Internet Engineering Task Force (IETF)
Request for Comments: 6650
Updates: 5965
Category: Standards Track
ISSN: 2070-1721

J. Falk
Return Path
M. Kucherawy, Ed.
Cloudmark

June 2012

Creation and Use of Email Feedback Reports: An Applicability Statement for the Abuse Reporting Format (ARF)

Abstract

RFC 5965 defines an extensible, machine-readable format intended for mail operators to report feedback about received email to other parties. This applicability statement describes common methods for utilizing this format for reporting both abuse and authentication failure events. Mailbox Providers of any size, mail-sending entities, and end users can use these methods as a basis to create procedures that best suit them. Some related optional mechanisms are also discussed.

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/rfc6650.

Table of Contents

1. Introduction

The Abuse Reporting Format (ARF) was initially developed for two very specific use cases. Initially, it was intended to be used for reporting feedback between large email operators, or from large email operators to end user network access operators, any of whom could be presumed to have automated abuse-handling systems. Secondarily, it is used by those same large mail operators to send those same reports to other entities, including those involved in sending bulk email for commercial purposes. In either case, the reports would be triggered by direct end user action such as clicking on a "report spam" button in their email client.

Though other uses for ARF as defined in [RFC5965] have been discussed (and may be documented similarly in the future), abuse reporting remains the primary application, with a small amount of adoption of extensions that enable authentication failure reporting.

This applicability statement provides direction for using ARF in both contexts. It also includes some statements about the use of ARF in conjunction with other email technologies.

The purpose for reporting abusive messages is to stop recurrences. The methods described in this document focus on automating abuse reporting as much as practical, so as to minimize the work of a site’s abuse team. There are further reasons why abuse feedback generation is worthwhile, such as instruction of mail filters or reputation trackers, or initiation of investigations of particularly egregious abuses. These other applications are not discussed in this memo.

Further introduction to this topic may be found in [RFC6449], which has more information about the general topic of abuse reporting. Many of the specific ARF guidelines in this document were taken from the principles presented in [RFC6449].

At the time of publication of this document, five feedback types are registered. This document only discusses two of them ("abuse" [RFC5965] and "auth-failure" [RFC6591]), as they are seeing sufficient use in practice that applicability statements can be made about them. The others, i.e., "fraud" [RFC5965], "other" [RFC5965], and "not-spam" [RFC6430], are either too new or too seldom used to be included here.

2. Definitions

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 [RFC2119] and are intended to replace the Requirement Levels described in Section 3.3 of [RFC2026].

Some of the terminology used in this document is taken from [RFC5598].

"Mailbox Provider" refers to an organization that accepts, stores, and offers access to [RFC5322] messages ("email messages") for end users. Such an organization has typically implemented SMTP [RFC5321] and might provide access to messages through IMAP [RFC3501], the Post Office Protocol (POP) [RFC1939], a proprietary interface designed for HTTP [RFC2616], or a proprietary protocol.

3. Solicited and Unsolicited Reports

The original, and still by far the most common, application of [RFC5965] is when two mail systems make a private agreement to exchange abuse reports -- usually reports due to recipients manually reporting messages as spam. We refer to these as solicited reports.

Other uses for ARF involve such reports sent between parties that don’t know each other. These unsolicited reports are sent without prior arrangement between the parties as to the context and meaning of the reports. Therefore, the constraints on how these unsolicited reports need to be structured such that they are likely to be useful to the recipient -- e.g., to what address(es) they can usefully be sent, what issues they can be used to report, and how they can be handled by the receiver of the report -- are very different.

The two cases are covered separately in the sections that follow.

4. Generating and Handling Solicited Abuse Reports

4.1. General Considerations for Feedback Providers

A Mailbox Provider receives reports of abusive or unwanted mail from its users, most often by providing a "report spam" button (or similar nomenclature) in the MUA (Mail User Agent). The method of transferring this message and any associated metadata from the MUA to the Mailbox Provider’s ARF processing system is not defined by any standards document but is discussed further in Section 3.2 of [RFC6449]. Policy concerns related to the collection of this data are discussed in Section 3.4 of [RFC6449].

To implement the recommendations of this memo, the reports are formatted per [RFC5965] and transmitted as an email message [RFC5322], typically using SMTP [RFC5321].

Ongoing maintenance of an ARF processing system is discussed in Section 3.6 of [RFC6449].

4.2. Where to Send Reports

The Mailbox Provider SHOULD NOT send reports to addresses that have not explicitly requested them. A valid deviation might be the result of local policy instructions. The process whereby such parties may request the reports is discussed in Section 3.5 of [RFC6449].

