Network Working Group
Internet-Draft
Intended status: Experimental
Expires: January 28, 2018

T. Akagiri
Regumi, Inc.
G. Yasutaka
Rakuten, Inc.
K. Okada
T. Hayashi
Lepidum Co. Ltd.
M. Kase
Individual Contributor
July 27, 2017

Mail Divide Framework
draft-akagiri-mail-divide-01

Abstract

Mail Divide Framework (MDF) is a recipient driven partitioning framework for E-Mail delivery. A protocol to divide mail delivery at the source of the message is defined in this draft. A mechanism called Reputation Service Provider is also introduced so that a third-party authority can assure senders' trust. With MDF, subdomaining is used for category-specific MTA designation. Senders decide which category the outgoing mail belongs. It then looks up DNS TXT record to find whether the recipient advertises a specific server for that category. The specified server puts the received mail into a corresponding per-category inbox for the user.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on January 28, 2018.

Table of Contents

1. Terminology

1.1. Key Words

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 [RFC2119].

1.2. Mail Divide Framework (MDF)

A recipient driven partitioning framework for E-Mail delivery. Receivers advertise that it prepares a separate delivery path, or receiving MTA, for a specific category of mail messages. MDF provides a mechanism to advertise and lookup category specific settings, and evaluate conformance of senders via RSPs.

1.3. Mail Category

The intended purpose of each mail message, such as communication, notification, etc. MDF requires that the definition of a Mail Category is agreed upon among senders, receivers, and RSPs.

1.4. Reputation Service Provider (RSP)

Reputation Service Provider keeps track of a white list of MDF-conforming senders. Receiving party can perform a query to see how a specific sender is conforming to MDF.

1.5. DIVIDE record

A DNS TXT resource record that advertises receiver's trust policy. DIVIDE record specifies that a mail message under a category is received by a specific subdomain.

1.6. Imported Definitions

ABNF (Augmented Backus-Naur Form) ABNF is defined in [RFC5234], as are the tokens "ALPHA", "DIGIT", and "SP" (space).

The tokens "Local-part", "Domain", "address-literal" and "Mailbox" are defined in [RFC5321].

"dot-atom", "quoted-string", "comment", "CFWS" (comment folded white space), "FWS" (folded white space), and "CRLF" (carriage-return/ line-feed) are defined in [RFC5322].

1.7. Message Handling Agent Definitions

This document is concerned with message delivery and handling. The following agents are defined in [RFC6409]:

- Message Submission Agent (MSA)
- Message Transfer Agent (MTA)
- Message User Agent (MUA)

Message Delivery Agent (MDA) receives messages and put them into users' mailbox. (non-normative reference [RFC5598])

2. Introduction

Current E-Mail traffic is flooded with Unsolicited Bulk E-Mail (UBE, aka spam). Traditional approaches against them were detecting and filtering them out from the network and user inboxes. In this document, another approach is presented. Instead of removing SPAM from the mail delivery network, we introduce a new partitioned delivery network for messages that are not SPAM.

It is possible to categorize E-Mail messages by their purposes. For example, communication messages usually expect replies. Typical communication messages thus show bi-directional exchange between peers. On the other hand, notification messages such as order confirmations or development activity updates are uni-directional.

E-Mail traffic in each categories may show different characteristics. For example, communication messages have problems like outbound bulk messages from a compromised account. Notification messages have risks of sender spoofing and phishing. Therefore, E-Mail abuse can be efficiently detected and filtered out if we have a different message delivery path per category.

This document defines a protocol by which domain owners may assign separate MTAs for each category of mail. This is done by subdomaining the receiving domain, while keeping the Local-part of the recipient. Subdomaining have an advantage that the separation can happen in transport layer. This effectively separates mail delivery paths at the source of the messages, as if a drainage divide does for water.

Compliant domain holders publish DIVIDE records that specify a subdomain for each mail category that it is willing to receive. DIVIDE records are defined as DNS TXT Resource Records similar to SPF [RFC7208] records. Compliant mail senders use the published DIVIDE records to find the destination MTA according to the category of the mail being sent. Receiver also specifies which method is used to authenticate the sender: DMARC [RFC7489], DKIM [RFC6376], SenderID [RFC4406], PRF [RFC4407], or SPF [RFC7208].

To make this framework effective, senders must label outgoing messages with correct categories. Senders that abusively categorize messages should be detected and removed from the network. A mechanism called Reputation Service Provider is also introduced so that a third-party authority can assure senders' trust. This enables per-category white-listing at the receivers' desired level of strictness.

