

Workgroup: DMARC

Internet-Draft: draft-ietf-dmarc-dmarcbis-00

Obsoletes: [7489](#) (if approved)

Published: 11 November 2020

Intended Status: Standards Track

Expires: 15 May 2021

Authors: E. Gustafsson (ed) T. Herr (ed) J. Levine (ed)
 Google Valimail Standcore LLC

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

Abstract

This document describes the Domain-based Message Authentication, Reporting, and Conformance (DMARC) protocol.

DMARC is a scalable mechanism by which a mail-originating organization can express domain-level policies and preferences for message validation, disposition, and reporting. Mail-receiving organizations can in turn use these expressions of policies and preferences to inform their mail handling decisions should they choose to do so.

This document obsoletes RFC 7489.

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 15 May 2021.

Copyright Notice

Copyright (c) 2020 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 Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- [1. Introduction](#)
- [2. Requirements](#)
 - [2.1. High-Level Goals](#)
 - [2.2. Out of Scope](#)
 - [2.3. Scalability](#)
 - [2.4. Anti-Phishing](#)
- [3. Terminology and Definitions](#)
 - [3.1. Identifier Alignment](#)
 - [3.1.1. DKIM-Authenticated Identifiers](#)
 - [3.1.2. SPF-Authenticated Identifiers](#)
 - [3.1.3. Alignment and Extension Technologies](#)
 - [3.2. Organizational Domain](#)
- [4. Overview](#)
 - [4.1. Authentication Mechanisms](#)
 - [4.2. Key Concepts](#)
 - [4.3. Flow Diagram](#)
- [5. Use of RFC5322.From](#)
- [6. Policy](#)
 - [6.1. DMARC Policy Record](#)
 - [6.2. DMARC URIs](#)
 - [6.3. General Record Format](#)
 - [6.4. Formal Definition](#)
 - [6.5. Domain Owner Actions](#)
 - [6.6. Mail Receiver Actions](#)
 - [6.6.1. Extract Author Domain](#)
 - [6.6.2. Determine Handling Policy](#)
 - [6.6.3. Policy Discovery](#)
 - [6.6.4. Message Sampling](#)
 - [6.6.5. Store Results of DMARC Processing](#)
 - [6.7. Policy Enforcement Considerations](#)
- [7. DMARC Feedback](#)
- [8. Minimum Implementations](#)
- [9. Other Topics](#)
 - [9.1. Issues Specific to SPF](#)
 - [9.2. DNS Load and Caching](#)
 - [9.3. Rejecting Messages](#)
 - [9.4. Identifier Alignment Considerations](#)
 - [9.5. Interoperability Issues](#)
- [10. IANA Considerations](#)
 - [10.1. Authentication-Results Method Registry Update](#)

- [10.2. Authentication-Results Result Registry Update](#)
- [10.3. Feedback Report Header Fields Registry Update](#)
- [10.4. DMARC Tag Registry](#)
- [10.5. DMARC Report Format Registry](#)
- [10.6. Underscored and Globally Scoped DNS Node Names Registry](#)
- [11. Security Considerations](#)
 - [11.1. Authentication Methods](#)
 - [11.2. Attacks on Reporting URIs](#)
 - [11.3. DNS Security](#)
 - [11.4. Display Name Attacks](#)
 - [11.5. External Reporting Addresses](#)
 - [11.6. Secure Protocols](#)
- [12. Normative References](#)
- [13. Informative References](#)
- [Appendix A. Technology Considerations](#)
 - [A.1. S/MIME](#)
 - [A.2. Method Exclusion](#)
 - [A.3. Sender Header Field](#)
 - [A.4. Domain Existence Test](#)
 - [A.5. Issues with ADSP in Operation](#)
 - [A.6. Organizational Domain Discovery Issues](#)
 - [A.6.1. Public Suffix Lists](#)
- [Appendix B. Examples](#)
 - [B.1. Identifier Alignment Examples](#)
 - [B.1.1. SPF](#)
 - [B.1.2. DKIM](#)
 - [B.2. Domain Owner Example](#)
 - [B.2.1. Entire Domain, Monitoring Only](#)
 - [B.2.2. Entire Domain, Monitoring Only, Per-Message Reports](#)
 - [B.2.3. Per-Message Failure Reports Directed to Third Party](#)
 - [B.2.4. Subdomain, Sampling, and Multiple Aggregate Report URIs](#)
 - [B.3. Mail Receiver Example](#)
 - [B.4. Processing of SMTP Time](#)
 - [B.5. Utilization of Aggregate Feedback: Example](#)
 - [B.6. mailto Transport Example](#)
- [Acknowledgements](#)
- [Authors' Addresses](#)

1. Introduction

RFC EDITOR: PLEASE REMOVE THE FOLLOWING PARAGRAPH BEFORE PUBLISHING:
The source for this draft is maintained in GitHub at: <https://github.com/ietf-wg-dmarc/draft-ietf-dmarc-dmarcbis>

The Sender Policy Framework ([RFC7208]) and DomainKeys Identified Mail ([RFC6376]) protocols provide domain-level authentication, and DMARC builds on these protocols. DMARC is designed to give Administrative Management Domains (ADMDs) that originate email the ability to publish in a DNS TXT record their email authentication

policies, specify preferred handling for mail that fails authentication checks, and request reports about mail purportedly originated by the ADMD, as determined by the RFC5322.From header in the message.

As with SPF and DKIM, DMARC authentication checks result in verdicts of "pass" or "fail". A DMARC pass verdict requires not only that SPF or DKIM pass for the message in question, but also that the domain validated by the SPF or DKIM check is aligned with the domain in the RFC5322.From header. In the DMARC protocol, two domains are said to be "in alignment" if they have the same Organizational Domain (a.k.a., relaxed alignment) or they are identical (a.k.a., strict alignment).

A DMARC pass verdict asserts only that the RFC5322.From domain has been authenticated in that message; there is no explicit or implied value assertion attributed to a message that receives such a verdict. A mail-receiving organization that performs DMARC validation checks on inbound mail can choose to use the results and the preferences expressed by the originating domain for message disposition to inform its mail handling decision for that message. For messages that pass DMARC validation checks, the mail-receiving organization can be confident in applying handling based on its known history for similarly authenticated messages, whereas messages that fail such checks cannot be reliably associated with a domain with a history of sending DMARC-validated messages.

DMARC also describes a reporting framework in which mail-receiving domains can generate regular reports containing data about messages seen that claim to be from domains that publish DMARC policies, and send those reports to the ADMD as requested by its DMARC policy record.

Experience with DMARC has revealed some issues of interoperability with email in general that require due consideration before deployment, particularly with configurations that can cause mail to be rejected. These are discussed in [Section 9](#).

2. Requirements

Specification of DMARC is guided by the following high-level goals, security dependencies, detailed requirements, and items that are documented as out of scope.

2.1. High-Level Goals

DMARC has the following high-level goals:

- *Allow Domain Owners to assert the preferred handling of authentication failures, for messages purporting to have authorship within the domain.
- *Allow Domain Owners to verify their authentication deployment.
- *Minimize implementation complexity for both senders and receivers, as well as the impact on handling and delivery of legitimate messages.
- *Reduce the amount of successfully delivered spoofed email.
- *Work at Internet scale.

2.2. Out of Scope

Several topics and issues are specifically out of scope for the initial version of this work. These include the following:

- *different treatment of messages that are not authenticated versus those that fail authentication;
- *evaluation of anything other than RFC5322.From;
- *multiple reporting formats;
- *publishing policy other than via the DNS;
- *reporting or otherwise evaluating other than the last-hop IP address;
- *attacks in the RFC5322.From field, also known as "display name" attacks;
- *authentication of entities other than domains, since DMARC is built upon SPF and DKIM, which authenticate domains; and
- *content analysis.

2.3. Scalability

Scalability is a major issue for systems that need to operate in a system as widely deployed as current SMTP email. For this reason, DMARC seeks to avoid the need for third parties or pre-sending agreements between senders and receivers. This preserves the positive aspects of the current email infrastructure.

Although DMARC does not introduce third-party senders (namely external agents authorized to send on behalf of an operator) to the email-handling flow, it also does not preclude them. Such third parties are free to provide services in conjunction with DMARC.

2.4. Anti-Phishing

DMARC is designed to prevent bad actors from sending mail that claims to come from legitimate senders, particularly senders of transactional email (official mail that is about business transactions). One of the primary uses of this kind of spoofed mail is phishing (enticing users to provide information by pretending to be the legitimate service requesting the information). Thus, DMARC is significantly informed by ongoing efforts to enact large-scale, Internet-wide anti-phishing measures.

Although DMARC can only be used to combat specific forms of exact-domain spoofing directly, the DMARC mechanism has been found to be useful in the creation of reliable and defensible message streams.

DMARC does not attempt to solve all problems with spoofed or otherwise fraudulent email. In particular, it does not address the use of visually similar domain names ("cousin domains") or abuse of the RFC5322.From human-readable <display-name>.

3. 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 encouraged to be familiar with the contents of [[RFC5598](#)]. In particular, that document defines various roles in the messaging infrastructure that can appear the same or separate in various contexts. For example, a Domain Owner could, via the messaging security mechanisms on which DMARC is based, delegate the ability to send mail as the Domain Owner to a third party with another role. This document does not address the distinctions among such roles; the reader is encouraged to become familiar with that material before continuing.

The following terms are also used:

Authenticated Identifiers: Domain-level identifiers that are validated using authentication technologies are referred to as

"Authenticated Identifiers". See [Section 4.1](#) for details about the supported mechanisms.

Author Domain: The domain name of the apparent author, as extracted from the RFC5322.From field.

Domain Owner: An entity or organization that owns a DNS domain. The term "owns" here indicates that the entity or organization being referenced holds the registration of that DNS domain. Domain Owners range from complex, globally distributed organizations, to service providers working on behalf of non-technical clients, to individuals responsible for maintaining personal domains. This specification uses this term as analogous to an Administrative Management Domain as defined in [[RFC5598](#)]. It can also refer to delegates, such as Report Receivers, when those are outside of their immediate management domain.

