DMARC
Internet-Draft
Obsoletes: 7489, 9091 (if approved)
Intended status: Standards Track
Expires: 6 October 2025

T. Herr (ed)
Valimail
J. Levine (ed)
Standcore LLC

4 April 2025

Domain-based Message Authentication, Reporting, and Conformance (DMARC)
draft-ietf-dmarc-dmarcbis-41

Abstract

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

DMARC permits the owner of an email's Author Domain to enable validation of the domain's use, to indicate the Domain Owner's or Public Suffix Operator's message handling preference regarding failed validation, and to request reports about the use of the domain name. Mail receiving organizations can use this information when evaluating handling choices for incoming mail.

This document obsoletes RFCs 7489 and 9091.

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 6 October 2025.

Copyright Notice

Table of Contents

1. Introduction

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

Abusive email often includes unauthorized and deceptive use of a domain name in the "From" header field defined in section 3.6.2 of [RFC5322] and referred to as RFC5322.From. The domain typically belongs to an organization expected to be known to - and presumably trusted by - the recipient. The Sender Policy Framework (SPF) [RFC7208] and DomainKeys Identified Mail (DKIM) [RFC6376] protocols

provide domain-level authentication but are not directly associated with the RFC5322.From domain, also known as the Author Domain (#author-domain). DMARC leverages these two protocols, providing a method for Domain Owners to publish a DNS TXT record describing the email authentication policies for the Author Domain and to request specific handling for messages using that domain that fail validation checks. These DNS records are called DMARC Policy Records (#dmarc-policy-record).

As with SPF and DKIM, DMARC validation results in a verdict of either "pass" or "fail". A DMARC result of "pass" requires not only an SPF or DKIM pass verdict for the email message, but also and more importantly that the domain associated with the SPF or DKIM pass be "aligned" with the Author Domain in one of two modes - "relaxed" or "strict". Domains are said to be in "relaxed alignment" if they have the same Organizational Domain (#organizational-domain); a domain's Organizational Domain is the domain at the top of the namespace hierarchy for that domain while having the same administrative authority as that domain. On the other hand, domains are in "strict alignment" if and only if they are identical. The choice of required alignment mode is left to the Domain Owner (#domain-owner) that publishes a DMARC Policy Record.

A DMARC pass for a message indicates only that the use of the Author Domain has been validated for that message as authorized by the Domain Owner. Such authorization does not carry an explicit or implicit value assertion about that message or about the Domain Owner, and so a DMARC pass by itself does not guarantee that delivery to the recipient's Inbox would be safe or desirable. For a mail-receiving organization participating in DMARC, a message that passes DMARC validation is part of a message stream reliably associated with the Author Domain. Therefore, reputation assessment of that stream by the mail-receiving organization can assume the use of that Author Domain is authorized by the Domain Owner.

On the other hand, a message that fails this validation is not necessarily associated with the Author Domain and so should not affect the Author Domain's reputation. The phrase "not necessarily associated" was purposely chosen here, as it is importatnt to understand that some messages making authorized use of the Author Domain can still fail DMARC validation checks. [RFC7960] and Section 7 of this document both discuss reasons why such failures may happen. Because of this, a mail-receiving organization that performs DMARC validation can choose to honor the Domain Owner's requested message handling for validation failures, but it is not required to do so. DMARC is commonly used as one input to more complex filtering decisions, and so the mail-receiving organization might choose different actions entirely.

DMARC, in the associated [I-D.ietf-dmarc-aggregate-reporting] and [I-D.ietf-dmarc-failure-reporting] documents, also specifies a reporting framework. Using it, a mail-receiving organization can generate regular reports about messages that use Author Domains for which a DMARC Policy Record exists; those reports are sent to the address(es) specified by the Domain Owner in the DMARC Policy Record. Domain Owners can use these reports, especially the aggregate reports, not only to identify sources of mail attempting to fraudulently use their domain, but also (and perhaps more importantly) to flag and fix gaps in their own authentication practices. However, as with honoring the Domain Owner's stated mail handling preference, a mail-receiving organization supporting DMARC is under no obligation to send requested reports, although it is recommended that they do send aggregate reports.

The use of DMARC creates some interoperability challenges that require due consideration before deployment, particularly with configurations that can cause mail to be rejected. These are discussed in Section 7.

2. Requirements

The following sections describe topics that guide the specification of DMARC.

2.1. High-Level Goals

DMARC has the following high-level goals:

* Allow Domain Owners (#domain-owner) and Public Suffix Operators (PSOs) (#public-suffix-operator) to validate their email authentication deployments.

* Allow Domain Owners and PSOs to assert their desired message handling for validation failures on messages purporting to have authorship within the domain.

* Minimize implementation complexity for both senders and receivers.

* Reduce the amount of successfully delivered spoofed emails.

* Work at Internet scale.

2.2. Anti-Phishing

DMARC is designed to prevent the unauthorized use of the Author Domain (#author-domain) of an email message, a technique known as "spoofing". Such unauthorized usage can frequently be found in messages impersonating a domain belonging to a business entity, messages that are meant to entice the recipient to provide sensitive information, such as usernames, passwords, and financial account information. These spoofed messages are commonly referred to as "phishing".

DMARC can only be used to combat specific forms of exact-domain spoofing directly. DMARC does not attempt to solve all problems with spoofed or otherwise fraudulent emails. In particular, it does not address the use of visually similar domain names or abuse of the RFC5322.From human-readable display-name, as defined in Section 3.4 of [RFC5322].

2.3. Scalability

Scalability is a significant issue for systems that need to operate in an environment 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. Out of Scope

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

* Different treatment of messages that are not authenticated (e.g., those that have no DKIM signature or those sent using an Author Domain (#author-domain) for which no DMARC Policy Record (#dmarc-policy-record) exists) versus those that fail validation;

* Evaluation of anything other than RFC5322.From header field;

* Multiple reporting formats;

* Publishing policy other than via the DNS;

* Reporting or otherwise evaluating other than the last-hop IP address;

* Attacks in the display-name portions of the RFC5322.From header 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.

3. Terminology and Definitions

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

3.1. Conventions Used in This 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 (#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.

3.2. Definitions

The following sections define the terms used in this document.

3.2.1. Authenticated Identifiers

Authenticated Identifiers are those domain-level identifiers for which authorized use is validated using a supported authentication mechanism (#authentication-mechanisms).

3.2.2. Author Domain

The domain name of the apparent author as extracted from the RFC5322.From header field.

3.2.3. DKIM Signing Domain

The domain name that is the value of the "d" tag in a validated DKIM-Signature header field in an email message.

3.2.4. SPF Domain

SPF, [RFC7208], can validate the uses of both the domain found in an SMTP [RFC5321] HELO/EHLO command (the HELO identity) and the domain found in an SMTP MAIL command (the MAIL FROM identity). DMARC relies solely on SPF validation of the MAIL FROM identity. Section 2.4 of [RFC7208] describes the determination of the MAIL FROM identity for cases in which the SMTP MAIL command has a null path, i.e., the mailbox composed of the local-part "postmaster" and the HELO identity.

The term "SPF Domain" when used in this document refers to an SPF validated MAIL FROM identity.

3.2.5. DMARC Policy Domain

The domain name at which an applicable DMARC Policy Record (#dmarc-policy-record) is discovered for the Author Domain (#author-domain) of an email message.

3.2.6. DMARC Policy Record

A DNS TXT record published by a Domain Owner (#domain-owner) or Public Suffix Operator (PSO) (#public-suffix-operator) to enable validation of an Author Domain’s (#author-domain) use, to indicate the Domain Owner’s or PSO’s message handling preference regarding failed validation, and optionally to request reports about the use of the Author Domain.

3.2.7. Domain Owner

An entity or organization that has control of a given DNS domain, usually by holding its registration. 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 Consumers when those are outside of their immediate management domain.

3.2.8. Domain Owner Assessment Policy

The message handling preference expressed in a DMARC Policy Record (#dmarc-policy-record) by the Domain Owner (#domain-owner) regarding failed validation of the Author Domain (#author-domain) is called the "Domain Owner Assessment Policy". Possible values are described in Section 4.7.

3.2.9. Enforcement

Enforcement describes a state where the existing Domain Owner Assessment Policy (#domain-owner-policy) for an Organizational Domain (#organizational-domain) and all subdomains below it is not "p=none". This state means that the Organizational Domain and its subdomains can only be used as Author Domains (#author-domain) if they are properly validated using the DMARC mechanism.

Historically, Domain Owner Assessment Policies of "p=quarantine" or "p=reject" have been higher value signals to Mail Receivers (#mail-receiver). Messages with Author Domains for which such policies exist that are not validated using the DMARC mechanism will not reach the inbox at Mail Receivers that participate in DMARC and honor the Domain Owner’s expressed handling preference.

3.2.10. Identifier Alignment

DMARC describes the concept of alignment between the Author Domain (#author-domain) and an Authenticated Identifier (#authenticated-identifiers), and requires such Identifier Alignment between the two for a message to achieve a DMARC pass. DMARC defines two states for alignment.

3.2.10.1. Relaxed Alignment

When the Author Domain (#author-domain) has the same Organizational Domain (#organizational-domain) as an Authenticated Identifier (#authenticated-identifier), the two are said to be in relaxed alignment.

3.2.10.2. Strict Alignment

When the Author Domain (#author-domain) is identical to an Authenticated Identifier (#authenticated-identifier), the two are said to be in strict alignment.

3.2.11. Mail Receiver

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

3.2.12. Monitoring Mode

Monitoring Mode describes a state where the existing Domain Owner Assessment Policy (#domain-owner-policy) for an Organizational Domain (#organizational-domain) and all subdomains below it is "p=none", and the Domain Owner (#domain-owner) is receiving aggregate reports for the Organizational Domain. While the use of the Organizational Domain and all its subdomains as Author Domains (#author-domain) can still be validated by a Mail Receiver (#mail-receiver) deploying the DMARC mechanism, the Domain Owner expresses no handling preference for messages that fail DMARC validation. The Domain Owner is, however, using the content of the DMARC aggregate reports to make any needed adjustments to the authentication practices for its mail streams.

3.2.13. Non-existent Domains

For DMARC purposes, a non-existent domain is consistent with the term’s meaning as described in [RFC8020]. That is, if the response code received for a query for a domain name is NXDOMAIN, then the domain name and any possible subdomains do not exist.

3.2.14. Organizational Domain

The Organizational Domain for any domain is akin to the ADMD described in [RFC5598]. A domain’s Organizational Domain is the domain at the top of the namespace hierarchy for that domain while having the same administrative authority as the domain. An Organizational Domain is determined by applying the algorithm found in Section 4.10.

3.2.15. Public Suffix Domain (PSD)

Some domains allow the registration of subdomains that are "owned" by independent organizations. Real-world examples of these domains are ".com", ".org", ".us", and ".co.uk", to name just a few. These domains are called "Public Suffix Domains" (PSDs). For example, "ietf.org" is a registered domain name, and ".org" is its PSD.

3.2.16. Public Suffix Operator (PSO)

A Public Suffix Operator is an organization that manages operations within a PSD, particularly the DNS records published for names at and under that domain name.

3.2.17. PSO Controlled Domain Names

PSO-Controlled Domain Names are names in the DNS that are managed by a PSO. PSO-controlled Domain Names may have one label (e.g., ".com") or more (e.g., ".co.uk"), depending on the PSD's policy.

3.2.18. Report Consumer

A Report Consumer is an operator that receives reports from another operator implementing the reporting mechanisms described in the documents [I-D.ietf-dmarc-aggregate-reporting] and [I-D.ietf-dmarc-failure-reporting]. This term applies collectively to the system components that receive and process these reports and the organizations that operate those components.

Report Consumers can receive reports concerning domains for which the Report Consumer is also the Domain Owner (#domain-owner) or PSO (#public-suffix-operator), or concerning domains that belong to another operator entirely. The DMARC mechanism permits a Domain Owner to act as a Report Consumer for its domain(s) and/or to designate third parties to so act. See Section 11.6 for further discussion of such designation.

4. Overview and Key Concepts

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

4.1. DMARC Basics

DMARC permits a Domain Owner (#domain-owner) or PSO (#public-suffix-operator) to enable validation of an Author Domain's (#author-domain) use in an email message, to indicate the Domain Owner's or PSO's message handling preference regarding failed validation, and to request reports about use of the Author Domain. A domain's DMARC Policy Record (#dmarc-policy-record) is published in the DNS as a TXT record at the name created by prepending the label "_dmarc" to the domain name and is retrieved through normal DNS queries.

DMARC's validation mechanism produces a "pass" result if a DMARC Policy Record exists for the Author Domain of an email message and the Author Domain is aligned (#identifier-alignment) with an

Authenticated Identifier (#authenticated-identifiers) from that message. When a DMARC Policy Record exists for the Author Domain and the DMARC mechanism does not produce a "pass" result, the Mail Receiver's (#mail-receiver) handling of that message can be influenced by the Domain Owner Assessment Policy (#domain-owner-policy) expressed in the DMARC Policy Record.

It is important to note that the authentication mechanisms employed by DMARC only validate the usage of a DNS domain in an email message. They do not validate the local-part of any email address identifier found in that message, nor do such validations carry an explicit or implicit value assertion about that message or about the Domain Owner.

DMARC's reporting component involves the collection of information about received messages using the Author Domain for periodic aggregate reports to the Domain Owner or PSO. The parameters and format for such reports are discussed in [I-D.ietf-dmarc-aggregate-reporting].

A Mail Receiver participating in DMARC might also generate per-message failure reports that contain information related to individual messages that fail DMARC validation checks. Per-message failure reports are a useful source of information when debugging deployments (if messages can be determined to be legitimate even though failing validation) or in analyzing attacks. The capability for such services is enabled by DMARC but defined in other referenced material such as [RFC6591] and [I-D.ietf-dmarc-failure-reporting]

4.2. 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. This field is the one used by end users to identify the source of the message, and so it has always been a prime target for abuse through such forgery and other means. That said, of all the identifiers that are part of the message itself, this is the only one required to be present. A message without a single, properly formed RFC5322.From header field does not comply with [RFC5322], and handling such a message is outside of the scope of this specification.

4.3. Authentication Mechanisms

The following mechanisms for determining Authenticated Identifiers (#authenticated-identifiers) are supported in this version of DMARC:

* DKIM, [RFC6376]. The DKIM Signing Domain (#dkim-signing-domain) from a validated DKIM-Signature header field is an Authenticated Identifier.

