Workgroup: DMARC Working Group
Internet-Draft:
draft-ietf-dmarc-failure-reporting-00
Obsoletes: 7489 (if approved)
Published: 12 November 2020
Intended Status: Standards Track
Expires: 16 May 2021

Authors: S. M. Jones, Ed. A. Vesely, Ed.
DMARC.org Tana

Domain-based Message Authentication, Reporting, and Conformance (DMARC) Failure Reporting

Abstract

Domain-based Message Authentication, Reporting, and Conformance (DMARC) is a scalable mechanism by which a domain owner can request feedback about email messages using their domain in the From: address field. This document describes "failure reports," or "failed message reports," which provide details about individual messages that failed to authenticate according to the DMARC mechanism.

Status of This Memo

Copyright Notice

Table of Contents

1. Introduction

Domain-based Message Authentication, Reporting, and Conformance (DMARC) [RFC7489] is a scalable mechanism by which a mail-originating organization can express domain-level policies and preferences for message validation, disposition, and reporting, that a mail-receiving organization can use to improve mail handling. This document focuses on one type of reporting that can be requested under DMARC.

"Failure reports," or "failed message reports," provide diagnostic information about messages that a Mail Receiver has determined do not pass the DMARC mechanism. These reports are generally sent at the time such messages are received and evaluated, to provide the Domain Owner with timely notification that such failures are occurring, and to provide information that may assist in diagnosing the cause of the failures.

2. Terminology and Definitions

This section defines terms used in the rest of the document.

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 BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

Readers are expected to be familiar with the contents of [RFC7489], specifically the terminology and definitions section.

3. Failure Reports

Providing Domain Owners with visibility into how Mail Receivers implement and enforce the DMARC mechanism in the form of feedback is critical to establishing and maintaining accurate authentication deployments. Failure reports can supply more detailed information about messages that failed to authenticate, enabling the Domain Owner to determine exactly what might be causing those specific failures.

Failure reports are normally generated and sent almost immediately after the Mail Receiver detects a DMARC failure. Rather than waiting for an aggregate report, these reports are useful for quickly notifying the Domain Owners when there is an authentication failure. Whether the failure is due to an infrastructure problem or the message is inauthentic, failure reports also provide more information about the failed message than is available in an aggregate report.

These reports SHOULD include any URI(s) from the message that failed authentication. These reports SHOULD include as much of the message and message header as is reasonable to support the Domain Owner's investigation into what caused the message to fail authentication and track down the sender.

When a Domain Owner requests failure reports for the purpose of forensic analysis, and the Mail Receiver is willing to provide such reports, the Mail Receiver generates and sends a message using the format described in [RFC6591]; this document updates that reporting format, as described in Section 3.1.

The destination(s) and nature of the reports are defined by the "ruf" and "fo" tags as defined in ([RFC7489] general-record-format).

Where multiple URIs are selected to receive failure reports, the report generator MUST make an attempt to deliver to each of them.

An obvious consideration is the denial-of-service attack that can be perpetrated by an attacker who sends numerous messages purporting to be from the intended victim Domain Owner but that fail both SPF and DKIM; this would cause participating Mail Receivers to send failure reports to the Domain Owner or its delegate in potentially huge volumes. Accordingly, participating Mail Receivers are encouraged to aggregate these reports as much as is practical, using the Incidents field of the Abuse Reporting Format ([RFC5965]). Various aggregation techniques are possible, including the following:

*only send a report to the first recipient of multi-recipient messages;

*store reports for a period of time before sending them, allowing detection, collection, and reporting of like incidents;

*apply rate limiting, such as a maximum number of reports per minute that will be generated (and the remainder discarded).

