MARF Working Group
Internet-Draft
Intended status: Standards Track
Expires: March 12, 2012

H. Fontana
eCert Inc.

September 9, 2011

Authentication Failure Reporting using the Abuse Report Format
draft-ietf-marf-authfailure-report-02

Abstract

This memo registers an extension report type to ARF to be used for reporting forensic information about messages that fail one or more message authentication schemes in use by the purported sender of the message.

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 March 12, 2012.

Copyright Notice

Table of Contents

1. Introduction

[ARF] defines a message format for sending reports of abuse in the messaging infrastructure, with an eye toward automating both the generating and consumption of those reports. This memo presents extensions to the Abuse Reporting Format (ARF) to allow for detailed reporting of message authentication failures.

2. Definitions

2.1. Keywords

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].

2.2. Imported Definitions

The ABNF token "qp-section" is imported from [MIME].

base64 is defined in [MIME].

3. Extension ARF Fields for Authentication Failure Reporting

The current report format defined in [ARF] lacks some specific features required to do effective sender authentication reporting. This section defines extensions to ARF to accommodate this requirement.

3.1. New ARF Feedback Type

A new feedback type of "auth-failure" is defined as an extension to Section 8.2 of [ARF]. See Section 3.3 for details.

A message that uses this feedback type has the following modified header field requirements for the second (machine-parseable) MIME part of the report:

Authentication-Results: This field MUST be formatted as defined in [AUTH-RESULTS], except that it MUST include explicit results for both DKIM and SPF. This field MUST appear at least once, and it is RECOMMENDED that the corresponding header fields be copied directly from the message about which a report is being generated.

Original-Envelope-Id: As specified in [ARF]. This field MUST appear exactly once.

Original-Mail-From: As specified in [ARF]. This field MUST appear exactly once.

Arrival-Date: As specified in [ARF]. This field MUST appear exactly once.

Source-IP: As specified in [ARF]. This field MUST appear exactly once. If this information is either not available at the time the report is generated, or the generating ADMD’s policy requires it be redacted, a value of 0.0.0.0 MUST be used.

Message-ID: As specified in [ARF]. This field MUST appear exactly once.

Reported-Domain: As specified in [ARF]. This field MUST appear exactly once.

Delivery-Result: As specified in Section 3.2.1. This field MUST NOT appear more than once. It SHOULD indicate the outcome of the message in some meaningful way, but might be redacted to ’other’ for local policy reasons.

The third MIME part of the message is either of type "message/rfc822"

(as defined in [MIME-TYPES]) or "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]).

For privacy reasons, report generators might need to redact portions of a reported message such as the end user whose complaint action resulted in the report. See Section 5 for a discussion of this.

3.2. New ARF Header Field Names

The following new ARF field names are defined as extensions to Section 3.1 of [ARF].

The values that are base64 encodings may contain FWS for formatting purposes as per the usual header field wrapping defined in [MAIL]. During decoding, any characters not in the base64 alphabet are ignored so that such line wrapping does not harm the value. The ABNF token "FWS" is defined in [DKIM].

3.2.1. Required For All Reports

Auth-Failure: Indicates the type of authentication failure that is being reported. The list of valid values is enumerated below.

Delivery-Result: The final message disposition that was enacted by the ADMD generating the report. Possible values are:

delivered: The message was delivered (not specific as to where).

spam: The message was delivered to the recipient’s spam folder (or equivalent).

policy: The message was not delivered to the intended inbox due to authentication failure. The specific action taken is not specified.

reject: The message was rejected.

other: The message had a final disposition not covered by one of the above values.

3.2.2. Required For DKIM Reports

DKIM-Canonicalized-Header: A base64 encoding of the canonicalized header of the message as generated by the verifier.

DKIM-Canonicalized-Body: A base64 encoding of the canonicalized body of the message as generated by the verifier.

DKIM-Domain: The domain that signed the message, taken from the "d=" tag of the signature.

DKIM-Identity: The identity of the signature that failed verification, taken from the "i=" tag of the signature.

