

Mail Abuse Reporting Format WG
Internet-Draft
Intended status: Standards Track
Expires: January 28, 2012

J. Falk
Return Path
July 27, 2011

A DNS TXT Record for Advertising and Discovering Willingness
to Provide or Receive ARF Reports
draft-ietf-marf-reporting-discovery-01

Abstract

This document defines a method for network operators to advertise their willingness to send feedback about received email to other parties, and for those other parties to advertise their willingness to receive such feedback.

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 28, 2012.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

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

Table of Contents

1. Introduction	3
2. Purpose	3
3. Requirements	4
4. Language	4
4.1. General	5
4.2. Email Specific	5
4.3. ARF Specific	5
5. Characteristics of a Feedback Reporting Advertisement	5
5.1. Feedback Consumers	6
5.1.1. Feedback Consumer Policies	6
5.1.2. Feedback Consumer re= Default	6
5.2. Feedback Generators	7
5.2.1. Feedback Generator Policies	7
5.3. Combining Generator and Consumer Tags in the Same Record	8
5.4. Note about URIs	8
5.5. Formal Definition	8
6. Example Records for Various Use Cases	8
6.1. Example Feedback Consumer Records	8
6.2. Example Feedback Generator Records	9
7. Authentication of Reported Message	10
7.1. DKIM signatures	10
7.2. SPF authorized sender	10
8. IANA Considerations	10
9. Security Considerations	10
9.1. Inherited from MARF-BASE	10
9.2. These Need Fleshing Out	10
9.3. Privacy considerations	11
10. Acknowledgements	11
11. Contributors	11
12. References	12
12.1. Normative References	12
12.2. Informative References	12
Appendix A. Public Discussion and Support	13
Appendix B. Document History & Open Issues	13
B.1. draft-jdfalk-marf-reporting-discovery-00	13
B.2. draft-jdfalk-marf-reporting-discovery-01	13
B.3. draft-jdfalk-marf-reporting-discovery-02	13
B.4. draft-jdfalk-marf-reporting-discovery-03	14
B.5. draft-ietf-marf-reporting-discovery-00	14
B.6. draft-ietf-marf-reporting-discovery-01	14
Author's Address	14

1. Introduction

As the spam problem continues to expand and potential solutions evolve, network operators are increasingly exchanging abuse reports among themselves and other parties. While [MARF-BASE] defines the Abuse Reporting Format (ARF) for these reports, it assumes that the operators will use some undefined method to discover each other and enter into any necessary agreements.

The advertisement method defined in this memo is intended to ease the process for potential ARF recipients to discover whether a particular Administrative Management Domain (ADMD) has the facility and willingness to generate ARF reports, and for ARF generators to discover whether a particular ADMD is able and willing (and authorized) to receive ARF reports.

While written primarily for initial discovery and configuration of feedback relationships, it is expected that these advertisements will also be useful for updating participants when parameters have changed.

Further, while this document only defines a DNS TXT record to contain these advertisements, other methods may be defined in the future.

This document only defines the process for advertisement and discovery of feedback recipients. Determination of when it is appropriate to send feedback or how trust may be established between report generators and report consumers is outside the scope of this document. It is assumed that best practices will continue to evolve over time, and will be codified in future documents.

Similarly, nothing in this draft is intended to preclude other methods that a Feedback Generator might use to determine where to send a report. For example, if the report were to be keyed off of the domain of an email address or a URL inside of the body of a message this discovery mechanism may be inappropriate.

2. Purpose

The reports defined in [MARF-BASE] are intended to inform mail operators about:

- o email abuse originating from their networks;
- o potential issues with the perceived quality of outbound mail, such as email service providers sending mail that attracts the attention of automated abuse detection systems.

To support these and other related purposes, this document addresses two primary use cases:

- o Any ADMD may advertise its willingness to receive reports from the internet at large, given particular criteria included in or referenced by the advertisement;
- o Any ADMD may advertise their willingness to provide reports to the Internet at large, given particular criteria included in or referenced by the advertisement;

Further, an ADMD which is generating reports may query the advertisement of an ADMD that wishes to receive reports, in order to confirm that an out-of-band request to receive reports is legitimate and/or to determine if any of the characteristics of the request have changed.

