad format.

For IPv6 addresses, the "i" macro expands to a dot-format address; it is intended for use in `%{ir}`. The "c" macro MAY expand to any of the hexadecimal colon-format addresses specified in [[RFC4291](#)], [Section 2.2](#). It is intended for humans to read.

The "p" macro expands to the validated domain name of `<ip>`. The procedure for finding the validated domain name is defined in [Section 5.5](#). If the `<domain>` is present in the list of validated domains, it SHOULD be used. Otherwise, if a subdomain of the `<domain>` is present, it SHOULD be used. Otherwise, any name from the list MAY be used. If there are no validated domain names or if a DNS error occurs, the string "unknown" is used. This macro is deprecated and SHOULD NOT be used.

The "r" macro expands to the name of the receiving MTA. This SHOULD be a fully qualified domain name, but if one does not exist (as when the checking is done by a MUA) or if policy restrictions dictate otherwise, the word "unknown" SHOULD be substituted. The domain name can be different from the name found in the MX record that the client MTA used to locate the receiving MTA.

The "t" macro expands to the decimal representation of the approximate number of seconds since the Epoch (Midnight, January 1, 1970, UTC). This is the same value as is returned by the POSIX `time()` function in most standards-compliant libraries.

When the result of macro expansion is used in a domain name query, if the expanded domain name exceeds 253 characters (the maximum length of a domain name), the left side is truncated to fit, by removing successive domain labels until the total length does not exceed 253 characters.

Uppercased macros expand exactly as their lowercased equivalents, and are then URL escaped. URL escaping must be performed for characters not in the "unreserved" set, which is defined in [[RFC3986](#)].

Note: Care must be taken so that macro expansion for legitimate email does not exceed the 63-character limit on DNS labels. The localpart of email addresses, in particular, can have more than 63 characters between dots.

Note: Domains should avoid using the "s", "l", "o", or "h" macros in conjunction with any mechanism directive. Although these macros are powerful and allow per-user records to be published, they severely limit the ability of implementations to cache results of `check_host()` and they reduce the effectiveness of DNS caches.

Implementations should be aware that if no directive processed during the evaluation of `check_host()` contains an "s", "l", "o", or "h" macro, then the results of the evaluation can be cached on the basis of <domain> and <ip> alone for as long as the shortest Time To Live (TTL) of all the DNS records involved.

8.2. Expansion Examples

The <sender> is strong-bad@email.example.com.

The IPv4 SMTP client IP is 192.0.2.3.

The IPv6 SMTP client IP is 2001:DB8::CB01.

The PTR domain name of the client IP is mx.example.org.

macro	expansion
-------	-----------

-----	-----
-------	-------

%{s}	strong-bad@email.example.com
------	------------------------------

%{o}	email.example.com
------	-------------------

%{d}	email.example.com
------	-------------------

%{d4}	email.example.com
-------	-------------------

%{d3}	email.example.com
-------	-------------------

%{d2}	example.com
-------	-------------

%{d1}	com
-------	-----

%{dr}	com.example.email
-------	-------------------

%{d2r}	example.email
--------	---------------

%{l}	strong-bad
------	------------

%{l-}	strong.bad
-------	------------

%{lr}	strong-bad
-------	------------

%{lr-}	bad.strong
--------	------------

%{l1r-}	strong
---------	--------

macro-string	expansion
--------------	-----------

-----	-----
-------	-------

%{ir}.%{v}._spf.%{d2}	3.2.0.192.in-addr._spf.example.com
-----------------------	------------------------------------

%{lr-}.lp._spf.%{d2}	bad.strong.lp._spf.example.com
----------------------	--------------------------------

%{lr-}.lp.%{ir}.%{v}._spf.%{d2}	bad.strong.lp.3.2.0.192.in-addr._spf.example.com
---------------------------------	--

%{ir}.%{v}.%{l1r-}.lp._spf.%{d2}	3.2.0.192.in-addr.strong.lp._spf.example.com
----------------------------------	--

%{d2}.trusted-domains.example.net	example.com.trusted-domains.example.net
-----------------------------------	---

IPv6:

%{ir}.%{v}._spf.%{d2}	1.0.B.C.0.0.0.0.
-----------------------	------------------

0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.0.8.B.D.0.1.0.0.2.ip6._spf.example.com	
--	--

9. Implications

This section outlines the major implications that adoption of this document will have on various entities involved in Internet email. It is intended to make clear to the reader where this document knowingly affects the operation of such entities. This section is not a "how-to" manual, or a "best practices" document, and it is not a comprehensive list of what such entities should do in light of this document.