Identifier Alignment: When the domain in the RFC5322.From address matches a domain validated by SPF or DKIM (or both), it has Identifier Alignment.

Mail Receiver: The entity or organization that receives and processes email. Mail Receivers operate one or more Internet-facing Mail Transport Agents (MTAs).

Organizational Domain: The domain that was registered with a domain name registrar. In the absence of more accurate methods, heuristics are used to determine this, since it is not always the case that the registered domain name is simply a top-level DNS domain plus one component (e.g., "example.com", where "com" is a top-level domain). The Organizational Domain is determined by applying the algorithm found in [Section 3.2](#).

Report Receiver: An operator that receives reports from another operator implementing the reporting mechanism described in this document. Such an operator might be receiving reports about its own messages, or reports about messages related to another operator. This term applies collectively to the system components that receive and process these reports and the organizations that operate them.

3.1. Identifier Alignment

Email authentication technologies authenticate various (and disparate) aspects of an individual message. For example, [[RFC6376](#)] authenticates the domain that affixed a signature to the message, while [[RFC7208](#)] can authenticate either the domain that appears in the RFC5321.MailFrom (MAIL FROM) portion of [[RFC5322](#)] or the RFC5321.EHLO/ HELO domain, or both. These may be different domains, and they are typically not visible to the end user.

DMARC authenticates use of the RFC5322.From domain by requiring that it match (be aligned with) an Authenticated Identifier. The RFC5322.From domain was selected as the central identity of the DMARC mechanism because it is a required message header field and therefore guaranteed to be present in compliant messages, and most Mail User Agents (MUAs) represent the RFC5322.From field as the originator of the message and render some or all of this header field's content to end users.

Thus, this field is the one used by end users to identify the source of the message and therefore is a prime target for abuse. Many high-profile email sources, such as email service providers, require that the sending agent have authenticated before email can be generated. Thus, for these mailboxes, the mechanism described in this document provides recipient end users with strong evidence that the message was indeed originated by the agent they associate with that mailbox, if the end user knows that these various protections have been provided.

Domain names in this context are to be compared in a case-insensitive manner, per [[RFC4343](#)].

It is important to note that Identifier Alignment cannot occur with a message that is not valid per [[RFC5322](#)], particularly one with a malformed, absent, or repeated RFC5322.From field, since in that case there is no reliable way to determine a DMARC policy that applies to the message. Accordingly, DMARC operation is predicated on the input being a valid RFC5322 message object, and handling of such non-compliant cases is outside of the scope of this specification. Further discussion of this can be found in [Section 6.6.1](#).

Each of the underlying authentication technologies that DMARC takes as input yields authenticated domains as their outputs when they succeed. From the perspective of DMARC, each can be operated in a "strict" mode or a "relaxed" mode. A Domain Owner would normally select strict mode if it wanted Mail Receivers to apply DMARC processing only to messages bearing an RFC5322.From domain exactly matching the domains those mechanisms will verify. Relaxed mode can be used when the operator also wishes to affect message flows bearing subdomains of the verified domains.

3.1.1. DKIM-Authenticated Identifiers

DMARC permits Identifier Alignment, based on the result of a DKIM authentication, to be strict or relaxed. (Note that these are not related to DKIM's "simple" and "relaxed" canonicalization modes.)

In relaxed mode, the Organizational Domains of both the [[RFC6376](#)]-authenticated signing domain (taken from the value of the "d=" tag in

the signature) and that of the RFC5322.From domain must be equal if the identifiers are to be considered aligned. In strict mode, only an exact match between both of the Fully Qualified Domain Names (FQDNs) is considered to produce Identifier Alignment.

To illustrate, in relaxed mode, if a validated DKIM signature successfully verifies with a "d=" domain of "example.com", and the RFC5322.From address is "alerts@news.example.com", the DKIM "d=" domain and the RFC5322.From domain are considered to be "in alignment". In strict mode, this test would fail, since the "d=" domain does not exactly match the FQDN of the address.

However, a DKIM signature bearing a value of "d=com" would never allow an "in alignment" result, as "com" should appear on all public suffix lists (see [Appendix A.6.1](#)) and therefore cannot be an Organizational Domain.

Identifier Alignment is required because a message can bear a valid signature from any domain, including domains used by a mailing list or even a bad actor. Therefore, merely bearing a valid signature is not enough to infer authenticity of the Author Domain.

Note that a single email can contain multiple DKIM signatures, and it is considered to be a DMARC "pass" if any DKIM signature is aligned and verifies.

3.1.2. SPF-Authenticated Identifiers

DMARC permits Identifier Alignment, based on the result of an SPF authentication, to be strict or relaxed.

In relaxed mode, the [[RFC3986](#)]-authenticated domain and RFC5322.From domain must have the same Organizational Domain. In strict mode, only an exact DNS domain match is considered to produce Identifier Alignment.

Note that the RFC5321.HELO identity is not typically used in the context of DMARC (except when required to "fake" an otherwise null reverse-path), even though a "pure SPF" implementation according to [[RFC7208](#)] would check that identifier.

For example, if a message passes an SPF check with an RFC5321.MailFrom domain of "cbg.bounces.example.com", and the address portion of the RFC5322.From field contains "payments@example.com", the Authenticated RFC5321.MailFrom domain identifier and the RFC5322.From domain are considered to be "in alignment" in relaxed mode, but not in strict mode.

3.1.3. Alignment and Extension Technologies

If in the future DMARC is extended to include the use of other authentication mechanisms, the extensions will need to allow for domain identifier extraction so that alignment with the RFC5322.From domain can be verified.

3.2. Organizational Domain

The Organizational Domain is determined using the following algorithm:

1. Acquire a "public suffix" list, i.e., a list of DNS domain names reserved for registrations. Some country Top-Level Domains (TLDs) make specific registration requirements, e.g., the United Kingdom places company registrations under ".co.uk"; other TLDs such as ".com" appear in the IANA registry of top-level DNS domains. A public suffix list is the union of all of these. [Appendix A.6.1](#) contains some discussion about obtaining a public suffix list.
2. Break the subject DNS domain name into a set of "n" ordered labels. Number these labels from right to left; e.g., for "example.com", "com" would be label 1 and "example" would be label 2.
3. Search the public suffix list for the name that matches the largest number of labels found in the subject DNS domain. Let that number be "x".
4. Construct a new DNS domain name using the name that matched from the public suffix list and prefixing to it the "x+1"th label from the subject domain. This new name is the Organizational Domain.

Thus, since "com" is an IANA-registered TLD, a subject domain of "a.b.c.d.example.com" would have an Organizational Domain of "example.com".

The process of determining a suffix is currently a heuristic one. No list is guaranteed to be accurate or current.

In addition to Mediators, mail that is sent by authorized, independent third parties might not be sent with Identifier Alignment, also preventing a "pass" result.

Issues specific to the use of policy mechanisms alongside DKIM are further discussed in [[RFC6377](#)], particularly Section 5.2.

4. Overview

This section provides a general overview of the design and operation of the DMARC environment.

4.1. Authentication Mechanisms

The following mechanisms for determining Authenticated Identifiers are supported in this version of DMARC:

*[\[RFC6376\]](#), which provides a domain-level identifier in the content of the "d=" tag of a validated DKIM-Signature header field.

*[\[RFC3986\]](#), which can authenticate both the domain found in an [\[RFC5322\]](#) HELO/EHLO command (the HELO identity) and the domain found in an SMTP MAIL command (the MAIL FROM identity). DMARC uses the result of SPF authentication of the MAIL FROM identity. Section 2.4 of [\[RFC7208\]](#) describes MAIL FROM processing for cases in which the MAIL command has a null path.

4.2. Key Concepts

DMARC policies are published by the Domain Owner, and retrieved by the Mail Receiver during the SMTP session, via the DNS.

DMARC's filtering function is based on whether the RFC5322.From field domain is aligned with (matches) an authenticated domain name from SPF or DKIM. When a DMARC policy is published for the domain name found in the RFC5322.From field, and that domain name is not validated through SPF or DKIM, the disposition of that message can be affected by that DMARC policy when delivered to a participating receiver.

It is important to note that the authentication mechanisms employed by DMARC authenticate only a DNS domain and do not authenticate the local-part of any email address identifier found in a message, nor do they validate the legitimacy of message content.

DMARC's feedback component involves the collection of information about received messages claiming to be from the Organizational Domain for periodic aggregate reports to the Domain Owner. The parameters and format for such reports are discussed in later sections of this document.

A DMARC-enabled Mail Receiver might also generate per-message reports that contain information related to individual messages that fail SPF and/or DKIM. Per-message failure reports are a useful source of information when debugging deployments (if messages can be determined to be legitimate even though failing authentication) or in analyzing

In essence, the steps are as follows:

1. Domain Owner constructs an SPF policy and publishes it in its DNS database as per [[RFC7208](#)]. Domain Owner also configures its system for DKIM signing as described in [[RFC6376](#)]. Finally, Domain Owner publishes via the DNS a DMARC message-handling policy.
2. Author generates a message and hands the message to Domain Owner's designated mail submission service.
3. Submission service passes relevant details to the DKIM signing module in order to generate a DKIM signature to be applied to the message.
4. Submission service relays the now-signed message to its designated transport service for routing to its intended recipient(s).
5. Message may pass through other relays but eventually arrives at a recipient's transport service.
6. Recipient delivery service conducts SPF and DKIM authentication checks by passing the necessary data to their respective modules, each of which requires queries to the Author Domain's DNS data (when identifiers are aligned; see below).
7. The results of these are passed to the DMARC module along with the Author's domain. The DMARC module attempts to retrieve a policy from the DNS for that domain. If none is found, the DMARC module determines the Organizational Domain and repeats the attempt to retrieve a policy from the DNS. (This is described in further detail in [Section 6.6.3](#).)
8. If a policy is found, it is combined with the Author's domain and the SPF and DKIM results to produce a DMARC policy result (a "pass" or "fail") and can optionally cause one of two kinds of reports to be generated (not shown).
9. Recipient transport service either delivers the message to the recipient inbox or takes other local policy action based on the DMARC result (not shown).
10. When requested, Recipient transport service collects data from the message delivery session to be used in providing feedback (see [Section 7](#)).