* SPF, [RFC7208]. The validated SPF Domain (#spf-domain) from the email message is the Authenticated Identifier.

4.4. Identifier Alignment Explained

DMARC validates the authorized use of the Author Domain (#author-domain) by requiring either that it have the same Organizational Domain (#organizational-domain) as an Authenticated Identifier (#authenticated-identifier) (a condition known as "Relaxed Alignment (#relaxed-alignment)") or that it be identical to the Authenticated Identifier (a condition known as "Strict Alignment (#strict-alignment)"). The choice of relaxed or strict alignment is left to the Domain Owner (#domain-owner) and is expressed in the domain's DMARC Policy Record (#dmarc-policy-record). In practice, nearly all Domain Owners have found relaxed alignment sufficient to meet their needs. Domain name comparisons in this context are case-insensitive, per [RFC4343].

The following table is meant to illustrate possible alignment conditions.

Authenticated Identifier	Author Domain	Identifier Alignment
foo.example.com	news.example.com	relaxed; the two have the same Organizational Domain, example.com
news.example.com	news.example.com	strict; the two are identical
foo.example.net	news.example.com	none; the two do not share a common Organizational Domain

Table 1: "Alignment Examples"

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 header field, since in that case there is no reliable way to determine a DMARC Policy Record (#dmarc-policy-record) that applies to the message. Accordingly,

DMARC operation is predicated on the input being a valid RFC5322 message object. For non-compliant cases, handling is outside of the scope of this specification. Further discussion of this can be found in Section 11.5.

4.4.1. DKIM-Authenticated Identifiers

DKIM permits a Domain Owner to claim some responsibility for a message by associating the domain to the message. This association is done by inserting the domain as the value of the "d" tag in a DKIM-Signature header field, and the assertion of responsibility is validated through a cryptographic signature in the header field. If the cryptographic signature validates, then the DKIM Signing Domain is the DKIM-Authenticated Identifier.

There is currently no generally accepted mechanism by which a Domain Owner may assert a list of third-party DKIM Signing Domains that are authorized to sign on behalf of a given Author Domain. Therefore, DMARC requires that Identifier Alignment is applied to the DKIM-Authenticated Identifier because a message can bear a valid signature from any domain, even one used by a bad actor. Only a DKIM-Authenticated Identifier that has Identifier Alignment with the Author Domain is enough to validate the authorized use of the Author Domain.

A single email can contain multiple DKIM signatures, and it is considered to produce a DMARC "pass" result if any DKIM-Authenticated Identifier aligns with the Author Domain.

4.4.2. SPF-Authenticated Identifiers

SPF can validate the uses of both the domain found in an SMTP HELO/EHLO command (the HELO identity) and the domain found in an SMTP MAIL command (the MAIL FROM identity). DMARC relies solely on SPF validation of the MAIL FROM identity. If the use of the domain in the MAIL FROM identity is validated by SPF, then that domain is the SPF-Authenticated Identifier.

There is currently no generally accepted mechanism by which a Domain Owner may assert a list of third-party domains that are authorized for use as the MAIL FROM identity for mail using a given Author Domain. Therefore, DMARC requires that Identifier Alignment is applied to the SPF-Authenticated Identifier because any Domain Owner, even a bad actor, can publish an SPF record for its domain and send email that will obtain an SPF pass result. Only an SPF-Authenticated Identifier that has Identifier Alignment with the Author Domain is enough to validate the authorized use of the Author Domain.

4.4.3. Alignment and Extension Technologies

If in the future DMARC is extended to include the use of other authentication mechanisms, the extensions MUST allow for the assignment of a domain as an Authenticated Identifier so that alignment with the Author Domain can be validated.

4.5. DMARC Policy Record Explained

A Domain Owner (#domain-owner) or PSO (#public-suffix-operator) advertises DMARC participation of one or more of its domains by publishing DMARC Policy Records (#dmarc-policy-record) that will apply to those domains. In doing so, Domain Owners and PSOs indicate their handling preference regarding failed validation for email messages using their domain in the RFC5322.From header field as well as their desire (if any) to receive feedback about such messages in the form of aggregate and/or failure reports.

DMARC Policy Records are stored as DNS TXT records with names starting with the label "_dmarc". For example, the Domain Owner of "example.com" would publish a DMARC Policy Record at the name "_dmarc.example.com", and a Mail Receiver (#mail-receiver) wishing to find the DMARC Policy Record for mail with an Author Domain (#author-domain) of "example.com" would issue a TXT query to the DNS for the name "_dmarc.example.com". A Domain Owner or PSO may choose not to participate in DMARC validation by Mail Receivers simply by not publishing a DMARC Policy Record for its Author Domain(s).

DMARC Policy Records can also apply to subdomains of the name at which they are published in the DNS, if the record is published at an Organizational Domain (#organizational-domain) for the subdomains. The Domain Owner Assessment Policy (#domain-owner-policy) that applies to the subdomains can be identical to the Domain Owner Assessment Policy that applies to the Organizational Domain or different, depending on the presence or absence of certain values in the DMARC Policy Record. See Section 4.7 for more details.

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 domain name targeted by a DNS query, specifically an isolated TXT record that is restricted to the DMARC context. Using 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 multiple "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.

A Domain Owner can choose not to have some underlying authentication mechanisms apply to DMARC evaluation of its Author Domain(s). For example, if a Domain Owner only wants to use DKIM as the underlying authentication mechanism, then the Domain Owner does not publish an SPF record that can produce Identifier Alignment between an SPF-Authenticated Identifier and the Author Domain. Alternatively, if the Domain Owner wishes to rely solely on SPF, then it can send email messages that have no DKIM-Signature header field that would produce Identifier Alignment between a DKIM-Authenticated Identifier and the Author Domain. However, it is RECOMMENDED that Domain Owners use both DKIM and SPF as underlying authentication mechanisms for DMARC.

A Mail Receiver implementing the DMARC mechanism gets the Domain Owner’s or PSO’s published Domain Owner Assessment Policy and can use it to inform its handling decisions for messages that undergo DMARC validation checks and do not produce a result of pass. Mail handling considerations based on Domain Owner Assessment Policy enforcement are discussed below in Section 5.4.

4.6. DMARC Reporting URIs

[RFC3986] defines a syntax for identifying a resource. The DMARC mechanism uses this as the format by which a Domain Owner (#domain-owner) or PSO (#public-suffix-organization) specifies the destination(s) for the two report types that are supported. The DMARC Policy Record format (#policy-record-format) allows for a list of these URIs to be provided, with each URI separated by commas (ASCII 0x2c).

A formal definition is provided in Section 4.8.

4.7. DMARC Policy Record Format

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

Section 9 creates a registry for known DMARC tags and registers the initial set defined in this document. Only tags defined in that registry are to be processed; unknown tags MUST be ignored.

The following tags are valid DMARC tags:

adkim: (plain-text; OPTIONAL; default is "r".) Indicates whether

the Domain Owner (#domain-owner) or PSO (#public-suffix-organization) requires strict or relaxed DKIM Identifier Alignment mode. See Section 4.4.1 for details. Valid values are as follows:

r: relaxed mode
s: strict mode

aspf: (plain-text; OPTIONAL; default is "r".) Indicates whether the Domain Owner or PSO requires strict or relaxed SPF Identifier Alignment mode. See Section 4.4.2 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 the 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. This tag can include one or more of the values shown here, with the exception that "0" and "1" are mutually exclusive. If more than one value is assigned to the tag, the list of values should be separated by colons (e.g., fo=0:d), and the values may appear in the list in any order. Valid values and their meanings are:

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 fails to produce 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].

np: Domain Owner Assessment Policy (#domain-owner-policy) for non-existent subdomains of the given Organizational Domain (plain-text; OPTIONAL). For this tag, the definition of "non-existent subdomain" is the same as that used for "Non-existent Domains" in Section 3.2.13. The policy expressed by this tag indicates the message handling preference of the Domain Owner or PSO for mail using non-existent subdomains of the prevailing Organizational Domain and not passing DMARC validation. It applies only to non-

existent subdomains of the Organizational Domain queried and not to either existing subdomains or the domain itself. Its syntax is identical to that of the "p" tag defined below. If the "np" tag is absent, the policy specified by the "sp" tag (if the "sp" tag is present) or the policy specified by the "p" tag, if the "sp" tag is not present, MUST be applied for non-existent subdomains.

p: Domain Owner Assessment Policy (#domain-owner-policy) (plain-text; RECOMMENDED for DMARC Policy Records). Indicates the message handling preference of the Domain Owner or PSO for mail using its domain but not passing DMARC validation. The policy applies to the domain queried and to subdomains, unless the subdomain policy is explicitly described using the "sp" or "np" tags. If this tag is not present in an otherwise syntactically valid DMARC Policy Record, then the record is treated as if it included "p=none" (see Section 4.10.1). This tag is not applicable for third-party reporting records (see [I-D.ietf-dmarc-aggregate-reporting] and [I-D.ietf-dmarc-failure-reporting]). Possible values are as follows:

none: The Domain Owner offers no expression of preference.
quarantine: The Domain Owner considers such mail to be suspicious. It is possible the mail is valid, although the failure creates a significant concern.
reject: The Domain Owner considers all such failures to be a clear indication that the use of the domain name is not valid. See Section 7.2 for some discussion of SMTP rejection methods and their implications.

psd: A flag indicating whether the domain is a PSD. (plain-text; OPTIONAL; default is "u"). Possible values are:

y: PSOs include this tag with a value of "y" to indicate that the domain is a PSD. If a record containing this tag with a value of "y" is found during policy discovery, this information will be used to determine the Organizational Domain and DMARC Policy Domain applicable to the message in question.
n: The DMARC Policy Record is published for a domain that is not a PSD, but it is the Organizational Domain for itself and its subdomains.
u: The default indicates that the DMARC Policy Record is published for a domain that is not a PSD, and may or may not be an Organizational Domain for itself and its subdomains. Use the mechanism described in Section 4.10 for determining the Organizational Domain for this domain.

rua: Addresses to which aggregate feedback reports are to be sent

(comma-separated plain-text list of DMARC Reporting URIs; OPTIONAL). If present, the Domain Owner is requesting Mail Receivers to send aggregate feedback reports as defined in [I-D.ietf-dmarc-aggregate-reporting] to the URIs listed. Any valid URI can be specified. A Mail Receiver that sends aggregate feedback reports MUST implement support for a "mailto:" URI, i.e., the ability to send a DMARC report via electronic mail. If the tag is not provided, Mail Receivers MUST NOT generate aggregate feedback reports for the domain. URIs involving schemes not supported by Mail Receivers MUST be ignored. [I-D.ietf-dmarc-aggregate-reporting] also discusses considerations that apply when the domain name of a URI differs from the domain publishing the DMARC Policy Record. See Section 11.6 for additional considerations.

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) to the URIs listed. Depending on the value of the "fo" tag, the format for such reports is described in [I-D.ietf-dmarc-failure-reporting], [RFC6651], or [RFC6652]. Any valid URI can be specified. A Mail Receiver sending failure reports MUST implement support for a "mailto:" URI, i.e., the ability to send message-specific failure information via electronic mail. If the tag is not provided, Mail Receivers MUST NOT generate failure reports for the domain. URIs involving schemes not supported by Mail Receivers MUST be ignored. [I-D.ietf-dmarc-aggregate-reporting] discusses considerations that apply when the domain name of a URI differs from that of the domain advertising the policy. See Section 11.6 for additional considerations.

sp: Domain Owner Assessment Policy for all subdomains of the given Organizational Domain (plain-text; OPTIONAL). Indicates the message handling preference of the Domain Owner or PSO for mail using an existing subdomain of the prevailing Organizational Domain for and not passing DMARC validation. It applies only to existing subdomains of the message’s Organizational Domain in the DNS hierarchy and not to the Organizational Domain itself. Its syntax is identical to that of the "p" tag defined above. If both the "sp" tag is absent, and the "np" tag is either absent or not applicable, the policy specified by the "p" tag MUST be applied for subdomains. Note that "sp" will be ignored for DMARC Policy Records published on subdomains of Organizational Domains and PSDs due to the effect of the DMARC Policy Discovery (#dmarc-policy-discovery).

t: DMARC policy test mode (plain-text; OPTIONAL; default is "n"). For the Author Domain to which the DMARC Policy Record applies, the "t" tag serves as a signal to the actor performing DMARC validation checks as to whether or not the Domain Owner wishes the Domain Owner Assessment Policy declared in the "p", "sp", and/or "np" tags to actually be applied. This tag does not affect the generation of DMARC reports, and it has no effect on any policy ("p", "sp", or "np") that is "none". See Appendix A.6 for further discussion of the use of this tag. Possible values are as follows:

y: A request that the actor performing the DMARC validation check not apply the policy, but instead apply any special handling rules it might have in place, such as rewriting the RFC5322.From header field (see Appendix A.6). The Domain Owner is currently testing its specified DMARC assessment policy, and has an expectation that the policy applied to any failing messages will be one level below the specified policy. That is, if the policy is "quarantine" and the value of the "t" tag is "y", a policy of "none" will be applied to failing messages; if the policy is "reject" and the value of the "t" tag is "y", a policy of "quarantine" will be applied to failing messages, irrespective of any other special handling rules that might be triggered by the "t" tag having a value of "y".

n: The default is a request to apply the Domain Owner Assessment Policy as specified to any message that produces a DMARC "fail" result.

v: Version (plain-text; REQUIRED). Identifies the record retrieved as a DMARC Policy Record. This tag MUST be the first tag in the list. The tag value is case sensitive, and the only possible value is "DMARC1". If the tag is not the first in the list, or the tag is absent, or the value is not "DMARC1", then the entire record MUST be ignored.

4.8. Formal Definition

A DMARC Policy Record MUST comply with the formal definition of same found in this section. Unknown tags MUST be ignored. Syntax errors in the remainder of the record MUST be discarded in favor of default values (if any) or ignored outright.

Because unknown tags MUST be ignored, the 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.