MDF provides the following advantages:

- a separate message delivery network per message category
- a separate message inbox per message category
- trust-based messaging
- receivers may advertise preferred sender authentication mechanism per category

o reputation based sender white-listing

o senders pay for trust, not receivers

3. Operational Overview

Figure 1 shows the overview of a mail transmission with Mail Divide Framework.

In this figure, solid lines indicate the flow of a message. Double lines indicate communications other than message deliveries (DNS queries, reputation queries over HTTPS).

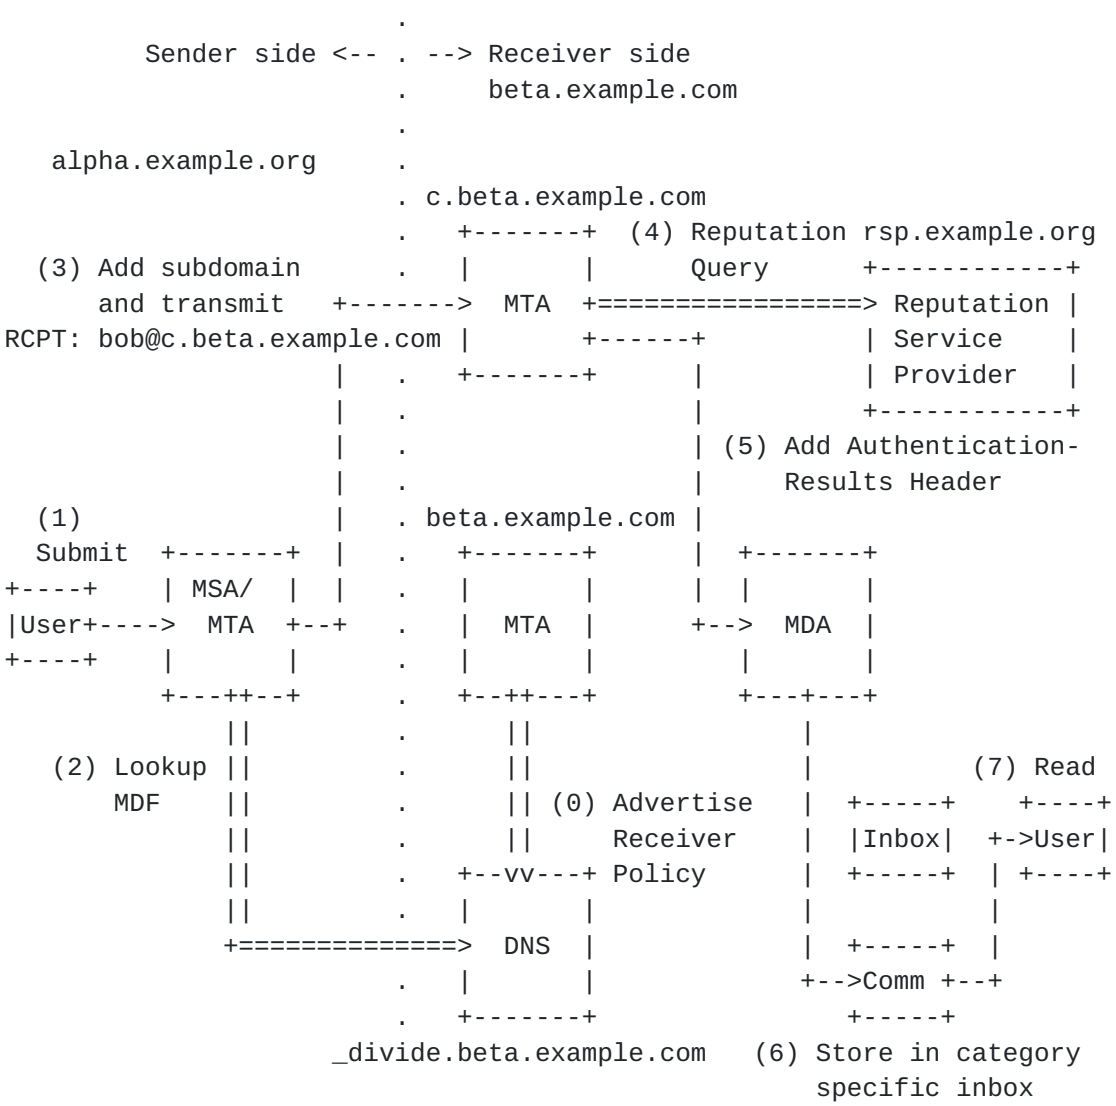

3.1. Preparation

3.1.1. Advertise Receiver Policy

Step (0): the receiver advertises its divide path per category with a DIVIDE record.