5. Use of RFC5322.From

One of the most obvious points of security scrutiny for DMARC is the choice to focus on an identifier, namely the RFC5322.From address, which is part of a body of data that has been trivially forged throughout the history of email.

Several points suggest that it is the most correct and safest thing to do in this context:

- *Of all the identifiers that are part of the message itself, this is the only one guaranteed to be present.

- *It seems the best choice of an identifier on which to focus, as most MUAs display some or all of the contents of that field in a manner strongly suggesting those data as reflective of the true originator of the message.

The absence of a single, properly formed RFC5322.From field renders the message invalid. Handling of such a message is outside of the scope of this specification.

Since the sorts of mail typically protected by DMARC participants tend to only have single Authors, DMARC participants generally operate under a slightly restricted profile of RFC5322 with respect to the expected syntax of this field. See [Section 6.6](#) for details.

6. Policy

DMARC policies are published by Domain Owners and applied by Mail Receivers.

A Domain Owner advertises DMARC participation of one or more of its domains by adding a DNS TXT record (described in [Section 6.1](#)) to those domains. In doing so, Domain Owners make specific requests of Mail Receivers regarding the disposition of messages purporting to be from one of the Domain Owner's domains and the provision of feedback about those messages.

A Domain Owner may choose not to participate in DMARC evaluation by Mail Receivers. In this case, the Domain Owner simply declines to advertise participation in those schemes. For example, if the results of path authorization checks ought not be considered as part of the overall DMARC result for a given Author Domain, then the Domain Owner does not publish an SPF policy record that can produce an SPF pass result.

A Mail Receiver implementing the DMARC mechanism SHOULD make a best-effort attempt to adhere to the Domain Owner's published DMARC policy when a message fails the DMARC test. Since email streams can be

complicated (due to forwarding, existing RFC5322.From domain-spoofing services, etc.), Mail Receivers MAY deviate from a Domain Owner's published policy during message processing and SHOULD make available the fact of and reason for the deviation to the Domain Owner via feedback reporting, specifically using the "PolicyOverride" feature of the aggregate report (see the DMARC reporting documents).

6.1. DMARC Policy Record

Domain Owner DMARC preferences are stored as DNS TXT records in subdomains named "_dmarc". For example, the Domain Owner of "example.com" would post DMARC preferences in a TXT record at "_dmarc.example.com". Similarly, a Mail Receiver wishing to query for DMARC preferences regarding mail with an RFC5322.From domain of "example.com" would issue a TXT query to the DNS for the subdomain of "_dmarc.example.com". The DNS-located DMARC preference data will hereafter be called the "DMARC record".

DMARC's use of the Domain Name Service is driven by DMARC's use of domain names and the nature of the query it performs. The query requirement matches with the DNS, for obtaining simple parametric information. It uses an established method of storing the information, associated with the target domain name, namely an isolated TXT record that is restricted to the DMARC context. Use of the DNS as the query service has the benefit of reusing an extremely well-established operations, administration, and management infrastructure, rather than creating a new one.

Per [[RFC1035](#)], a TXT record can comprise several "character-string" objects. Where this is the case, the module performing DMARC evaluation MUST concatenate these strings by joining together the objects in order and parsing the result as a single string.

6.2. DMARC URIs

[[RFC3986](#)] defines a generic syntax for identifying a resource. The DMARC mechanism uses this as the format by which a Domain Owner specifies the destination for the two report types that are supported.

The place such URIs are specified (see [Section 6.3](#)) allows a list of these to be provided. A report is normally sent to each listed URI in the order provided by the Domain Owner. Receivers MAY impose a limit on the number of URIs to which they will send reports but MUST support the ability to send to at least two. The list of URIs is separated by commas (ASCII 0x2C).

Each URI can have associated with it a maximum report size that may be sent to it. This is accomplished by appending an exclamation point

(ASCII 0x21), followed by a maximum-size indication, before a separating comma or terminating semicolon.

Thus, a DMARC URI is a URI within which any commas or exclamation points are percent-encoded per [[RFC3986](#)], followed by an OPTIONAL exclamation point and a maximum-size specification, and, if there are additional reporting URIs in the list, a comma and the next URI.

For example, the URI "mailto:reports@example.com!50m" would request that a report be sent via email to "reports@example.com" so long as the report payload does not exceed 50 megabytes.

A formal definition is provided in [Section 6.4](#).

6.3. General Record Format

DMARC records follow the extensible "tag-value" syntax for DNS-based key records defined in DKIM [[RFC6376](#)].

[Section 10](#) creates a registry for known DMARC tags and registers the initial set defined in this document. Only tags defined in this document or in later extensions, and thus added to that registry, are to be processed; unknown tags MUST be ignored.

The following tags are introduced as the initial valid DMARC tags:

adkim: (plain-text; OPTIONAL; default is "r".) Indicates whether strict or relaxed DKIM Identifier Alignment mode is required by the Domain Owner. See [Section 3.1.1](#) for details. Valid values are as follows:

r: relaxed mode

s: strict mode

aspf: (plain-text; OPTIONAL; default is "r".) Indicates whether strict or relaxed SPF Identifier Alignment mode is required by the Domain Owner. See [Section 3.1.2](#) for details. Valid values are as follows:

r: relaxed mode

s: strict mode

fo: Failure reporting options (plain-text; OPTIONAL; default is "0") Provides requested options for generation of failure reports. Report generators MAY choose to adhere to the requested options. This tag's content MUST be ignored if a "ruf" tag (below) is not also specified. The value of this tag is a colon-separated list of characters that indicate failure reporting options as follows:

- 0:** Generate a DMARC failure report if all underlying authentication mechanisms fail to produce an aligned "pass" result.
- 1:** Generate a DMARC failure report if any underlying authentication mechanism produced something other than an aligned "pass" result.
- d:** Generate a DKIM failure report if the message had a signature that failed evaluation, regardless of its alignment. DKIM-specific reporting is described in [[RFC6651](#)].
- s:** Generate an SPF failure report if the message failed SPF evaluation, regardless of its alignment. SPF-specific reporting is described in [[RFC6652](#)].
- p:** Requested Mail Receiver policy (plain-text; REQUIRED for policy records). Indicates the policy to be enacted by the Receiver at the request of the Domain Owner. Policy applies to the domain queried and to subdomains, unless subdomain policy is explicitly described using the "sp" tag. This tag is mandatory for policy records only, but not for third-party reporting records (as discussed in the document(s) that discuss DMARC reporting in more detail). Possible values are as follows:
- none:** The Domain Owner requests no specific action be taken regarding delivery of messages.
- quarantine:** The Domain Owner wishes to have email that fails the DMARC mechanism check be treated by Mail Receivers as suspicious. Depending on the capabilities of the Mail Receiver, this can mean "place into spam folder", "scrutinize with additional intensity", and/or "flag as suspicious".
- reject:** The Domain Owner wishes for Mail Receivers to reject email that fails the DMARC mechanism check. Rejection SHOULD occur during the SMTP transaction. See [Section 9.3](#) for some discussion of SMTP rejection methods and their implications.
- pct:** (plain-text integer between 0 and 100, inclusive; OPTIONAL; default is 100). Percentage of messages from the Domain Owner's mail stream to which the DMARC policy is to be applied. However, this MUST NOT be applied to the DMARC-generated reports, all of which must be sent and received unhindered. The purpose of the "pct" tag is to allow Domain Owners to enact a slow rollout enforcement of the DMARC mechanism. The prospect of "all or nothing" is recognized as preventing many organizations from experimenting with strong authentication-based mechanisms. See

[Section 6.6.4](#) for details. Note that random selection based on this percentage, such as the following pseudocode, is adequate:

```
if (random mod 100) < pct then selected = true else selected =  
false
```

rf: Format to be used for message-specific failure reports (colon-separated plain-text list of values; OPTIONAL; default is "afrf"). The value of this tag is a list of one or more report formats as requested by the Domain Owner to be used when a message fails both [\[RFC3986\]](#) and [\[RFC6376\]](#) tests to report details of the individual failure. The values MUST be present in the registry of reporting formats defined in [Section 10](#); a Mail Receiver observing a different value SHOULD ignore it or MAY ignore the entire DMARC record. For this version, only "afrf" (the auth-failure report type defined in [\[RFC6591\]](#)) is presently supported. See the DMARC reporting documents for details. For interoperability, the Authentication Failure Reporting Format (AFRF) MUST be supported.

ri: Interval requested between aggregate reports (plain-text 32-bit unsigned integer; OPTIONAL; default is 86400). Indicates a request to Receivers to generate aggregate reports separated by no more than the requested number of seconds. DMARC implementations MUST be able to provide daily reports and SHOULD be able to provide hourly reports when requested. However, anything other than a daily report is understood to be accommodated on a best- effort basis.

rua: Addresses to which aggregate feedback is to be sent (comma-separated plain-text list of DMARC URIs; OPTIONAL). A comma or exclamation point that is part of such a DMARC URI MUST be encoded per Section 2.1 of [\[RFC3986\]](#) so as to distinguish it from the list delimiter or an OPTIONAL size limit. The DMARC reporting documents discuss considerations that apply when the domain name of a URI differs from that of the domain advertising the policy. See [Section 11.5](#) for additional considerations. Any valid URI can be specified. A Mail Receiver MUST implement support for a "mailto:" URI, i.e., the ability to send a DMARC report via electronic mail. If not provided, Mail Receivers MUST NOT generate aggregate feedback reports. URIs not supported by Mail Receivers MUST be ignored. The aggregate feedback report format is described in the DMARC reporting documents.