This section is non-normative. [[RFC5598](#)] describes the Internet email architecture. This section is organized based on the different segments of the architecture.

9.1. Sending Domains

Originating ADMDs (ADministrative Management Domains - [[RFC5598](#) [Section 2.2.1](#) and [Section 2.3](#)]) that wish to be compliant with this specification will need to determine the list of relays ([\[RFC5598\] Section 2.2.2](#)) that they allow to use their domain name in the "HELO" and "MAIL FROM" identities when relaying to other ADMDs. It is recognized that forming such a list is not just a simple technical exercise, but involves policy decisions with both technical and administrative considerations.

9.1.1. DNS Resource Considerations

Minimizing the DNS resources required for SPF lookups can be done by choosing directives that require less DNS information and placing lower-cost mechanisms earlier in the SPF record.

For example, consider a domain set up as follows:

```
example.com.      IN MX    10 mx.example.com.
mx.example.com.   IN A      192.0.2.1
a.example.com.    IN TXT    "v=spf1 mx:example.com -all"
b.example.com.    IN TXT    "v=spf1 a:mx.example.com -all"
c.example.com.    IN TXT    "v=spf1 ip4:192.0.2.1 -all"
```

Evaluating `check_host()` for the domain "a.example.com" requires the MX records for "example.com", and then the A records for the listed hosts. Evaluating for "b.example.com" requires only the A records. Evaluating for "c.example.com" requires none.

[Section 4.6.4](#) specifies the limits receivers have to use. It is essential to publish records that do not exceed these requirements. As an example, if you have more than 10 MX records, do not use the 'MX' mechanism to describe them in your SPF record.

9.1.2. Administrator's Considerations

There might be administrative considerations: using "a" over "ip4" or "ip6" allows hosts to be renumbered easily. Using "mx" over "a" allows the set of mail hosts to be changed easily. Unless such changes are common, it is better to use the less resource intensive mechanisms like "ip4" and "ip6".

Validating correct deployment is difficult. [RFC6652] describes one mechanism for soliciting feedback on SPF failures. Another approach that can be helpful to publish records that include a "tracking exists:" mechanism. By looking at the name server logs, a rough list can then be generated. For example:

```
v=spf1 exists:_h.%{h}._l.%{l}._o.%{o}._i.%{i}._spf.%{d} ?all
```

Regardless of the method used, understanding the ADMD's outbound mail architecture is essential to effective deployment.

9.2. Mediators

Broadly speaking, there are two types of mediating ADMDs that can affect SPF deployment of other ADMDs: mailing lists (see [RFC5598] [Section 5.3](#)) and ReSenders ([RFC5598] [Section 5.2](#)).

9.2.1. Mailing Lists

Mailing lists must be aware of how they re-inject mail that is sent to the list. Mailing lists MUST comply with the requirements in [RFC5321], [Section 3.10](#), and [RFC1123], [Section 5.3.6](#), that say that the reverse-path MUST be changed to be the mailbox of a person or other entity who administers the list. Whereas the reasons for changing the reverse-path are many and long-standing, SPF adds enforcement to this requirement.

In practice, almost all mailing list software in use already complies with this requirement. Mailing lists that do not comply might encounter problems depending on how access to the list is restricted. Such lists that are entirely internal to a domain (only people in the domain can send to or receive from the list) are not affected.

9.2.2. Forwarding Services and Aliases

Forwarding services take mail that is received at a mailbox and direct it to some external mailbox. At the time of this writing, the near-universal practice of such services is to use the original "MAIL FROM" of a message when re-injecting it for delivery to the external mailbox. [RFC1123] and [RFC5321] describe this action as an "alias"

rather than a "mail list". This means that the external mailbox's MTA sees all such mail in a connection from a host of the forwarding service, and so the "MAIL FROM" identity will not, in general, pass authorization.

There are three places that techniques can be used to ameliorate this problem.

1. The beginning, when email is first sent (Originating ADMDs).

1. "Neutral" results could be given for IP addresses that might be forwarders, instead of "fail" results. For example:

```
"v=spf1 mx -exists:%{ir}.sbl.spamhaus.example.org ?all"
```

This would cause a lookup on an anti-spam DNS blacklist (DNSBL) and cause a result of "fail" only for email coming from listed sources. All other email, including email sent through forwarders, would receive a "neutral" result. By checking the DNSBL after the known good sources, problems with incorrect listing on the DNSBL are greatly reduced.