4.3. What to Put in Reports

The reports SHOULD use "Feedback-Type: abuse" for the report type. Although a Mailbox Provider generating the reports can use other types appropriate to the nature of the abuse being reported, the operator receiving the reports might not treat different feedback types differently.

The following fields are optional in [RFC5965] but SHOULD be used in this context when their corresponding values are available: Original-Mail-From, Arrival-Date, Source-IP, and Original-Rcpt-To. Other optional fields can be included as deemed appropriate by the implementer.

User-identifiable data MAY be obscured as described in [RFC6590].

4.4. General Considerations for Feedback Consumers

ARF report streams are established proactively between Feedback Providers and Feedback Consumers. Recommendations for preparing to request feedback are discussed in Section 4.1 of [RFC6449].

Operators MUST be able to accept ARF [RFC5965] reports as email messages [RFC5322] over SMTP [RFC5321]. These messages, and other types of email messages that can be received, are discussed in Section 4.2 of [RFC6449].

Recipients of feedback reports that are part of formal feedback arrangements have to be capable of handling large volumes of reports. This could require automation of report processing as discussed in Section 4.4 of [RFC6449].

4.5. What to Expect

The list of valid Feedback-Types is defined in [RFC5965], which created an IANA registry for valid values to allow for extensions. However, to allow for handling of new types that are not yet supported, an automated report processing system MUST NOT reject (in the SMTP sense) a report based solely on an unknown Feedback-Type. The automated system can simply set reports of unknown types aside for manual handling. However, Mailbox Providers might only make use of the "abuse" Feedback-Type. Therefore, report receivers might be required to do additional analysis to separate different types of abuse reports after receipt if they do not have prior specific knowledge of the sender of the report.

Report receivers MUST accept reports that have obscured their user-identifiable data as described in [RFC6590]. That document also discusses the handling of such reports. This technique is also discussed in Section 4.4 of [RFC6449].

4.6. What to Do with Reports

Section 4.3 of [RFC6449] discusses actions that mail operators might take upon receiving a report (or multiple reports).

5. Generating and Handling Unsolicited Abuse Reports

5.1. General Considerations

It is essential for report recipients to be capable of throttling reports being sent to avoid damage to their own installations. Therefore, Feedback Providers MUST provide a way for report recipients to request that no further reports be sent. Unfortunately, no standardized mechanism for such requests exists to date, and all existing mechanisms for meeting this requirement are out-of-band.

Message authentication is generally a good idea, but it is especially important to encourage credibility of, and thus response to, unsolicited reports. Therefore, as with any other message, Feedback Providers sending unsolicited reports SHOULD send reports that they expect will pass the Sender Policy Framework (SPF) [RFC4408] and/or DomainKeys Identified Mail (DKIM) [RFC6376] checks.

5.2. When to Generate Reports

Handling of unsolicited reports has a significant cost to the report receiver. Senders of unsolicited reports, especially those sending large volumes of them automatically, SHOULD NOT send reports that cannot be used as a basis for action by the recipient, whether this is due to the report being sent about an incident that is not abuse-related, the report being sent to an email address that won't result in action, or the content or format of the report being hard for the recipient to read or use.

Feedback Providers SHOULD NOT report all mail sent from a particular sender merely because some of it is determined to be abusive.

Mechanical reports of mail that "looks like" spam, based solely on the results of inline content analysis tools, SHOULD NOT be sent since, because of their subjective nature, they are unlikely to provide a basis for the recipient to take action. Complaints generated by end users about mail that is determined by them to be abusive, or mail delivered to "spam trap" or "honeypot" addresses, are far more likely to be accurate and MAY be sent.

If a Feedback Provider applies SPF [RFC4408] to arriving messages, a report SHOULD NOT be generated to the RFC5321.MailFrom domain if the SPF evaluation produced a "Fail", "SoftFail", "TempError", or "PermError" report, as no reliable assertion or assumption can be made that use of the domain was authorized. A valid exception would be specific knowledge that the SPF result is not definitive for that domain under those circumstances (for example, a message that is also signed using DKIM [RFC6376] by the same domain, and that signature validates).

5.3. Where to Send Reports

Rather than generating feedback reports themselves, MUAs SHOULD create abuse reports and send these reports back to their Mailbox Providers so that they can generate and send ARF messages on behalf of end users (see Section 3.2 of [RFC6449]). This allows centralized processing and tracking of reports, and provides training input to filtering systems. There is, however, no standard mechanism for this signaling between MUAs and Mailbox Providers to trigger abuse reports.

