

Internet Engineering Task Force (IETF)
Request for Comments: 5965
Category: Standards Track
ISSN: 2070-1721

Y. Shafranovich
ShafTek Enterprises
J. Levine
Taughannock Networks
M. Kucherawy
Cloudmark
August 2010

An Extensible Format for Email Feedback Reports

Abstract

This document defines an extensible format and MIME type that may be used by mail operators to report feedback about received email to other parties. This format is intended as a machine-readable replacement for various existing report formats currently used in Internet email.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc5965>.

Copyright Notice

Copyright (c) 2010 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
1.1. Purpose	3
1.2. Requirements	4
1.3. Definitions	4
1.3.1. General	4
1.3.2. Email Specific	4
2. Format of Email Feedback Reports	4
3. The 'message/feedback-report' Content Type	5
3.1. Required Fields	6
3.2. Optional Fields Appearing Once	6
3.3. Optional Fields Appearing Multiple Times	7
3.4. Notes about URIs	8
3.5. Formal Definition	8
4. Handling Malformed Reports	10
5. Transport Considerations	10
6. Extensibility	10
7. IANA Considerations	11
7.1. MIME Type Registration of 'message/feedback-report'	11
7.2. Feedback Report Header Fields	12
7.3. Feedback Report Type Values	15
8. Security Considerations	17
8.1. Inherited from RFC 3462	17
8.2. Interpretation	17
8.3. Attacks against Authentication Methods	17
8.4. Intentionally Malformed Reports	18
8.5. Omitting Data from ARF Reports	18
8.6. Automatically Generated ARF Reports	18
8.7. Attached Malware	18
8.8. The User-Agent Field	18
8.9. Malformed Messages	19
9. References	19
9.1. Normative References	19
9.2. Informative References	20
Appendix A. Acknowledgements	22
Appendix B. Sample Feedback Reports	22
B.1. Simple Report for Email Abuse without Optional Headers	22
B.2. Full Report for Email Abuse with All Headers	23

1. Introduction

As the spam problem continues to expand and potential solutions evolve, mail operators are increasingly exchanging abuse reports among themselves and other parties. However, different operators have defined their own formats, and thus the receivers of these reports are forced to write custom software to interpret each of them. In addition, many operators use various other report formats to provide non-abuse-related feedback about processed email. This memo uses the "multipart/report" content type defined in [REPORT], and in that context defines a standard extensible format by creating the "message/feedback-report" [MIME] type for these reports.

While there has been previous work in this area (e.g., [STRADS-BCP] and [ASRG-ABUSE]), none of it has yet been successful. It is hoped that this document will have a better fate.

This format is intended primarily as an Abuse Reporting Format (ARF) for reporting email abuse but also includes support for direct feedback via end user mail clients, reports of some types of virus activity, and some similar issues. This memo also contains provision for extensions should other specific types of reports be desirable in the future.

This document only defines the format and [MIME] content type to be used for these reports. Determination of where these reports should be sent, validation of their contents, and how trust among report generators and report recipients is established are outside the scope of this document. It is assumed that best practices will evolve over time, and will be codified in future documents.

1.1. Purpose

The reports defined in this document 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.

Please note that while the parent "multipart/report" content type defined in [REPORT] is used for all kinds of administrative messages, this format is intended specifically for communications among providers regarding email abuse and related issues, and SHOULD NOT be used for other reports.

1.2. Requirements

The following requirements are necessary for feedback reports (the actual specification is defined later in this document):

- o They must be both human and machine readable;
- o A copy of the original email message (both body and header) or the message header must be enclosed in order to allow the receiver to handle the report properly;
- o The machine-readable section must provide ability for the report generators to share meta-data with receivers;
- o The format must be extensible.

1.3. Definitions

This section defines various terms used throughout this document.

1.3.1. 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].

1.3.2. 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.

2. Format of Email Feedback Reports

To satisfy the requirements, an email feedback report is defined as a [MIME] message with a top-level MIME content type of "multipart/report" (as defined in [REPORT]). The following apply:

- a. The "report-type" parameter of the "multipart/report" type is set to "feedback-report";
- b. The first MIME part of the message contains a human-readable description of the report and MUST be included.