2. The "MAIL FROM" identity could have additional information in the localpart that cryptographically identifies the mail as coming from an authorized source. In this case, such an SPF record could be used:

```
"v=spf1 mx exists:%{l}._spf_verify.%{d} -all"
```

Then, a specialized DNS server can be set up to serve the _spf_verify subdomain that validates the localpart. Although this requires an extra DNS lookup, this happens only when the email would otherwise be rejected as not coming from a known good source.

Note that due to the 63-character limit for domain labels, this approach only works reliably if the localpart signature scheme is guaranteed either to only produce localparts with a maximum of 63 characters or to gracefully handle truncated localparts.

3. Similarly, a specialized DNS server could be set up that will rate-limit the email coming from unexpected IP addresses.

```
"v=spf1 mx exists:%{ir}._spf_rate.%{d} -all"
```

4. SPF allows the creation of per-user policies for special cases. For example, the following SPF record and appropriate wildcard DNS records can be used:

`"v=spf1 mx redirect=%{l1r+}._at_.%{o}._spf.%{d}"`

2. The middle, when email is forwarded (Mediating ADMDs).
 1. Forwarding services can solve the problem by rewriting the "MAIL FROM" to be in their own domain. This means that mail rejected from the external mailbox will have to be forwarded back to the original sender by the forwarding service. Various schemes to do this exist though they vary widely in complexity and resource requirements on the part of the forwarding service.
 2. Several popular MTAs can be forced from "alias" semantics to "mailing list" semantics by configuring an additional alias with "owner-" prepended to the original alias name (e.g., an alias of "friends: george@example.com, fred@example.org" would need another alias of the form "owner-friends: localowner").
3. The end, when email is received (Receiving ADMDs).
 1. If the owner of the external mailbox wishes to trust the forwarding service, he can direct the external mailbox's MTA to skip SPF tests when the client host belongs to the forwarding service.
 2. Tests against other identities, such as the "HELO" identity, MAY be used to override a failed test against the "MAIL FROM" identity.
 3. For larger domains, it might not be possible to have a complete or accurate list of forwarding services used by the owners of the domain's mailboxes. In such cases, whitelists of generally-recognized forwarding services could be employed.

9.3. Mail Services

MSPs (Mail Service Providers - [\[RFC5598\] Section 2.3](#)) that offer mail services to third-party domains, such as sending of bulk mail, might want to adjust their setup in light of the authorization check described in this document. If the "MAIL FROM" identity used for such email uses the domain of the service provider, then the provider needs only to ensure that its sending host is authorized by its own SPF record, if any.

If the "MAIL FROM" identity does not use the mail service provider's domain, then extra care must be taken. The SPF record format has

several options for the third-party domain to authorize the service provider's MTAs to send mail on its behalf. For mail service providers, such as ISPs, that have a wide variety of customers using the same MTA, steps should be taken to prevent cross-customer forgery (see [Section 10.4](#)).

9.4. MTA Relays

Relays are described in [\[RFC5598\] Section 2.2.2](#). The authorization check generally precludes the use of arbitrary MTA relays between sender and receiver of an email message.

Within an organization, MTA relays can be effectively deployed. However, for purposes of this document, such relays are effectively transparent. The SPF authorization check is a check between border MTAs of different ADMDs.

For mail senders, this means that published SPF records must authorize any MTAs that actually send across the Internet. Usually, these are just the border MTAs as internal MTAs simply forward mail to these MTAs for delivery.

The receiving ADMD will generally want to perform the authorization check at the boundary MTAs, specifically including all secondary MXs. This allows mail that fails to be rejected during the SMTP session rather than sending a failure report. Internal MTAs then do not perform the authorization test. To perform the authorization test other than at the boundary, the host that first transferred the message to the organization must be determined, which can be difficult to extract from the message header. Testing other than at the boundary is likely to produce unreliable results.