This specification inherits from [MARF-BASE] that it is intended specifically for communications among providers regarding email abuse and related issues, and SHOULD NOT be used for other reports unless those feedback-types specifically mention this document. For example, the [DKIM-REPORTING] extension includes its own ARF recipient discovery method that should not be confused with the method defined in this memo.

3. Requirements

The advertisement and discovery process must be easily accessible to the software involved in providing email service, preferably using concepts and technologies an email operator can be assumed to be familiar with. Thus, following the examples of [DKIM] and [DKIM-REPORTING], the advertisement is in the form of a [DNS] TXT record. While this may provide challenges for offline processing, this is outweighed by the advantages of security and maintainability.

In order to reflect current usage, advertisements must also provide the ability to reference complex "terms of service" or other documents outside of the scope of a simple discovery method. This is accomplished through the inclusion of a URI.

And finally, the advertisement must be readable by humans (assuming they have access to this RFC) as well as software specifically written for the purpose.

4. Language

This section defines various terms used throughout this document.

4.1. General

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

4.2. Email Specific

[EMAIL-ARCH] introduces several terms and concepts that are used in this memo, and thus readers are advised to become familiar with it as well.

4.3. ARF Specific

[MARF-BASE] introduces terms and concepts that are necessary for a full understanding of this memo, and thus readers are advised to read it before continuing.

5. Characteristics of a Feedback Reporting Advertisement

An advertisement of willingness to generate or receive feedback is accomplished by publishing a TXT record in the [DNS] using the name '_report' within the given DNS domain.

This record will contain a sequence of "tag=value" pairs, separated by semicolons. The tags and possible values are described in the next two sections, followed by the precise ABNF grammar.

In the case of a feedback consumer, the advertisement should be published in the DNS domain matching the [DKIM] 'd=' value used on outgoing signatures, and/or in the DNS domain matching the one present in the [SMTP] MAIL commands it issues when sending mail, and/or in the DNS domain referenced by the DNS PTR record (sometimes called "reverse DNS") of the IP address of the border MTA used to transfer the message.

In the case of a feedback generator, to inquire whether or not an ADMD wishes to receive feedback reports, the DNS domain to which the report should be sent is determined (using the [DKIM] d= string) and then a TXT record query to the above name is issued. For example, if a report generator wishes to generate a report about a message bearing DNS domain 'example.com', the generator would issue a TXT record query for '_report.example.com'.

Feedback generators SHOULD NOT send reports to ADMDs that are not in any way responsible for the reported message, for both security and efficiency reasons. Responsibility can be ascertained as described

in Section 7, and by applying local policy.

In the case of a feedback generator who wishes to advertise that reports are available, the TXT record is placed by the DNS domain at which they receive mail. For example, to advertise reports regarding mail received at example.net, the advertisement is placed in _report.example.net.

5.1. Feedback Consumers

A "Feedback Consumer" is an entity which wishes to receive feedback. In most cases, the Feedback Consumer will be within the same ADMD as the identifier (domain name or IP address) used to determine the source of a message.

The following tags are defined for Feedback Consumers:

- r the address to which reports should be sent. Required; there is no default. This address MUST be able to respond to an emailed subscription verification request; see Security Considerations below.
- rf the format of the report requested; currently only "ARF" ([MARF-BASE]) is supported. Optional; defaults to ARF.
- ri requested report interval; may not be supported by all implementors. Optional; if omitted, all reports may be sent.
- rt colon-separated list of ARF ([MARF-BASE]) feedback types for which reports are requested. Optional; if omitted, all report types may be sent.
- re email address of a person or role account responsible for handling any issues related to receiving reports. Optional, but SHOULD be defined; defaults to abuse@ the DNS domain.
- rp stated policy, as listed below. Optional; defaults to "o".
- ru URI for additional contact information. Optional, but SHOULD be defined; there is no default value.

5.1.1. Feedback Consumer Policies

Policies are listed in the "rp" tag, described above.

- o open to reports from all sources. This is the default.
- c closed; no reports are requested. This option is intended for testing purposes, or for feedback arrangements which have been set up using methods outside of the scope of this document.