- c. The second MIME part of the message is a machine-readable section with the content type of "message/feedback-report" (defined later in this memo) and MUST be included. This section is intended to convey meta-data about the report in question that may not be readily available from the included email message itself.
- d. The third MIME part of the message is either of type "message/rfc822" (as defined in [MIME-TYPES]) and contains the original message in its entirety OR is of type "text/rfc822-headers" (as defined in [REPORT]) and contains a copy of the entire header block from the original message. This part MUST be included (contrary to [REPORT]). While some operators may choose to modify or redact this portion for privacy or legal reasons, it is RECOMMENDED that the entire original email message be included without any modification as such modifications can impede forensic work by the recipient of this report. See Section 8 for further discussion.
- e. Except as discussed below, each feedback report MUST be related to only a single email message. Summary and aggregate formats are outside of the scope of this specification.
- f. The Subject header field of the feedback report SHOULD be the same as the included email message about which the report is being generated. If it differs, the difference MUST be limited to only a typical forwarding prefix used by Mail User Agents (MUAs) such as "FW:". (Many smaller operators using MUAs for abuse handling rely on the subject lines for processing.)
- g. The primary evidence of the abuse being reported is found in the third part of the report, which contains the original message. The second part contains additional derived data that may help the receiver, but in terms of selecting actionable report data, report recipients SHOULD use the content of the third part first, then data from the second part. The first part is meant to contain explanatory text for human use but is not itself a part of the report, and SHOULD NOT be used if it is in conflict with the other parts.

3. The 'message/feedback-report' Content Type

A new [MIME] content type called "message/feedback-report" is defined. This content type provides a machine-readable section intended to let the report generator convey meta-data to the report receiver. The intent of this section is to convey information that may not be obvious or may not be easily extracted from the original email message body or header.

The body of this content type consists of multiple "fields" formatted according to the ABNF of [MAIL] header fields. This section defines the initial set of fields provided by this specification. Additional fields may be registered according to the procedure described later in this memo. Although these fields have a syntax similar to those of mail message header fields, they are semantically distinct; hence, they SHOULD NOT be repeated as header fields of the message containing the report. Note that these fields represent information that the receiver is asserting about the report in question, but are not necessarily verifiable. Report receivers MUST NOT assume that these assertions are always accurate.

Note that the above limitation in no way restricts the use of message header fields that are registered in the IANA header field registry with the same field names.

3.1. Required Fields

The following report header fields MUST appear exactly once:

- o "Feedback-Type" contains the type of feedback report (as defined in the corresponding IANA registry and later in this memo). This is intended to let report parsers distinguish among different types of reports.
- o "User-Agent" indicates the name and version of the software program that generated the report. The format of this field MUST follow section 14.43 of [HTTP]. This field is for documentation only; there is no registry of user agent names or versions, and report receivers SHOULD NOT expect user agent names to belong to a known set.
- o "Version" indicates the version of specification that the report generator is using to generate the report. The version number in this specification is set to "1".

3.2. Optional Fields Appearing Once

The following header fields are optional and MUST NOT appear more than once:

- o "Original-Envelope-Id" contains the envelope ID string used in the original [SMTP] transaction (see section 2.2.1 of [DSN]).
- o "Original-Mail-From" contains a copy of the email address used in the MAIL FROM portion of the original SMTP transaction. The format of this field is defined in section 4.1.2 of [SMTP] as "Reverse-path".

- o "Arrival-Date" indicates the date and time at which the original message was received by the Mail Transfer Agent (MTA) of the generating ADMD (Administrative Management Domain). This field MUST be formatted as per section 3.3 of [MAIL].
- o "Reporting-MTA" indicates the name of the MTA generating this feedback report. This field is defined in section 2.2.2 of [DSN], except that it is an optional field in this report.
- o "Source-IP" contains an IPv4 or IPv6 address of the MTA from which the original message was received. Addresses MUST be formatted as per section 4.1.3 of [SMTP].
- o "Incidents" contains an unsigned 32-bit integer indicating the number of incidents this report represents. The absence of this field implies the report covers a single incident.

The historic field "Received-Date" SHOULD also be accepted and interpreted identically to "Arrival-Date". However, if both are present, the report is malformed and SHOULD be treated as described in Section 4.