10. Security Considerations

10.1. Processing Limits

As with most aspects of email, there are a number of ways that malicious parties could use the protocol as an avenue for a Denial-of-Service (DoS) attack. The processing limits outlined in [Section 4.6.4](#) are designed to prevent attacks such as the following:

- o A malicious party could create an SPF record with many references to a victim's domain and send many emails to different SPF clients; those SPF clients would then create a DoS attack. In effect, the SPF clients are being used to amplify the attacker's bandwidth by using fewer bytes in the SMTP session than are used by the DNS queries. Using SPF clients also allows the attacker to hide the true source of the attack.
- o Whereas implementations of `check_host()` are supposed to limit the number of DNS lookups, malicious domains could publish records that exceed these limits in an attempt to waste computation effort at their targets when they send them mail. Malicious domains could also design SPF records that cause particular implementations to use excessive memory or CPU usage, or to trigger bugs.
- o Malicious parties could send a large volume of mail purporting to come from the intended target to a wide variety of legitimate mail hosts. These legitimate machines would then present a DNS load on the target as they fetched the relevant records.

Of these, the case of a third party referenced in the SPF record is the easiest for a DoS attack to effectively exploit. As a result, limits that might seem reasonable for an individual mail server can still allow an unreasonable amount of bandwidth amplification. Therefore, the processing limits need to be quite low.

MTAs or other processors SHOULD impose a limit on the maximum amount of elapsed time to evaluate `check_host()`. Such a limit SHOULD allow at least 20 seconds. If such a limit is exceeded, the result of authorization SHOULD be "temperror".

10.2. SPF-Authorized Email May Contain Other False Identities

The "MAIL FROM" and "HELO" identity authorizations must not be construed to provide more assurance than they do. It is entirely possible for a malicious sender to inject a message using his own domain in the identities used by SPF, to have that domain's SPF record authorize the sending host, and yet the message can easily

list other identities in its header. Unless the user or the MUA takes care to note that the authorized identity does not match the other more commonly-presented identities (such as the From: header field), the user might be lulled into a false sense of security.

10.3. Spoofed DNS and IP Data

There are two aspects of this protocol that malicious parties could exploit to undermine the validity of the `check_host()` function:

- o The evaluation of `check_host()` relies heavily on DNS. A malicious attacker could attack the DNS infrastructure and cause `check_host()` to see spoofed DNS data, and then return incorrect results. This could include returning "pass" for an <ip> value where the actual domain's record would evaluate to "fail". See [\[RFC3833\]](#) for a description of DNS weaknesses.
- o The client IP address, <ip>, is assumed to be correct. A malicious attacker could spoof TCP sequence numbers to make mail appear to come from a permitted host for a domain that the attacker is impersonating.

10.4. Cross-User Forgery

By definition, SPF policies just map domain names to sets of authorized MTAs, not whole email addresses to sets of authorized users. Although the "l" macro ([Section 8](#)) provides a limited way to define individual sets of authorized MTAs for specific email addresses, it is generally impossible to verify, through SPF, the use of specific email addresses by individual users of the same MTA.

It is up to mail services and their MTAs to directly prevent cross-user forgery: based on SMTP AUTH ([\[RFC4954\]](#)), users should be restricted to using only those email addresses that are actually under their control (see [\[RFC6409, Section 6.1\]](#)). Another means to verify the identity of individual users is message cryptography such as PGP ([\[RFC4880\]](#)) or S/MIME ([\[RFC5751\]](#)).

10.5. Untrusted Information Sources

An SPF compliant receiver gathers information from the SMTP commands it receives and from the published DNS records of the sending domain holder, (e.g., "HELO" domain name, the "MAIL FROM" address from the envelope, and SPF DNS records published by the domain holder).

This information, passed to the receiver in the Received-SPF: trace fields, may be returned to the client MTA as an SMTP rejection message. If such an SMTP rejection message is generated, the

information from the trace fields must be checked for such problems as invalid characters and excessively long lines.

When the authorization check fails, an explanation string could be included in the reject response. Both the sender and the rejecting receiver need to be aware that the explanation was determined by the publisher of the SPF record checked and, in general, not the receiver. The explanation can contain malicious URLs, or it might be offensive or misleading.

Explanations returned to sender domains due to exp modifiers, ([Section 6.2](#)), were generated by the sender policy published by the domain holders themselves. As long as messages are only returned with non-delivery notification to domains publishing the explanation strings from their own DNS SPF records, the only affected parties are the original publishers of the domain's SPF records.

In practice, such non-delivery notifications can be misdirected, such as when an MTA accepts an email and only later generates the notification to a forged address, or when an email forwarder does not direct the bounce back to the original sender.