ruf: Addresses to which message-specific failure information is to be reported (comma-separated plain-text list of DMARC URIs; OPTIONAL). If present, the Domain Owner is requesting Mail Receivers to send detailed failure reports about messages that fail the DMARC evaluation in specific ways (see the "fo" tag above). The format of the message to be generated MUST follow the

format specified for the "rf" tag. The DMARC reporting documents discuss considerations that apply when the domain name of a URI differs from that of the domain advertising the policy. A Mail Receiver MUST implement support for a "mailto:" URI, i.e., the ability to send a DMARC report via electronic mail. If not provided, Mail Receivers MUST NOT generate failure reports. See [Section 11.5](#) for additional considerations.

sp: Requested Mail Receiver policy for all subdomains (plain-text; OPTIONAL). Indicates the policy to be enacted by the Receiver at the request of the Domain Owner. It applies only to subdomains of the domain queried and not to the domain itself. Its syntax is identical to that of the "p" tag defined above. If absent, the policy specified by the "p" tag MUST be applied for subdomains. Note that "sp" will be ignored for DMARC records published on subdomains of Organizational Domains due to the effect of the DMARC policy discovery mechanism described in [Section 6.6.3](#).

v: Version (plain-text; REQUIRED). Identifies the record retrieved as a DMARC record. It MUST have the value of "DMARC1". The value of this tag MUST match precisely; if it does not or it is absent, the entire retrieved record MUST be ignored. It MUST be the first tag in the list.

A DMARC policy record MUST comply with the formal specification found in [Section 6.4](#) in that the "v" tag MUST be present and MUST appear first. Unknown tags MUST be ignored. Syntax errors in the remainder of the record SHOULD be discarded in favor of default values (if any) or ignored outright.

Note that given the rules of the previous paragraph, addition of a new tag into the registered list of tags does not itself require a new version of DMARC to be generated (with a corresponding change to the "v" tag's value), but a change to any existing tags does require a new version of DMARC.

6.4. Formal Definition

The formal definition of the DMARC format, using [[RFC5234](#)], is as follows:

[FIXTHIS: Reference to [RFC3986] in code block]

```

dmarc-uri      = URI [ "!" 1*DIGIT [ "k" / "m" / "g" / "t" ] ]
                  ; "URI" is imported from [RFC3986]; commas (ASCII
                  ; 0x2C) and exclamation points (ASCII 0x21)
                  ; MUST be encoded; the numeric portion MUST fit
                  ; within an unsigned 64-bit integer

dmarc-record   = dmarc-version dmarc-sep
                  [dmarc-request]
                  [dmarc-sep dmarc-srequest]
                  [dmarc-sep dmarc-auri]
                  [dmarc-sep dmarc-furi]
                  [dmarc-sep dmarc-adkim]
                  [dmarc-sep dmarc-aspf]
                  [dmarc-sep dmarc-ainterval]
                  [dmarc-sep dmarc-fo]
                  [dmarc-sep dmarc-rfmt]
                  [dmarc-sep dmarc-percent]
                  [dmarc-sep]
                  ; components other than dmarc-version and
                  ; dmarc-request may appear in any order

dmarc-version  = "v" *WSP "=" *WSP %x44 %x4d %x41 %x52 %x43 %x31

dmarc-sep      = *WSP %x3b *WSP

dmarc-request  = "p" *WSP "=" *WSP
                  ( "none" / "quarantine" / "reject" )

dmarc-srequest = "sp" *WSP "=" *WSP
                  ( "none" / "quarantine" / "reject" )

dmarc-auri     = "rua" *WSP "=" *WSP
                  dmarc-uri *( *WSP "," *WSP dmarc-uri )

dmarc-furi     = "ruf" *WSP "=" *WSP
                  dmarc-uri *( *WSP "," *WSP dmarc-uri )

dmarc-adkim    = "adkim" *WSP "=" *WSP
                  ( "r" / "s" )

dmarc-aspf     = "aspf" *WSP "=" *WSP
                  ( "r" / "s" )

dmarc-ainterval = "ri" *WSP "=" *WSP 1*DIGIT

dmarc-fo       = "fo" *WSP "=" *WSP
                  ( "0" / "1" / "d" / "s" )
                  *( *WSP ":" *WSP ( "0" / "1" / "d" / "s" ) )

dmarc-rfmt     = "rf" *WSP "=" *WSP Keyword *( *WSP ":" Keyword )

```

```
                ; registered reporting formats only  
  
dmarc-percent  = "pct" *WSP "=" *WSP  
                1*3DIGIT
```

"Keyword" is imported from Section 4.1.2 of [\[RFC5321\]](#).

A size limitation in a dmarc-uri, if provided, is interpreted as a count of units followed by an OPTIONAL unit size ("k" for kilobytes, "m" for megabytes, "g" for gigabytes, "t" for terabytes). Without a unit, the number is presumed to be a basic byte count. Note that the units are considered to be powers of two; a kilobyte is 2¹⁰, a megabyte is 2²⁰, etc.

6.5. Domain Owner Actions

To implement the DMARC mechanism, the only action required of a Domain Owner is the creation of the DMARC policy record in the DNS. However, in order to make meaningful use of DMARC, a Domain Owner must at minimum either establish an address to receive reports, or deploy authentication technologies and ensure Identifier Alignment. Most Domain Owners will want to do both.

DMARC reports will be of significant size, and the addresses that receive them are publicly visible, so we encourage Domain Owners to set up dedicated email addresses to receive and process reports, and to deploy abuse countermeasures on those email addresses as appropriate.

Authentication technologies are discussed in [\[RFC6376\]](#) (see also [\[RFC5585\]](#) and [\[RFC5863\]](#)) and [\[RFC7208\]](#).

6.6. Mail Receiver Actions

This section describes receiver actions in the DMARC environment.

6.6.1. Extract Author Domain

The domain in the RFC5322.From field is extracted as the domain to be evaluated by DMARC. If the domain is encoded with UTF-8, the domain name must be converted to an A-label, as described in Section 2.3 of [\[RFC5890\]](#), for further processing.

In order to be processed by DMARC, a message typically needs to contain exactly one RFC5322.From domain (a single From: field with a single domain in it). Not all messages meet this requirement, and handling of them is outside of the scope of this document. Typical

exceptions, and the way they have been historically handled by DMARC participants, are as follows:

*Messages with multiple RFC5322.From fields are typically rejected, since that form is forbidden under RFC 5322 [[RFC5322](#)];

*Messages bearing a single RFC5322.From field containing multiple addresses (and, thus, multiple domain names to be evaluated) are typically rejected because the sorts of mail normally protected by DMARC do not use this format;

*Messages that have no RFC5322.From field at all are typically rejected, since that form is forbidden under RFC 5322 [[RFC5322](#)];

*Messages with an RFC5322.From field that contains no meaningful domains, such as RFC 5322 [[RFC5322](#)]'s "group" syntax, are typically ignored.

The case of a syntactically valid multi-valued RFC5322.From field presents a particular challenge. The process in this case is to apply the DMARC check using each of those domains found in the RFC5322.From field as the Author Domain and apply the most strict policy selected among the checks that fail.

6.6.2. Determine Handling Policy

To arrive at a policy for an individual message, Mail Receivers MUST perform the following actions or their semantic equivalents. Steps 2-4 MAY be done in parallel, whereas steps 5 and 6 require input from previous steps.

The steps are as follows:

1. Extract the RFC5322.From domain from the message (as above).
2. Query the DNS for a DMARC policy record. Continue if one is found, or terminate DMARC evaluation otherwise. See [Section 6.6.3](#) for details.
3. Perform DKIM signature verification checks. A single email could contain multiple DKIM signatures. The results of this step are passed to the remainder of the algorithm and MUST include the value of the "d=" tag from each checked DKIM signature.
4. Perform SPF validation checks. The results of this step are passed to the remainder of the algorithm and MUST include the domain name used to complete the SPF check.
5. Conduct Identifier Alignment checks. With authentication checks and policy discovery performed, the Mail Receiver checks to see

if Authenticated Identifiers fall into alignment as described in [Section 3](#). If one or more of the Authenticated Identifiers align with the RFC5322.From domain, the message is considered to pass the DMARC mechanism check. All other conditions (authentication failures, identifier mismatches) are considered to be DMARC mechanism check failures.

6. Apply policy. Emails that fail the DMARC mechanism check are disposed of in accordance with the discovered DMARC policy of the Domain Owner. See [Section 6.3](#) for details.

Heuristics applied in the absence of use by a Domain Owner of either SPF or DKIM (e.g., [[Best-Guess-SPF](#)]) SHOULD NOT be used, as it may be the case that the Domain Owner wishes a Message Receiver not to consider the results of that underlying authentication protocol at all.

DMARC evaluation can only yield a "pass" result after one of the underlying authentication mechanisms passes for an aligned identifier. If neither passes and one or both of them fail due to a temporary error, the Receiver evaluating the message is unable to conclude that the DMARC mechanism had a permanent failure; they therefore cannot apply the advertised DMARC policy. When otherwise appropriate, Receivers MAY send feedback reports regarding temporary errors.

Handling of messages for which SPF and/or DKIM evaluation encounter a permanent DNS error is left to the discretion of the Mail Receiver.

6.6.3. Policy Discovery

As stated above, the DMARC mechanism uses DNS TXT records to advertise policy. Policy discovery is accomplished via a method similar to the method used for SPF records. This method, and the important differences between DMARC and SPF mechanisms, are discussed below.

To balance the conflicting requirements of supporting wildcarding, allowing subdomain policy overrides, and limiting DNS query load, the following DNS lookup scheme is employed:

1. Mail Receivers MUST query the DNS for a DMARC TXT record at the DNS domain matching the one found in the RFC5322.From domain in the message. A possibly empty set of records is returned.
2. Records that do not start with a "v=" tag that identifies the current version of DMARC are discarded.
3. If the set is now empty, the Mail Receiver MUST query the DNS for a DMARC TXT record at the DNS domain matching the

Organizational Domain in place of the RFC5322.From domain in the message (if different). This record can contain policy to be asserted for subdomains of the Organizational Domain. A possibly empty set of records is returned.

