

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: 18 September 2026

A. Brotman
Comcast, Inc
T. Corbett
Iterable
E. Gustafsson
Google
17 March 2026

Aggregate Performance Reporting
draft-brotman-aggregate-performance-reporting-00

Abstract

Definition of an aggregate performance report format for email messaging, the means to discover target destinations, and a specified delivery method.

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 <https://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 18 September 2026.

Copyright Notice

Copyright (c) 2026 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 (<https://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 Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	2
2. Terminology	3
3. Glossary	3
4. Destination Discovery via DNS Record	3
4.1. Record Attributes	4
4.1.1. DNS Record ABNF	4
4.2. Signer-Defined Identifiers (SDI)	4
4.2.1. Note about Usage of SDI	5
4.3. DNS Record Samples	5
5. Report Format & Contents	5
5.1. Report Time Period	5
5.2. Document Contents	5
5.2.1. Header	5
5.2.2. Body	7
5.3. Empty SDI	7
6. Classification	7
7. Engagement	8
8. Report Delivery	8
8.1. Compression	9
9. Destination Validation	9
10. Security Considerations	9
10.1. Data Leakage	9
10.1.1. Personal Data	10
10.1.2. Report Data	10
11. IANA Considerations	10
12. Appendix	10
12.1. Definition of "Performance"	10
12.2. Report Samples	10
12.2.1. Report Sample 1	10
12.2.2. Report Sample 2	12
13. Normative References	12
14. Informative References	13
Authors' Addresses	13

1. Introduction

In the Email industry, there is a great amount of emphasis put upon various types of reputation. Mailbox Providers (MBP) typically use these reputation systems to govern placement, permitted volume, or other characteristics related to delivery. In many cases, the entities sending the messages have little insight into these reputation systems that are directly impacted by their actions. That lack of usable information directly impacts their ability to take future corrective actions.

A report which contains some metrics impacting reputation could allow these sending entities to have more insight into how their actions impact relevant metrics.

Proposed below will be a document format, as well as methods for report destination discovery and delivery methods. The reporting data contained within will focus on domain-based, specifically DKIM, metrics.

2. Terminology

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

3. Glossary

DKIM - DomainKeys Identified Mail MBP - Mailbox Provider RUA - The destination for reports, a list of email addresses SDI - Signer-Defined Identifier

4. Destination Discovery via DNS Record

In order to discover where the reports will be delivered, the report generator will perform a DNS lookup against the domain used with the valid DKIM [RFC6376] signatures. The lookup will also leverage the selector attached to the signature. An example signature might be:

```
DKIM-Signature: v=1;a=rsa-sha256;c=relaxed/  
relaxed;s=sell;d=foo.example.org;...
```

And the resulting TXT record lookup would be:

```
sell._aprf._domainkey.foo.example.org
```

Effectively: <s value>._aprf._domainkey.<d value>

There is the option to use a wildcard [RFC1034] to the left of the '_aprf' label. This would use the same record for all selectors, unless specifically stated in the DNS system.

A signing entity can opt to mix wildcard and explicit selector definitions. As defined with DNS, the explicit definition gets precedence over the wildcard result. In the absence of an explicit selector-based record, the wildcard record would then be used.

4.1. Record Attributes

v: This MUST always exist, and the value must always be "APRFv1". Failure to do so should be treated as an invalid record, and the record MUST be ignored.

rua: This is the destination address for the report data. The value must be prefaced with a "mailto:", and then include a valid 5321 address as the destination. If there is more than one destination, they should be separated with a comma, and each should have its own "mailto:". Absence of this attribute suggests the signing entity has no interest in receiving reports, and the record MUST be ignored.

sdi: An optional attribute which helps segment the data. The contents are a header and separator character, separated by a comma. Defined in the section below. The separator character MUST NOT be any of ';', '=', or ','. If both the header name and the separator are not defined, the 'sdi' declaration MUST be ignored.

4.1.1. DNS Record ABNF

TODO

4.2. Signer-Defined Identifiers (SDI)

There is an attribute ('sdi') by which the DKIM domain holder can share a header name which can help segment the data contained within the report.