[10.6](#). Privacy Exposure

Checking SPF records causes DNS queries to be sent to the domain owner. These DNS queries, especially if they are caused by the "exists" mechanism, can contain information about who is sending email and likely to which MTA the email is being sent. This can introduce some privacy concerns, which are more or less of an issue depending on local laws and the relationship between the domain owner and the person sending the email.

11. Contributors and Acknowledgements

This document is largely based on the work of Meng Weng Wong, Mark Lentczner, and Wayne Schlitt. Although, as this section acknowledges, many people have contributed to this document, a very large portion of the writing and editing are due to Meng, Mark, and Wayne.

This design owes a debt of parentage to [\[RMX\]](#) by Hadmut Danisch and to [\[DMP\]](#) by Gordon Fecyk. The idea of using a DNS record to check the legitimacy of an email address traces its ancestry further back through messages on the namedroppers mailing list by Paul Vixie [\[Vixie\]](#) (based on suggestion by Jim Miller) and by David Green [\[Green\]](#).

Philip Gladstone contributed the concept of macros to the specification, multiplying the expressiveness of the language and making per-user and per-IP lookups possible.

The authors of both this document and [\[RFC4408\]](#) would also like to thank the literally hundreds of individuals who have participated in the development of this design. They are far too numerous to name, but they include the following:

- The participants in the SPFbis working group.
- The folks on the spf-discuss mailing list.
- The folks on the SPAM-L mailing list.
- The folks on the IRTF ASRG mailing list.
- The folks on the IETF MARID mailing list.
- The folks on #perl.

12. IANA Considerations

12.1. The SPF DNS Record Type

Per [[RFC4408](#)], the IANA assigned the Resource Record Type and Qtype from the DNS Parameters Registry for the SPF RR type with code 99. The format of this type is identical to the TXT RR [[RFC1035](#)]. The character content of the record is encoded as [[US-ASCII](#)]. Use of this record type is obsolete for SPF Version 1.

IANA is requested to add (to be changed to "has added" upon publication) an annotation to the SPF RRTYPE saying "(OBSOLETE - use TXT)" in the DNS Parameters registry.

12.2. The Received-SPF Mail Header Field

Per [[RFC3864](#)], the "Received-SPF:" header field is added to the IANA Permanent Message Header Field Registry. The following is the registration template:

Header field name: Received-SPF
Applicable protocol: mail ([[RFC5322](#)])
Status: Standards Track
Author/Change controller: IETF
Specification document(s): [[RFC4408](#)], RFC XXXX

13. References

13.1. Normative References

- [RFC1035] Mockapetris, P., "Domain names - implementation and specification", STD 13, [RFC 1035](#), November 1987.
- [RFC1123] Braden, R., "Requirements for Internet Hosts - Application and Support", STD 3, [RFC 1123](#), October 1989.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3463] Vaudreuil, G., "Enhanced Mail System Status Codes", [RFC 3463](#), January 2003.
- [RFC3864] Klyne, G., Nottingham, M., and J. Mogul, "Registration Procedures for Message Header Fields", [BCP 90](#), [RFC 3864](#), September 2004.
- [RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, [RFC 3986](#), January 2005.
- [RFC4291] Hinden, R. and S. Deering, "IP Version 6 Addressing Architecture", [RFC 4291](#), February 2006.
- [RFC5234] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, [RFC 5234](#), January 2008.
- [RFC5321] Klensin, J., "Simple Mail Transfer Protocol", [RFC 5321](#), October 2008.
- [RFC5322] Resnick, P., Ed., "Internet Message Format", [RFC 5322](#), October 2008.
- [US-ASCII]
American National Standards Institute (formerly United States of America Standards Institute), "USA Code for Information Interchange, X3.4", 1968.

ANSI X3.4-1968 has been replaced by newer versions with slight modifications, but the 1968 version remains definitive for the Internet.