4. Records that do not start with a "v=" tag that identifies the current version of DMARC are discarded.
5. If the remaining set contains multiple records or no records, policy discovery terminates and DMARC processing is not applied to this message.
6. If a retrieved policy record does not contain a valid "p" tag, or contains an "sp" tag that is not valid, then:
 1. if a "rua" tag is present and contains at least one syntactically valid reporting URI, the Mail Receiver SHOULD act as if a record containing a valid "v" tag and "p=none" was retrieved, and continue processing;
 2. otherwise, the Mail Receiver applies no DMARC processing to this message.

If the set produced by the mechanism above contains no DMARC policy record (i.e., any indication that there is no such record as opposed to a transient DNS error), Mail Receivers SHOULD NOT apply the DMARC mechanism to the message.

Handling of DNS errors when querying for the DMARC policy record is left to the discretion of the Mail Receiver. For example, to ensure minimal disruption of mail flow, transient errors could result in delivery of the message ("fail open"), or they could result in the message being temporarily rejected (i.e., an SMTP 4yx reply), which invites the sending MTA to try again after the condition has possibly cleared, allowing a definite DMARC conclusion to be reached ("fail closed").

6.6.4. Message Sampling

If the "pct" tag is present in the policy record, the Mail Receiver MUST NOT enact the requested policy ("p" tag or "sp" tag") on more than the stated percent of the totality of affected messages. However, regardless of whether or not the "pct" tag is present, the Mail Receiver MUST include all relevant message data in any reports produced.

If email is subject to the DMARC policy of "quarantine", the Mail Receiver SHOULD quarantine the message. If the email is not subject to the "quarantine" policy (due to the "pct" tag), the Mail Receiver SHOULD apply local message classification as normal.

If email is subject to the DMARC policy of "reject", the Mail Receiver SHOULD reject the message (see [Section 9.3](#)). If the email is not subject to the "reject" policy (due to the "pct" tag), the Mail Receiver SHOULD treat the email as though the "quarantine" policy applies. This behavior allows Domain Owners to experiment with progressively stronger policies without relaxing existing policy.

Mail Receivers implement "pct" via statistical mechanisms that achieve a close approximation to the requested percentage and provide a representative sample across a reporting period.

6.6.5. Store Results of DMARC Processing

The results of Mail Receiver-based DMARC processing should be stored for eventual presentation back to the Domain Owner in the form of aggregate feedback reports. [Section 6.3](#) and the DMARC reporting documents discuss aggregate feedback.

6.7. Policy Enforcement Considerations

Mail Receivers MAY choose to reject or quarantine email even if email passes the DMARC mechanism check. The DMARC mechanism does not inform Mail Receivers whether an email stream is "good". Mail Receivers are encouraged to maintain anti-abuse technologies to combat the possibility of DMARC-enabled criminal campaigns.

Mail Receivers MAY choose to accept email that fails the DMARC mechanism check even if the Domain Owner has published a "reject" policy. Mail Receivers need to make a best effort not to increase the likelihood of accepting abusive mail if they choose not to comply with a Domain Owner's reject, against policy. At a minimum, addition of the Authentication-Results header field (see [[RFC8601](#)]) is RECOMMENDED when delivery of failing mail is done. When this is done, the DNS domain name thus recorded MUST be encoded as an A-label.

Mail Receivers are only obligated to report reject or quarantine policy actions in aggregate feedback reports that are due to DMARC policy. They are not required to report reject or quarantine actions that are the result of local policy. If local policy information is exposed, abusers can gain insight into the effectiveness and delivery rates of spam campaigns.

Final disposition of a message is always a matter of local policy. An operator that wishes to favor DMARC policy over SPF policy, for example, will disregard the SPF policy, since enacting an SPF-determined rejection prevents evaluation of DKIM; DKIM might otherwise pass, satisfying the DMARC evaluation. There is a trade-off to doing so, namely acceptance and processing of the entire message body in exchange for the enhanced protection DMARC provides.

DMARC-compliant Mail Receivers typically disregard any mail-handling directive discovered as part of an authentication mechanism (e.g., Author Domain Signing Practices (ADSP), SPF) where a DMARC record is also discovered that specifies a policy other than "none". Deviating from this practice introduces inconsistency among DMARC operators in terms of handling of the message. However, such deviation is not proscribed.

To enable Domain Owners to receive DMARC feedback without impacting existing mail processing, discovered policies of "p=none" SHOULD NOT modify existing mail disposition processing.

Mail Receivers SHOULD also implement reporting instructions of DMARC, even in the absence of a request for DKIM reporting [[RFC6651](#)] or SPF reporting [[RFC6652](#)]. Furthermore, the presence of such requests SHOULD NOT affect DMARC reporting.

7. DMARC Feedback

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. When Domain Owners can see what effect their policies and practices are having, they are better willing and able to use quarantine and reject policies.

The details of this feedback are described in a separate document.

8. Minimum Implementations

A minimum implementation of DMARC has the following characteristics:

- *Is able to send and/or receive reports at least daily;
- *Is able to send and/or receive reports using "mailto" URIs;
- *Other than in exceptional circumstances such as resource exhaustion, can generate or accept a report up to ten megabytes in size;
- *If acting as a Mail Receiver, fully implements the provisions of [Section 6.6](#).

9. Other Topics

This section discusses some topics regarding choices made in the development of DMARC, largely to commit the history to record.

9.1. Issues Specific to SPF

Though DMARC does not inherently change the semantics of an SPF policy record, historically lax enforcement of such policies has led many to publish extremely broad records containing many large network ranges. Domain Owners are strongly encouraged to carefully review their SPF records to understand which networks are authorized to send on behalf of the Domain Owner before publishing a DMARC record.

Some receiver architectures might implement SPF in advance of any DMARC operations. This means that a "-" prefix on a sender's SPF mechanism, such as "-all", could cause that rejection to go into effect early in handling, causing message rejection before any DMARC processing takes place. Operators choosing to use "-all" should be aware of this.

9.2. DNS Load and Caching

DMARC policies are communicated using the DNS and therefore inherit a number of considerations related to DNS caching. The inherent conflict between freshness and the impact of caching on the reduction of DNS-lookup overhead should be considered from the Mail Receiver's point of view. Should Domain Owners publish a DNS record with a very short TTL, Mail Receivers can be provoked through the injection of large volumes of messages to overwhelm the Domain Owner's DNS. Although this is not a concern specific to DMARC, the implications of a very short TTL should be considered when publishing DMARC policies.

Conversely, long TTLs will cause records to be cached for long periods of time. This can cause a critical change to DMARC parameters advertised by a Domain Owner to go unnoticed for the length of the TTL (while waiting for DNS caches to expire). Avoiding this problem can mean shorter TTLs, with the potential problems described above. A balance should be sought to maintain responsiveness of DMARC preference changes while preserving the benefits of DNS caching.

9.3. Rejecting Messages

This proposal calls for rejection of a message during the SMTP session under certain circumstances. This is preferable to generation of a Delivery Status Notification ([[RFC3464](#)]), since fraudulent messages caught and rejected using DMARC would then result in annoying generation of such failure reports that go back to the RFC5321.MailFrom address.

This synchronous rejection is typically done in one of two ways:

- *Full rejection, wherein the SMTP server issues a 5xy reply code as an indication to the SMTP client that the transaction failed; the

SMTP client is then responsible for generating notification that delivery failed (see Section 4.2.5 of [[RFC5321](#)]).

*A "silent discard", wherein the SMTP server returns a 2xy reply code implying to the client that delivery (or, at least, relay) was successfully completed, but then simply discarding the message with no further action.

Each of these has a cost. For instance, a silent discard can help to prevent backscatter, but it also effectively means that the SMTP server has to be programmed to give a false result, which can confound external debugging efforts.

Similarly, the text portion of the SMTP reply may be important to consider. For example, when rejecting a message, revealing the reason for the rejection might give an attacker enough information to bypass those efforts on a later attempt, though it might also assist a legitimate client to determine the source of some local issue that caused the rejection.

In the latter case, when doing an SMTP rejection, providing a clear hint can be useful in resolving issues. A receiver might indicate in plain text the reason for the rejection by using the word "DMARC" somewhere in the reply text. Many systems are able to scan the SMTP reply text to determine the nature of the rejection. Thus, providing a machine-detectable reason for rejection allows the problems causing rejections to be properly addressed by automated systems. For example:

```
550 5.7.1 Email rejected per DMARC policy for example.com
```

If a Mail Receiver elects to defer delivery due to inability to retrieve or apply DMARC policy, this is best done with a 4xy SMTP reply code.

9.4. Identifier Alignment Considerations

The DMARC mechanism allows both DKIM and SPF-authenticated identifiers to authenticate email on behalf of a Domain Owner and, possibly, on behalf of different subdomains. If malicious or unaware users can gain control of the SPF record or DKIM selector records for a subdomain, the subdomain can be used to generate DMARC-passing email on behalf of the Organizational Domain.

For example, an attacker who controls the SPF record for "evil.example.com" can send mail with an RFC5322.From field containing "foo@example.com" that can pass both authentication and the DMARC check against "example.com".

The Organizational Domain administrator should be careful not to delegate control of subdomains if this is an issue, and to consider using the "strict" Identifier Alignment option if appropriate.

9.5. Interoperability Issues

DMARC limits which end-to-end scenarios can achieve a "pass" result.

Because DMARC relies on [\[RFC7208\]](#) and/or [\[RFC6376\]](#) to achieve a "pass", their limitations also apply.

Additional DMARC constraints occur when a message is processed by some Mediators, such as mailing lists. Transiting a Mediator often causes either the authentication to fail or Identifier Alignment to be lost. These transformations may conform to standards but will still prevent a DMARC "pass".