DKIM-Selector: The selector of the signature that failed verification, taken from the "s=" tag of the signature.

3.3. Authentication Failure Types

The list of defined authentication failure types, used in the "Auth-Failure:" header field (defined above), is as follows:

adsp: The message did not conform to the sender’s published [ADSP] signing practises. The DKIM-ADSP-DNS field MUST be included in the report.

bodyhash: The body hash in the signature and the body hash computed by the verifier did not match. The DKIM-Canonicalized-Body field SHOULD be included in the report.

granularity: The DKIM key referenced by the signature on the message was not authorized for use by the sender. The DKIM-Domain and DKIM-Selector fields MUST be included in the report, and the DKIM-Identity field SHOULD be included.

revoked: The DKIM key referenced by the signature on the message has been revoked. The DKIM-Domain and DKIM-Selector fields MUST be included in the report.

signature: The DKIM signature on the message did not successfully verify against the header hash and public key. The DKIM-Domain, DKIM-Selector and DKIM-Canonicalized-Header fields MUST be included in the report.

spf: The evaluation of the sending domain’s SPF record produced a "fail", "softfail", "temperror" or "permerror" result.

Supplementary data MAY be included in the form of [MAIL]-compliant comments. For example, "Auth-Failure: adsp" could be augmented by a comment to indicate that the failed message was rejected because it was not signed when it should have been. See Appendix B for examples.

4. Syntax For Added ARF Header Fields

The ABNF definitions for the new fields are as follows:

```
auth-failure = "Auth-Failure:" [CFWS] token [CFWS] CRLF
  ; "token" must be a registered authentication failure type
  ; as specified elsewhere in this memo

delivery-result = "Delivery-Result:" [CFWS]
                  ( "delivered" / "spam" /"policy" /
                    "reject" / "other" ) [CFWS] CRLF

dkim-header = "DKIM-Canonicalized-Header:" [CFWS]
              base64string CRLF
  ; "base64string" is imported from [DKIM]

dkim-domain = "DKIM-Domain:" [CFWS] domain [CFWS] CRLF

dkim-identity = "DKIM-Identity:" [CFWS] [ local-part ] "@"
                domain-name [CFWS] CRLF
  ; "local-part" is imported from [MAIL]

dkim-selector = "DKIM-Selector:" [CFWS] token [CFWS] CRLF

dkim-adsp-dns = "DKIM-ADSP-DNS:" [CFWS]
                quoted-string [CFWS] CRLF
  ; "quoted-string" is imported from [MAIL]

dkim-body = "DKIM-Canonicalized-Body:" [CFWS]
            base64string CRLF

dkim-selector-dns = "DKIM-Selector-DNS:" [CFWS]
                    quoted-string [CFWS] CRLF

spf-dns = "SPF-DNS:" [CFWS] quoted-string [CFWS] CRLF
```

5. Redacting Data

For privacy considerations it might be the policy of a report generator to redact, or obscure, portions of the report that might identify an end user that caused the report to be generated. Precisely how this is done is unspecified in [ARF] as it will generally be a matter of local policy. That specification does admonish generators against being overly zealous with this practice, as obscuring too much data makes the report inactionable.

Previous redaction practices, such as replacing local-parts of addresses with a uniform string like "xxxxxxxx", often frustrated any kind of prioritizing or grouping of reports.

Generally, it is assumed that the recipient fields of a message (i.e. those containing recipient addresses), when copied into a report, are to be obscured to protect the identify of an end user that submitted a complaint about a message. However, it is also presumed that other data will be left intact, data that could be correlated against logs to determine the source of the message that drew a complaint.

See [I-D.IETF-MARF-REDACTION] for further details.

6. IANA Considerations

As required by [IANA-CONSIDERATIONS], this section contains registry information for the new tag, and the extension to [ARF].