Feedback Providers SHOULD NOT send reports to recipients that are uninvolved or only peripherally involved. For example, they SHOULD NOT send reports to the operator of every Autonomous System in the path between the apparent originating system and the operator

generating the report. Instead, they need to send reports to recipients that are both responsible for the messages and able to do something about them.

Deciding where to send an unsolicited report will typically rely on heuristics. Abuse addresses in WHOIS [RFC3912] records of the IP address relaying the subject message and/or of the domain name found in the results of a PTR ("reverse lookup") query on that address are likely reasonable candidates, as is the abuse@domain role address (see [RFC2142]) of related domains. Unsolicited reports SHOULD NOT be sent to email addresses that are not clearly intended to handle abuse reports. Legitimate candidates include those found in WHOIS records or on a web site that either are explicitly described as an abuse contact or are of the form "abuse@domain".

Where an abusive message is authenticated using a domain-level authentication technology such as DKIM [RFC6376] or SPF [RFC4408], the domain that has been verified by the authentication mechanism is often a reasonable candidate for receiving feedback about the message. For DKIM, though, while the authenticated domain has some responsibility for the mail sent, it can be a poor contact point for abuse issues (for example, it could represent the message's author but not its sender, it could identify the bad actor responsible for the message, or it could refer to a domain that cannot receive mail at all).

Often, unsolicited reports will have no meaning if sent to abuse reporting addresses belonging to the abusive parties themselves. In fact, it is possible that such reports might reveal information about complainants. Reports SHOULD NOT be sent to such addresses if they can be identified beforehand, except where the abusive party is known to be responsive to such reports.

5.4. What to Put in Reports

Reports SHOULD use "Feedback-Type: abuse" but can use other types as appropriate. However, the Mailbox Provider generating the reports cannot assume that the operator receiving the reports will treat different Feedback-Types differently.

Reports SHOULD include the following optional fields whenever their corresponding values are available and applicable to the report: Original-Mail-From, Arrival-Date, Source-IP, and Original-Rcpt-To. Other optional fields can be included as deemed appropriate by the implementer.

Experience suggests that the use of ARF is advisable in most contexts. Automated recipient systems can handle abuse reports sent in ARF at least as well as any other format such as plain text, with or without a copy of the message attached. That holds even for systems that did not request ARF reports, assuming such reports are generated considering the possibility of recipients that don't use automated ARF parsing. Anyone sending unsolicited reports in ARF can legitimately presume that some recipients will only be able to access the human-readable (first, text/plain) part of it and SHOULD include all information needed also in this part. Further, they SHOULD ensure that the report is readable when viewed as plain text, to give low-end ticketing systems as much assistance as possible. In extreme cases, failure to take these steps may result in the report being discarded or ignored.

5.5. What to Do with Reports

Receivers of unsolicited reports can take advantage of the standardized parts of ARF to automate processing. Independent of the sender of the report, they can improve processing by separating valid reports from invalid reports by, for example, looking for references to IP address ranges, domains, and mailboxes for which the recipient organization is responsible in the copy of the reported message, and by correlating multiple reports of similar messages to identify bulk email senders.

Per Section 4.4 of [RFC6449], a network service provider MAY use ARF data for automated forwarding of feedback messages to the originating customer.

Published abuse mailbox addresses SHOULD NOT reject non-ARF messages based solely on the format, as generation of ARF messages can occasionally be unavailable or not applicable. Deviation from this requirement could be done due to local policy decisions regarding other message criteria.

Although [RFC6449] suggests that replying to feedback is not useful, in the case of receipt of ARF reports where no feedback arrangement has been established, a non-automated reply might be desirable to indicate what action resulted from the complaint, heading off more severe filtering by the Feedback Provider. In addition, using an address that cannot receive replies precludes any requests for additional information and increases the likelihood that further reports will be discarded or blocked. Thus, a Feedback Provider sending unsolicited reports SHOULD NOT generate reports for which a reply cannot be received. Where an unsolicited report results in the establishment of contact with a responsible and responsive party, this data can be saved for future complaint handling and possible

establishment of a formal (solicited) feedback arrangement. See Section 3.5 of [RFC6449] for a discussion of establishment of feedback arrangements.

6. Generating Automatic Authentication Failure Reports

There are some cases where report generation is caused by automation rather than user requests. A specific example of this is reporting, using ARF (or extensions to it), of messages that fail particular message authentication checks. Examples of this include [RFC6651] and [RFC6652]. The considerations presented below apply in those cases.