The attribute is defined in two parts, separated by a comma (','). The parts MUST be the [RFC5322] Header Name, and then a single character separator. The separator MUST be a printable ASCII [RFC20] character, and MUST NOT be ';', '=', or ','.

The header field MUST be DKIM signed by the related signature.

The information contained within is expected to be increasing in granularity as it moves to the right, with a maximum of four fields. There MUST be a limit of four parts within the header.

Example:

```
sell._aprf._domainkey.example.org TXT "v=ARPFv1;...;sdi=Signer-Info,^(")
```

Where the header name is "Signer-Info", and the separator is '^('.)

An example header:

Signer-

Info:SenderCommonName^BrandNameRegionalDistinction^CampaignName

Use of this by the report generator is optional. The report generator MAY ignore this attribute.

4.2.1. Note about Usage of SDI

It's quite possible that two signing entities could attempt to use the same header field. This is not explicitly forbidden, however, if each entity uses a different separator string, the results will likely be undesirable.

It is very much suggested that signers choose a header name that is likely unique to their entity.

Additionally, a signer should take caution when creating segment names. These segment names may create a data leakage. Those names could appear in reports.

4.3. DNS Record Samples

```
v=APRFv1;rua=mailto:reports@example.org (mailto:reports@example.org);
```

```
v=APRFv1;rua=mailto:reports@example.org,mailto:reports2@example2.net  
(mailto:reports@example.org,mailto:reports2@example2.net);
```

```
v=APRFv1;rua=mailto:reports@example.org;sdi=MsgInfo,^  
(mailto:reports@example.org;sdi=MsgInfo,^);
```

If one of the destinations does not align with the sending RFC5322.From domain, there should be destination validation, discussed below.

5. Report Format & Contents

The report format MUST be valid JSON [RFC8259].

5.1. Report Time Period

The report MUST contain data for one UTC day. It MUST begin at 0000UTC, and MUST end at 2359.59UTC on that day.

5.2. Document Contents

5.2.1. Header

The "header" portion of the report will include data about the entity creating the report. The fields will be:

version: This is a string provided by the report generator. This allows the MBP to notate when there is a deviation from previous versions as it relates to the data contained within. An example might be that the MBP adds or removes some action from the "positive" feedback field, and therefore alters the version string to illustrate the demarcation.

source: The common name of the reporting entity. If a company is reporting on behalf of a MBP, that MBP name should be in this field.

dkim_domain: The domain which is being reported on. This should be from the d= field of the signature. If the report is going to roll up data to the higher level domain, the field should show this as an asterisk (e.g., "*.example.org"). The domain may not be aligned with the 5322.From domain as the report is based on the signing domain.

dkim_selector: The selector which is being reported on. If this is an aggregate report that is rolling up data, the reporting entity should make this value an asterisk ("*").

report_start: The time where the data within the report begins, noted in the epoch format.

report_end: The time where the data within the report ends, noted in the epoch format.

contact_info: An email address that may be used to contact the report generator with questions. This field is optional.

sdi_used: Demonstrate which SDI was used to create the report data. The string MUST be in the same format retrieved from DNS. If the report generator is not using the declared 'sdi', this value MUST be "N/A". If the value is not in the record, or is not valid, the value here MUST be "N/F".

extra_info: Optional field. This could be used by the report generator to provide additional information for the report recipient. This could include information such as how to better interpret the reports, contact information, information about data points that may influence reputation.

5.2.2. Body

The second portion is called the "Body", and will include the data. Within the body, there are a list of segments. Each segment contains a "classification", and "engagement" part in addition to a sender defined identifier, if applicable. The "classification" section is meant to disclose placement information, while "engagement" is meant to disclose what happens after the message is classified, and include "positive", "negative", and "neutral" data points.

Each of the "classification" and "engagement" portions of the report are optional, and the decision to include that data will be made by the report generator. It may be that the report generator only wishes to share some portion of the data, or the data is simply not available.