The formal definition of the DMARC Policy Record format, using [RFC5234] and [RFC7405], is as follows:

```
  dmarc-uri     = URI
                ; "URI" is imported from [RFC3986];
                ; commas (ASCII 0x2C) and exclamation
                ; points (ASCII 0x21) MUST be
                ; encoded

  obs-dmarc-uri = dmarc-uri obs-dmarc-report-size
                ; Obsolete syntax, reporters should ignore the
                ; obs-dmarc-report-size if it is found in a DMARC Policy Record.

  obs-dmarc-report-size = "!" 1*DIGIT [ "k" / "m" / "g" / "t" ]

  dmarc-sep     = *WSP ";" *WSP

  equals        = *WSP "=" *WSP

  dmarc-record  = dmarc-version *(dmarc-sep dmarc-tag) [dmarc-sep]

  dmarc-tag     = 1*ALPHA equals 1*dmarc-value

  ; any printing characters but semicolon
  dmarc-value   = %x20-3A / %x3C-7E

  dmarc-version = "v" equals %s"DMARC1" ; case sensitive

  ; specialized syntax of DMARC values
  dmarc-request = "none" / "quarantine" / "reject"

  dmarc-yorn    = "y" / "n"

  dmarc-psd     = "y" / "n" / "u"

  dmarc-rors    = "r" / "s"

  dmarc-urilist = (dmarc-uri / obs-dmarc-uri) *(*WSP "," *WSP (dmarc-uri / obs-dmarc-uri)
)

  dmarc-fo      = ("0" / "1") *(":" dmarc-afrf)
                / dmarc-afrf [":" ("0" / "1")] [":" dmarc-afrf]
                / *(dmarc-afrf ":") ("0" / "1")

  dmarc-afrf    = "d" / "s"
                ; each may appear at most once in dmarc-fo
```

In each dmarc-tag, the dmarc-value has a syntax that depends on the tag name. The ABNF rule for each dmarc-value is specified in the following table:

Tag Name	Value Rule
p	dmarc-request
t	dmarc-yorn
psd	dmarc-psd
np	dmarc-request
sp	dmarc-request
adkim	dmarc-rors
aspf	dmarc-rors
rua	dmarc-urilist
ruf	dmarc-urilist
fo	dmarc-fo

Table 2: "Tag Names and Values"

4.9. Flow Diagram

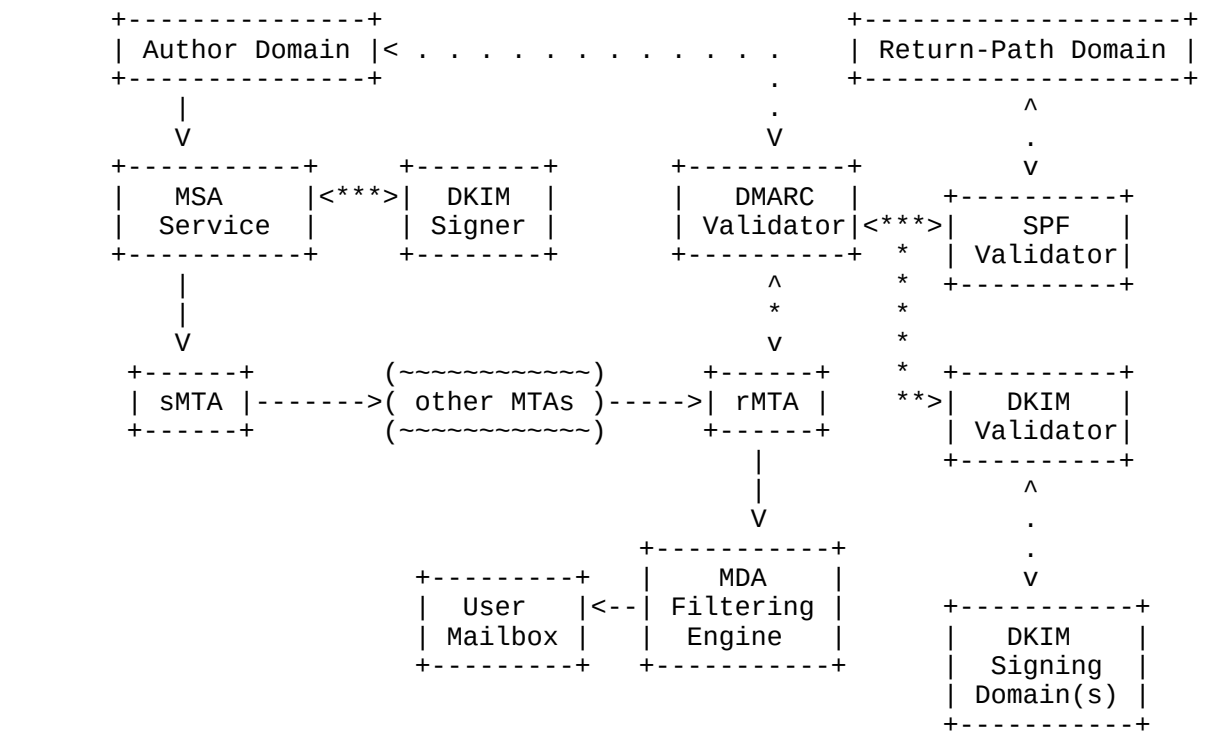

The above diagram shows a typical flow of messages through a DMARC-aware system. Dashed lines (e.g., -->) denote the actual message flow, dotted lines (e.g., < . . >) represent DNS queries used to retrieve message policy related to the supported message authentication schemes, and starred lines (e.g., <**>) indicate data exchange between message-handling modules and message authentication modules. "sMTA" is the sending MTA, and "rMTA" is the receiving MTA.

Put simply, when a message reaches a DMARC-aware rMTA, a DNS query will be initiated to determine if a DMARC Policy Record exists that applies to the Author Domain. If a DMARC Policy Record is found, the rMTA will use the results of SPF and DKIM validation checks to determine DMARC validation status. The DMARC validation status can then factor into the message handling decision made by the recipient’s mail system.

More details on specific actions for the parties involved can be found in Section 5.1 and Section 5.3.

4.10. DNS Tree Walk

An Organizational Domain (#organizational-domain) serves two different purposes, depending on the context:

* The Organizational Domain of the Author Domain (#author-domain) establishes the DMARC Policy Record (#dmarc-policy-record) for that domain when no DMARC Policy Record is published specifically for the Author Domain. (see Section 4.10.1)

* The Organizational Domains of an Authenticated Identifier (#authenticated-identifiers) and the Author Domain are used in determining Identifier Alignment between the two. (see Section 4.10.2).

[RFC7489] defined an Organizational Domain as "The domain that was registered with a domain name registrar." RFC 7489 discussed using a "public suffix" list (PSL) as the authoritative list of the parent domains for Organizational Domains, and further described a method for determining the Organizational Domain of an Author Domain or an Authenticated Identifier. However, RFC 7489 mandated no requirement for a specific PSL for Mail Receivers to use (though it did suggest the one found at https://publicsuffix.org/ (https://publicsuffix.org/)) nor did it provide any guidance for the frequency of regular retrieval of the PSL by Mail Receivers participating in DMARC. RFC 7489 acknowledged the possibility of interoperability issues caused by Mail Receivers choosing different PSLs, and even suggested that if a more reliable and secure method for determining the Organizational Domain could be created, that method should replace reliance on a public suffix list.

This update to DMARC offers more flexibility to Domain Owners, especially those with large, complex organizations that might want to apply decentralized management to their DNS and their DMARC Policy Records. Rather than just using a public suffix list to help identify an Organizational Domain, this update defines a discovery technique known colloquially as the "DNS Tree Walk". The target of any DNS Tree Walk is discovery of a valid DMARC Policy Record, and its use in determining an Organizational Domain allows for publishing DMARC Policy Records at multiple points in the namespace.

This flexibility comes at a possible cost, however. Since the DNS Tree Walk relies on the Mail Receiver making a series of DNS queries, the potential exists for an ill-intentioned Domain Owner to send mail with Author Domains with tens or even hundreds of labels for the purpose of executing a Denial of Service Attack on the Mail Receiver. To guard against such abuse of the DNS, a shortcut is built into the process so that Author Domains with more than eight labels do not

result in more than eight DNS queries. Observed data at the time of publication showed that Author Domains with up to seven labels were in usage, and so eight was chosen as the query limit to allow for some future expansion of the name space that did not require updating this document.

The generic steps for a DNS Tree Walk are as follows:

1. Query the DNS for a TXT record that matches the format of a DMARC Policy Record at the starting point for the Tree Walk. The starting point for the DNS Tree Walk will depend on the ultimate target of the DNS Tree Walk. Section 4.10.1 and Section 4.10.2 describe the possible starting points. 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. If multiple DMARC Policy Records are returned for a single target, they are all discarded. If a single record remains and it contains a "psd=n" or "psd=y" tag, stop.

3. Break the subject DNS domain name into a set of ordered labels. Assign the count of labels to "x", and number the labels from right to left; e.g., for "a.mail.example.com", "x" would be assigned the value 4, "com" would be label 1, "example" would be label 2, "mail" would be label 3, and so forth.

4. If x < 8, remove the left-most (highest-numbered) label from the subject domain. If x >= 8, remove the left-most (highest-numbered) labels from the subject domain until 7 labels remain. The resulting DNS domain name is the new target for the next lookup.

5. Query the DNS for a DMARC Policy Record at the DNS domain name matching this new target. A possibly empty set of records is returned.

6. Records that do not start with a "v" tag that identifies the current version of DMARC are discarded. If multiple DMARC Policy Records are returned for a single target, they are all discarded. If a single record remains and it contains a "psd=n" or "psd=y" tag, stop.

7. Determine the target for the next query by removing the left-most label from the target of the previous query. Repeat steps 5, 6, and 7 until the process stops or there are no more labels remaining.

To illustrate, for a message with the arbitrary Author Domain of "a.b.c.d.e.f.g.h.i.j.mail.example.com", a full DNS Tree Walk would require the following eight queries to potentially locate the DMARC Policy Record or Organizational Domain:

* _dmarc.a.b.c.d.e.f.g.h.i.j.mail.example.com
* _dmarc.g.h.i.j.mail.example.com
* _dmarc.h.i.j.mail.example.com
* _dmarc.i.j.mail.example.com
* _dmarc.j.mail.example.com
* _dmarc.mail.example.com
* _dmarc.example.com
* _dmarc.com

4.10.1. DMARC Policy Discovery

The DMARC Policy Record to be applied to an email message will be the record found at any of the following locations, listed from highest preference to lowest:

* The Author Domain
* The Organizational Domain of the Author Domain
* The Public Suffix Domain of the Author Domain

Policy discovery starts first with a query for a valid DMARC Policy Record at the name created by prepending the label "_dmarc" to the Author Domain of the message being evaluated. If a valid DMARC Policy Record is found there, then this is the DMARC Policy Record to be applied to the message; however, this does not necessarily mean that the Author Domain is the Organizational Domain to be used in Identifier Alignment checks. Whether this is also the Organizational Domain is dependent on the value of the "psd" tag, if present, or some conditions described in Section 4.10.2.

If no valid DMARC Policy Record is found by the first query, then perform a DNS Tree Walk to find the Author Domain’s Organizational Domain or its Public Suffix Domain. The starting point for this DNS Tree Walk is determined as follows:

* If the Author Domain has eight or fewer labels, the starting point will be the immediate parent domain of the Author Domain.
* Otherwise, the starting point will be the name produced by shortening the Author Domain as described starting in step 3 of Section 4.10.

If the DMARC Policy Record to be applied is that of the Author Domain, then the Domain Owner Assessment Policy is taken from the "p" tag of the record.

If the DMARC Policy Record to be applied is that of either the Organizational Domain or the Public Suffix Domain and the Author Domain is a subdomain of that domain, then the Domain Owner Assessment Policy is taken from the "sp" tag (if any) if the Author Domain exists, or the "np" tag (if any) if the Author Domain does not exist. In the absence of applicable "sp" or "np" tags, the "p" tag policy is used for subdomains.

If a retrieved DMARC Policy Record does not contain a valid "p" tag, or contains an "sp" or "np" tag that is not valid, then:

* If a "rua" tag is present and contains at least one syntactically valid reporting URI, the Mail Receiver MUST act as if a record containing "p=none" was retrieved and continue processing;

* Otherwise, the Mail Receiver applies no DMARC processing to this message.

If the set produced by the DNS Tree Walk contains no DMARC Policy Record (i.e., any indication that there is no such record as opposed to a transient DNS error), Mail Receivers MUST 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").

Note: PSD policy is not used for Organizational Domains that have published a DMARC Policy Record. Specifically, this is not a mechanism to provide feedback addresses (rua/ruf) when an Organizational Domain has declined to do so.

4.10.2. Identifier Alignment Evaluation

It may be necessary to perform multiple DNS Tree Walks to determine if an Authenticated Identifier and an Author Domain are in alignment, meaning that they have either the same Organizational Domain (relaxed alignment) or that they’re identical (strict alignment). DNS Tree Walks done to discover an Organizational Domain for use in Identifier Alignment Evaluation might start at any of the following locations:

* The Author Domain of the message being evaluated.

* The SPF-Authenticated Identifier if there is an SPF pass result for the message being evaluated.
* Any DKIM-Authenticated Identifier if one or more DKIM pass results exist for the message being evaluated.

Note: There is no need to perform Identifier Alignment Evaluations under any of the following conditions:

* The Author Domain and the Authenticated Identifier(s) are all the same domain, and there is a DMARC Policy Record published for that domain. In this case, this common domain is treated as the Organizational Domain. For example, if the common domain in question is "mail.example.com", and there is a valid DMARC Policy Record published at "_dmarc.mail.example.com", then "mail.example.com" is the Organizational Domain.
* No applicable DMARC Policy Record is discovered for the Author Domain. In this case, the DMARC mechanism does not apply to the message in question.
* The DMARC Policy Record for the Author Domain indicates strict alignment. In this case, a simple string comparison of the Author Domain and the Authenticated Identifier(s) is all that is required.

To discover the Organizational Domain for a domain, perform the DNS Tree Walk described in Section 4.10 as needed for any of the domains in question.

For each Tree Walk that retrieved valid DMARC Policy Records, select the Organizational Domain from the domains for which valid DMARC Policy Records were retrieved from the longest to the shortest:

1. If a valid DMARC Policy Record contains the "psd" tag set to "n" ("psd=n"), this is the Organizational Domain, and the selection process is complete.

2. If a valid DMARC Policy Record, other than the one for the domain where the tree walk started, contains the "psd" tag set to "y" ("psd=y"), the Organizational Domain is the domain one label below this one in the DNS hierarchy, and the selection process is complete. For example, if in the course of a tree walk a DMARC Policy Record is queried for at first "_dmarc.mail.example.com" and then "_dmarc.example.com", and a valid DMARC Policy Record containing the "psd" tag set to "y" is found at "_dmarc.example.com", then "mail.example.com" is the domain one label below "example.com" in the DNS hierarchy and is thus the Organizational Domain.