3.3. Optional Fields Appearing Multiple Times

The following set of header fields are optional and may appear any number of times as appropriate:

- o "Authentication-Results" indicates the result of one or more authentication checks run by the report generator. The format of this field is defined in [AUTH-RESULTS]. Report receivers should note that this field only indicates an assertion made by the report generator.
- o "Original-Rcpt-To" includes a copy of the email address used in the RCPT TO portion of the original [SMTP] transaction. The format of this field is a "Reverse-path" defined in section 4.1.2 of that memo. This field SHOULD be repeated for every SMTP recipient seen by the report generator.
- o "Reported-Domain" includes a domain name that the report generator believes to be relevant to the report, e.g., the domain whose apparent actions provoked the generation of the report. It is unspecified how the report generator determines this information, and thus the report receiver cannot be certain how it was chosen. It is often used as a means of suggesting to the report receiver how this report might be handled. In cases where the derivation

is not obvious, the report generator is encouraged to clarify in the text section of the report. Domain format is defined in section 2.3.1 of [DNS].

- o "Reported-URI" indicates a URI that the report generator believes to be relevant to the report, e.g., a suspect URI that was found in the message that caused the report to be generated. The same caveats about the origin of the value of "Reported-Domain" apply to this field. The URI format is defined in [URI].

3.4. Notes about URIs

Implementors should be aware that the Reported-URI field 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.

Furthermore, it is outside the scope of this standard whether the data carried in this field implies any additional information. Implementors may negotiate their own agreements surrounding the interpretation of this data.

3.5. Formal Definition

The formal definition of the contents of a "message/feedback-report" media type using [ABNF] is as follows:

```
feedback-report = *( feedback-type / user-agent / version )
                  opt-fields-once
                  *( opt-fields-many )
                  *( ext-field )
```

```
feedback-type = "Feedback-Type:" [CFWS] token [CFWS] CRLF
                ; the "token" must be a registered feedback type as
                ; described elsewhere in this document
```

```
user-agent = "User-Agent:" [CFWS] product *( CFWS product )
            [CFWS] CRLF
```

```
version = "Version:" [CFWS] %x31-39 *DIGIT [CFWS] CRLF
         ; as described above
```

```
opt-fields-once = [ arrival-date ]
                  [ incidents ]
                  [ original-envelope-id ]
                  [ original-mail-from ]
                  [ reporting-mta ]
                  [ source-ip ]
```



```

arrival-date = "Arrival-Date:" [CFWS] date-time CRLF

incidents = "Incidents:" [CFWS] 1*DIGIT [CFWS] CRLF
           ; must be a 32-bit unsigned integer

original-envelope-id = "Original-Envelope-Id:" [CFWS]
                     envelope-id [CFWS] CRLF

original-mail-from = "Original-Mail-From:" [CFWS]
                   reverse-path [CFWS] CRLF

reporting-mta = "Reporting-MTA:" [CFWS] mta-name-type [CFWS] ";"
              [CFWS] mta-name [CFWS] CRLF

source-ip = "Source-IP:" [CFWS]
           ( IPv4-address-literal /
             IPv6-address-literal ) [CFWS] CRLF

opt-fields-many = [ authres-header ]
                 [ original-rcpt-to ]
                 [ reported-domain ]
                 [ reported-uri ]

original-rcpt-to = "Original-Rcpt-To:" [CFWS]
                 forward-path [CFWS] CRLF

reported-domain = "Reported-Domain:" [CFWS]
                 domain [CFWS] CRLF

reported-uri = "Reported-URI:" [CFWS] URI [CFWS] CRLF

ext-field = field-name ":" unstructured

```

A set of fields satisfying this ABNF may appear in the transmitted message in any order.

"CRLF" and "DIGIT" are imported from [ABNF].

"token" is imported from [MIME].

"product" is imported from [HTTP].

"field-name", "unstructured", "CFWS", "date-time", and "domain" are imported from [MAIL].

"envelope-id", "mta-name-type", and "mta-name" are imported from [DSN].

"reverse-path", "forward-path", "local-part", "IPv4-address-literal", and "IPv6-address-literal" are imported from [SMTP].