Sample default segment (no signer defined identifier):

```
... { "classification": { ... }, "engagement": { ... } }, ...
```

Sample segment with signer defined identifier:

```
... { "segment": "ExampleCustomerID", "classification": { ... },  
  "engagement": { ... } }, ...
```

5.3. Empty SDI

In the event of an absent SDI, either by the Report Generator or the Signer, then the report will essentially be "flat". There will be just one stanza, without a reference a segment. A sample is provided below.

NOTE:specify which sample below, currently sample #2

6. Classification

Each segment disclosed in the report should share information about placement of the message. This could include values such as: inbox, unwanted, forwarded, promotional, or some other placement information. These are scalar values, and recommended to be created as "buckets". While the recommendation is to use buckets, the decision to do so is left to the report generator. Additionally, when buckets are used during reporting, it's suggested that the report generator will use the upper bound of the bucket for values. The size (10, 10k, etc) of the buckets will be a decision made by the report generator.

These metrics should pertain to the reporting period, and measure the number of messages in each category that were received during that time.

Sample segment:

```
... "classification": { "inbox": 10000, "unwanted": 500 }, ...
```

The report generator MAY provide information about these categories in the "extra_info" header.

7. Engagement

This segment is to disclose data around user engagement, and how that may impact the reputation.

These metrics should pertain to the reporting period, and measure the number of engagement actions in each category, regardless of when the messages were received. The values are meant to be bucketed scalar values, similar to the metrics in the prior section.

Sample segment:

```
... "engagement": { "positive": 100, "neutral": 500, "negative": 50 }, ...
```

The report generator MAY provide information about these categories in the "extra_info" header.

8. Report Delivery

The reports must be attached to the messages as a standard attachment.

When using the date below, the date used should be the ending date of the report. For a single UTC day report, the ending date should be that date.

The attachment name MUST be in the form:

```
<yyyymmdd>_<DKIM domain>_<DKIM selector>_<source>
```

If the "source" has any spaces, those should be removed.

The suffix for the file MUST be ".json".

The subject MUST be in the form:

ARPF: <yyyymmdd>_<DKIM domain>_<DKIM selector>_<source>

If the "source" has any spaces, those should be removed.

The Content-Type for the message MUST be "multipart/report". The Content-Type for the report attachment MUST be "application/json". If there is a plain-text portion of the report, the Content-Type MUST be "text/plain".

Reports will be delivered via SMTP to the destination address.

8.1. Compression

A report receiver SHOULD be able to receive messages that are compressed using gzip [RFC1952], as a report generator MAY opt to compress the attachment for the report. If the attachment is to be compressed, it MUST have the Content-Type of "application/gzip", and a file extension of ".json.gz".

NOTE: If a later version shows up with the same date period, does it overwrite, discard? Should there be a flag showing it to be an update?

9. Destination Validation

When the report destination and DKIM domain are not aligned, there is a method by which the report generator SHOULD validate the report destination before attempting delivery.

<selector>.<DKIM domain>._aprf.<report destination domain>

And the value MUST be "v=APRFv1;" or "v=APRFv1"

The <selector> or <DKIM domain> may be a wildcard entry. This would allow all portions to the left to go to that destination.

This is similar to the validation performed for DMARC [RFC7489].

TODO: proper ABNF

10. Security Considerations

10.1. Data Leakage

10.1.1. Personal Data

At a sufficient low volume, it may be possible to expose some information that would otherwise be "lost in the noise" about individuals. One suggestion may be that the MBP that is supplying the reports set a threshold for a number of messages or distinct recipients in order to better obfuscate that information.

10.1.2. Report Data

Report generators may not wish to send reports to all of the destinations requested. The decision on destinations could be related to 5322.From alignment, domain reputation, or other considerations. A report generator is not required to send reports to every entity requesting these reports.

11. IANA Considerations

12. Appendix

12.1. Definition of "Performance"

This report is meant to allow the MBP/report generator to share information about metrics relating to performance and reputation. The metrics that are noted above are positive/neutral/negative. As one may have noticed, there is no formal definition of these metrics or what they may mean. Each MBP/report generator may have different actions that belong in each of these categories.