13.2. Informative References

- [DMP] Fecyk, G., "Designated Mailers Protocol".
Work In Progress
- [Green] Green, D., "Domain-Authorized SMTP Mail", 2002.
- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, [RFC 1034](#), November 1987.
- [RFC1983] Malkin, G., "Internet Users' Glossary", [RFC 1983](#), August 1996.
- [RFC2308] Andrews, M., "Negative Caching of DNS Queries (DNS NCACHE)", [RFC 2308](#), March 1998.
- [RFC3696] Klensin, J., "Application Techniques for Checking and Transformation of Names", [RFC 3696](#), February 2004.
- [RFC3833] Atkins, D. and R. Austein, "Threat Analysis of the Domain Name System (DNS)", [RFC 3833](#), August 2004.
- [RFC4408] Wong, M. and W. Schlitt, "Sender Policy Framework (SPF) for Authorizing Use of Domains in E-Mail, Version 1", [RFC 4408](#), April 2006.
- [RFC4880] Callas, J., Donnerhacke, L., Finney, H., Shaw, D., and R. Thayer, "OpenPGP Message Format", [RFC 4880](#), November 2007.
- [RFC4954] Siemborski, R. and A. Melnikov, "SMTP Service Extension for Authentication", [RFC 4954](#), July 2007.
- [RFC5598] Crocker, D., "Internet Mail Architecture", [RFC 5598](#), July 2009.
- [RFC5751] Ramsdell, B. and S. Turner, "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2 Message Specification", [RFC 5751](#), January 2010.
- [RFC6409] Gellens, R. and J. Klensin, "Message Submission for Mail", STD 72, [RFC 6409](#), November 2011.
- [RFC6647] Kucherawy, M. and D. Crocker, "Email Greylisting: An Applicability Statement for SMTP", [RFC 6647](#), June 2012.
- [RFC6652] Kitterman, S., "Sender Policy Framework (SPF) Authentication Failure Reporting Using the Abuse Reporting

Format", [RFC 6652](#), June 2012.

[RMX] Danisch, H., "The RMX DNS RR Type for light weight sender authentication".

Work In Progress

[Vixie] Vixie, P., "Repudiating MAIL FROM", 2002.

[[draft-ietf-spfbis-experiment](#)]

Kucherawy, M., "Resolution of The SPF and Sender ID Experiments", [draft-ietf-spfbis-experiment](#) (work in progress), 2012.

Appendix A. Collected ABNF

This section is normative and any discrepancies with the ABNF fragments in the preceding text are to be resolved in favor of this grammar.

See [\[RFC5234\]](#) for ABNF notation. Please note that as per this ABNF definition, literal text strings (those in quotes) are case-insensitive. Hence, "mx" matches "mx", "MX", "mX", and "Mx".

```

record          = version terms *SP
version         = "v=spf1"

terms           = *( 1*SP ( directive / modifier ) )

directive       = [ qualifier ] mechanism
qualifier       = "+" / "-" / "?" / "~"
mechanism       = ( all / include
                  / A / MX / PTR / IP4 / IP6 / exists )

all             = "all"
include         = "include" ":" domain-spec
A               = "a"      [ ":" domain-spec ] [ dual-cidr-length ]
MX              = "mx"     [ ":" domain-spec ] [ dual-cidr-length ]
PTR             = "ptr"    [ ":" domain-spec ]
IP4             = "ip4"    ":" ip4-network   [ ip4-cidr-length ]
IP6             = "ip6"    ":" ip6-network   [ ip6-cidr-length ]
exists          = "exists" ":" domain-spec

modifier        = redirect / explanation / unknown-modifier
redirect        = "redirect" "=" domain-spec
explanation      = "exp" "=" domain-spec
unknown-modifier = name "=" macro-string
                  ; where name is not any known modifier

ip4-cidr-length = "/" 1*DIGIT
ip6-cidr-length = "/" 1*DIGIT
dual-cidr-length = [ ip4-cidr-length ] [ "/" ip6-cidr-length ]

ip4-network     = qnum "." qnum "." qnum "." qnum
qnum            = DIGIT ; 0-9
                  / %x31-39 DIGIT ; 10-99
                  / "1" 2DIGIT ; 100-199
                  / "2" %x30-34 DIGIT ; 200-249
                  / "25" %x30-35 ; 250-255
                  ; conventional dotted quad notation. e.g., 192.0.2.0
ip6-network     = <as per \[RFC 4291\], section 2.2>
                  ; e.g., 2001:DB8::CD30