The applicability statement for this use case is somewhat smaller, as many of the issues associated with abuse reports are not relevant to reports about authentication failures.

Automatic feedback generators MUST select actual message recipients based on data provided by willing report receivers. In particular, recipients MUST NOT be selected using heuristics.

If the message under evaluation by the Verifier is an ARF [RFC5965] message, a report MUST NOT be automatically generated.

The message for a new report sent via SMTP MUST be constructed so as to avoid amplification attacks, deliberate or otherwise. The envelope sender address of the report MUST be chosen so that these reports will not generate mail loops. Similar to Section 2 of [RFC3464], the envelope sender address of the report MUST be chosen to ensure that no feedback reports will be issued in response to the report itself. Therefore, when an SMTP transaction is used to send a report, the MAIL FROM command SHOULD use the NULL reverse-path, i.e., "MAIL FROM:<>". An exception to this would be the use of a reverse-path selected such that SPF checks on the report will pass; in such cases, the operator will need to make provisions to avoid the amplification attack or mail loop via other means.

Reports SHOULD use "Feedback-Type: auth-failure" but MAY use other types as appropriate. However, the Mailbox Provider generating the reports cannot assume that the operator receiving the reports will treat different Feedback-Types differently.

These reports SHOULD include the following fields, although they are optional in [RFC5965], whenever their corresponding values are available: Original-Mail-From, Arrival-Date, Source-IP, and Original-Rcpt-To. Other optional fields can be included as deemed appropriate by the implementer.

7. Security Considerations

7.1. Security Considerations in Other Documents

Implementers are strongly urged to review, at a minimum, the Security Considerations sections of [RFC5965] and [RFC6449].

7.2. Forgeries

Feedback Providers that relay user complaints directly, rather than by reference to a stored message (e.g., IMAP or POP), could be duped into sending a complaint about a message that the complaining user never actually received, as an attack on the purported originator of the falsified message. Feedback Providers need to be resilient to such attack methods.

Also, these reports may be forged as easily as ordinary Internet electronic mail. User agents and automatic mail handling facilities (such as mail distribution list exploders) that wish to make automatic use of reports of any kind should take appropriate precautions to minimize the potential damage from denial-of-service attacks.

Perhaps the simplest means of mitigating this threat is to assert that these reports should themselves be signed with something like DKIM and/or authorized by something like SPF. Note, however, that if there is a problem with the email infrastructure at either end, DKIM and/or SPF may result in reports that aren’t trusted or even accepted by their intended recipients, so it is important to make sure those components are properly configured. The use of both technologies in tandem can resolve this concern to a degree, since they generally have disjoint failure modes.

7.3. Amplification Attacks

Failure to comply with the recommendations regarding selection of the envelope sender can lead to amplification denial-of-service attacks. This is discussed in Section 6 as well as in [RFC3464].

7.4. Automatic Generation

ARF [RFC5965] reports have historically been generated individually as a result of some kind of human request, such as someone clicking a "Report Abuse" button in a mail reader. In contrast, the mechanisms described in some extension documents (i.e., [RFC6651] and [RFC6652]) are focused around automated reporting. This obviously implies the

potential for much larger volumes or higher frequency of messages, and thus greater mail system load (both for Feedback Providers and report receivers).

Those mechanisms are primarily intended for use in generating reports to aid implementers of DKIM [RFC6376], Author Domain Signing Practices (ADSP) [RFC5617], and SPF [RFC4408], and other related protocols during development and debugging. They are not generally intended for prolonged forensic use, specifically because of these load concerns. However, extended use is possible by ADministrative Management Domains (ADMDs) that want to keep a close watch for fraud or infrastructure problems. It is important to consider the impact of doing so on both Feedback Providers and the requesting ADMDs.

A sender requesting these reports can cause its mail servers to be overwhelmed if it sends out signed messages whose signatures fail to verify for some reason, provoking a large number of reports from Feedback Providers. Similarly, a Feedback Provider could be overwhelmed by a large volume of messages requesting reports whose signatures fail to validate, as the Feedback Provider now needs to send reports back to the Signer.

Limiting the rate of generation of these messages may be appropriate but threatens to inhibit the distribution of important and possibly time-sensitive information.

In general ARF feedback loop terms, it is often suggested that Feedback Providers only create these (or any) ARF reports after an out-of-band arrangement has been made between two parties. These extension mechanisms provide ways to adjust parameters of an authorized abuse report feedback loop that is configured and activated by private agreement. The alternative (sending reports automatically based solely on data found in the messages) may have unintended consequences.