"URI" is imported from [URI].

"authres-header" is imported from [AUTH-RESULTS].

"ext-field" refers to extension fields, which are discussed in Section 6.

4. Handling Malformed Reports

When an agent that accepts and handles ARF messages receives a message that purports (by MIME type) to be an ARF message but syntactically deviates from this specification, that agent SHOULD ignore or reject the message. Where rejection is performed, the rejection notice (either via an [SMTP] reply or generation of a [DSN]) SHOULD identify the specific cause for the rejection.

See Section 8.9 for further discussion.

5. Transport Considerations

[DSN] requires that its reports be sent with the empty [SMTP] envelope sender to avoid bounce loops. A similar requirement was considered for this specification, but it seems unlikely that an ARF report would be generated in response to receipt of an ARF report, and furthermore such a requirement would prevent an ARF generator from ever determining that an ARF report was not actually received.

On the other hand, if an ARF report is generated without the empty envelope sender and is sent to an address that actually does not work, then the generating address can also be overwhelmed by DSNs as a denial-of-service attack (see Section 8.6).

This specification therefore makes no requirement related to the envelope sender of a generated report. Operators will have to consider what envelope sender to use within the context of their own installations.

6. Extensibility

Like many other formats and protocols, this format may need to be extended over time to fit the ever-changing landscape of the Internet. Therefore, extensibility is provided via two IANA registries: one for feedback types and a second for report header fields. The feedback type registry is to be used in conjunction with the "Feedback-Type" field above. The header name registry is

intended for registration of new meta-data fields to be used in the machine-readable portion (part 2) of this format. Please note that version numbers do not change with new field registrations unless a new specification of this format is published. Also, note that all new field registrations may only be registered as optional fields. Any new required fields REQUIRE a new version of this specification to be published.

In order to encourage extensibility and interoperability of this format, implementors MUST ignore any fields or report types they do not explicitly support.

Additional report types (extension report types) or report header fields might be defined in the future by later revisions to this specification, or by registrations as described above. Such types and fields MUST be registered as described above and published in an Open Specification such as an RFC.

Experimental report types and report header fields MUST only be used between ADMDs that have explicitly consented to use them. These names and the parameters associated with them are not documented in RFCs. Therefore, they are subject to change at any time and are not suitable for general use.

7. IANA Considerations

IANA has registered a new [MIME] type and created two new registries, as described below.

7.1. MIME Type Registration of 'message/feedback-report'

This section provides the media type registration application from [MIME-REG] for processing by IANA:

To: ietf-types@iana.org

Subject: Registration of media type message/feedback-report

Type name: message

Subtype name: feedback-report

Required parameters: none

Optional parameters: none

Encoding considerations: "7bit" encoding is sufficient and MUST be used to maintain readability when viewed by non-MIME mail readers.

Security considerations: See Section 8 of [RFC5965].

Interoperability considerations: Implementors MUST ignore any fields they do not support.

Published specification: [RFC5965]

Applications that use this media type: Abuse helpdesk software for ISPs, mail service bureaus, mail certifiers, and similar organizations

Additional information: none

Person and email address to contact for further information:

Yakov Shafranovich <ietf@shaftek.org>

Murray S. Kucherawy <msk@cloudmark.com>

Intended usage: COMMON

Author:

Yakov Shafranovich

John R. Levine

Murray S. Kucherawy

Change controller: IESG

7.2. Feedback Report Header Fields

IANA has created the "Feedback Report Header Fields" registry. This registry contains header fields for use in feedback reports, as defined by this memo.

New registrations or updates MUST be published in accordance with the "Specification Required" guidelines as described in [IANA]. Any new field thus registered is considered optional by this specification unless a new version of this memo is published.

New registrations and updates MUST contain the following information:

1. Name of the field being registered or updated
2. Short description of the field

3. Whether the field can appear more than once
4. To which feedback type(s) this field applies (or "any")
5. The document in which the specification of the field is published
6. New or updated status, which MUST be one of:

current: The field is in current use

deprecated: The field is in current use but its use is discouraged

historic: The field is no longer in current use

An update may make a notation on an existing registration indicating that a registered field is historic or deprecated if appropriate.