```



```

domain-spec      = macro-string domain-end
domain-end      = ( "." toplevel [ "." ] ) / macro-expand

toplevel        = ( *alphanum ALPHA *alphanum ) /
                  ( 1*alphanum "-" *( alphanum / "-" ) alphanum )
                  ; LDH rule plus additional TLD restrictions
                  ; (see \[RFC3696\], Section 2 for background)

alphanum        = ALPHA / DIGIT

explain-string   = *( macro-string / SP )

macro-string     = *( macro-expand / macro-literal )
macro-expand     = ( "%{" macro-letter transformers *delimiter "}" )
                  / "%%" / "%_" / "%-"

macro-literal    = %x21-24 / %x26-7E
                  ; visible characters except "%"

macro-letter     = "s" / "l" / "o" / "d" / "i" / "p" / "h" /
                  "c" / "r" / "t" / "v"

transformers     = *DIGIT [ "r" ]
delimiter        = "." / "-" / "+" / "," / "/" / "_" / "="

name            = ALPHA *( ALPHA / DIGIT / "-" / "_" / "." )

header-field     = "Received-SPF:" [CFWS] result FWS [comment FWS]
                  [ key-value-list ] CRLF

result          = "pass" / "fail" / "softfail" / "neutral" /
                  "none" / "temperror" / "permerror"

key-value-list   = key-value-pair *( ";" [CFWS] key-value-pair )
                  [ ";" ]

key-value-pair   = key [CFWS] "=" ( dot-atom / quoted-string )

key             = "client-ip" / "envelope-from" / "helo" /
                  "problem" / "receiver" / identity /
                  mechanism / name

identity         = "mailfrom"      ; for the "MAIL FROM" identity
                  / "helo"         ; for the "HELO" identity
                  / name           ; other identities

ALPHA            = <A-Z / a-z as per \[RFC5234\]>
DIGIT            = <0-9 as per \[RFC5234\]>
SP              = <space character as per \[RFC5234\]>
dot-atom        = <unquoted word as per \[RFC5322\]>
quoted-string    = <quoted string as per \[RFC5322\]>
comment         = <comment string as per \[RFC5322\]>

```


CFWS = <comment or folding white space as per [[RFC5322](#)]>
FWS = <folding white space as per [[RFC5322](#)]>
CRLF = <standard end-of-line token as per [[RFC5322](#)]>

[Appendix B](#). Extended Examples

These examples are based on the following DNS setup:

```
; A domain with two mail servers, two hosts
; and two servers at the domain name
$ORIGIN example.com.
```

```
@           MX  10 mail-a
            MX  20 mail-b
            A   192.0.2.10
            A   192.0.2.11
amy         A   192.0.2.65
bob         A   192.0.2.66
mail-a     A   192.0.2.129
mail-b     A   192.0.2.130
www        CNAME example.com.
```

```
; A related domain
$ORIGIN example.org.
```

```
@           MX  10 mail-c
mail-c     A   192.0.2.140
```

```
; The reverse IP for those addresses
```

```
$ORIGIN 2.0.192.in-addr.arpa.
10        PTR example.com.
11        PTR example.com.
65        PTR amy.example.com.
66        PTR bob.example.com.
129       PTR mail-a.example.com.
130       PTR mail-b.example.com.
140       PTR mail-c.example.org.
```

```
; A rogue reverse IP domain that claims to be
```

```
; something it's not
```

```
$ORIGIN 0.0.10.in-addr.arpa.
4         PTR bob.example.com.
```

[B.1](#). Simple Examples

These examples show various possible published records for example.com and which values if <ip> would cause check_host() to return "pass". Note that <domain> is "example.com".

```
v=spf1 +all
```



```
-- any <ip> passes

v=spf1 a -all
  -- hosts 192.0.2.10 and 192.0.2.11 pass

v=spf1 a:example.org -all
  -- no sending hosts pass since example.org has no A records

v=spf1 mx -all
  -- sending hosts 192.0.2.129 and 192.0.2.130 pass

v=spf1 mx:example.org -all
  -- sending host 192.0.2.140 passes

v=spf1 mx mx:example.org -all
  -- sending hosts 192.0.2.129, 192.0.2.130, and 192.0.2.140 pass

v=spf1 mx/30 mx:example.org/30 -all
  -- any sending host in 192.0.2.128/30 or 192.0.2.140/30 passes

v=spf1 ptr -all
  -- sending host 192.0.2.65 passes (reverse DNS is valid and is in
  -- example.com)
  -- sending host 192.0.2.140 fails (reverse DNS is valid, but not
  -- in example.com)
  -- sending host 10.0.0.4 fails (reverse IP is not valid)

v=spf1 ip4:192.0.2.128/28 -all
  -- sending host 192.0.2.65 fails
  -- sending host 192.0.2.129 passes
```

B.2. Multiple Domain Example

These examples show the effect of related records:

```
example.org: "v=spf1 include:example.com include:example.net -all"
```

This record would be used if mail from example.org actually came through servers at example.com and example.net. Example.org's designated servers are the union of example.com's and example.net's designated servers.

```
la.example.org: "v=spf1 redirect=example.org"
ny.example.org: "v=spf1 redirect=example.org"
sf.example.org: "v=spf1 redirect=example.org"
```

These records allow a set of domains that all use the same mail system to make use of that mail system's record. In this way, only

the mail system's record needs to be updated when the mail setup changes. These domains' records never have to change.