7.5. Reporting Multiple Incidents

If it is known that a particular host generates abuse reports upon certain incidents, an attacker could forge a high volume of messages that will trigger such a report. The recipient of the report could then be inundated with reports. This could easily be extended to a distributed denial-of-service attack by finding a number of report-generating servers.

The incident count referenced in ARF [RFC5965] provides a limited form of mitigation. The host that generates reports can elect to send reports only periodically, with each report representing a number of identical or nearly identical incidents. One might even do

something inverse-exponentially, sending reports for each of the first ten incidents, then every tenth incident up to 100, then every 100th incident up to 1000, etc., until some period of relative quiet after which the limitation resets.

The use of this technique for "nearly identical" incidents in particular causes a degradation in reporting quality, however. If for example a large number of pieces of spam arrive from one attacker, a reporting agent could decide only to send a report about a fraction of those messages. While this averts a flood of reports to a system administrator, the precise details of each incident are similarly not sent.

Other rate-limiting provisions might be considered, such as detecting a temporary failure response from the report destination and thus halting report generation to that destination for some period, or simply imposing or negotiating a hard limit on the number of reports to be sent to a particular receiver in a given time frame.

8. Acknowledgements

The author and editor wish to thank Steve Atkins, John Levine, Shmuel Metz, S. Moonesamy, and Alessandro Vesely for their contributions to this memo.

All of the best practices referenced by this document are found in [RFC6449], written within the Collaboration Committee of the Messaging Anti-Abuse Working Group (MAAWG).

Finally, the original author wishes to thank the doctors and staff at the University of Texas MD Anderson Cancer Center for doing what they do.

9. References

9.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC5321] Klensin, J., "Simple Mail Transfer Protocol", RFC 5321, October 2008.

[RFC5322] Resnick, P., Ed., "Internet Message Format", RFC 5322, October 2008.

[RFC5598] Crocker, D., "Internet Mail Architecture", RFC 5598, July 2009.

[RFC5965] Shafranovich, Y., Levine, J., and M. Kucherawy, "An Extensible Format for Email Feedback Reports", RFC 5965, August 2010.

[RFC6591] Fontana, H., "Authentication Failure Reporting Using the Abuse Reporting Format", RFC 6591, April 2012.

9.2. Informative References

[RFC1939] Myers, J. and M. Rose, "Post Office Protocol - Version 3", STD 53, RFC 1939, May 1996.

[RFC2026] Bradner, S., "The Internet Standards Process -- Revision 3", BCP 9, RFC 2026, October 1996.

[RFC2142] Crocker, D., "Mailbox Names for Common Services, Roles and Functions", RFC 2142, May 1997.

[RFC2616] 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.

[RFC3464] Moore, K. and G. Vaudreuil, "An Extensible Message Format for Delivery Status Notifications", RFC 3464, January 2003.

[RFC3501] Crispin, M., "INTERNET MESSAGE ACCESS PROTOCOL - VERSION 4rev1", RFC 3501, March 2003.

[RFC3912] Daigle, L., "WHOIS Protocol Specification", RFC 3912, September 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.

[RFC5617] Allman, E., Fenton, J., Delany, M., and J. Levine, "DomainKeys Identified Mail (DKIM) Author Domain Signing Practices (ADSP)", RFC 5617, August 2009.

[RFC6376] Crocker, D., Ed., Hansen, T., Ed., and M. Kucherawy, Ed., "DomainKeys Identified Mail (DKIM) Signatures", RFC 6376, September 2011.

[RFC6430] Li, K. and B. Leiba, "Email Feedback Report Type Value: not-spam", RFC 6430, November 2011.

[RFC6449] Falk, J., Ed., "Complaint Feedback Loop Operational Recommendations", RFC 6449, November 2011.

[RFC6590] Falk, J., Ed., and M. Kucherawy, Ed., "Redaction of Potentially Sensitive Data from Mail Abuse Reports", RFC 6590, April 2012.

[RFC6651] Kucherawy, M., "Extensions to DomainKeys Identified Mail (DKIM) for Failure Reporting", RFC 6651, June 2012.

[RFC6652] Kitterman, S., "Sender Policy Framework (SPF) Authentication Failure Reporting Using the Abuse Reporting Format", RFC 6652, June 2012.

Authors' Addresses

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

URI: http://www.returnpath.net/

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

EMail: superuser@gmail.com