The initial registry contains these values:

Field Name: Arrival-Date
Description: date/time the original message was received
Multiple Appearances: No
Related "Feedback-Type": any
Published in: [RFC5965]
Status: current

Field Name: Authentication-Results
Description: results of authentication check(s)
Multiple Appearances: Yes
Related "Feedback-Type": any
Published in: [RFC5965]
Status: current

Field Name: Feedback-Type
Description: registered feedback report type
Multiple Appearances: No
Related "Feedback-Type": N/A
Published in: [RFC5965]
Status: current

Field Name: Incidents
Description: expression of how many similar incidents are
represented by this report
Multiple Appearances: No
Related "Feedback-Type": any
Published in: [RFC5965]
Status: current

Field Name: Original-Mail-From
Description: email address used in the MAIL FROM portion of the
original SMTP transaction
Multiple Appearances: No
Related "Feedback-Type": any
Published in: [RFC5965]
Status: current

Field Name: Original-Rcpt-To
Description: email address used in the RCPT TO portion of the
original SMTP transaction
Multiple Appearances: Yes
Related "Feedback-Type": any
Published in: [RFC5965]
Status: current

Field Name: Received-Date
Description: date/time the original message was received
(replaced by "Arrival-Date")
Multiple Appearances: No
Related "Feedback-Type": any
Published in: [RFC5965]
Status: historic

Field Name: Reported-Domain
Description: a domain name the report generator considers to
be key to the message about which a report is
being generated
Multiple Appearances: Yes
Related "Feedback-Type": any
Published in: [RFC5965]
Status: current

Field Name: Reported-URI
Description: a URI the report generator considers to be key
to the message about which a report is being
generated
Multiple Appearances: Yes
Related "Feedback-Type": any
Published in: [RFC5965]
Status: current

Field Name: Reporting-MTA
Description: MTA generating this report
Multiple Appearances: No
Related "Feedback-Type": any
Published in: [RFC5965]
Status: current

Field Name: Source-IP
Description: IPv4 or IPv6 address from which the original message
was received
Multiple Appearances: No
Related "Feedback-Type": any
Published in: [RFC5965]
Status: current

Field Name: User-Agent
Description: name and version of the program generating the
report
Multiple Appearances: No
Related "Feedback-Type": any
Published in: [RFC5965]
Status: current

Field Name: Version
Description: version of specification used
Multiple Appearances: No
Related "Feedback-Type": any
Published in: [RFC5965]
Status: current

7.3. Feedback Report Type Values

IANA has created the "Feedback Report Type Values" registry. This registry contains feedback types for use in feedback reports, defined by this memo.

New registrations or updates MUST be published in accordance with the "Specification Required" guidelines as described in [IANA]. Any new field thus registered is considered optional by this specification unless a new version of this memo is published.

New registrations MUST contain the following information:

1. Name of the feedback type being registered
2. Short description of the feedback type
3. The document in which the specification of the field is published
4. New or updated status, which MUST be one of:

current: The field is in current use

deprecated: The field is in current use but its use is discouraged

historic: The field is no longer in current use

The initial registry contains these values:

Feedback Type Name: abuse

Description: unsolicited email or some other kind of email abuse

Published in: [RFC5965]

Status: current

Feedback Type Name: fraud

Description: indicates some kind of fraud or phishing activity

Published in: [RFC5965]

Status: current

Feedback Type Name: other

Description: any other feedback that does not fit into other registered types

Published in: [RFC5965]

Status: current

Feedback Type Name: virus

Description: report of a virus found in the originating message

Published in: [RFC5965]

Status: current

8. Security Considerations

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

8.1. Inherited from RFC 3462

All of the Security Considerations from [REPORT] are inherited here.

8.2. Interpretation

This specification describes a report format. The authentication and validity of the content of the report SHOULD be established through other means. The content of an unvetted report could be wrong, incomplete or deliberately false, including the alleged abuse incident in the third part, derived data in the second part or the human-readable first part.

There will be some desire to perform some actions in an automated fashion in order to enact timely responses to common feedback reports. Caution must be taken, however, as there is no substantial security around the content of these reports. An attacker could craft a report meant to generate undesirable actions on the part of a report recipient.