6.1. Updates to ARF Feedback Types

The following feedback type is added to the Feedback Report Feedback Type Registry:

Feedback Type: auth-failure
Description: sender authentication failure report
Registration: (this document)

6.2. Updates to ARF Header Field Names

The following headers are added to the Feedback Report Header Names Registry:

Field Name: Auth-Failure
Description: Type of authentication failure
Multiple Appearances: No
Related "Feedback-Type": auth-failure

Field Name: Delivery-Result
Description: Final disposition of the subject message
Multiple Appearances: No
Related "Feedback-Type": auth-failure

Field Name: DKIM-ADSP-DNS
Description: Retrieved DKIM ADSP record
Multiple Appearances: No
Related "Feedback-Type": auth-failure

Field Name: DKIM-Canonicalized-Body
Description: Canonicalized body, per DKIM
Multiple Appearances: No
Related "Feedback-Type": auth-failure

Field Name: DKIM-Canonicalized-Header
Description: Canonicalized header, per DKIM
Multiple Appearances: No
Related "Feedback-Type": auth-failure

Field Name: DKIM-Domain
Description: DKIM signing domain from "d=" tag
Multiple Appearances: No
Related "Feedback-Type": auth-failure

Field Name: DKIM-Identity
Description: Identity from DKIM signature
Multiple Appearances: No
Related "Feedback-Type": auth-failure

Field Name: DKIM-Selector
Description: Selector from DKIM signature
Multiple Appearances: No
Related "Feedback-Type": auth-failure

Field Name: DKIM-Selector-DNS
Description: Retrieved DKIM key record
Multiple Appearances: No
Related "Feedback-Type": auth-failure

Field Name: SPF-DNS
Description: Retrieved SPF record
Multiple Appearances: No
Related "Feedback-Type": auth-failure

7. Security Considerations

Security issues with respect to these reports are similar to those found in [DSN].

7.1. Inherited Considerations

Implementors are advised to consider the Security Considerations sections of [DKIM], [ADSP] [SPF] and [ARF].

7.2. Forgeries

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 DSNs of any kind should take appropriate precautions to minimize the potential damage from denial-of-service attacks.

Security threats related to forged DSNs include the sending of:

a. A falsified authentication failure notification when the message was in fact delivered to the indicated recipient;

b. Falsified signature information, such as selector, domain, etc.

Perhaps the simplest means of mitigating this threat is to assert that these reports should themselves be signed with something like DKIM. On the other hand, if there’s a problem with the DKIM infrastructure at the verifier, signing DKIM failure reports may produce reports that aren’t trusted or even accepted by their intended recipients.

7.3. Automatic Generation

Automatic generation of these reports by verifying agents can cause a denial-of-service attack when a large volume of e-mail is sent that causes sender authentication failures for whatever reason.

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 suggested that report generators only create these (or any) ARF reports after an out-of-band arrangement has been made between two parties. This mechanism then becomes a way to adjust parameters of an authorized abuse report feedback loop that is configured and activated by private agreement rather than starting to send them automatically based solely on

discovered data in the DNS.

7.4. Envelope Sender Selection

In the case of transmitted reports in the form of a new message, it is necessary to construct the message so as to avoid amplification attacks, deliberate or otherwise. Thus, per Section 2 of [DSN], the envelope sender address of the report SHOULD be chosen to ensure that no delivery status reports will be issued in response to the report itself, and MUST be chosen so that these reports will not generate mail loops. Whenever an [SMTP] transaction is used to send a report, the MAIL FROM command MUST use a NULL return address, i.e. "MAIL FROM:<>".

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 innundated 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] provides a limited form of mitigation. The host generating reports may elect to send reports only periodically, with each report representing a number of identical or near-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 for "near-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 may 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.

8. References

8.1. Normative References

[ADSP] Allman, E., Delany, M., Fenton, J., and J. Levine, "DKIM Sender Signing Practises", RFC 5617, August 2009.

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

[AUTH-RESULTS] Kucherawy, M., "Message Header Field for Indicating Message Authentication Status", RFC 5451, April 2009.