5.1.2. Feedback Consumer re= Default

As described above, the default re= value is abuse@ the DNS domain. Report Generators are cautioned that re= is the address of a responsible person; this address may not be equipped to parse ARF

reports. Instead, ARF reports MUST only be sent to the address in the r= value.

Feedback Consumers are similarly cautioned that deciding not to publish an advertisement in accordance with this specification does not guarantee that all Report Generators won't send unsolicited ARF reports to your abuse@ or postmaster@ address. Thus, all domain owners are encouraged to publish explicit information even in case it happens to agree with the default values: in addition to better clarity, it may also improve caching.

5.2. Feedback Generators

A "Feedback Generator" is an entity which generates feedback reports. Often, the Feedback Generator is within the same ADMD as the mail server which received the message.

The following tags are defined for Feedback Generators:

gf the format of reports offered; currently only "ARF" ([MARF-BASE]) is supported. Optional; defaults to "ARF".
gt colon-separated list of ARF ([MARF-BASE]) feedback types for which reports are available. (Optional; if omitted, any report types may be generated.)
ge email address of a person (or role account) responsible for handling any issues related to receiving reports. Optional, but SHOULD be defined; defaults to postmaster@ the DNS domain.
gp stated policy, as listed below. Optional; defaults to "o".
gu URI for additional information. This field SHOULD be defined for a policy of "o" or "c", and MUST be defined when the policy is "r". Otherwise, the field is optional; there is no default.

5.2.1. Feedback Generator Policies

Policies are listed in the "gp" tag, described above.

- o open to providing reports to any consumer. This option is the default policy.
- r open to providing reports only after the prospective consumer has completed an application process, which may be found at the URI defined by the "gu" tag above.
- c closed; no reports are available. This option is intended for testing purposes, or for feedback arrangements which have been set up using methods outside of the scope of this document.

5.3. Combining Generator and Consumer Tags in the Same Record

It is common for a Feedback Generator to also act as a Feedback Consumer. When this happens, they MAY include both types of tags in the same TXT record.

When parsing these records, implementors MUST accept both types of tags in the same record, and MUST NOT expect the tags to appear in any particular order. Implementors MUST ignore any unfamiliar tags or other unexpected text.

5.4. Note about URIs

While this memo assumes that advertisements will contain `http://` or similar URIs, implementors should be aware that the URI-related fields can carry many different types of data depending on the URI scheme used. For more information, please consult the URI Schemes registry maintained by IANA.

5.5. Formal Definition

The formal definition using [ABNF] is TBD.

6. Example Records for Various Use Cases

While (in the author's mind) these examples address many of the most common use cases, implementers MUST NOT assume that only these configurations will ever be seen on the real live internet.

6.1. Example Feedback Consumer Records

Perhaps the most common scenario today is where a sender of bulk commercial email wishes to receive any complaints about messages originating from the servers under their control. These servers are identified by IP address or range. The PTR record for one of the IP addresses in question is "outmail5.example.com", so the advertisement could be:

```
_report.outmail5.example.com TXT "r=complaints@example.com; rf=ARF;
rt=abuse,fraud,virus,other; re=isprelations@example.com;"
```

This would need to be repeated for each PTR record, perhaps outmail11.example.com through outmail7.example.com. (NOTE: per [DNS], you cannot have both PTR and TXT associated with the same DNS record. Instead, the TXT record is associated with the `_destination_` of the PTR pointer.)

Or, if they sign all of their outbound mail with a [DKIM] d= string of "example.com", they could place the same advertisement in a TXT record at _report.example.com.

```
example.com TXT "r=complaints@example.com; rf=ARF;
rt=abuse,fraud,virus,other; re=isprelations@example.com;"
```

Another common scenario is of a mailbox or access provider who wishes to receive complaints about mail sent by their users. That record would be very similar, and would again be applied to PTR and/or d= domains:

```
example.com TXT "r=abuse@example.net; rf=ARF; rt=abuse,fraud,virus;
re=postmaster@example.net;"
```