It is suggested that the origin of an ARF report be vetted, such as by using common message authentication schemes like [SMIME], [DKIM], [SPF], or [SENDERID], prior to the undertaking of any kind of automated action in response to receipt of the report. In particular, S/MIME offers the strongest authentication and the cost of key exchange is assumed in the process of establishing a bilateral reporting relationship that uses this specification; however, it is not as transparent as the others and thus will interfere with the parsing capabilities of code that is designed specifically to handle multipart/report messages.

The details of the required validation to achieve this are a matter of local policy and are thus outside the scope of this specification.

8.3. Attacks against Authentication Methods

If an attack becomes known against an authentication method, clearly then the agent verifying that method can be fooled into thinking an inauthentic message is authentic, and thus the value of this header field can be misleading. It follows that any attack against an authentication method that might be used to protect the authenticity of an abuse report is also a security consideration here.

8.4. Intentionally Malformed Reports

It is possible for an attacker to generate an ARF message field that is extraordinarily large or otherwise malformed in an attempt to discover or exploit weaknesses in recipient parsing code. Implementors SHOULD thoroughly verify all such messages and be robust against intentionally as well as unintentionally malformed messages.

8.5. Omitting Data from ARF Reports

The sending of these reports can reveal possibly private information about the person sending the report. For example, such a report sent in response to a mailing list posting will reveal to the report recipient a valid email address on the list that might otherwise have remained hidden.

For this reason, report generators might wish to redact portions of the report to conceal private information. Doing so could be necessary where privacy trumps operational necessity, but, as mentioned in Section 2, it might impede a timely or meaningful response from the report recipient.

8.6. Automatically Generated ARF Reports

Systems have been implemented that generate ARF reports automatically in response to an event. For example, software monitoring a honeypot email address might generate an ARF report immediately upon delivery of any message to it. An attacker that becomes aware of such a configuration can exploit it to attack an ARF recipient with automatically generated ARF reports.

8.7. Attached Malware

As this format is sometimes used to automatically report malware, ARF processors (human or otherwise) SHOULD ensure that attachments are processed in a manner appropriate for unverified and potentially hostile data.

8.8. The User-Agent Field

Further to Section 8.2, the User-Agent field is an assertion of the generating software and is neither specified in this memo nor derived from the message represented in the third part of the report. It is intended for documentation and debugging, and since it is trivially forged by a malicious agent, it SHOULD NOT be interpreted by recipients.

8.9. Malformed Messages

Further to the discussion in Section 4, there might be cases where an ARF processing agent elects to accept messages not consistent with this specification, such as during transition periods where some fields are moving toward "historic" or "deprecated" status, or the introduction of new non-standard extension or experimental fields. Such choices need to be implemented with extreme caution; where two different fields have related meaning (e.g., "Received-Date", which is historic, and "Arrival-Date", which is current), an attacker could craft a report that makes a confusing claim in an attempt to exploit such liberal parsing logic.

9. References

9.1. Normative References

- [ABNF] Crocker, D. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, January 2008.
- [AUTH-RESULTS] Kucherawy, M., "Message Header Field for Indicating Message Authentication Status", RFC 5451, April 2009.
- [DNS] Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, November 1987.
- [DSN] Moore, K. and G. Vaudreuil, "An Extensible Message Format for Delivery Status Notifications", RFC 3464, January 2003.
- [HTTP] Fielding, R., Gettys, J., Mogul, J., Frystyk, H., Masinter, L., Leach, P., and T. Berners-Lee, "Hypertext Transfer Protocol -- HTTP/1.1", RFC 2616, June 1999.
- [KEYWORDS] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [MAIL] Resnick, P., Ed., "Internet Message Format", RFC 5322, October 2008.
- [MIME] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies", RFC 2045, November 1996.

- [MIME-REG] Freed, N. and J. Klensin, "Media Type Specifications and Registration Procedures", BCP 13, RFC 4288, December 2005.
- [MIME-TYPES] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part Two: Media Types", RFC 2046, November 1996.
- [REPORT] Vaudreuil, G., "The Multipart/Report Content Type for the Reporting of Mail System Administrative Messages", RFC 3462, January 2003.
- [SMTP] Klensin, J., "Simple Mail Transfer Protocol", RFC 5321, October 2008.
- [URI] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, January 2005.