3. Otherwise, select the DMARC Policy Record found at the name with the fewest number of labels. This is the Organizational Domain and the selection process is complete.

If this process does not determine the Organizational Domain, then the initial target domain is the Organizational Domain.

For example, given the starting domain "a.mail.example.com", a search for the Organizational Domain would require a series of DNS queries for DMARC Policy Records starting with "_dmarc.a.mail.example.com" and finishing with "_dmarc.com". If there are DMARC Policy Records published at "_dmarc.mail.example.com" and "_dmarc.example.com", but not at "_dmarc.a.mail.example.com" or "_dmarc.com", then the Organizational Domain for this domain would be "example.com".

As another example, given the starting domain "a.mail.example.com", if a search for the Organizational Domain yields a DMARC Policy Record at "_dmarc.mail.example.com" with the "psd" tag set to "n", then the Organizational Domain for this domain would be "mail.example.com".

As a last example, given the starting domain "a.mail.example.com", if a search for the Organizational Domain only yields a DMARC Policy Record at "_dmarc.com" and that record contains the tag "psd=y", then the Organizational Domain for this domain would be "example.com".

5. DMARC Participation

This section describes the actions for participating in DMARC for each of three unique entities - Domain Owners, PSOs, and Mail Receivers.

5.1. Domain Owner Actions

A Domain Owner (#domain-owner) wishing to fully participate in DMARC will publish a DMARC Policy Record (#dmarc-policy-record) to cover each Author Domain (#author-domain) and corresponding Organizational Domain (#organizational-domain) to which DMARC validation should apply, send email that produces at least one, and preferably two, Authenticated Identifiers (#authenticated-identifiers) that align with the Author Domain, will receive and monitor the content of DMARC aggregate reports, and will correct any authentication shortcomings in mail making authorized use of its domains.

The following sections describe how to achieve this.

5.1.1. Publish an SPF Record for an Aligned Domain

To configure SPF for DMARC, the Domain Owner MUST send mail that has an RFC5321.MailFrom domain that will produce an SPF-Authenticated Identifier (#spf-identifiers) that has Identifier Alignment (#identifier-alignment-explained) with the Author Domain.

5.1.2. Configure Sending System for DKIM Signing Using an Aligned Domain

To configure DKIM for DMARC, the Domain Owner MUST send mail that has a DKIM Signing Domain (#dkim-signing-domain) that will produce a DKIM-Authenticated Identifier (#dkim-identifiers) that has Identifier Alignment (#identifier-alignment-explained) with the Author Domain.

5.1.3. Set Up a Mailbox to Receive Aggregate Reports

Proper consumption and analysis of DMARC aggregate reports are essential to any successful DMARC deployment for a Domain Owner. DMARC aggregate reports, which are defined in [I-D.ietf-dmarc-aggregate-reporting], contain valuable data for the Domain Owner, showing sources of mail using the Author Domain.

5.1.4. Publish a DMARC Policy Record for the Author Domain and Organizational Domain

Once SPF, DKIM, and the aggregate reports mailbox are all in place, it’s time to publish a DMARC Policy Record. For best results, Domain Owners usually start with "p=none", (see Section 5.1.5) with the "rua" tag containing a URI that references the mailbox created in the previous step. This is commonly referred to as putting the Author Domain into Monitoring Mode (#monitoring-mode). If the Organizational Domain is different from the Author Domain, a record also needs to be published for the Organizational Domain.

5.1.5. Collect and Analyze Reports

The reason for starting at "p=none" is to ensure that nothing’s been missed in the initial SPF and DKIM deployments. In all but the most trivial setups, a Domain Owner can overlook a server here or be unaware of a third party sending agreement there. Starting at "p=none", therefore, takes advantage of DMARC’s aggregate reporting function, with the Domain Owner using the reports to audit its own mail streams’ authentication configurations.

While it is possible for a human to read aggregate reports, they are formatted in such a way that it is recommended that they be machine-parsed, so setting up a mailbox involves more than just the physical

creation of that mailbox. Many third-party services exist that will process DMARC aggregate reports or the Domain Owner can create its own set of tools. No matter which method is chosen, the ability to consume these reports and parse the data contained in them will go a long way to ensuring a successful deployment.

5.1.6. Remediate Unaligned or Unauthenticated Mail Streams

DMARC aggregate reports can reveal to the Domain Owner mail streams using the Author Domain but not passing DMARC validation checks. These mail streams may be a combination of illegitimate uses of the domain, such as spoofing or other attempted abuse, and legitimate uses, as in the case of a mail stream created by an agent of the Domain Owner but one which is not passing is due to Authenticated Identifiers being unaligned or missing entirely. For such legitimate uses, these shortcomings MUST be addressed prior to any attempt by the Domain Owner to publish a Domain Owner Assessment Policy (#domain-owner-policy) of Enforcement (#enforcement) for the Author Domain.

5.1.7. Decide Whether to Update Domain Owner Assessment Policy to Enforcement

Once the Domain Owner is satisfied that it is properly authenticating all of its mail, then it is time to decide if it is appropriate to change its Domain Owner Assessment Policy to Enforcement (#enforcement). Depending on its cadence for sending mail, it may take many months of consuming DMARC aggregate reports before a Domain Owner reaches the point where it is sure that it is properly authenticating all of its mail, and the decision on which "p" value to use will depend on its needs.

In making this decision it is important to understand the interoperability issues involved and problems that can result for mailing lists and for delivery of legitimate mail. Those issues are discussed in detail in Section 7.4

5.1.8. A Note on Large, Complex Organizations and Decentralized DNS Management

Large, complex organizations frequently adopt a decentralized model for DNS management, whereby management of a subtree of the name space is delegated to a local department by the central IT organization. In such situations, the "psd" tag makes it possible for those local departments to declare any arbitrary node in their subtree as an Organizational Domain. This would be accomplished by publishing a DMARC Policy Record at that node with the "psd" tag set to "n". The reasons that departments might declare their own Organizational

Domains include a desire to have different policy settings or reporting URIs than the DMARC Policy Record published for the apex domain.

Such configurations would work in theory, and they might involve domain names with many labels, reflecting the structure of the organization, for example:

* Apex domain (DMARC Policy Record published here): example.com
* Zone cut domain (DMARC Policy Record with "psd=n" published here): b.c.d.e.f.g.example.com
* Author Domain: mail.a.b.c.d.e.f.g.example.com

However, Domain Owners should be aware that due to the anti-abuse protections built into the DNS Tree Walk (#dns-tree-walk), the DMARC Policy Record published at the zone cut domain in this example will never be discovered. A Mail Receiver performing a Tree Walk would only perform queries for these names:

* _dmarc.mail.a.b.c.d.e.f.g.example.com
* _dmarc.c.d.e.f.g.example.com
* _dmarc.d.e.f.g.example.com
* _dmarc.e.f.g.example.com
* _dmarc.f.g.example.com
* _dmarc.g.example.com
* _dmarc.example.com
* _dmarc.com

To avoid this circumstance, Domain Owners wishing to have a specific DMARC Policy Record applied to a given Author Domain (#author-domain) longer than eight labels MUST publish a DMARC Policy Record at that domain’s location in the DNS namespace, as such records are always queried by Mail Receivers that participate in DMARC before the Tree Walk begins. In the above example, this would mean publishing a DMARC Policy Record at the name "_dmarc.mail.a.b.c.d.e.f.g.example.com.".

5.2. PSO Actions

In addition to the DMARC Domain Owner actions, if a PSO (#public-suffix-operator) publishes a DMARC Policy Record it MUST include the "psd" tag (see Section 4.7) with a value of "y" ("psd=y").

5.3. Mail Receiver Actions

Mail Receivers (#mail-receiver) wishing to fully participate in DMARC will apply the DMARC mechanism to inbound email messages when a DMARC Policy Record (#dmarc-policy-record) exists that applies to the Author Domain (#author-domain), and will send aggregate feedback reports to Domain Owners that request them. Mail Receivers might also send failure reports to Domain Owners that request them. The following sections describe how to achieve this.

The steps for applying the DMARC mechanism to an email message can take place during the SMTP transaction, and should do so if the Mail Receiver plans to honor Domain Owner Assessment Policies (#domain-owner-policy) that are at the Enforcement (#enforcement) state.

Many Mail Receivers perform one or both of the underlying Authentication Mechanisms (#authentication-mechanisms) on inbound messages even in cases where no DMARC Policy Record exists for the Author Domain of a given message, or where the Mail Receiver is not participating in DMARC. Nothing in this section is intended to imply that the underlying Authentication Mechanisms should only be performed by Mail Receivers participating in DMARC.

5.3.1. Extract Author Domain

Once the email message has been transmitted to the Mail Receiver, the Mail Receiver extracts the domain in the RFC5322.From header field as the Author Domain. If the domain is a U-label, the domain MUST be converted to an A-label, as described in Section 2.3 of [RFC5890], for further processing.

If zero or more than one domain is extracted from the RFC5322.From header field, then DMARC validation is not possible and the process terminates. In the case where more than one domain is retrieved, the Mail Receiver MAY choose to go forward with DMARC validation anyway. See Section 11.5 for further discussion.

5.3.2. Determine If The DMARC Mechanism Applies

If precisely one Author Domain exists for the message, then perform the step described in DMARC Policy Discovery (#dmarc-policy-discovery) to determine if the DMARC mechanism applies. If a DMARC Policy Record (#dmarc-policy-record) is not discovered during this step, then the DMARC mechanism does not apply and DMARC validation terminates for the message.

5.3.3. Determine If Authenticated Identifiers Exist

For each Authentication Mechanism underlying DMARC, perform the required check to determine if an Authenticated Identifier (#authenticated-identifier) exists for the message if such check has not already been performed. Results from each check must be preserved for later use as follows:

* For SPF, the preserved results MUST include "pass" or "fail", and if "fail", SHOULD include information about the reasons for failure if available. The results MUST further include the domain name used to complete the SPF check.
* For DKIM signature validation checks, for each signature checked, the results MUST include "pass" or "fail", and if "fail", SHOULD include information about the reasons for failure. The results MUST further include the value of the "d" and "s" tags from each checked DKIM signature.

5.3.4. Conduct Identifier Alignment Checks If Necessary

For each Authenticated Identifier found in the message, the Mail Receiver checks to see if the Authenticated Identifier is aligned (#identifier-alignment-evaluation) with the Author Domain.

5.3.5. Determine DMARC "Pass" or "Fail"

If one or more of the Authenticated Identifiers align with the Author Domain, the message is considered to pass the DMARC mechanism check.

If no Authenticated Identifiers exist for the domain, or none of the Authenticated Identifiers align with the Author Domain, the message is considered to fail the DMARC mechanism check.

5.3.6. Apply Policy If Appropriate

Email messages that fail the DMARC mechanism check are handled in accordance with the Mail Receiver’s local policies. These local policies may take into account the Domain Owner Assessment Policy for the Author Domain at the Mail Receiver’s discretion.

If one or more DNS queries required to perform DMARC validation on the message do not complete due to temporary or permanent DNS errors, the message cannot be considered to pass or fail the DMARC mechanism check. In such cases, the Domain Owner Assessment Policy cannot be applied to the message, and any other handling decisions for the message are left to the discretion of the Mail Receiver.

See Section 7.2 for further discussion of topics regarding rejecting messages.

5.3.7. Store Results of DMARC Processing

If the Mail Receiver intends to send aggregate feedback reports and/or failure reports, then results obtained from the application of the DMARC mechanism by the Mail Receiver MUST be preserved for eventual presentation back to the Domain Owner in the form of such reports. Section 4.7 and [I-D.ietf-dmarc-aggregate-reporting] discuss aggregate feedback reports.

5.3.8. Send Aggregate Reports

To ensure maximum usefulness for DMARC across the email ecosystem, Mail Receivers SHOULD generate and send aggregate reports with a frequency of at least once every 24 hours. Such reports provide Domain Owners with insight into all mail streams using Author Domains under the Domain Owner’s control, and aid the Domain Owner in determining whether and when to transition from Monitoring Mode (#monitoring-mode) to Enforcement (#enforcement).

The most common reasons for a Mail Receiver to opt out of sending aggregate reports include resource constraints, local policy against sharing data, and concerns about user privacy.

5.3.9. Optionally Send Failure Reports

Per-message failure reports can be a useful source of information for a Domain Owner, either for debugging deployments or in analyzing attacks, and so Mail Receivers MAY choose to send them. Experience has shown, however, that Mail Receivers rightly concerned about protecting user privacy have either chosen to heavily redact the information in such reports (which can hinder their usefulness) or not send them at all. See [I-D.ietf-dmarc-failure-reporting] for further information.

5.4. Policy Enforcement Considerations

The final handling of any message is always a matter of local policy and is left to the discretion of the Mail Receiver.

A DMARC pass for a message indicates only that the use of the Author Domain (#author-domain) has been validated for that message as authorized by the Domain Owner (#domain-owner). Such authorization does not carry an explicit or implicit value assertion about that message or the Domain Owner, and Mail Receivers MAY choose to reject or quarantine a message even if it passes the DMARC validation check. Mail Receivers are encouraged to maintain anti-abuse technologies to combat the possibility of DMARC-enabled abuse.

Mail Receivers MAY choose to accept email that fails the DMARC validation check even if the published Domain Owner Assessment Policy is "reject". In particular, because of the considerations discussed in [RFC7960] and in Section 7.4 of this document, it is important that Mail Receivers SHOULD NOT reject messages solely because of a published policy of "reject", but that they apply other knowledge and analysis to avoid situations such as rejection of legitimate messages sent in ways that DMARC cannot describe, harm to the operation of mailing lists, and similar.

If a Mail Receiver chooses not to honor the published Domain Owner Assessment Policy to improve interoperability among mail systems, it may increase the likelihood of accepting abusive mail. At a minimum, Mail Receivers SHOULD add the Authentication-Results header field (see [RFC8601]), and it is RECOMMENDED when delivering messages that fail the DMARC validation check.

When Mail Receivers deviate from a published Domain Owner Assessment Policy during message processing they 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 defined in [I-D.ietf-dmarc-aggregate-reporting].