6.2. Example Feedback Generator Records

The most common scenario today is where a large mailbox provider offers feedback to qualified Feedback Consumers who have filled out an application form. While the full application process cannot be adequately represented inside of a single DNS record, the "gp" and "gu" fields permit advertisement of this policy:

```
example.net TXT "gf=ARF; gt=abuse; ge=postmaster@example.net; gp=r;
gu=http://postmaster.example.net/fbl/"
```

However, it is rare for a Feedback Generator -- most often a mailbox provider or similar service -- to not also be a Feedback Consumer. If they wish to place an advertisement using the PTR method, there is usually no conflict. However, if their outbound mail is signed with [DKIM] and the d= value is the same domain at which they receive mail, then the advertisement will need to contain both generator and consumer tags:

```
example.net TXT "gf=ARF; gt=abuse; ge=postmaster@example.net; gp=r;
gu=http://postmaster.example.net/fbl/"; rf=ARF;
r=abuse+arf@example.net; rt=abuse,fraud,other;
re=postmaster@example.net;"
```

Or, they MAY split the Generator and Consumer tags into separate TXT records under the same domain.

Were a Feedback Generator to only offer authentication results, they could advertise thusly:

```
"gf=ARF; gt=auth; ge=postmaster@example.net; gp=o;
gu=http://postmaster.example.net/auth/"
```

7. Authentication of Reported Message

[AUTH-METHODS] and its extensions define a number of authentication methods that allow the recipient of a message to determine which ADMD is responsible for the message. A Report Generator MUST utilize one or more of these methods to ensure that reports are being sent to a Report Consumer within the correct ADMD, and furthermore it is RECOMMENDED that the Report Consumer utilize these same methods to ensure that a message reported to them was indeed sent by who they think it was sent by. For example:

7.1. DKIM signatures

Any valid DKIM signature in the reported message provides a domain name, either the signature "d" tag, or the domain-part of the signature "i" tag.

7.2. SPF authorized sender

When the `check_host()` function described in [SPF] results in a "pass" for the envelope sender, the domain part of that address. Otherwise, if a "pass" is obtained for the HELO/EHLO name, that domain name is used.

8. IANA Considerations

This memo will create an IANA registry of MARF discovery record tags and their legal values and defining documents; this section is yet to be written.

9. Security Considerations

The following security considerations apply when generating or processing a feedback report:

9.1. Inherited from MARF-BASE

All of the Security Considerations from [MARF-BASE] are inherited here.

9.2. These Need Fleshing Out

Subscription verification requests.

[DNSSEC] SHOULD be used to ensure authenticity of all DNS requests and replies.

Reference the data redaction discussion that will appear in an updated dkim-reporting draft soon.

Feedback destination mailboxes are potentially sinks for private information, and should be secured accordingly. The precise methods of securing mailbox files are outside of the scope of this document.

Additional security considerations are likely, and TBD.

9.3. Privacy considerations

[Should these be moved to a BCP?]

When recipients report received messages as spam, they deny any involvement with whoever may turn out to be responsible of authoring or forwarding that mail. Users submit that spam to the community at large in the hope that a corrective action may be taken. In that respect, the only privacy issue is to avoid to further divulge the email addresses of reporting users, or let them be profiled. Suitably redacted email addresses are an acceptable remedy. Responsible senders are advised to include opaque tokens as necessary, in order to reconstruct needed data.

However, users may also err. Unwanted smear may result in case a message containing confidential data is erroneously reported as spam. Authentication and policies are concerned with avoiding such circumstances. Recipients of special addresses at an ADMD that handled a given message could have intercepted the message when they signed or relayed it, if they wanted to, therefore it is less of a concern to forward abuse reports there.

Contact addresses for a given domain name should bear such name as their domain parts. Generators may treat non-conforming addresses with great suspicion, possibly avoiding to relay to them.

10. Acknowledgements

This document was heavily influenced by discussions on the topic within the IRTF Anti-Spam Research Group, collected at [ASRG-ABUSE].

11. Contributors

Many thanks to Murray Kucherawy, Alessandro Vesely, Todd Herr, Jacob Rideout, Derek Digeet, Yakov Shafranovitch, Barry Leiba, Tim Draegen, Michael Adkins, and Tanguy Ortolto for their suggestions and contributions.