In addition to Mediators, mail that is sent by authorized, independent third parties might not be sent with Identifier Alignment, also preventing a "pass" result.

Issues specific to the use of policy mechanisms alongside DKIM are further discussed in [\[RFC6377\]](#), particularly Section 5.2.

10. IANA Considerations

This section describes actions completed by IANA.

10.1. Authentication-Results Method Registry Update

IANA has added the following to the "Email Authentication Methods" registry:

Method: dmarc

Defined: RFC 7489

ptype: header

Property: from

Value: the domain portion of the RFC5322.From field

Status: active

Version: 1

10.2. Authentication-Results Result Registry Update

IANA has added the following in the "Email Authentication Result Names" registry:

Code: none

Existing/New Code: existing

Defined: [[RFC8601](#)]

Auth Method: dmarc (added)

Meaning: No DMARC policy record was published for the aligned identifier, or no aligned identifier could be extracted.

Status: active

Code: pass

Existing/New Code: existing

Defined: [[RFC8601](#)]

Auth Method: dmarc (added)

Meaning: A DMARC policy record was published for the aligned identifier, and at least one of the authentication mechanisms passed.

Status: active

Code: fail

Existing/New Code: existing

Defined: [[RFC8601](#)]

Auth Method: dmarc (added)

Meaning: A DMARC policy record was published for the aligned identifier, and none of the authentication mechanisms passed.

Status: active

Code: temperror

Existing/New Code: existing

Defined: [[RFC8601](#)]

Auth Method: dmarc (added)

Meaning: A temporary error occurred during DMARC evaluation. A later attempt might produce a final result.

Status: active

Code: permerror

Existing/New Code: existing

Defined: [[RFC8601](#)]

Auth Method: dmarc (added)

Meaning: A permanent error occurred during DMARC evaluation, such as encountering a syntactically incorrect DMARC record. A later attempt is unlikely to produce a final result.

Status: active

10.3. Feedback Report Header Fields Registry Update

The following has been added to the "Feedback Report Header Fields" registry:

Field Name: Identity-Alignment

Description: indicates whether the message about which a report is being generated had any identifiers in alignment as defined in RFC 7489

Multiple Appearances: No

Related "Feedback-Type": auth-failure

Reference: RFC 7489

Status: current

10.4. DMARC Tag Registry

A new registry tree called "Domain-based Message Authentication, Reporting, and Conformance (DMARC) Parameters" has been created. Within it, a new sub-registry called the "DMARC Tag Registry" has been created.

Names of DMARC tags must be registered with IANA in this new sub-registry. New entries are assigned only for values that have been documented in a manner that satisfies the terms of Specification

Required, per [\[RFC8126\]](#). Each registration must include the tag name; the specification that defines it; a brief description; and its status, which must be one of "current", "experimental", or "historic". The Designated Expert needs to confirm that the provided specification adequately describes the new tag and clearly presents how it would be used within the DMARC context by Domain Owners and Mail Receivers.

To avoid version compatibility issues, tags added to the DMARC specification are to avoid changing the semantics of existing records when processed by implementations conforming to prior specifications.

The initial set of entries in this registry is as follows:

Tag Name	Reference	Status	Description
adkim	RFC 7489	current	DKIM alignment mode
aspf	RFC 7489	current	SPF alignment mode
fo	RFC 7489	current	Failure reporting options
p	RFC 7489	current	Requested handling policy
pct	RFC 7489	current	Sampling rate
rf	RFC 7489	current	Failure reporting format(s)
ri	RFC 7489	current	Aggregate Reporting interval
rua	RFC 7489	current	Reporting URI(s) for aggregate data
ruf	RFC 7489	current	Reporting URI(s) for failure data
sp	RFC 7489	current	Requested handling policy for subdomains
v	RFC 7489	current	Specification version

Table 1: "DMARC Tag Registry"

10.5. DMARC Report Format Registry

Also within "Domain-based Message Authentication, Reporting, and Conformance (DMARC) Parameters", a new sub-registry called "DMARC Report Format Registry" has been created.

Names of DMARC failure reporting formats must be registered with IANA in this registry. New entries are assigned only for values that satisfy the definition of Specification Required, per [\[RFC8126\]](#). In addition to a reference to a permanent specification, each registration must include the format name; a brief description; and its status, which must be one of "current", "experimental", or "historic". The Designated Expert needs to confirm that the provided specification adequately describes the report format and clearly presents how it would be used within the DMARC context by Domain Owners and Mail Receivers.

The initial entry in this registry is as follows:

Format Name	Reference	Status	Description
afrf	RFC 7489	current	Authentication Failure Reporting Format (see [RFC6591])

Table 2: "DMARC Report Format Registry"

10.6. Underscored and Globally Scoped DNS Node Names Registry

Per [\[!@RFC8552\]](#), please add the following entry to the "Underscored and Globally Scoped DNS Node Names" registry:

RR Type	_NODE NAME	Reference
TXT	_dmarc	RFC 7489

Table 3: "Underscored and Globally Scoped DNS Node Names" registry

11. Security Considerations

This section discusses security issues and possible remediations (where available) for DMARC.

11.1. Authentication Methods

Security considerations from the authentication methods used by DMARC are incorporated here by reference.

11.2. 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.

11.3. 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.

11.4. Display Name Attacks

A common attack in messaging abuse is the presentation of false information in the display-name portion of the RFC5322.From field. For example, it is possible for the email address in that field to be an arbitrary address or domain name, while containing a well-known name (a person, brand, role, etc.) in the display name, intending to fool the end user into believing that the name is used legitimately. The attack is predicated on the notion that most common MUAs will show the display name and not the email address when both are available.

Generally, display name attacks are out of scope for DMARC, as further exploration of possible defenses against these attacks needs to be undertaken.

There are a few possible mechanisms that attempt mitigation of these attacks, such as the following:

*If the display name is found to include an email address (as specified in [RFC5322]), execute the DMARC mechanism on the domain name found there rather than the domain name discovered originally. However, this addresses only a very specific attack space, and spoofer can easily circumvent it by simply not using an email address in the display name. There are also known cases of legitimate uses of an email address in the display name with a domain different from the one in the address portion, e.g.,

From: "user@example.org via Bug Tracker" support@example.com

*In the MUA, only show the display name if the DMARC mechanism succeeds. This too is easily defeated, as an attacker could

arrange to pass the DMARC tests while fraudulently using another domain name in the display name.

*In the MUA, only show the display name if the DMARC mechanism passes and the email address thus validated matches one found in the receiving user's list of known addresses.

11.5. 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, The DMARC reporting documents describe 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 the DMARC reporting documents is currently not mandatory ("MUST") but strongly recommended ("SHOULD"). It is possible that it would be elevated to a "MUST" by later security review.

11.6. 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.

12. 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>>.
- [RFC4343] Eastlake 3rd, D., "Domain Name System (DNS) Case Insensitivity Clarification", RFC 4343, DOI 10.17487/RFC4343, January 2006, <<https://www.rfc-editor.org/info/rfc4343>>.
- [RFC5234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, DOI 10.17487/RFC5234, January 2008, <<https://www.rfc-editor.org/info/rfc5234>>.
- [RFC5321] Klensin, J., "Simple Mail Transfer Protocol", RFC 5321, DOI 10.17487/RFC5321, October 2008, <<https://www.rfc-editor.org/info/rfc5321>>.
- [RFC5322] Resnick, P., Ed., "Internet Message Format", RFC 5322, DOI 10.17487/RFC5322, October 2008, <<https://www.rfc-editor.org/info/rfc5322>>.
- [RFC5890] Klensin, J., "Internationalized Domain Names for Applications (IDNA): Definitions and Document Framework", RFC 5890, DOI 10.17487/RFC5890, August 2010, <<https://www.rfc-editor.org/info/rfc5890>>.
- [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>>.
- [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>>.
- [RFC6651] Kucherawy, M., "Extensions to DomainKeys Identified Mail (DKIM) for Failure Reporting", RFC 6651, DOI 10.17487/

RFC6651, June 2012, <<https://www.rfc-editor.org/info/rfc6651>>.

[RFC6652] Kitterman, S., "Sender Policy Framework (SPF) Authentication Failure Reporting Using the Abuse Reporting Format", RFC 6652, DOI 10.17487/RFC6652, June 2012, <<https://www.rfc-editor.org/info/rfc6652>>.

[RFC7208] Kitterman, S., "Sender Policy Framework (SPF) for Authorizing Use of Domains in Email, Version 1", RFC 7208, DOI 10.17487/RFC7208, April 2014, <<https://www.rfc-editor.org/info/rfc7208>>.

13. Informative References

[Best-Guess-SPF] Kitterman, S., "Sender Policy Framework: Best guess record (FAQ entry)", May 2010, <http://www.openspf.org/FAQ/Best_guess_record>.

[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>>.

[RFC3464] Moore, K. and G. Vaudreuil, "An Extensible Message Format for Delivery Status Notifications", RFC 3464, DOI 10.17487/RFC3464, January 2003, <<https://www.rfc-editor.org/info/rfc3464>>.

[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>>.

[RFC5585] Hansen, T., Crocker, D., and P. Hallam-Baker, "DomainKeys Identified Mail (DKIM) Service Overview", RFC 5585, DOI 10.17487/RFC5585, July 2009, <<https://www.rfc-editor.org/info/rfc5585>>.

[RFC5598] Crocker, D., "Internet Mail Architecture", RFC 5598, DOI 10.17487/RFC5598, July 2009, <<https://www.rfc-editor.org/info/rfc5598>>.

[RFC5617] Allman, E., Fenton, J., Delany, M., and J. Levine, "DomainKeys Identified Mail (DKIM) Author Domain Signing Practices (ADSP)", RFC 5617, DOI 10.17487/RFC5617, August 2009, <<https://www.rfc-editor.org/info/rfc5617>>.