3.1. Reporting Format Update

Operators implementing this specification also implement an augmented version of [RFC6591] as follows:

1. A DMARC failure report includes the following ARF header fields, with the indicated normative requirement levels:

   *Identity-Alignment (REQUIRED; defined below)

   *Delivery-Result (OPTIONAL)

   *DKIM-Domain, DKIM-Identity, DKIM-Selector (REQUIRED if the message was signed by DKIM)

   *DKIM-Canonicalized-Header, DKIM-Canonicalized-Body (OPTIONAL if the message was signed by DKIM)

   *SPF-DNS (REQUIRED)

2. The "Identity-Alignment" field is defined to contain a comma-separated list of authentication mechanism names that produced an aligned identity, or the keyword "none" if none did. ABNF:

```
id-align     = "Identity-Alignment:" [CFWS]
               ( "none" /
                 dmarc-method *( [CFWS] "," [CFWS] dmarc-method ) )
               [CFWS]

dmarc-method = ( "dkim" / "spf" )
               ; each may appear at most once in an id-align
```

3. Authentication Failure Type "dmarc" is defined, which is to be used when a failure report is generated because some or all of the authentication mechanisms failed to produce aligned

identifiers. Note that a failure report generator MAY also independently produce an AFRF message for any or all of the underlying authentication methods.

3.2. Verifying External Destinations

It is possible to specify destinations for the different reports that are outside the authority of the Domain Owner making the request. This allows domains that do not operate mail servers to request reports and have them go someplace that is able to receive and process them.

Without checks, this would allow a bad actor to publish a DMARC policy record that requests that reports be sent to a victim address, and then send a large volume of mail that will fail both DKIM and SPF checks to a wide variety of destinations; the victim will in turn be flooded with unwanted reports. Therefore, a verification mechanism is included.

When a Mail Receiver discovers a DMARC policy in the DNS, and the Organizational Domain at which that record was discovered is not identical to the Organizational Domain of the host part of the authority component of a [RFC3986] specified in the "rua" or "ruf" tag, the following verification steps are to be taken:

1. Extract the host portion of the authority component of the URI. Call this the "destination host", as it refers to a Report Receiver.

2. Prepend the string "_report._dmarc".

3. Prepend the domain name from which the policy was retrieved, after conversion to an A-label if needed.

4. Query the DNS for a TXT record at the constructed name. If the result of this request is a temporary DNS error of some kind (e.g., a timeout), the Mail Receiver MAY elect to temporarily fail the delivery so the verification test can be repeated later.

5. For each record returned, parse the result as a series of "tag=value" pairs, i.e., the same overall format as the policy record (see ([RFC7489] formal-definition)). In particular, the "v=DMARC1;" tag is mandatory and MUST appear first in the list. Discard any that do not pass this test.

6. If the result includes no TXT resource records that pass basic parsing, a positive determination of the external reporting relationship cannot be made; stop.

7. If at least one TXT resource record remains in the set after parsing, then the external reporting arrangement was authorized by the Report Receiver.

8. If a "rua" or "ruf" tag is thus discovered, replace the corresponding value extracted from the domain's DMARC policy record with the one found in this record. This permits the Report Receiver to override the report destination. However, to prevent loops or indirect abuse, the overriding URI MUST use the same destination host from the first step.

For example, if a DMARC policy query for "blue.example.com" contained "rua=mailto:reports@red.example.net", the host extracted from the latter ("red.example.net") does not match "blue.example.com", so this procedure is enacted. A TXT query for "blue.example.com._report._dmarc.red.example.net" is issued. If a single reply comes back containing a tag of "v=DMARC1;", then the relationship between the two is confirmed. Moreover, "red.example.net" has the opportunity to override the report destination requested by "blue.example.com" if needed.

Where the above algorithm fails to confirm that the external reporting was authorized by the Report Receiver, the URI MUST be ignored by the Mail Receiver generating the report. Further, if the confirming record includes a URI whose host is again different than the domain publishing that override, the Mail Receiver generating the report MUST NOT generate a report to either the original or the override URI.

A Report Receiver publishes such a record in its DNS if it wishes to receive reports for other domains.

A Report Receiver that is willing to receive reports for any domain can use a wildcard DNS record. For example, a TXT resource record at "*._report._dmarc.example.com" containing at least "v=DMARC1;" confirms that example.com is willing to receive DMARC reports for any domain.