9.2. 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.
- [EMAIL-ARCH] Crocker, D., "Internet Mail Architecture", RFC 5598, July 2009.
- [IANA] Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.
- [SENDERID] Lyon, J. and M. Wong, "Sender ID: Authenticating E-Mail", RFC 4406, April 2006.
- [SMIME] Ramsdell, B. and S. Turner, "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.2 Message Specification", RFC 5751, January 2010.
- [SPF] Wong, M. and W. Schlitt, "Sender Policy Framework (SPF) for Authorizing Use of Domains in E-Mail, Version 1", RFC 4408, April 2006.

[STRADS-BCP] Crissman, G., "Proposed Spam Reporting BCP Document",
May 2005.

Appendix A. Acknowledgements

The authors would like to thank many of the members of the email community who provided helpful comments and suggestions for this document including many of the participants in ASRG, IETF, and MAAWG activities, and all of the members of the abuse-feedback-report public mailing list.

Appendix B. Sample Feedback Reports

This section presents some examples of the use of this message format to report feedback about an arriving message.

B.1. Simple Report for Email Abuse without Optional Headers

Simple report:

```
From: <abusedesk@example.com>
Date: Thu, 8 Mar 2005 17:40:36 EDT
Subject: FW: Earn money
To: <abuse@example.net>
MIME-Version: 1.0
Content-Type: multipart/report; report-type=feedback-report;
    boundary="part1_13d.2e68ed54_boundary"
```

```
--part1_13d.2e68ed54_boundary
Content-Type: text/plain; charset="US-ASCII"
Content-Transfer-Encoding: 7bit
```

This is an email abuse report for an email message received from IP 192.0.2.1 on Thu, 8 Mar 2005 14:00:00 EDT. For more information about this format please see <http://www.mipassoc.org/arf/>.

```
--part1_13d.2e68ed54_boundary
Content-Type: message/feedback-report
```

```
Feedback-Type: abuse
User-Agent: SomeGenerator/1.0
Version: 1
```

```
--part1_13d.2e68ed54_boundary
Content-Type: message/rfc822
Content-Disposition: inline
```

```
Received: from mailserver.example.net
    (mailserver.example.net [192.0.2.1])
    by example.com with ESMTP id M63d4137594e46;
    Thu, 08 Mar 2005 14:00:00 -0400
```


Reporting-MTA: dns; mail.example.com
Source-IP: 192.0.2.1
Authentication-Results: mail.example.com;
 spf=fail smtp.mail=somespammer@example.com
Reported-Domain: example.net
Reported-Uri: http://example.net/earn_money.html
Reported-Uri: mailto:user@example.com
Removal-Recipient: user@example.com

--part1_13d.2e68ed54_boundary
Content-Type: message/rfc822
Content-Disposition: inline

From: <somespammer@example.net>
Received: from mailserver.example.net (mailserver.example.net
 [192.0.2.1]) by example.com with ESMTTP id M63d4137594e46;
 Thu, 08 Mar 2005 14:00:00 -0400

To: <Undisclosed Recipients>
Subject: Earn money
MIME-Version: 1.0
Content-type: text/plain
Message-ID: 8787KJKJ3K4J3K4J3K4J3K4J3.mail@example.net
Date: Thu, 02 Sep 2004 12:31:03 -0500

Spam Spam Spam
Spam Spam Spam
Spam Spam Spam
Spam Spam Spam
--part1_13d.2e68ed54_boundary--

Example 1: Generic abuse report with maximum returned information

A contrived example in which the report generator has returned all possible information about an abuse incident.

Authors' Addresses

Yakov Shafranovich
ShafTek Enterprises
4014 Labyrinth Rd.
Baltimore, MD 21215
US

E-Mail: ietf@shaftek.org
URI: <http://www.shaftek.org>

John R. Levine
Taughannock Networks
PO Box 727
Trumansburg, NY 14886
US

Phone: +1 831 480 2300
E-Mail: standards@taugh.com

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

Phone: +1 415 946 3800
E-Mail: msk@cloudmark.com