[RFC5863] Hansen, T., Siegel, E., Hallam-Baker, P., and D. Crocker, "DomainKeys Identified Mail (DKIM) Development, Deployment, and Operations", RFC 5863, DOI 10.17487/

RFC5863, May 2010, <<https://www.rfc-editor.org/info/rfc5863>>.

- [RFC6377] Kucherawy, M., "DomainKeys Identified Mail (DKIM) and Mailing Lists", BCP 167, RFC 6377, DOI 10.17487/RFC6377, September 2011, <<https://www.rfc-editor.org/info/rfc6377>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [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>>.
- [RFC8601] Kucherawy, M., "Message Header Field for Indicating Message Authentication Status", RFC 8601, DOI 10.17487/RFC8601, May 2019, <<https://www.rfc-editor.org/info/rfc8601>>.

Appendix A. Technology Considerations

This section documents some design decisions that were made in the development of DMARC. Specifically, addressed here are some suggestions that were considered but not included in the design. This text is included to explain why they were considered and not included in this version.

A.1. S/MIME

S/MIME, or Secure Multipurpose Internet Mail Extensions, is a standard for encryption and signing of MIME data in a message. This was suggested and considered as a third security protocol for authenticating the source of a message.

DMARC is focused on authentication at the domain level (i.e., the Domain Owner taking responsibility for the message), while S/MIME is really intended for user-to-user authentication and encryption. This alone appears to make it a bad fit for DMARC's goals.

S/MIME also suffers from the heavyweight problem of Public Key Infrastructure, which means that distribution of keys used to verify signatures needs to be incorporated. In many instances, this alone is a showstopper. There have been consistent promises that PKI usability and deployment will improve, but these have yet to materialize. DMARC can revisit this choice after those barriers are addressed.

S/MIME has extensive deployment in specific market segments (government, for example) but does not enjoy similar widespread

deployment over the general Internet, and this shows no signs of changing. DKIM and SPF both are deployed widely over the general Internet, and their adoption rates continue to be positive.

Finally, experiments have shown that including S/MIME support in the initial version of DMARC would neither cause nor enable a substantial increase in the accuracy of the overall mechanism.

A.2. Method Exclusion

It was suggested that DMARC include a mechanism by which a Domain Owner could tell Message Receivers not to attempt validation by one of the supported methods (e.g., "check DKIM, but not SPF").

Specifically, consider a Domain Owner that has deployed one of the technologies, and that technology fails for some messages, but such failures don't cause enforcement action. Deploying DMARC would cause enforcement action for policies other than "none", which would appear to exclude participation by that Domain Owner.

The DMARC development team evaluated the idea of policy exception mechanisms on several occasions and invariably concluded that there was not a strong enough use case to include them. The specific target audience for DMARC does not appear to have concerns about the failure modes of one or the other being a barrier to DMARC's adoption.

In the scenario described above, the Domain Owner has a few options:

1. Tighten up its infrastructure to minimize the failure modes of the single deployed technology.
2. Deploy the other supported authentication mechanism, to offset the failure modes of the first.
3. Deploy DMARC in a reporting-only mode.

A.3. Sender Header Field

It has been suggested in several message authentication efforts that the Sender header field be checked for an identifier of interest, as the standards indicate this as the proper way to indicate a re-mailing of content such as through a mailing list. Most recently, it was a protocol-level option for DomainKeys, but on evolution to DKIM, this property was removed.

The DMARC development team considered this and decided not to include support for doing so, for the following reasons:

1. The main user protection approach is to be concerned with what the user sees when a message is rendered. There is no consistent

behavior among MUAs regarding what to do with the content of the Sender field, if present. Accordingly, supporting checking of the Sender identifier would mean applying policy to an identifier the end user might never actually see, which can create a vector for attack against end users by simply forging a Sender field containing some identifier that DMARC will like.

2. Although it is certainly true that this is what the Sender field is for, its use in this way is also unreliable, making it a poor candidate for inclusion in the DMARC evaluation algorithm.
3. Allowing multiple ways to discover policy introduces unacceptable ambiguity into the DMARC evaluation algorithm in terms of which policy is to be applied and when.

A.4. Domain Existence Test

A common practice among MTA operators, and indeed one documented in [[RFC5617](#)], is a test to determine domain existence prior to any more expensive processing. This is typically done by querying the DNS for MX, A, or AAAA resource records for the name being evaluated and assuming that the domain is nonexistent if it could be determined that no such records were published for that domain name.

The original pre-standardization version of this protocol included a mandatory check of this nature. It was ultimately removed, as the method's error rate was too high without substantial manual tuning and heuristic work. There are indeed use cases this work needs to address where such a method would return a negative result about a domain for which reporting is desired, such as a registered domain name that never sends legitimate mail and thus has none of these records present in the DNS.

A.5. Issues with ADSP in Operation

DMARC has been characterized as a "super-ADSP" of sorts.

Contributors to DMARC have compiled a list of issues associated with ADSP, gained from operational experience, that have influenced the direction of DMARC:

1. ADSP has no support for subdomains, i.e., the ADSP record for example.com does not explicitly or implicitly apply to subdomain.example.com. If wildcarding is not applied, then spammers can trivially bypass ADSP by sending from a subdomain with no ADSP record.
2. Nonexistent subdomains are explicitly out of scope in ADSP. There is nothing in ADSP that states receivers should simply reject mail from NXDOMAINS regardless of ADSP policy (which of

course allows spammers to trivially bypass ADSP by sending email from nonexistent subdomains).

3. ADSP has no operational advice on when to look up the ADSP record.
4. ADSP has no support for using SPF as an auxiliary mechanism to DKIM.
5. ADSP has no support for a slow rollout, i.e., no way to configure a percentage of email on which the receiver should apply the policy. This is important for large-volume senders.
6. ADSP has no explicit support for an intermediate phase where the receiver quarantines (e.g., sends to the recipient's "spam" folder) rather than rejects the email.
7. The binding between the "From" header domain and DKIM is too tight for ADSP; they must match exactly.

A.6. Organizational Domain Discovery Issues

Although protocols like ADSP are useful for "protecting" a specific domain name, they are not helpful at protecting subdomains. If one wished to protect "example.com" by requiring via ADSP that all mail bearing an RFC5322.From domain of "example.com" be signed, this would "protect" that domain; however, one could then craft an email whose RFC5322.From domain is "security.example.com", and ADSP would not provide any protection. One could use a DNS wildcard, but this can undesirably interfere with other DNS activity; one could add ADSP records as fraudulent domains are discovered, but this solution does not scale and is a purely reactive measure against abuse.

The DNS does not provide a method by which the "domain of record", or the domain that was actually registered with a domain registrar, can be determined given an arbitrary domain name. Suggestions have been made that attempt to glean such information from SOA or NS resource records, but these too are not fully reliable, as the partitioning of the DNS is not always done at administrative boundaries.

When seeking domain-specific policy based on an arbitrary domain name, one could "climb the tree", dropping labels off the left end of the name until the root is reached or a policy is discovered, but then one could craft a name that has a large number of nonsense labels; this would cause a Mail Receiver to attempt a large number of queries in search of a policy record. Sending many such messages constitutes an amplified denial-of-service attack.

The Organizational Domain mechanism is a necessary component to the goals of DMARC. The method described in [Section 3.2](#) is far from

perfect but serves this purpose reasonably well without adding undue burden or semantics to the DNS. If a method is created to do so that is more reliable and secure than the use of a public suffix list, DMARC should be amended to use that method as soon as it is generally available.

A.6.1. Public Suffix Lists

A public suffix list for the purposes of determining the Organizational Domain can be obtained from various sources. The most common one is maintained by the Mozilla Foundation and made public at <http://publicsuffix.org>. License terms governing the use of that list are available at that URI.

Note that if operators use a variety of public suffix lists, interoperability will be difficult or impossible to guarantee.

Appendix B. Examples

This section illustrates both the Domain Owner side and the Mail Receiver side of a DMARC exchange.

B.1. Identifier Alignment Examples

The following examples illustrate the DMARC mechanism's use of Identifier Alignment. For brevity's sake, only message headers are shown, as message bodies are not considered when conducting DMARC checks.

B.1.1. SPF

The following SPF examples assume that SPF produces a passing result.

Example 1: SPF in alignment:

```
MAIL FROM: <sender@example.com>

From: sender@example.com
Date: Fri, Feb 15 2002 16:54:30 -0800
To: receiver@example.org
Subject: here's a sample
```

In this case, the RFC5321.MailFrom parameter and the RFC5322.From field have identical DNS domains. Thus, the identifiers are in alignment.

Example 2: SPF in alignment (parent):

MAIL FROM: <sender@child.example.com>

From: sender@example.com
Date: Fri, Feb 15 2002 16:54:30 -0800
To: receiver@example.org
Subject: here's a sample

In this case, the RFC5322.From parameter includes a DNS domain that is a parent of the RFC5321.MailFrom domain. Thus, the identifiers are in alignment if relaxed SPF mode is requested by the Domain Owner, and not in alignment if strict SPF mode is requested.

Example 3: SPF not in alignment:

MAIL FROM: <sender@example.net>

From: sender@child.example.com
Date: Fri, Feb 15 2002 16:54:30 -0800
To: receiver@example.org
Subject: here's a sample

In this case, the RFC5321.MailFrom parameter includes a DNS domain that is neither the same as nor a parent of the RFC5322.From domain. Thus, the identifiers are not in alignment.

B.1.2. DKIM

The examples below assume that the DKIM signatures pass verification. Alignment cannot exist with a DKIM signature that does not verify.

Example 1: DKIM in alignment:

DKIM-Signature: v=1; ...; d=example.com; ...
From: sender@example.com
Date: Fri, Feb 15 2002 16:54:30 -0800
To: receiver@example.org
Subject: here's a sample

In this case, the DKIM "d=" parameter and the RFC5322.From field have identical DNS domains. Thus, the identifiers are in alignment.