If the Report Receiver is overcome by volume, it can simply remove the confirming DNS record. However, due to positive caching, the change could take as long as the time-to-live (TTL) on the record to go into effect.

A Mail Receiver might decide not to enact this procedure if, for example, it relies on a local list of domains for which external reporting addresses are permitted.

4. Privacy Considerations

This section discusses issues specific to private data that may be included in the DMARC reporting functions.

4.1. Data Exposure Considerations

Failed-message reporting provides message-specific details pertaining to authentication failures. Individual reports can contain message content as well as trace header fields. Domain Owners are able to analyze individual reports and attempt to determine root causes of authentication mechanism failures, gain insight into misconfigurations or other problems with email and network infrastructure, or inspect messages for insight into abusive practices.

These reports may expose sender and recipient identifiers (e.g., RFC5322.From addresses), and although the [RFC6591] format used for failed-message reporting supports redaction, failed-message reporting is capable of exposing the entire message to the report recipient.

Domain Owners requesting reports will receive information about mail claiming to be from them, which includes mail that was not, in fact, from them. Information about the final destination of mail where it might otherwise be obscured by intermediate systems will therefore be exposed.

When message-forwarding arrangements exist, Domain Owners requesting reports will also receive information about mail forwarded to domains that were not originally part of their messages' recipient lists. This means that destination domains previously unknown to the Domain Owner may now become visible.

Disclosure of information about the messages is being requested by the entity generating the email in the first place, i.e., the Domain Owner and not the Mail Receiver, so this may not fit squarely within existing privacy policy provisions. For some providers, failed-message reporting is viewed as a function similar to complaint reporting about spamming or phishing and is treated similarly under the privacy policy. Report generators (i.e., Mail Receivers) are encouraged to review their reporting limitations under such policies before enabling DMARC reporting.

4.2. Report Recipients

A DMARC record can specify that reports should be sent to an intermediary operating on behalf of the Domain Owner. This is done when the Domain Owner contracts with an entity to monitor mail streams for abuse and performance issues. Receipt by third parties of such data may or may not be permitted by the Mail Receiver's privacy policy, terms of use, or other similar governing document. Domain Owners and Mail Receivers should both review and understand if their own internal policies constrain the use and transmission of DMARC reporting.

Some potential exists for report recipients to perform traffic analysis, making it possible to obtain metadata about the Receiver's traffic. In addition to verifying compliance with policies, Receivers need to consider that before sending reports to a third party.

5. Security Considerations

This section discusses security issues related to DMARC reporting, and possible remediations.

5.1. Attacks on Reporting URIs

URIs published in DNS TXT records are well-understood possible targets for attack. Specifications such as [RFC1035] and [RFC2142] either expose or cause the exposure of email addresses that could be flooded by an attacker, for example; MX, NS, and other records found in the DNS advertise potential attack destinations; common DNS names such as "www" plainly identify the locations at which particular services can be found, providing destinations for targeted denial-of-service or penetration attacks.

Thus, Domain Owners will need to harden these addresses against various attacks, including but not limited to:

- high-volume denial-of-service attacks;

- deliberate construction of malformed reports intended to identify or exploit parsing or processing vulnerabilities;

- deliberate construction of reports containing false claims for the Submitter or Reported-Domain fields, including the possibility of false data from compromised but known Mail Receivers.

5.2. DNS Security

The DMARC mechanism and its underlying technologies (SPF, DKIM) depend on the security of the DNS. To reduce the risk of subversion of the DMARC mechanism due to DNS-based exploits, serious consideration should be given to the deployment of DNSSEC in parallel with the deployment of DMARC by both Domain Owners and Mail Receivers.

Publication of data using DNSSEC is relevant to Domain Owners and third-party Report Receivers. DNSSEC-aware resolution is relevant to Mail Receivers and Report Receivers.

5.3. External Reporting Addresses

To avoid abuse by bad actors, reporting addresses generally have to be inside the domains about which reports are requested. In order to accommodate special cases such as a need to get reports about domains that cannot actually receive mail, Section 3.2 describes a DNS-based mechanism for verifying approved external reporting.