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

6. DMARC Feedback

DMARC Feedback is described in [I-D.ietf-dmarc-aggregate-reporting]

As an operational note for Public Suffix Operators, feedback for non-existent domains can be desirable and useful, just as it can be for Organizational Domains. Therefore, both such entities should consider including "rua=" tags in any DMARC Policy Records they publish for themselves. See Section 10 for discussion of Privacy Considerations.

7. Other Topics

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

7.1. 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 extensive network ranges. Domain Owners (#domain-owner) 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 Policy Record. Furthermore, Domain Owners should periodically review their SPF records to ensure that the authorization conveyed by the records matches the domain’s current needs.

SPF was intended to be implemented early in the SMTP transaction, meaning it’s possible for a message to fail SPF validation prior to any message content being transmitted, and so some Mail Receiver architectures might implement SPF in advance of any DMARC operations. This means that an SPF hard fail ("-") prefix on a sender’s SPF mechanism, such as "-all", could cause a message to be rejected early in the SMTP transaction, before any DMARC processing takes place, if the message fails SPF authentication checks. Domain Owners choosing to use "-all" to terminate SPF records should be aware of this, and should understand that messages that might otherwise pass DMARC due to an aligned DKIM-Authenticated Identifier (#dkim-identifiers) could be rejected solely due to an SPF fail. Moreover, messages rejected early in the SMTP transaction will never appear in aggregate DMARC reports, as the transaction will never proceed to the DATA phase and so the RFC5322.From domain will never be revealed and its DMARC policy will never be discovered. Domain Owners and Mail Receivers (#mail-receiver) can consult [M3SPF] and [M3AUTH] for more discussion of the topic and best practices regarding publishing SPF records and when to reject based solely on SPF failure:

7.2. Rejecting Messages

The DMARC mechanism 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 the DMARC mechanism would then result in the 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 to the DATA command as an indication to the SMTP client that the transaction failed; the SMTP client is then responsible for generating a 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 discards 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 Mail Receiver (#mail-receiver) might indicate in plain text the reason for the rejection by using the word "DMARC" somewhere in the reply text. For example:

550 5.7.1 Email rejected per DMARC policy for example.com

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.

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

7.3. Interoperability Issues

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

Because DMARC relies on SPF [RFC7208] and/or DKIM [RFC6376] to achieve a "pass", their limitations also apply.

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

Mail that is sent by authorized, independent third parties might not be sent with Identifier Alignment, also preventing a "pass" result. A Domain Owner can use DMARC aggregate reports to identify this mail and take steps to address authentication shortcomings.

7.4. Interoperability Considerations

As discussed in "Interoperability Issues between DMARC and Indirect Email Flows" [RFC7960], use of "p=reject" can be incompatible with and cause interoperability problems to indirect message flows such as "alumni forwarders", role-based email aliases, and mailing lists across the Internet.

As an example of this, a bank might send only targeted messages to account holders. Those account holders might have given their bank addresses such as "jones@alumni.example.edu" (an address that relays the messages to another address with a real mailbox) or "finance@association.example" (a role-based address that does similar relaying for the current head of finance at the association). When such mail is delivered to the actual recipient mailbox, it will most likely fail SPF checks unless the RFC5321.MailFrom address is rewritten by the relaying MTA, as the incoming IP address will be that of "example.edu" or "association.example", and not an IP address authorized by the originating RFC5321.MailFrom domain. DKIM signatures will generally remain valid in these relay situations.

| It is therefore critical that domains that publish "p=reject" MUST NOT rely solely on SPF to secure a DMARC pass, and MUST apply valid DKIM signatures to their messages.

In the case of domains that have general users who send routine email, those that publish a Domain Owner Assessment Policy (#domain-owner-policy) of "p=reject" are likely to create significant interoperability issues. In particular, if users in such domains post messages to mailing lists on the Internet, those messages can cause significant operational problems for the mailing lists and for the subscribers to those lists, as explained below and in [RFC7960].

> | It is therefore critical that domains that host users who might
> | post messages to mailing lists SHOULD NOT publish Domain Owner
> | Assessment Policies of "p=reject". Any such domains wishing to
> | publish "p=reject" SHOULD first take advantage of DMARC aggregate
> | report data for their domain to determine the possible impact to
> | their users, first by publishing "p=none" for at least a month,
> | followed by publishing "p=quarantine" for an equally long period
> | of time, and comparing the message disposition results. Domains
> | that choose to publish "p=reject" SHOULD either implement policies
> | that their users not post to Internet mailing lists and/or inform
> | their users that their participation in mailing lists may be
> | hindered.

As noted in Section 5.4, Mail Receivers (#mail-receivers) need to apply more analysis than just DMARC validation in their disposition of incoming messages. An example of the consequences of honoring a Domain Owner Assessment Policy of "p=reject" without further analysis is that rejecting messages that have been relayed by a mailing list can cause the Mail Receiver’s users to have their subscriptions to that mailing list canceled by the list software’s automated handling of such rejections - it looks to the list manager as though the recipient’s email address is no longer working, so the address is automatically unsubscribed. An example of this scenario, albeit with DKIM Author Domain Signing Practices (ADSP) rather than DMARC, can be found in Section 5.2 of [RFC6377].

> | It is therefore critical that Mail Receivers MUST NOT reject
> | incoming messages solely on the basis of a "p=reject" policy by
> | the sending domain. Mail Receivers must use the DMARC policy as
> | part of their disposition decision, along with other knowledge and
> | analysis. "Other knowledge and analysis" here might refer to
> | observed sending patterns for properly-authenticated mail using
> | the sending domain, content filtering, etc. In the absence of
> | other knowledge and analysis, Mail Receivers MUST treat such
> | failing mail as if the policy were "p=quarantine" rather than
> | "p=reject".

Failure to understand and abide by these considerations can cause legitimate, sometimes important email to be rejected, can cause operational damage to mailing lists throughout the Internet, and can result in trouble-desk calls and complaints from the Mail Receiver’s employees, customers, and clients.

In practice, despite this advice, few Mail Receivers apply any mitigation techniques when receiving indirect mail flows, few organizations consider the effect of DMARC policies on their users' indirect mail, and it is unlikely that any advice in this document will change that. As a result, mail forwarded through mailing lists with unmodified From: header lines is frequently rejected due to a p=reject policy.

In the ten years since large consumer mail systems started publishing p=reject policies, mailing list software has all adopted workarounds to make the From: header line DMARC aligned. Some simply use the list's address, while others do per-address modifications intended to be reversible or to allow mail to be forwarded back to the original author, e.g., bob@example.com turned into bob=example.com@user.somelist.example. While these workarounds are far from ideal, they are firmly established and list operators treat them as a fact of life.

Mail developers have been trying for a decade to invent technical methods to allow mailing lists to continue to work without modifying the From: header line, with a prominent example being the Authenticated Received Chain (ARC) protocol described in [RFC8617]. While work continues, as of this document's publication, none of the methods have become widely used. Should such a technical method achieve widespread adoption in the future, this document can be updated to reflect that.

8. Conformance Requirements for Full DMARC Participation

This document describes the DMARC mechanism, and allows Domain Owners and Mail Receivers some leeway in deciding which parts of the mechanism to implement. This section summarizes the requirements for full participation in DMARC, either by Domain Owners or by Mail Receivers.

In order to fully participate in DMARC, Domain Owners:

* MUST send mail so it produces an SPF-Authenticated identifier that has Identifier Alignment with the Author Domain
* MUST send mail that has a DKIM Signing Domain that will produce a DKIM-Authenticated Identifier that has Identifier Alignment with the Author Domain
* MUST set up a mailbox to receive aggregate reports and collect and analyze those reports
* MUST publish a DMARC Policy Record for the Author Domain and the Organizational Domain, if it differs from the Author Domain
* MUST NOT rely solely on SPF for a DMARC pass if the DMARC policy for the Author Domain is "p=reject"

In order to fully participate in DMARC, Mail Receivers

* MUST check for the existence of a DMARC Policy Record for the Author Domain of an inbound mail message to determine if the DMARC mechanism applies to that message.
* MUST determine if Authenticated Identifiers exist for the message and preserve the results of those checks for future use in reportging if the DMARC mechanism applies to the message
* MUST conduct necessary Identifier Alignmeent checks if the DMARC mechanism applies for the message and Authenticated Identifiers exist
* MUST use the information from the checks for Authenticated Identifiers to determine if the DMARC validation result is "pass" or "fail" for the message.
* MUST support the "mailto:" URI for sending requested reports
* SHOULD send aggregate reports on at least a daily basis
* MUST NOT reject messages solely on the basis of a "p=reject" policy for the Author Domain

9. IANA Considerations

This section describes actions to be completed by IANA.

9.1. Email Authentication Methods Registry Update

A registry group called "Email Authentication Parameters" exists, and within it a registry group called "Email Authentication Methods" exists and needs to be updated in the manner specified in this section.

The properties of an email message to be evaluated by an email authentication method are registered with IANA in this registry. Entries are assigned only for values that have been documented in a manner that satisfies the terms of Specification Required, per [RFC8126]. Each registration includes the authentication method; the specification that defines the authentication method; the property type (ptype), which is one of the ptype values from the entries in the "Email Authentication Property Types" registry in this same registry group; the property; the value for that property; the status of the property, which is one of "active" or "deprecated"; and its version. The Designated Expert needs to confirm that the provided specification adequately describes the property and the method for its evaluation and clearly presents how they would be used within the DMARC context by Domain Owners and Mail Receivers.

The set of entries to be updated in this registry is as follows:

Method	Defined	ptype	Property	Value	Status	Version
dmarc	[this document]	header	from	The domain portion of the RFC5322.From header field	active	1
dmarc	[this document]	policy	dmarc	The evaluated DMARC policy applied/to be applied after policy options have been processed. Must be "none", "quarantine", or "reject".	active	1

Table 3: "Email Authentication Methods Registry Update"

9.2. Email Authentication Result Names Registry Update

Also within the registry group "Email Authentication Parameters" a registry called "Email Authentication Result Names" exists and should be updated to reference this section of this document.

Result codes for DMARC are registered with IANA in this registry. Entries are assigned only for values that have been documented in a manner that satisfies the terms of Specification Required, per [RFC8126]. Each registration includes the auth method; the code; the specification that defines the code; and the code’s status, which is one of "active" or "deprecated". The "Description" field is included here solely for the reader’s reference, and does not appear in the IANA registry. The Designated Expert needs to confirm that the provided specification adequately describes the result code and clearly presents how it would be used within the DMARC context by Domain Owners and Mail Receivers.

The set of entries to be updated in this registry is as follows:

Auth Method	Code	Specification	Status	Description
dmarc	fail	[this document]	active	A DMARC Policy Record exists for the Author Domain, but no Authenticated Identifier with Identifier Alignment exists
dmarc	none	[this document]	active	No DMARC Policy Record exists for the Author Domain
dmarc	pass	[this document]	active	A DMARC Policy Record exists for the Author Domain, and an Authenticated Identifier with Identifier Alignment exists
dmarc	permerror	[this document]	active	An error occurred during DMARC evaluation that is unrecoverable, such as the retrieval of an improperly formatted DMARC Policy Record. A later attempt is unlikely to produce a final result
dmarc	temperror	[this document]	active	An error occurred during DMARC evaluation that is likely transient in nature, such as a DNS server being temporarily

				unreachable. A later attempt might produce a final result

Table 4: "Email Authentication Result Names Registry Update"

9.3. DMARC Tags Registry Update

A registry group called "Domain-based Message Authentication, Reporting, and Conformance (DMARC)" exists, and within it, a registry called "DMARC Tags" exists. That registry should be updated as described in this section.

Names of tags used in DMARC Policy Records as part of "tag-value" pairs are registered with IANA in this registry. Entries are assigned only for values that have been documented in a manner that satisfies the terms of Specification Required, per [RFC8126]. Each registration includes the tag name, the specification that defines the tag, the status of the tag, and a brief description of the tag. The Designated Expert needs to confirm that the provided specification adequately describes the tag and clearly presents how it would be used within the DMARC context by Domain Owners and Mail Receivers. The "status" column is one of the following:

* "active", meaning the tag is in use in current implementations, and its specifications is expected to be stable
* "experimental", meaning the tag is relatively new and may be in use in some current implementations but not in others, and its specification is not expected to be stable
* "historic", meaning the tag is considered deprecated and is not expected to be in use in any current implementation

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 set of entries to be updated in this registry is as follows:

Tag Name	Reference	Status	Description
adkim	[this document]	active	DKIM Identifier Alignment mode
aspf	[this	active	SPF Identifier Alignment

	document]		mode
fo	[this document]	active	Failure reporting options
np	[this document]	active	Requested Domain Owner Assessment Policy for non-existent subdomains
p	[this document]	active	Requested Domain Owner Assessment Policy
pct	[RFC7489]	historic	Sampling rate
psd	[this document]	active	Indicates whether the DMARC Policy Record is published by a Public Suffix Domain
rf	[RFC7489]	historic	Failure reporting format(s)
ri	[RFC7489]	historic	Aggregate Reporting interval
rua	[this document]	active	Reporting URI(s) for DMARC aggregate feedback reports
ruf	[this document]	active	Reporting URI(s) for message-specific DMARC failure reports
sp	[this document]	active	Requested Domain Owner Assessment Policy for subdomains
t	[this document]	active	DMARC policy test mode
v	[this document]	active	DMARC specification version

Table 5: "DMARC Tags Registry Updatee"

9.4. DMARC Report Formats Registry Update

Also within the registry group "Domain-based Message Authentication, Reporting, and Conformance (DMARC)" a registry called "DMARC Report Formats" exists and should be updated to reference this document.

Names of DMARC failure reporting formats are registered with IANA in this registry. Entries are assigned only for values that have been documented in a manner that satisfies the terms of Specification Required, per [RFC8126]. In addition to a reference to a permanent specification, each registration includes the format name, the format's status, and a brief description of the format. 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 "status" column is one of the following:

* "active", meaning the format is in use in current implementations, and its specifications is expected to be stable
* "experimental", meaning the format is relatively new and may be in use in some current implementations but not in others, and its specification is not expected to be stable
* "historic", meaning the format is considered deprecated and is not expected to be in use in any current implementation

The entry to be updated in this registry is as follows:

Format Name	Reference	Status	Description
afrf	[this document]	active	Authentication Failure Reporting Format (see [RFC6591])

Table 6: "DMARC Report Formats Registry Update"

9.5. Underscored and Globally Scoped DNS Node Names Registry Update

A registry group called "Domain Name System (DNS) Parameters" exists, and within it, a registry called "Underscored and Globally Scoped DNS Node Names" exists, and that registry should be updated to reference this document.