[DKIM] Allman, E., Callas, J., Delany, M., Libbey, M., Fenton, J., and M. Thomas, "DomainKeys Identified Mail (DKIM) Signatures", RFC 4871, May 2007.

[I-D.IETF-MARF-REDACTION] Falk, JD., "Redaction of Potentially Sensitive Data from Mail Abuse Reports", I-D draft-ietf-marf-redaction, March 2011.

[IANA-CONSIDERATIONS] Alvestrand, H. and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", RFC 5226, May 2008.

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

[MAIL] Resnick, P., "Internet Message Format", RFC 2822, April 2001.

[MIME] Freed, N. and N. Borenstein, "Multipurpose Internet Mail Extensions (MIME) Part One: Format of Internet Message Bodies", RFC 2045, November 1996.

[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.

[SPF] Wong, M. and W. Schlitt, "Sender Policy Framework (SPF) for Authorizing Use of Domains in E-Mail, Version 1", RFC 4408, April 2006.

8.2. Informative References

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

Appendix A. Acknowledgements

The authors wish to acknowledge the following for their review and constructive criticism of this proposal: Mike Markley, Kelly Wanser and Murray Kucherawy.

Appendix B. Examples

This section contains examples of the use of each the extension defined by this memo.

B.1. Example Use of ARF Extension Headers

An ARF-formatted report using some of the proposed ARF extension fields:

```
From: arf-daemon@example.com
To: recipient@example.net
Subject: This is a test
Date: Wed, 14 Apr 2010 12:17:45 -0700 (PDT)
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 Wed, 14 Apr 2010 12:15:31 PDT. For more
information about this format please see
http://www.mipassoc.org/arf/.

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

Feedback-Type: auth-failure
User-Agent: SomeDKIMFilter/1.0
Version: 1.0
Original-Mail-From: <randomuser@example.net>
Original-Rcpt-To: <user@example.com>
Received-Date: Wed, 14 Apr 2010 12:15:31 -0700 (PDT)
Source-IP: 192.0.2.1
Authentication-Results: mail.example.com; dkim=fail
    header.d=example.net
Reported-Domain: example.net
DKIM-Domain: example.net
DKIM-Failure: bodyhash

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

DKIM-Signature: v=1; c=relaxed/simple; a=rsa-sha256;
    s=testkey; d=example.net; h=From:To:Subject:Date;
```

```
          bh=2jUSOH9NhtVGCQWNr9BrIAPreKQjO6Sn7XIkfJVOzv8=;
          b=AuUoFEfDxTDkHlLXSZEpZj79LICEps6eda7W3deTVFOk4yAUoqOB
           4nujc7YopdG5dWLSdNg6xNAZpOPr+kHxt1IrE+NahM6L/LbvaHut
           KVdkLLkpVaVVQPzeRDI009SO2Il5Lu7rDNH6mZckBdrIx0orEtZV
           4bmp/YzhwvcubU4=
      Received: from smtp-out.example.net by mail.example.com
          with SMTP id o3F52gxO029144;
          Wed, 14 Apr 2010 12:15:31 -0700 (PDT)
      Received: from internal-client-001.example.com
          by mail.example.com
          with SMTP id o3F3BwdY028431;
          Wed, 14 Apr 2010 12:12:09 -0700 (PDT)
      From: randomuser@example.net
      To: user@example.com
      Date: Wed, 14 Apr 2010 12:12:09 -0700 (PDT)
      Subject: This is a test

      Hi, just making sure DKIM is working!

      --part1_13d.2e68ed54_boundary--
```

Example 3: Example ARF report using these extensions

This example ARF message is making the following assertion:

- DKIM verification of the signature added within "example.net" failed when it was processed on arrival at "mail.example.com".

- The cause for the verification failure was a mismatch between the body contents observed at the verifier and the body hash contained in the signature.

Author's Address

Hilda L. Fontana
eCert Inc.
One Market Street Suite 3600
San Francisco, CA 94107
US

Phone: +1 626 676 8852
Email: hfontana@ecertsystems.com