The obvious consideration here is an increased DNS load against domains that are claimed as external recipients. Negative caching will mitigate this problem, but only to a limited extent, mostly dependent on the default TTL in the domain's SOA record.

Where possible, external reporting is best achieved by having the report be directed to domains that can receive mail and simply having it automatically forwarded to the desired external destination.

Note that the addresses shown in the "ruf" tag receive more information that might be considered private data, since it is possible for actual email content to appear in the failure reports. The URIs identified there are thus more attractive targets for intrusion attempts than those found in the "rua" tag. Moreover, attacking the DNS of the subject domain to cause failure data to be routed fraudulently to an attacker's systems may be an attractive prospect. Deployment of [RFC4033] is advisable if this is a concern.

The verification mechanism presented in Section 3.2 is currently not mandatory ("MUST") but strongly recommended ("SHOULD"). It is possible that it would be elevated to a "MUST" by later security review.

5.4. Secure Protocols

This document encourages use of secure transport mechanisms to prevent loss of private data to third parties that may be able to monitor such transmissions. Unencrypted mechanisms should be avoided.

In particular, a message that was originally encrypted or otherwise secured might appear in a report that is not sent securely, which could reveal private information.

6. Normative References

[RFC1035]

Mockapetris, P., "Domain names - implementation and specification", STD 13, RFC 1035, DOI 10.17487/RFC1035, November 1987, <https://www.rfc-editor.org/info/rfc1035>.

[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>.

[RFC3986] Berners-Lee, T., Fielding, R., and L. Masinter, "Uniform Resource Identifier (URI): Generic Syntax", STD 66, RFC 3986, DOI 10.17487/RFC3986, January 2005, <https://www.rfc-editor.org/info/rfc3986>.

[RFC6591] Fontana, H., "Authentication Failure Reporting Using the Abuse Reporting Format", RFC 6591, DOI 10.17487/RFC6591, April 2012, <https://www.rfc-editor.org/info/rfc6591>.

[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>.

7. Informative References

[RFC2142] Crocker, D., "Mailbox Names for Common Services, Roles and Functions", RFC 2142, DOI 10.17487/RFC2142, May 1997, <https://www.rfc-editor.org/info/rfc2142>.

[RFC4033] Arends, R., Austein, R., Larson, M., Massey, D., and S. Rose, "DNS Security Introduction and Requirements", RFC 4033, DOI 10.17487/RFC4033, March 2005, <https://www.rfc-editor.org/info/rfc4033>.

[RFC5965] Shafranovich, Y., Levine, J., and M. Kucherawy, "An Extensible Format for Email Feedback Reports", RFC 5965, DOI 10.17487/RFC5965, August 2010, <https://www.rfc-editor.org/info/rfc5965>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

Appendix A. Examples

This section presents some examples related to the use of DMARC reporting functions.

A.1. Entire Domain, Monitoring Only, Per-Message Reports

The Domain Owner from the previous example has used the aggregate reporting to discover some messaging systems that had not yet implemented DKIM correctly, but they are still seeing periodic authentication failures. In order to diagnose these intermittent problems, they wish to request per-message failure reports when authentication failures occur.

Not all Receivers will honor such a request, but the Domain Owner feels that any reports it does receive will be helpful enough to justify publishing this record. The default per-message report format ([RFC6591]) meets the Domain Owner's needs in this scenario.