The names of DNS Resource Records beginning with an underscore character that are globally scoped (as per [RFC8552]) are registered with IANA in this registry. In addition to a reference to a permanent specification, each registration contains the DNS Resource Record (RR) type and Node Name. The Designated Expert needs to confirm that the provided specification adequately describes the Node Name and clearly presents how it would be used within the DMARC context by Domain Owners and Mail Receivers.

The entry to be updated in this registry is as follows:

RR Type	_NODE NAME	Reference
TXT	_dmarc	[this document]

Table 7: "Underscored and Globally Scoped DNS Node Names Registry Update"

10. Privacy Considerations

This section discusses issues specific to private data that may be included if DMARC reports are requested. Issues associated with sending aggregate reports and failure reports are addressed in [I-D.ietf-dmarc-aggregate-reporting] and [I-D.ietf-dmarc-failure-reporting] respectively.

10.1. Aggregate Report Considerations

Aggregate reports may, particularly for small organizations, provide some limited insight into email sending patterns. As an example, in a small organization, an aggregate report from a particular domain may be sufficient to make the Report Consumer aware of sensitive personal or business information. If setting an "rua" tag in a DMARC Policy Record, the reporting address needs controls appropriate to the organizational requirements to mitigate any risk associated with receiving and handling reports.

In the case of "rua" requests for multi-organizational PSDs, additional information leakage considerations exist. Multi-organizational PSDs that do not mandate DMARC use by registrants risk exposure of private data of registrant domains if they include the "rua" tag in their DMARC Policy Record.

10.2. Failure Report Considerations

Failure reports do provide insight into email sending patterns, including specific users. If requesting failure reports, data management controls are needed to support appropriate management of this information. The additional detail available through failure reports (relative to aggregate reports) can drive a need for additional controls. As an example, a company may be legally restricted from receiving data related to a specific subsidiary. Before requesting failure reports, any such data spillage risks have to be addressed through data management controls or publishing DMARC Policy Records for relevant subdomains to prevent reporting on data related to their emails.

Due to the nature of the email contents which may be shared through Failure Reports, most Mail Receivers refuse to send them out of privacy concerns. Out of band agreements between Report Consumers and Mail Receivers may be required to address these concerns.

DMARC Policy Records for multi-organizational PSDs MUST NOT include the "ruf" tag.

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.

Both of the email authentication methods that underlie DMARC provide some assurance that an email was transmitted by an MTA which is authorized to do so. SPF policies map domain names to sets of authorized MTAs (see Section 11.4 of [RFC7208]). Validated DKIM signatures indicate that an email was transmitted by an MTA with access to a private key that matches the published DKIM key record.

Whenever mail is sent, there is a risk that an overly permissive source may send mail that will receive a DMARC pass result that was not, in fact, intended by the Domain Owner. These results may lead to issues when systems interpret DMARC pass results to indicate a message is in some way authentic. They also allow such unauthorized senders to evade the Domain Owner’s intended message handling for DMARC validation failures.

To avoid this risk one must ensure that no unauthorized source can add DKIM signatures to the domain’s mail or transmit mail which will evaluate as SPF pass. If, nonetheless, a Domain Owner wishes to include a permissive source in a domain’s SPF record, the source can be excluded from DMARC consideration by using the "?" qualifier on the SPF record mechanism associated with that source. The DMARC working group had a lively discussion about possibly eliminating SPF entirely as an underlying Authentication Mechanism for DMARC, but consensus was not reached, and the suggestion to use the "?" qualifier for permissive sources is presented here instead.

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. Records found in the DNS such as MX, NS, and others 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. This all means that 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 Authentication Mechanisms (SPF and DKIM) depend on the security of the DNS. Examples of how hostile parties can have an adverse impact on DNS traffic include:

* If they can snoop on DNS traffic, they can get an idea of who is receiving mail using the domain(s) in question.

* If they can block outgoing or reply DNS messages, they can prevent systems from discovering senders’ DMARC policies.

* If they can send forged response packets, they can make aligned mail appear unaligned or vice-versa.

None of these threats are unique to DMARC, and they can be addressed using a variety of techniques, including, but not limited to:

* Signing DNS records with Domain Name System Security Extensions (DNSSEC) [RFC9364], which enables recipients to validate the integrity of DNS data and detect and discard forged responses.

* DNS over TLS [RFC7858] or DNS over HTTPS [RFC8484] can mitigate snooping and forged responses.

11.4. Display Name Attacks

An increasingly common attack in messaging abuse is the presentation of false information in the display-name portion of the RFC5322.From header 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.

Such attacks, known as display name attacks, are out of scope for DMARC.

11.5. Denial of DMARC Processing Attacks

The declaration in Section 5.3.1 and elsewhere in this document that messages that do not contain precisely one RFC5322.From domain are outside the scope of this document exposes an attack vector that must be taken into consideration.

Because such messages are outside the scope of this document, an attacker can craft messages with multiple RFC5322.From domains, including the spoofed domain, in an effort to bypass DMARC validation and get the fraudulent message to be displayed by the victim’s MUA with the spoofed domain successfully shown to the victim. In those cases where such messages are not rejected due to other reasons (for example, many such messages would violate RFC5322’s requirement that there be precisely one From: header field), care must be taken by the Mail Receiver to recognize such messages as the threats they might be and handle them appropriately.

The case of a syntactically valid multi-valued RFC5322.From field presents a particular challenge. Experience has shown that most such messages are abusive and/or unwanted by their recipients, and given this fact, a Mail Receiver may make a negative disposition decision for the message prior to and instead of its being subjected to DMARC processing. However, in a case where a Mail Receiver requires that the message be subject to DMARC validation, a recommended approach as per [RFC7489] is to apply the DMARC mechanism to each domain found in the RFC5322.From field as the Author Domain and apply the most strict policy selected among the checks that fail. Such an approach might prove useful for a small number of Author Domains, but it is possible that applying such logic to messages with a large number of domains (where "large" is defined by each Mail Receiver) will expose the Mail Receiver to a form of denial of service attack. Limiting the number of Author Domains processed will avoid this risk. If not all Author Domains are processed, then the DMARC evaluation is incomplete.

11.6. External Reporting Addresses

To avoid abuse by bad actors, reporting addresses generally have to be inside the domains about which reports are requested. To accommodate special cases such as a need to get reports about domains that cannot actually receive mail, Section 3 of [I-D.ietf-dmarc-aggregate-reporting] describes a DNS-based mechanism for validating 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 DNSSEC [RFC9364] is advisable if this is a concern.

11.7. Secure Protocols

This document encourages the use of secure transport mechanisms to prevent the 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.

11.8. Relaxed Alignment Considerations

The DMARC mechanism allows both DKIM- and SPF-Authenticated Identifiers (#identifier-alignment-explained) to validate authorized use of an Author Domain (#author-domain) on behalf of a Domain Owner (#domain-owner). If malicious or unaware users can gain control of the SPF record or DKIM selector records for a subdomain of the Organizational Domain, the subdomain can be used to generate email that achieves a DMARC pass on behalf of the Organizational Domain.

A scenario such as this could occur under the following conditions:

* A DMARC Policy Record exists for the domain "example.com", such that "example.com" is an Organizational Domain
* An attacker controls DNS for the domain "evil.example.com" and publishes an SPF record for that domain
* The attacker sends email with RFC5322.From header field containing "foo@example.com" and an SPF-Authenticated Identifier of "evil.example.com"

Although this email was not authorized by the Domain Owner, it can produce a DMARC pass because the SPF-Authenticated Identifier ("evil.example.com") has Identifier Alignment with the Author Domain ("example.com").

The Organizational Domain Owner should be careful not to delegate control of subdomains if this is an issue, and consider using the Strict Alignment (#strict-alignment) option if appropriate.

DMARC evaluation for relaxed alignment is also highly sensitive to errors in determining the Organizational Domain if the Author Domain does not have a published DMARC Policy Record (#dmarc-policy-record). If an incorrectly selected Organizational Domain is a parent of the correct Organizational Domain, then relaxed alignment could potentially allow a malicious sender to send mail that achieves a DMARC pass verdict. This potential exists for both the legacy [RFC7489] and current methods for determining the organizational domain, the latter described in Section 4.10.2.

The following example illustrates this possibility:

* Mail is sent with an Author Domain of "a.mail.example.com" and Authenticated Identifiers of "mail.example.com"
* There is no DMARC Policy Record published at "_dmarc.a.mail.example.com"
* There is one published at "_dmarc.mail.example.com" and this is intended to be the Organizational Domain for this message
* There is also a DMARC Policy Record published at "_dmarc.example.com", with default alignment (relaxed)
* An attacker is able to send mail with the Author Domain of "evil.example.com" and an Authenticated Identifier of "mail.example.com"

In this scenario, if a Mail Receiver incorrectly determines the Organizational Domain to be "example.com", then the attacker’s mail will pass DMARC validation checks.

This issue is entirely avoided by the use of Strict Alignment and publishing explicit DMARC Policy Records for all Author Domains used in an organization’s email.

For cases where Strict Alignment is not appropriate, this issue can be mitigated by the Domain Owner periodically (perhaps weekly, or whatever frequency might be appropriate for a given organization’s operational needs) checking the DMARC Policy Records, if any, of PSDs (#public-suffix-domain) above the Organizational Domain in the DNS tree and (for legacy [RFC7489] checking that appropriate PSL entries remain present). If a PSD publishes a DMARC Policy Record without the appropriate "psd=y" tag, Organizational Domain owners can add "psd=n" to their Organizational Domain’s DMARC Policy Record so that the PSD’s DMARC Policy Record will not be incorrectly interpreted to indicate that the PSD is the Organizational Domain.

12. References

12.1. Normative References

[I-D.ietf-dmarc-aggregate-reporting]
Brotman, A., "Domain-based Message Authentication, Reporting, and Conformance (DMARC) Aggregate Reporting", Work in Progress, Internet-Draft, draft-ietf-dmarc-aggregate-reporting-32, 17 March 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-dmarc-aggregate-reporting-32>.

[I-D.ietf-dmarc-failure-reporting]
Jones, S. M. and A. Vesely, "Domain-based Message Authentication, Reporting, and Conformance (DMARC) Failure Reporting", Work in Progress, Internet-Draft, draft-ietf-dmarc-failure-reporting-12, 9 January 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-dmarc-failure-reporting-12>.

[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>.

[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>.

[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>.

[RFC7405] Kyzivat, P., "Case-Sensitive String Support in ABNF", RFC 7405, DOI 10.17487/RFC7405, December 2014, <https://www.rfc-editor.org/info/rfc7405>.

[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>.

12.2. Informative References

[M3AUTH] "M3AAWG Email Authentication Recommended Best Practices", <https://www.m3aawg.org/sites/default/files/m3aawg-email-authentication-recommended-best-practices-09-2020.pdf>.

[M3SPF] "M3AAWG Best Practices for Managing SPF Records", <https://www.m3aawg.org/Managing-SPF-Records>.

[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>.

[RFC2308] Andrews, M., "Negative Caching of DNS Queries (DNS NCACHE)", RFC 2308, DOI 10.17487/RFC2308, March 1998, <https://www.rfc-editor.org/info/rfc2308>.

[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>.

[RFC4870] Delany, M., "Domain-Based Email Authentication Using Public Keys Advertised in the DNS (DomainKeys)", RFC 4870, DOI 10.17487/RFC4870, May 2007, <https://www.rfc-editor.org/info/rfc4870>.

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

[RFC7489] Kucherawy, M., Ed. and E. Zwicky, Ed., "Domain-based Message Authentication, Reporting, and Conformance (DMARC)", RFC 7489, DOI 10.17487/RFC7489, March 2015, <https://www.rfc-editor.org/info/rfc7489>.

[RFC7858] Hu, Z., Zhu, L., Heidemann, J., Mankin, A., Wessels, D., and P. Hoffman, "Specification for DNS over Transport Layer Security (TLS)", RFC 7858, DOI 10.17487/RFC7858, May 2016, <https://www.rfc-editor.org/info/rfc7858>.

[RFC7960] Martin, F., Ed., Lear, E., Ed., Draegen, T., Ed., Zwicky, E., Ed., and K. Andersen, Ed., "Interoperability Issues between Domain-based Message Authentication, Reporting, and Conformance (DMARC) and Indirect Email Flows", RFC 7960, DOI 10.17487/RFC7960, September 2016, <https://www.rfc-editor.org/info/rfc7960>.

[RFC8020] Bortzmeyer, S. and S. Huque, "NXDOMAIN: There Really Is Nothing Underneath", RFC 8020, DOI 10.17487/RFC8020, November 2016, <https://www.rfc-editor.org/info/rfc8020>.

[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>.

[RFC8484] Hoffman, P. and P. McManus, "DNS Queries over HTTPS (DoH)", RFC 8484, DOI 10.17487/RFC8484, October 2018, <https://www.rfc-editor.org/info/rfc8484>.

[RFC8551] Schaad, J., Ramsdell, B., and S. Turner, "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 4.0 Message Specification", RFC 8551, DOI 10.17487/RFC8551, April 2019, <https://www.rfc-editor.org/info/rfc8551>.

[RFC8552] Crocker, D., "Scoped Interpretation of DNS Resource Records through "Underscored" Naming of Attribute Leaves", BCP 222, RFC 8552, DOI 10.17487/RFC8552, March 2019, <https://www.rfc-editor.org/info/rfc8552>.

[RFC8617] Andersen, K., Long, B., Ed., Blank, S., Ed., and M. Kucherawy, Ed., "The Authenticated Received Chain (ARC) Protocol", RFC 8617, DOI 10.17487/RFC8617, July 2019, <https://www.rfc-editor.org/info/rfc8617>.

[RFC9091] Kitterman, S. and T. Wicinski, Ed., "Experimental Domain-Based Message Authentication, Reporting, and Conformance (DMARC) Extension for Public Suffix Domains", RFC 9091, DOI 10.17487/RFC9091, July 2021, <https://www.rfc-editor.org/info/rfc9091>.

[RFC9364] Hoffman, P., "DNS Security Extensions (DNSSEC)", BCP 237, RFC 9364, DOI 10.17487/RFC9364, February 2023, <https://www.rfc-editor.org/info/rfc9364>.

Appendix A. Technology Considerations

This section documents some design decisions made in the development of DMARC. Specifically addressed here are some suggestions that were considered but not included in the design, with explanatory text regarding the decision.