B.3. DNSBL Style Example

Imagine that, in addition to the domain records listed above, there are these:

```
$ORIGIN _spf.example.com.
mary.mobile-users          A 127.0.0.2
fred.mobile-users          A 127.0.0.2
15.15.168.192.joel.remote-users  A 127.0.0.2
16.15.168.192.joel.remote-users  A 127.0.0.2
```

The following records describe users at example.com who mail from arbitrary servers, or who mail from personal servers.

example.com:

```
v=spf1 mx
      include:mobile-users._spf.{d}
      include:remote-users._spf.{d}
      -all
```

mobile-users._spf.example.com:

```
v=spf1 exists:%{l1r+}.{d}
```

remote-users._spf.example.com:

```
v=spf1 exists:%{ir}.{d}
```

B.4. Multiple Requirements Example

Say that your sender policy requires both that the IP address is within a certain range and that the reverse DNS for the IP matches. This can be done several ways, including the following:

```
example.com.          SPF ( "v=spf1 "
                          "-include:ip4._spf.{d} "
                          "-include:ptr._spf.{d} "
                          "+all" )
ip4._spf.example.com.  SPF "v=spf1 -ip4:192.0.2.0/24 +all"
ptr._spf.example.com.  SPF "v=spf1 -ptr +all"
```

This example shows how the "-include" mechanism can be useful, how an SPF record that ends in "+all" can be very restrictive, and the use of De Morgan's Law.

Appendix C. Change History

Changes since [RFC 4408](#) (to be removed prior to publication)

Moved to standards track

Authors updated

IESG Note regarding experimental use replaced with discussion of results

Process errata:

Add %v macro to ABNF grammar

Replace "uric" by "unreserved"

Recommend an SMTP reply code for optional permerror rejections

Correct syntax in Received-SPF examples

Fix unknown-modifier clause is too greedy in ABNF

Correct use of empty domain-spec on exp modifier

Fix minor typo errata

Convert to spfbis working group draft,
[draft-ietf-spfbis-4408bis-00](#)

Addressed Ticket #1, [RFC 4408 Section 2.5.6](#) - Temporary errors by giving the option to turn repeated SERVFAIL into permerror and adding [RFC 2308](#) reference.

Clarified text about IPv4 mapped addresses to resolve test suite ambiguity

Clarified ambiguity about result when more than 10 "mx" or "ptr" records are returned for lookup to specify permerror. This resolves one of the test suite ambiguities

Made all references to result codes lower case per issue #7

Adjusted [section 2.2](#) Requirement to check mail from per issue #15

Added missing "v" element in macro-letter in the collected ABNF per issue #16 - [section 8.1](#) was already fixed in the pre-WG draft

Marked ptr and "p" macro deprecated/SHOULD NOT use per issue #27

Expunged lower case may from the draft per issue #8

Expunged "x-" name as an obsolete concept

Updated obsolete references: [RFC2821](#) to [RFC5321](#), [RFC2822](#) to [RFC5322](#), and [RFC4234](#) to [RFC5234](#)

Refer to [RFC6647](#) to describe greylisting instead of trying to describe it directly.

Updated informative references to the current versions.

Added definition for deprecated since there are questions.

Start to rework [section 9](#) with some [RFC5598](#) terms.

Added mention of [RFC 6552](#) feedback reports in [section 9](#).

Added [draft-ietf-spfbis-experiment](#) as an informational reference.

Drop Type SPF.

Try and clarify informational nature of [RFC3696](#)

Fix ABNF nits and add missing definitions per Bill's ABNF checker.

Make DNS lookup time limit SHOULD instead of MAY.

Reorganize and clarify processing limits. Move hard limits to new [section 4.6.4](#), Evaluation Limits. Move advice to non-normative [section 9](#).

Removed paragraph in [section 10.1](#) about limiting total data volumes as it is unused (and removable per the charter) and serves no purpose (it isn't something that actually can be implemented in any reasonable way).

Author's Address

Scott Kitterman
Agari
3611 Scheel Dr
Ellicott City, MD 21042
United States of America

Email: scott@kitterman.com