The Domain Owner accomplishes this by adding the following to its policy record from ([RFC7489] domain-owner-example):

- Per-message failure reports should be sent via email to the address "auth-reports@example.com" ("ruf=mailto:auth-reports@example.com")

The DMARC policy record might look like this when retrieved using a common command-line tool (the output shown would appear on a single line but is wrapped here for publication):

```
% dig +short TXT _dmarc.example.com.
"v=DMARC1; p=none; rua=mailto:dmarc-feedback@example.com;
 ruf=mailto:auth-reports@example.com"
```

To publish such a record, the DNS administrator for the Domain Owner might create an entry like the following in the appropriate zone file (following the conventional zone file format):

```
; DMARC record for the domain example.com

_dmarc  IN TXT ( "v=DMARC1; p=none; "
                 "rua=mailto:dmarc-feedback@example.com; "
                 "ruf=mailto:auth-reports@example.com" )
```

A.2. Per-Message Failure Reports Directed to Third Party

The Domain Owner from the previous example is maintaining the same policy but now wishes to have a third party receive and process the per-message failure reports. Again, not all Receivers will honor this request, but those that do may implement additional checks to validate that the third party wishes to receive the failure reports for this domain.

The Domain Owner needs to alter its policy record from Appendix A.1 as follows:

*Per-message failure reports should be sent via email to the address "auth-reports@thirdparty.example.net" ("ruf=mailto:auth-reports@thirdparty.example.net")

The DMARC policy record might look like this when retrieved using a common command-line tool (the output shown would appear on a single line but is wrapped here for publication):

```
% dig +short TXT _dmarc.example.com.
"v=DMARC1; p=none; rua=mailto:dmarc-feedback@example.com;
 ruf=mailto:auth-reports@thirdparty.example.net"
```

To publish such a record, the DNS administrator for the Domain Owner might create an entry like the following in the appropriate zone file (following the conventional zone file format):

```
; DMARC record for the domain example.com

_dmarc IN TXT ( "v=DMARC1; p=none; "
                "rua=mailto:dmarc-feedback@example.com; "
                "ruf=mailto:auth-reports@thirdparty.example.net" )
```

Because the address used in the "ruf" tag is outside the Organizational Domain in which this record is published, conforming Receivers will implement additional checks as described in Section 3.2 of this document. In order to pass these additional checks, the third party will need to publish an additional DNS record as follows:

*Given the DMARC record published by the Domain Owner at "_dmarc.example.com", the DNS administrator for the third party will need to publish a TXT resource record at "example.com._report._dmarc.thirdparty.example.net" with the value "v=DMARC1;".

The resulting DNS record might look like this when retrieved using a common command-line tool (the output shown would appear on a single line but is wrapped here for publication):

```
% dig +short TXT example.com._report._dmarc.thirdparty.example.net
"v=DMARC1;"
```

To publish such a record, the DNS administrator for example.net might create an entry like the following in the appropriate zone file (following the conventional zone file format):

```
  ; zone file for thirdparty.example.net
  ; Accept DMARC failure reports on behalf of example.com

  example.com._report._dmarc   IN   TXT    "v=DMARC1;"
```

Intermediaries and other third parties should refer to Section 3.2 for the full details of this mechanism.

Acknowledgements

DMARC and the draft version of this document submitted to the Independent Submission Editor were the result of lengthy efforts by an informal industry consortium: DMARC.org (see http://dmarc.org). Participating companies included Agari, American Greetings, AOL, Bank of America, Cloudmark, Comcast, Facebook, Fidelity Investments, Google, JPMorgan Chase & Company, LinkedIn, Microsoft, Netease, PayPal, ReturnPath, The Trusted Domain Project, and Yahoo!. Although the contributors and supporters are too numerous to mention, notable individual contributions were made by J. Trent Adams, Michael Adkins, Monica Chew, Dave Crocker, Tim Draegen, Steve Jones, Franck Martin, Brett McDowell, and Paul Midgen. The contributors would also like to recognize the invaluable input and guidance that was provided early on by J.D. Falk.

Additional contributions within the IETF context were made by Kurt Anderson, Michael Jack Assels, Les Barstow, Anne Bennett, Jim Fenton, J. Gomez, Mike Jones, Scott Kitterman, Eliot Lear, John Levine, S. Moonesamy, Rolf Sonneveld, Henry Timmes, and Stephen J. Turnbull.

Authors' Addresses

Steven M Jones (editor)
DMARC.org

Email: smj@dmarc.org

Alessandro Vesely (editor)
Tana

Email: vesely@tana.it