A.1. S/MIME

S/MIME, or Secure Multipurpose Internet Mail Extensions [RFC8551], is a standard for encrypting and signing 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 validate 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 are both 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 instruct Mail 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 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 [RFC4870], 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 the 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 validation algorithm in terms of which policy is to be applied and when.

A.4. Domain Existence Test

The presence of the "np" tag in this specification seemingly implies that there would be an agreed-upon standard for determining a domain’s existence.

Since the DMARC mechanism is focused on email, one might think that the definition of "resolvable" in [RFC5321] applies. Using that definition, only names that resolve to MX Resource Records (RRs), A RRs, or AAAA RRs are deemed to be resolvable and to exist in the DNS. This is a common practice among Mail Receivers to determine whether or not to accept a mail message before performing other more expensive processing.

The DMARC mechanism makes no such requirement for the existence of specific DNS RRs in order for a domain to exist; instead, if any RR exists for a domain, then the domain exists. To use the terminology from [RFC2308], an "NXDOMAIN" response (rcode "Name Error") to a DNS query means that the domain name does not exist, while a "NODATA" response (rcode "NOERROR") means that the given resource record type queried for does not exist, but the domain name does.

Furthermore, in keeping with [RFC8020], if a query for a name returns NXDOMAIN, then not only does the name not exist, every name below it in the DNS hierarchy also does not exist.

A.5. Organizational Domain Discovery Issues

An earlier informational version of the DMARC mechanism [RFC7489] noted that 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. That version further mentioned suggestions that 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.

That previous version posited that one could "climb the tree" to find the Organizational Domain, but expressed concern that an attacker could exploit this for a denial-of-service attack through sending a high number of messages each with a relatively large number of nonsense labels, causing a Mail Receiver to perform a large number of DNS queries in search of a DMARC Policy Record. This version defines a method for performing a DNS Tree Walk (#dns-tree-walk), and further mitigates the risk of the denial-of-service attack by expressly limiting the number of DNS queries to execute regardless of the number of labels in the domain name.

Readers curious about the previous method for Organizational Domain Discovery are directed to Section 3.2 of [RFC7489].

A.6. Removal of the "pct" Tag

An earlier informational version of the DMARC mechanism [RFC7489] included a "pct" tag and specified all integers from 0 to 100 inclusive as valid values for the tag. The intent of the tag was to provide domain owners with a method to gradually change their preferred Domain Owner Assessment Policy (the "p" tag) from "none" to "quarantine" or from "quarantine" to "reject" by requesting the stricter treatment for just a percentage of messages that produced DMARC results of "fail".

Operational experience showed that the pct tag was usually not accurately applied, unless the value specified was either 0 or 100 (the default), and the inaccuracies with other values varied widely from one implementation to another. The default value was easily implemented, as it required no special processing on the part of the Mail Receiver, while the value of 0 took on unintended significance as a value used by some intermediaries and mailbox providers as an indicator to deviate from standard handling of the message, usually by rewriting the RFC5322.From header field in an effort to avoid DMARC failures downstream.

These custom actions when the "pct" tag was set to 0 proved valuable to the email community. In particular, header field rewriting by an intermediary meant that a Domain Owner’s aggregate reports could reveal to the Domain Owner how much of its traffic was routing through intermediaries that don’t rewrite the RFC5322.From header field. Such information wasn’t explicit in the aggregate reports received; rather, sussing it out required work on the part of the Domain Owner to compare aggregate reports from before and after the "p" value was changed and "pct=0" was included in the DMARC Policy Record, but the data was there. Consequently, knowing how much mail

was subject to possible DMARC failure due to a lack of RFC5322.From header field rewriting by intermediaries could assist the Domain Owner in choosing whether to move from Monitoring Mode (#monitoring-mode) to Enforcement (#enforcement). Armed with this knowledge, the Domain Owner could make an informed decision regarding subjecting its mail traffic to possible DMARC failures based on the Domain Owner's tolerance for such things.

Because of the value provided by "pct=0" to Domain Owners, it was logical to keep this functionality in the protocol; at the same time, it didn't make sense to support a tag named "pct" that had only two valid values. This version of the DMARC mechanism, therefore, introduces the "t" tag as shorthand for "testing", with the valid values of "y" and "n", which are meant to be analogous in their application by mailbox providers and intermediaries to the "pct" tag values "0" and "100", respectively.

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 header fields and relevant SMTP commands 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. Alignment cannot exist if SPF does not produce a passing result.

Example 1: SPF in Strict 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 domain and the Author Domain are identical. Thus, the identifiers are in Strict Alignment.

Example 2: SPF in Relaxed Alignment:

```
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 Author Domain (example.com) is a parent of the RFC5321.MailFrom domain. Thus, the identifiers are in relaxed alignment because they both have the same Organizational Domain (example.com).

Example 3: No SPF Identifier 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 domain that is neither the same as, a parent of, nor a child of the Author Domain. Thus, the identifiers are not in alignment.

B.1.2. DKIM

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

Example 1: DKIM in Strict 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" tag and the Author Domain have identical DNS domains. Thus, the identifiers are in Strict Alignment.

Example 2: DKIM in Relaxed Alignment:

```
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" tag includes a DNS domain that is a parent of the Author Domain. Thus, the identifiers are in relaxed alignment, as they have the same Organizational Domain (example.com).

Example 3: No DKIM Identifier Alignment:

```
DKIM-Signature: v=1; ...; d=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 DKIM signature's "d" tag includes a DNS domain that is neither the same as, a parent of, nor a child of the Author 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 DMARC Policy Record for the Domain Owner's Organizational Domain.

B.2.1. Entire Domain, Monitoring Mode

The Domain Owner for "example.com" has deployed SPF and DKIM on its messaging infrastructure. The Domain Owner wishes to begin using DMARC with a policy that will solicit aggregate feedback from Mail Receivers without affecting how the messages are processed in order to:

* Confirm that its legitimate messages are authenticating correctly
* Validate that all authorized message sources have implemented authentication measures
* Determine how many messages from other sources would be affected by publishing a Domain Owner Assessment Policy at Enforcement

The Domain Owner accomplishes this by constructing a DMARC Policy Record indicating that:

* The version of DMARC being used is "DMARC1" ("v=DMARC1;")
* Mail Receivers should not alter how they treat these messages because of this DMARC Policy Record ("p=none")

* Aggregate feedback reports are 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 "t" tag present, so the default of "n" applies).

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 Policy 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 Mode, Per-Message Failure 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. To diagnose these intermittent problems, they wish to request per-message failure reports when authentication failures occur.

Not all Mail 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 failure report format ([RFC6591]) meets the Domain Owner’s needs in this scenario.

The Domain Owner accomplishes this by adding the following to its DMARC Policy Record from Appendix B.2.1:

* Per-message failure reports are sent via email to the address "auth-reports@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 Policy 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 serve as a Report Consumer. Again, not all Mail Receivers will honor this request, but those that do MUST implement additional checks to validate that the third party authorizes reception of failure reports on behalf of this domain.

The Domain Owner needs to alter its DMARC Policy Record from Appendix B.2.2 as follows:

* Per-message failure reports are sent via email to the address "auth-reports@thirdparty.example.net" ("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 Policy 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 Mail Receivers MUST implement additional checks as described in Section 3 of [I-D.ietf-dmarc-aggregate-reporting]. To pass these additional checks, the Report Consumer’s Domain Owner will need to publish an additional DMARC Policy Record as follows:

* Given the DMARC Policy Record published by the Domain Owner at "_dmarc.example.com", the DNS administrator for the Report Consumer will need to publish a TXT resource record at "example.com._report._dmarc.thirdparty.example.net" with the value "v=DMARC1;" to authorize receipt of the reports.

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 reports on behalf of example.com
example.com._report._dmarc   IN   TXT    "v=DMARC1;"
```

B.2.4. Overriding destination addresses

The third party Report Consumer can also publish "rua" and "ruf" tags in order to override the specific address published by example.com with a different address in the same third party domain. This may be necessary if the third party Report Consumer has changed its email address, or want to guard against typos in the DMARC Policy Record of the Author Domain. Intermediaries and other third parties should refer to Section 3 of [I-D.ietf-dmarc-aggregate-reporting] for the full details of this mechanism.

The third party Report Consumer accomplishes this by adding the following to its DMARC Policy Record from Appendix B.2.3:

* The override address for aggregate reports is "aggregate-reports@thirdparty.example.net" ("rua=mailto:aggregate-reports@thirdparty.example.net")
* The override address for failure reports is "failure-reports@thirdparty.example.net" ("ruf=mailto:failure-reports@thirdparty.example.net")

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 reports on behalf of example.com
; Override destination mailboxes
example.com._report._dmarc   IN   TXT    (
        "v=DMARC1; "
        "rua=mailto:aggregate-reports@thirdparty.example.net; "
        "ruf=mailto:failure-reports@thirdparty.example.net" )
```

In this case only the "ruf" tag is actually overridden, because, in the previous example, failure reporting is the only reporting type that was directed to the third party Report Consumer.

B.2.5. Subdomain, Testing, 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 express a handling preference for messages from this subdomain that fail DMARC validation to indicate to participating Mail Receivers that use of this domain is not valid.

As a first step, it will express that it considers messages using this subdomain that fail DMARC validation to be suspicious. The goal here will be to enable examination of messages sent to mailboxes

hosted by participating Mail Receivers as a method for troubleshooting any existing authentication issues. Aggregate feedback reports will be sent to a mailbox within the Organizational Domain, and to a mailbox at a Report Consumer selected and authorized to receive them by the Domain Owner.

The Domain Owner will accomplish this by constructing a DMARC Policy Record indicating that:

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

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

* Mail Receivers are advised that the Domain Owner considers messages that fail to authenticate to be suspicious ("p=quarantine")

* Aggregate feedback reports are sent via email to the addresses "dmarc-feedback@example.com" and "example-tld-test@thirdparty.example.net" ("rua=mailto:dmarc-feedback@example.com, mailto:example-tld-test@thirdparty.example.net")

* The Domain Owner desires only that an actor performing a DMARC validation check apply any special handling rules it might have in place, such as rewriting the RFC53322.From header field; the Domain Owner is testing its setup at this point and so does not want the Domain Owner Assessment Policy to be applied. ("t=y")

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 Policy Record for the domain test.example.com
_dmarc IN  TXT  ( "v=DMARC1; p=quarantine; "
                  "rua=mailto:dmarc-feedback@example.com,"
                  "mailto:tld-test@thirdparty.example.net; "
                  "t=y" )
```

Once enough time has passed to allow for collecting enough reports to give the Domain Owner confidence that all authorized email sent using the subdomain is properly authenticating and passing DMARC validation checks, then the Domain Owner can update the DMARC Policy Record to indicate that it considers validation failures to be a clear indication that use of the subdomain is not valid. It would do this by altering the record to advise Mail Receivers of its position on such messages ("p=reject") and removing the testing flag ("t=y").

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 Policy Record for the domain test.example.com
_dmarc IN  TXT  ( "v=DMARC1; p=reject; "
                  "rua=mailto:dmarc-feedback@example.com,"
                  "mailto:tld-test@thirdparty.example.net" )
```

B.3. Mail Receiver Example

A Mail Receiver that wants to participate in 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 Consumers).

B.3.1. SMTP Session Example

An optimal DMARC-enabled Mail Receiver performs validation and Identifier Alignment checking during the SMTP [RFC5321] conversation.

Before 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 Record lookup.

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

Given a DMARC Policy 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 Policy 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 Policy Record:
  "v=DMARC1; p=reject; aspf=r;
   rua=mailto:dmarc-feedback@example.com"
```

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

If no Authenticated Identifiers align with the Author Domain, then the Mail Receiver applies the Domain Owner Assessment Policy. However, before this action is taken, the Mail Receiver can consult external information to override the Domain Owner Assessment 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 Assessment 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 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 Policy Record requests AFRF-based reporting. If so, then the Mail Receiver can emit an AFRF to the reporting address supplied in the DMARC Policy 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 unnecessary if the Domain Owner has not requested aggregate reports, i.e., no "rua" tag was found in the policy record.

B.4. Organizational and Policy Domain Tree Walk Examples

If an Author Domain has no DMARC Policy Record, a Mail Receiver uses a tree walk to find the DMARC Policy.

If the DMARC Policy Record allows relaxed alignment and the SPF- or DKIM-Authenticated Identifiers are different from the Author Domain, a Mail Receiver uses a tree walk to discover the respective Organizational Domains to determine Identifier Alignment.

B.4.1. Simple Organizational and Policy Example

A Mail Receiver receives an email with:

* Author Domain: example.com
* RFC5321.MailFrom Domain: example.com
* DKIM-Authenticated Identifier: signing.example.com

In this example, "_dmarc.example.com" and "_dmarc.signing.example.com" both have DMARC Policy Records while "_dmarc.com" does not. If SPF or DKIM yield pass results, they still have to be aligned to support a DMARC pass. Since not all domains are the same, if the alignment is relaxed then the tree walk is performed to determine the Organizational Domain for each.

To determine the Organizational Domain for the Author Domain, query "_dmarc.example.com" and "_dmarc.com"; "example.com" is the last element of the DNS tree with a DMARC Policy Record, so it is the Organizational Domain for "example.com".

For the RFC5321.MailFrom domain, the Organizational Domain already found for "example.com" is "example.com", so SPF is aligned.

To determine the Organizational Domain for the DKIM-Authenticated Identifier, query "_dmarc.signing.example.com", "_dmarc.example.com", and "_dmarc.com". Both "signing.example.com" and "example.com" have DMARC Policy Records, but "example.com" is the highest element in the tree with a DMARC Policy Record (it has the fewest labels), so "example.com" is the Organizational Domain. Since this is also the Organizational Domain for the Author Domain, DKIM is aligned for relaxed alignment.

Since both SPF and DKIM are aligned, they can be used to determine if the message has a DMARC pass result. If the result is not pass, then the policy domain’s DMARC Policy Record is used to determine the appropriate policy. In this case, since the RFC5322.From domain has a DMARC Policy Record, that is the policy domain.

B.4.2. Deep Tree Walk Example

A Mail Receiver receives an email with:

* Author Domain: a.b.c.d.e.f.g.h.i.j.k.example.com
* RFC5321.MailFrom Domain: example.com
* DKIM-Authenticated Identifier: signing.example.com

Both "_dmarc.example.com" and "_dmarc.signing.example.com" have DMARC Policy Records, while "_dmarc.com" does not. If SPF or DKIM yield pass results, they still have to be aligned to support a DMARC pass. Since not all domains are the same, if the alignment is relaxed then the tree walk is performed to determine the Organizational Domain for each.