12. References

12.1. Normative References

- [ABNF] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", RFC 5234, January 2008.
- [AUTH-METHODS] Kucherawy, M., "Message Header Field for Indicating Message Authentication Status", RFC 5451, April 2009.
- [DNS] Mockapetris, P., "Domain Names -- Implementation and Specification", RFC 1035, November 1987.
- [KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", RFC 2119, March 1997.
- [MAILBOX-NAMES] Crocker, D., "MAILBOX NAMES FOR COMMON SERVICES, ROLES AND FUNCTIONS", RFC 2142, May 1997.
- [MARF-BASE] Shafranovich, Y., Levine, J., and M. Kucherawy, "An Extensible Format for Email Feedback Reports", RFC 5965, April 2010.
- [SMTP] Klensin, J., "Simple Mail Transfer Protocol", RFC 5321, October 2008.
- [SPF] Wong, M. and W. Schlitt, "Sender Policy Framework (SPF) for Authorizing Use of Domains in E-Mail, Version 1", RFC 4408, April 2006.

12.2. Informative References

- [ASRG-ABUSE] Anti-Spam Research Group (ASRG) of the Internet Research Task Force (IRTF), "Abuse Reporting Standards Subgroup of the ASRG", May 2005.
- [DKIM] Allman, E., Callas, J., Delany, M., Libbey, M., Fenton, J., and M. Thomas, "DomainKeys Identified Mail (DKIM) Signatures", RFC 4871, May 2007.
- [DKIM-REPORTING] Kucherawy, M., "Reporting of DKIM Verification Failures", April 2010, <<http://tools.ietf.org/html/>

draft-ietf-marf-dkim-reporting>.

[EMAIL-ARCH]

Crocker, D., "Internet Mail Architecture", RFC 5598,
July 2009.

Appendix A. Public Discussion and Support

[REMOVE BEFORE PUBLICATION]

Public discussion of this proposed specification is handled via the marf@ietf.org mailing list. The list is open. Access to subscription forms and to list archives can be found at <http://www.ietf.org/mail-archive/web/marf/current/maillist.html>

Appendix B. Document History & Open Issues

[REMOVE BEFORE PUBLICATION]

B.1. draft-jdfalk-marf-reporting-discovery-00

- o NEEDED: WG input, references cleanup, ABNF and other formal definitions, and probably lots of other stuff.
- o QUESTION: should this include IODEF as a format?

B.2. draft-jdfalk-marf-reporting-discovery-01

- o Changed from "Experimental" to "Standards Track".
- o Various non-normative textual & formatting improvements, most importantly changing "hangtext" to "hangText" because xml2rfc is (surprisingly) case-sensitive.
- o Added the PTR record as another place to look for advertisements published by ARF consumers. (This may require additional clarifications later in the text.)
- o Moved a few references from normative to informative.
- o Questions & needs listed for version 00 remain valid.

B.3. draft-jdfalk-marf-reporting-discovery-02

- o Added authentication-related and privacy considerations sections written by Alessandro Vesely.
- o Questions & needs listed for versions 00 and 01 remain valid.

B.4. draft-jdfalk-marf-reporting-discovery-03

- o Removed references to end-user report submissions; those should be left for a separate document, and probably be discoverable in the same way as MUA settings.
- o Rewrote some confusing parts.
- o Added the examples section.
- o QUESTION: should the advertisement be published somewhere other than DNS?
- o Questions & needs listed for earlier versions remain valid.
- o QUESTION: should this document include a description of how the complaint feedback signup process works today, and how this proposal would change it? Or is that a separate Informational document?
- o QUESTION: should the current Privacy Considerations section be moved to a separate BCP?

B.5. draft-ietf-marf-reporting-discovery-00

- o This draft now belongs to the MARF WG. Yay!
- o Added "Feedback Consumer re= Default" section.

B.6. draft-ietf-marf-reporting-discovery-01

- o Minor wording change to clarify DKIM d= domain.
- o Removed unused references.

Author's Address

J.D. Falk
Return Path
100 Mathilda Place, Suite 100
Sunnyvale, CA 94086
US

Email: ietf@cybernothing.org
URI: <http://www.returnpath.net/>