Example 2: DKIM in alignment (parent):

DKIM-Signature: v=1; ...; d=example.com; ...
From: sender@child.example.com
Date: Fri, Feb 15 2002 16:54:30 -0800
To: receiver@example.org
Subject: here's a sample

In this case, the DKIM signature's "d=" parameter includes a DNS domain that is a parent of the RFC5322.From domain. Thus, the identifiers are in alignment for relaxed mode, but not for strict mode.

Example 3: DKIM not in alignment:

```
DKIM-Signature: v=1; ...; d=sample.net; ...
From: sender@child.example.com
Date: Fri, Feb 15 2002 16:54:30 -0800
To: receiver@example.org
Subject: here's a sample
```

In this case, the DKIM signature's "d=" parameter includes a DNS domain that is neither the same as nor a parent of the RFC5322.From domain. Thus, the identifiers are not in alignment.

B.2. Domain Owner Example

A Domain Owner that wants to use DMARC should have already deployed and tested SPF and DKIM. The next step is to publish a DNS record that advertises a DMARC policy for the Domain Owner's Organizational Domain.

B.2.1. Entire Domain, Monitoring Only

The owner of the domain "example.com" has deployed SPF and DKIM on its messaging infrastructure. The owner wishes to begin using DMARC with a policy that will solicit aggregate feedback from receivers without affecting how the messages are processed, in order to:

- *Confirm that its legitimate messages are authenticating correctly
- *Verify that all authorized message sources have implemented authentication measures
- *Determine how many messages from other sources would be affected by a blocking policy

The Domain Owner accomplishes this by constructing a policy record indicating that:

- *The version of DMARC being used is "DMARC1" ("v=DMARC1;")
- *Receivers should not alter how they treat these messages because of this DMARC policy record ("p=none")
- *Aggregate feedback reports should be sent via email to the address "dmarc-feedback@example.com" ("rua=mailto:dmarc-feedback@example.com")

*All messages from this Organizational Domain are subject to this policy (no "pct" tag present, so the default of 100% applies)

The DMARC policy record might look like this when retrieved using a common command-line tool:

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

To publish such a record, the DNS administrator for the Domain Owner creates 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" )
```

B.2.2. 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 [Appendix B.2](#):

```
*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" )
```

B.2.3. 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 B.2.2](#) 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 the DMARC reporting documents. 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 the DMARC reporting documents for the full details of this mechanism.

B.2.4. Subdomain, Sampling, and Multiple Aggregate Report URIs

The Domain Owner has implemented SPF and DKIM in a subdomain used for pre-production testing of messaging services. It now wishes to request that participating receivers act to reject messages from this subdomain that fail to authenticate.

As a first step, it will ask that a portion (1/4 in this example) of failing messages be quarantined, enabling examination of messages sent to mailboxes hosted by participating receivers. Aggregate feedback reports will be sent to a mailbox within the Organizational Domain, and to a mailbox at a third party selected and authorized to receive same by the Domain Owner. Aggregate reports sent to the third party are limited to a maximum size of ten megabytes.

The Domain Owner will accomplish this by constructing a policy record indicating that:

*The version of DMARC being used is "DMARC1" ("v=DMARC1;")

*It is applied only to this subdomain (record is published at "_dmarc.test.example.com" and not "_dmarc.example.com")

*Receivers should quarantine messages from this Organizational Domain that fail to authenticate ("p=quarantine")

*Aggregate feedback reports should be sent via email to the addresses "dmarc-feedback@example.com" and "example-tld-

test@thirdparty.example.net", with the latter subjected to a maximum size limit ("rua=mailto:dmARC-feedback@example.com,mailto:tld-test@thirdparty.example.net!10m")

*25% of messages from this Organizational Domain are subject to action based on this policy ("pct=25")

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.test.example.com
"v=DMARC1; p=quarantine; rua=mailto:dmARC-feedback@example.com,
mailto:tld-test@thirdparty.example.net!10m; pct=25"
```

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

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

```
_dmarc IN TXT ( "v=DMARC1; p=quarantine; "
                "rua=mailto:dmARC-feedback@example.com,"
                "mailto:tld-test@thirdparty.example.net!10m; "
                "pct=25" )
```

B.3. Mail Receiver Example

A Mail Receiver that wants to use DMARC should already be checking SPF and DKIM, and possess the ability to collect relevant information from various email-processing stages to provide feedback to Domain Owners (possibly via Report Receivers).

B.4. Processing of SMTP Time

An optimal DMARC-enabled Mail Receiver performs authentication and Identifier Alignment checking during the [[RFC5322](#)] conversation.

Prior to returning a final reply to the DATA command, the Mail Receiver's MTA has performed:

1. An SPF check to determine an SPF-authenticated Identifier.
2. DKIM checks that yield one or more DKIM-authenticated Identifiers.
3. A DMARC policy lookup.

The presence of an Author Domain DMARC record indicates that the Mail Receiver should continue with DMARC-specific processing before returning a reply to the DATA command.

Given a DMARC record and the set of Authenticated Identifiers, the Mail Receiver checks to see if the Authenticated Identifiers align with the Author Domain (taking into consideration any strict versus relaxed options found in the DMARC record).

For example, the following sample data is considered to be from a piece of email originating from the Domain Owner of "example.com":

```
Author Domain: example.com
SPF-authenticated Identifier: mail.example.com
DKIM-authenticated Identifier: example.com
DMARC record:
  "v=DMARC1; p=reject; aspf=r;
  rua=mailto:dmarc-feedback@example.com"
```

In the above sample, both the SPF-authenticated Identifier and the DKIM-authenticated Identifier align with the Author Domain. The Mail Receiver considers the above email to pass the DMARC check, avoiding the "reject" policy that is to be applied to email that fails to pass the DMARC check.

If no Authenticated Identifiers align with the Author Domain, then the Mail Receiver applies the DMARC-record-specified policy. However, before this action is taken, the Mail Receiver can consult external information to override the Domain Owner's policy. For example, if the Mail Receiver knows that this particular email came from a known and trusted forwarder (that happens to break both SPF and DKIM), then the Mail Receiver may choose to ignore the Domain Owner's policy.

The Mail Receiver is now ready to reply to the DATA command. If the DMARC check yields that the message is to be rejected, then the Mail Receiver replies with a 5xy code to inform the sender of failure. If the DMARC check cannot be resolved due to transient network errors, then the Mail Receiver replies with a 4xy code to inform the sender as to the need to reattempt delivery later. If the DMARC check yields a passing message, then the Mail Receiver continues on with email processing, perhaps using the result of the DMARC check as an input to additional processing modules such as a domain reputation query.

Before exiting DMARC-specific processing, the Mail Receiver checks to see if the Author Domain DMARC record requests AFRF-based reporting. If so, then the Mail Receiver can emit an AFRF to the reporting address supplied in the DMARC record.

At the exit of DMARC-specific processing, the Mail Receiver captures (through logging or direct insertion into a data store) the result of

DMARC processing. Captured information is used to build feedback for Domain Owner consumption. This is not necessary if the Domain Owner has not requested aggregate reports, i.e., no "rua" tag was found in the policy record.

B.5. Utilization of Aggregate Feedback: Example

Aggregate feedback is consumed by Domain Owners to verify a Domain Owner's understanding of how the Domain Owner's domain is being processed by the Mail Receiver. Aggregate reporting data on emails that pass all DMARC-supporting authentication checks is used by Domain Owners to verify that authentication practices remain accurate. For example, if a third party is sending on behalf of a Domain Owner, the Domain Owner can use aggregate report data to verify ongoing authentication practices of the third party.

Data on email that only partially passes underlying authentication checks provides visibility into problems that need to be addressed by the Domain Owner. For example, if either SPF or DKIM fails to pass, the Domain Owner is provided with enough information to either directly correct the problem or understand where authentication-breaking changes are being introduced in the email transmission path. If authentication-breaking changes due to email transmission path cannot be directly corrected, then the Domain Owner at least maintains an understanding of the effect of DMARC-based policies upon the Domain Owner's email.

Data on email that fails all underlying authentication checks provides baseline visibility on how the Domain Owner's domain is being received at the Mail Receiver. Based on this visibility, the Domain Owner can begin deployment of authentication technologies across uncovered email sources. Additionally, the Domain Owner may come to an understanding of how its domain is being misused.

B.6. mailto Transport Example

A DMARC record can contain a "mailto" reporting address, such as:

`mailto:dmarc-feedback@example.com`

A sample aggregate report from the Mail Receiver at `mail.receiver.example` follows:

DKIM-Signature: v=1; ...; d=mail.receiver.example; ...
From: dmarc-reporting@mail.receiver.example
Date: Fri, Feb 15 2002 16:54:30 -0800
To: dmarc-feedback@example.com
Subject: Report Domain: example.com
 Submitter: mail.receiver.example
 Report-ID: <2002.02.15.1>
MIME-Version: 1.0
Content-Type: multipart/alternative;
 boundary="-----_NextPart_000_024E_01CC9B0A.AFE54C00"
Content-Language: en-us

This is a multipart message in MIME format.

-----_NextPart_000_024E_01CC9B0A.AFE54C00
Content-Type: text/plain; charset="us-ascii"
Content-Transfer-Encoding: 7bit

This is an aggregate report from mail.receiver.example.

-----_NextPart_000_024E_01CC9B0A.AFE54C00
Content-Type: application/gzip
Content-Transfer-Encoding: base64
Content-Disposition: attachment;
 filename="mail.receiver.example!example.com!
 1013662812!1013749130.gz"

<gzipped content of report>

-----_NextPart_000_024E_01CC9B0A.AFE54C00--

Not shown in the above example is that the Mail Receiver's feedback should be authenticated using SPF. Also, the value of the "filename" MIME parameter is wrapped for printing in this specification but would normally appear as one continuous string.

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

Emil Gustafsson
Google

Email: emgu@google.com

Todd M. Herr
Valimail

Email: todd.herr@valimail.com

John Levine
Standcore LLC

Email: standards@standore.com