To determine the Organizational Domain For the Author Domain, query "_dmarc.a.b.c.d.e.f.g.h.i.j.k.example.com", then query "_dmarc.g.h.i.j.k.example.com" (skipping the intermediate names), then query "_dmarc.h.i.j.k.example.com", "_dmarc.i.j.k.example.com", "_dmarc.j.k.example.com", "_dmarc.k.example.com", "_dmarc.example.com", and "_dmarc.com". None of "a.b.c.d.e.f.g.h.i.j.k.example.com", "g.h.i.j.k.example.com", "h.i.j.k.example.com", "i.j.k.example.com", "j.k.example.com", or "k.example.com" have a DMARC Policy Record.

Since "example.com" is the last element of the DNS tree with a DMARC Policy Record, it is the Organizational Domain for "a.b.c.d.e.f.g.h.i.j.k.example.com".

For the RFC5321.MailFrom domain, the Organizational domain already found for "example.com" is "example.com". SPF is aligned.

For the DKIM-Authenticated Identifier, query "_dmarc.signing.example.com", "_dmarc.example.com", and "_dmarc.com". Both "signing.example.com" and "example.com" have DMARC Policy Records, but "example.com" is the highest element in the tree with a DMARC Policy Record, so "example.com" is the Organizational Domain. Since this is also the Organizational Domain for the Author Domain, DKIM is aligned for relaxed alignment.

Since both SPF and DKIM are aligned, they can be used to determine if the message has a DMARC pass result. If the results for both are not pass, then the policy domain’s DMARC Policy Record is used to determine the appropriate policy. In this case, the Author Domain does not have a DMARC Policy Record, so the policy domain is the highest element in the DNS tree with a DMARC Policy Record, example.com.

B.4.3. Example with a PSD DMARC Policy Record

In rare cases, a PSD publishes a DMARC Policy Record with a psd tag, which the tree walk must take into account.

A Mail Receiver receives an email with:

* Author Domain: giant.bank.example
* RFC5321.MailFrom Domain: mail.giant.bank.example
* DKIM-Authenticated Identifier: mail.mega.bank.example

In this case, "_dmarc.bank.example" has a DMARC Policy Record which includes the "psd=y" tag, and "_dmarc.example" does not have a DMARC Policy Record. While "_dmarc.giant.bank.example" has a DMARC Policy Record without a "psd" tag, "_dmarc.mega.bank.example" and "_dmarc.mail.mega.bank.example" have no DMARC Policy Records.

Since the three domains are all different, tree walks find their Organizational Domains to see which are aligned.

For the Author Domain "giant.bank.example", the tree walk finds the DMARC Policy Record at "_dmarc.giant.bank.example", then the DMARC Policy Record at "_dmarc.bank.example", and stops because of the "psd=y" flag. The Organizational Domain is "giant.bank.example" because it is the domain directly below the one with "psd=y". Since the Organizational Domain has a DMARC Policy Record, it is also the policy domain.

For the RFC5321.MailFrom domain "mail.giant.bank.example", the tree walk finds no DMARC Policy Record at "_dmarc.mail.giant.bank.example", but does find both the DMARC Policy Record at "_dmarc.giant.bank.example" and then the DMARC Policy Record at "_dmarc.bank.example", and stops because of the "psd=y" flag. Again the Organizational Domain is "giant.bank.example" because it is the domain directly below the one with "psd=y". Since this is the same Organizational Domain as the Author Domain, SPF is aligned.

For the DKIM-Authenticated Identifier "mail.mega.bank.example", the tree walk finds no DMARC Policy Records at "_dmarc.mail.mega.bank.example" or "_dmarc.mega.bank.example", then finds the DMARC Policy Record at "_dmarc.bank.example" and stops because of the "psd=y" flag. The Organizational Domain is "mega.bank.example", so DKIM is not aligned.

Since SPF is aligned, it can be used to determine if the message has a DMARC pass result. If the result is not pass, then the policy domain's DMARC Policy Record is used to determine the appropriate policy.

B.5. Utilization of Aggregate Feedback: Example

Aggregate feedback is consumed by Domain Owners to enable their understanding of how a given domain is being processed by the Mail Receiver. Aggregate reporting data on emails that pass all underlying authentication checks is used by Domain Owners to validate that their 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 validate 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 produce an Authenticated Identifier, 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, if the mail that is failing the checks was generated by or on behalf of the Domain Owner. Data regarding failing authentication checks can also allow the Domain Owner to come to an understanding of how its domain is being misused.

Appendix C. Changes from RFC 7489

This document is intended to render [RFC7489] obsolete. As one might guess, that means there are significant differences between RFC 7489 and this document. This section will summarize those changes.

C.1. Informational vs. Standards Track

RFC 7489 was not the product of any IETF work stream, but was instead published into the RFC series by the Independent Submissions Editor and is classified as an Informational RFC.

This document, by contrast, is intended to be Internet Standards Track.

C.2. Changes to Terminology and Definitions

The following changes were made to the Terminology and Definitions section.

C.2.1. Terms Added

These terms were added:

* Domain Owner Assessment Policy
* Enforcement
* Monitoring Mode
* Non-existent Domains
* Public Suffix Domain (PSD)
* Public Suffix Operator (PSO)
* PSO Controlled Domain Names

C.2.2. Definitions Updated

These definitions were updated:

* Organizational Domain
* Report Receiver (renamed to Report Consumer)

C.3. Policy Discovery and Organizational Domain Determination

The algorithms for DMARC policy discovery and for determining the Organizational Domain have been changed. Specifically, reliance on a Public Suffix List (PSL) has been replaced by a technique called a "DNS Tree Walk", and the methodology for the DNS Tree Walk is explained in detail in this document.

The DNS Tree Walk also incorporates PSD policy discovery, which was introduced in [RFC9091]. That RFC was an Experimental RFC, and the results of that experiment were that the RFC was not implemented as written. Instead, this document redefines the algorithm for PSD policy discovery, and thus obsoletes [RFC9091]. Specifically, the DNS Tree Walk defined in this document obviates the need for a PSD DMARC registry, and that PSD DMARC registry is what made RFC 9091 an Experimental RFC.

These algorithm changes introduce the possibility of interoperability issues where a Domain Owner expects a DMARC Policy Record or an Organizational Domain to be identified by the Tree Walk process, but a Mail Receiver using an RFC 7489-based implementation of DMARC and relying on a PSL might arrive at a different answer.

This issue is entirely avoided by the use of Strict Alignment and publishing explicit DMARC Policy Records for all Author Domains used in an organization’s email.

C.4. Reporting

Discussion of both aggregate and failure reporting have been moved to separate documents dedicated to the topics.

In addition, the ability to specify a maximum report size in the DMARC URI has been removed.

C.5. Tags

Several tags have been added to the "DMARC Policy Record Format" section of this document since RFC 7489 was published, and at the same time, several others were removed.

C.5.1. Tags Added

* np - Policy for non-existent domains (Imported from [RFC9091])
* psd - Flag indicating whether a domain is a Public Suffix Domain
* t - Replacement for some pct tag functionality. See Appendix A.6 for further discussion

C.5.2. Tags Removed

* pct - Tag requesting application of DMARC policy to only a percentage of messages. See Appendix A.6 for discussion
* rf - Tag specifying requested format of failure reports
* ri - Tag specifying requested interval between aggregate reports

C.6. Expansion of Domain Owner Actions Section

RFC 7489 had just two paragraphs in its Domain Owner Actions section, and while the content of those paragraphs was correct, it was minimalist in its approach to providing guidance to domain owners on just how to implement DMARC.

This document provides much more detail and explanatory text to a Domain Owner, focusing not just on what to do to implement DMARC, but also on the reasons for each step and the repercussions of each decision.

In particular, this document makes explicit that domains for general-purpose email SHOULD NOT deploy a DMARC policy of p=reject. See Section 7.4 for further discussion of this topic.

C.7. Report Generator Recommendations

In the cases where a DMARC Policy Record specifies multiple destinations for either aggregate reports or failure reports, RFC 7489 stated:

> 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.

This document in Section 4.6 says:

> A report **SHOULD** be sent to each listed URI provided in the DMARC Policy Record.

C.8. Removal of RFC 7489 Appendix A.5

One of the appendices in RFC 7489, specifically Appendix A.5, has been removed from the text with this update. The appendix was titled "Issues with ADSP in Operation" and it contained a list of issues associated with ADSP that influenced the direction of DMARC. The ADSP protocol was moved to "Historic" status in 2013 and working group consensus was that such a discussion of ADSP’s influence on DMARC was no longer relevant.

C.9. RFC 7489 Errata Summary

This document and its companion documents ([I-D.ietf-dmarc-aggregate-reporting] and [I-D.ietf-dmarc-failure-reporting]) address the following errata filed against [RFC7489] since that document’s publication in March, 2015. More details on each of these can be found at https://www.rfc-editor.org/errata_search.php?rfc=7489 (https://www.rfc-editor.org/errata_search.php?rfc=7489)

C.9.1. RFC Errata, Erratum ID 5365, RFC 7489, Section 7.2.1.1 (https://www.rfc-editor.org/errata/eid5365)

Addressed in [I-D.ietf-dmarc-aggregate-reporting].

C.9.2. RFC Errata, Erratum ID 5371, RFC 7489, Section 7.2.1.1 (https://www.rfc-editor.org/errata/eid5371)

Addressed in [I-D.ietf-dmarc-aggregate-reporting].

C.9.3. RFC Errata, Erratum ID 5440, RFC 7489, Sections 7.1, B.2.1, B.2.3, and B.2.4 (https://www.rfc-editor.org/errata/eid5440)

This erratum references several mentions in RFC 7489 of the "v=" tag from the Domain Owner Assessment Policy and/or its value, specifically mentions that were not, but should have been, "v=DMARC1;". Some of those mentions are preserved in this document and those mentions have been addressed as per the erratum. The rest have moved to [I-D.ietf-dmarc-aggregate-reporting] and are addressed there.

C.9.4. RFC Errata, Erratum ID 6439, RFC 7489, Section 7.1 (https://www.rfc-editor.org/errata/eid6439)

Addressed in [I-D.ietf-dmarc-aggregate-reporting].

C.9.5. RFC Errata, Erratum ID 5221, RFC 7489, Appendix C (https://www.rfc-editor.org/errata/eid5221)

The regular expression pattern for IP addresses has been removed from this document and from [I-D.ietf-dmarc-aggregate-reporting].

C.9.6. RFC Errata, Erratum ID 5229, RFC 7489, Appendix C (https://www.rfc-editor.org/errata/eid5229)

Addressed in [I-D.ietf-dmarc-aggregate-reporting].

C.9.7. RFC Errata, Erratum 5495, RFC 7489, Section 6.6.3 (https://www.rfc-editor.org/errata/eid5495)

This erratum is in reference to the description of the process documented in RFC 7489 for the applicable DMARC policy for an email message. The process for doing this has drastically changed in DMARCbis, and so the text identified in this erratum no longer exists.

C.9.8. RFC Errata, Erratum ID 6485, RFC 7489, Section 7.2.1.1 (https://www.rfc-editor.org/errata/eid6485)

Addressed in [I-D.ietf-dmarc-aggregate-reporting].

C.9.9. RFC Errata, Erratum ID 6729, RFC 7489, Section 3.2 (https://www.rfc-editor.org/errata/eid6729)

This erratum is in reference to a search of the Public Suffix List (PSL) as part of finding a DMARC Policy Record (a.k.a., Domain Owner Assessment Policy). The PSL is no longer relied upon for this practice, and the text at issue has been removed from this document.

C.9.10. RFC Errata, Erratum ID 7099, RFC 7489, Section 7.2.1.1 (https://www.rfc-editor.org/errata/eid7099)

Addressed in [I-D.ietf-dmarc-aggregate-reporting].

C.9.11. RFC Errata, Erratum ID 7100, RFC 7489, Section 7.2.1.1 (https://www.rfc-editor.org/errata/eid7100)

Addressed in [I-D.ietf-dmarc-aggregate-reporting].

C.9.12. RFC Errata, Erratum ID 7835, RFC 7489, Section 6.6.3 (https://www.rfc-editor.org/errata/eid7835)

This erratum is in reference to the description of the process documented in RFC 7489 for the applicable DMARC policy for an email message. The process for doing this has drastically changed in DMARCbis, and so the text identified in this erratum no longer exists.

C.9.13. RFC Errata, Erratum ID 7865, RFC 7489, Appendix C (https://www.rfc-editor.org/errata/eid7865)

The regular expression pattern for IP addresses has been removed from this document and from [I-D.ietf-dmarc-aggregate-reporting].

C.9.14. RFC Errata, Erratum ID 5151, RFC 7489, Section 1 (https://www.rfc-editor.org/errata/eid5151)

This erratum is in reference to the Introduction section of RFC 7489. That section has been substantially rewritten in DMARCbis, and the text at issue for this erratum no longer exists.

C.9.15. RFC Errata, Erratum ID 5774, RFC 7489, Appendix C (https://www.rfc-editor.org/errata/eid5774)

Addressed in [I-D.ietf-dmarc-aggregate-reporting].

C.10. General Editing and Formatting

A great deal of the content from RFC 7489 was preserved in this document, but much of it was subject to either minor editing, re-ordering of sections, and/or both.

Acknowledgements

This reworking of the DMARC mechanism specified in [RFC7489] is the result of contributions from many participants in the IETF Working Group dedicated to this effort. Although the contributors are too numerous to mention, significant contributions were made by Kurt Andersen, Laura Atkins, Seth Blank, Alex Brotman, Dave Crocker, Douglas E. Foster, Ned Freed, Mike Hammer, Steven M. Jones, Scott Kitterman, Murray S. Kucherawy, Barry Leiba, Alessandro Vesely, and Tim Wicinski.

The authors and contributors also recognize that this document would not have been possible without the work done by those who had a hand in producing [RFC7489]. The Acknowledgements section from that document is preserved in full below.

Acknowledgements - RFC 7489

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 https://dmarc.org (https://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 Andersen, 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

Todd M. Herr
Valimail
Email: todd@someguyinva.com

John Levine
Standcore LLC
Email: standards@standcore.com