Some examples could be:

positive: open or click engagement, False positive ("not-spam")
report neutral: filing message into a folder, forwarding
negative: marking as spam, deletion without opening, unsubscribe

These could be considered "secret sauce", so it is likely not advantageous for the MBP to disclose precisely what each category consists of.

A report generator MAY choose to divulge some or all of this information via the "extra_info" field in the report header.

12.2. Report Samples

12.2.1. Report Sample 1

```
[
  {
    "header": {
      "version": 42,
      "source": "Receiver MBP, Inc.",
      "dkim_domain": "example.com",
      "dkim_selector": "selector1",
      "report_start": 1709164800,
      "report_end": 1709251199,
      "contact_info": "reports@mbp.net",
      "extra_info": "TBD",
      "sdi_used": "UniqueHeaderName,^"
    },
    "body": [
      {
        "segment": ["Seg1", "Seg2", "Seg3"],
        "classification": {
          "inbox": 10000,
          "unwanted": 500
        },
        "engagement": {
          "positive": 300,
          "negative": 200,
          "neutral": 50
        }
      },
      {
        "segment": ["Seg1", "Seg2", "Other"],
        "classification": {
          "inbox": 200,
          "unwanted": 50
        },
        "engagement": {
          "positive": 50,
          "negative": 20,
          "neutral": 0
        }
      },
      {
        "segment": ["Seg1", "Other"],
        "classification": {
          "inbox": 50,
          "unwanted": 50
        }
      }
    ]
  }
]
```

```
    },
    "engagement":
    {
        "positive": 50,
        "negative": 50,
        "neutral": 50
    }
    ]
}
]
```

12.2.2. Report Sample 2

```
[
{
    "header": {
        "version": 4,
        "source": "Receiver MBP, Inc.",
        "dkim_domain": "example.com",
        "dkim_selector": "sell",
        "report_start": 1709164800,
        "report_end": 1709251199,
        "contact_info": "reports@mbp.net",
        "extra_info": "TBD",
        "sdi_used": "N/F"
    },
    "body": [
    {
        "classification":
        {
            "inbox": 10000,
            "unwanted": 100
        },
        "engagement":
        {
            "positive": 200,
            "negative": 100,
            "neutral": 20
        }
    }
    ]
}
]
```

13. Normative References

- [RFC6376] Crocker, D., Ed., Hansen, T., Ed., and M. Kucherawy, Ed., "DomainKeys Identified Mail (DKIM) Signatures", STD 76, RFC 6376, DOI 10.17487/RFC6376, September 2011, <<https://www.rfc-editor.org/info/rfc6376>>.
- [RFC8259] Bray, T., Ed., "The JavaScript Object Notation (JSON) Data Interchange Format", STD 90, RFC 8259, DOI 10.17487/RFC8259, December 2017, <<https://www.rfc-editor.org/info/rfc8259>>.

14. Informative References

- [RFC1034] Mockapetris, P., "Domain names - concepts and facilities", STD 13, RFC 1034, DOI 10.17487/RFC1034, November 1987, <<https://www.rfc-editor.org/info/rfc1034>>.
- [RFC20] Cerf, V., "ASCII format for network interchange", STD 80, RFC 20, DOI 10.17487/RFC0020, October 1969, <<https://www.rfc-editor.org/info/rfc20>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC5322] Resnick, P., Ed., "Internet Message Format", RFC 5322, DOI 10.17487/RFC5322, October 2008, <<https://www.rfc-editor.org/info/rfc5322>>.
- [RFC7489] Kucherawy, M., Ed. and E. Zwicky, Ed., "Domain-based Message Authentication, Reporting, and Conformance (DMARC)", RFC 7489, DOI 10.17487/RFC7489, March 2015, <<https://www.rfc-editor.org/info/rfc7489>>.

Authors' Addresses

Alex Brotman
Comcast, Inc
Email: alex_brotman@comcast.com

Tom Corbett
Iterable
Email: tom.corbett@iterable.com

Emil Gustafsson
Google

Internet-Draft

APR

March 2026

Email: emgu@google.com