Administrator of the mail-receiving domain designs per-category path partition. For example, "beta.example.com" separates communication and notification to "c.beta.example.com" and "n.beta.example.com", respectively. All other categories should go to "beta.example.com". "beta.example.com" builds a DNS TXT Resource Record to express these, as described in Section 5.1. It puts the record in its DNS under "_divide.beta.example.com".

```
v=DIVIDE1\; a=DMARC p=comm:c rsp=rsp.example.org;
            a=DMARC p=notif:n rsp=reputation.example.com
```

The "rsp=" part specifies an RSP associated for each category. Sender's reputation should be managed by this RSP so that the receiver can decide whether it trusts the sender.

3.2. Sending in MDF

3.2.1. Submission with category

Step (1): a user submits a mail message.

MUA/MSA assigns a category to the message according to the context. If the user is a notification sender system, assign "notification". If the user is a human and the message is a reply, assign "communication".

For example, when "alice@alpha.example.org" sends a communication message to "bob@beta.example.com", the MSA of "alice@alpha.example.org" assigns "communication" to the message.

3.2.2. Looking Up DIVIDE Records

Step (2): Sender's MTA looks up MDF policy of the recipient domain.

The final MTA in "alpha.example.org" is going to transmit the message to "beta.example.com".

It first looks up DNS under "_divide.beta.example.com" and finds a DNS TXT Resource Record with "v=DIVIDE1". It now knows that beta.example.com uses Mail Divide Framework.

3.2.3. Subdomaining Recipient Domain

Step (3): Set destination to the divided mail server.

The record has entries for "p=comm:c" and "p=notif:n". This specifies messages in category "communication" should be sent to recipient's subdomain "c", namely, "c.beta.example.com"; similarly category "notification" to "n.beta.example.com".

Since the message from alice to bob has the category "communication", sender's MTA SHOULD choose "c" as target subdomain. It creates a new envelope RCPT address (as defined in [RFC5321]) "bob@c.beta.example.com".

Note that the header To: (as defined in [RFC5322]) MUST stay intact.

3.2.4. Transmitting Mail

Now the mail is sent from "alpha.example.org" to "c.beta.example.com". This is done with ordinary mail transfer protocol, SMTP [RFC5321].

The sender's MTA authenticates itself with DMARC, in this example, according to the DIVIDE record specifies.

3.3. Receiving in MDF

When "c.beta.example.com" receives the mail, it verifies the sender's identity and reputation. The result of the verification is added to the message as Authentication-Results header.

3.3.1. Sender Authentication

Sender's identity is verified by DMARC, DKIM, SPF, etc. according to the authentication method specified in the DIVIDE record.

3.3.2. Reputation Lookup

Step (4): Reputation Lookup

The recipient MTA "c.beta.example.com" is specifically configured to receive messages that are categorized as "communication" according to MDF. It should verify whether the sender complies with MDF, i.e. not sending spam mail under a category label "communication".

The recipient MTA makes a query to a reputation server as defined in Repute protocol [RFC7072]. New assertion-types are introduced to specify MDF mail categories. If the obtained reputation rate is

acceptable, the recipient MTA continue processing the message. Otherwise it should reject the message and return a 5xx status.

3.3.3. Headers and Envelope Handling

Step (5): Add Headers and revert RCPT

At this point, the receiver MTA has verified that the sender conforms to MDF. The mail message is transmitted to the MDA with an Authentication-Results header, which is defined in [RFC7410]. MDF specific parameters are added to the Authentication-Results header.

```
Authentication-Results: c.beta.example.com;
   dkim=pass (good signature) header.d=alpha.example.org;
   divide=pass policy.category=communication
```

Upon forwarding the mail message to the MDA, the receiver MTA MAY remove the category subdomain from the envelope RCPT. This reverts the final recipient to "bob@beta.example.com".

3.3.4. Deliver to Specific Inbox

Step (6): Put the message into a specific inbox for the category

MDA looks at Authentication-Results header of the mail message and will find "divide=pass" field that indicates this mail has been transported via MDF-conformed partitioned delivery path. The MDA puts the message into a separate inbox for the user. In this example, it is "Comm" folder in the user bob's IMAP server.

3.3.5. Read the mail

Step (7): Find the mail as partitioned

The user reads the newly received mail in the "Comm" folder. The MUA looks at the Authentication-Results header to know this is a partitioned mail. It displays a prominent sign to the user that the sender is trusted.

4. Mail Categories

For the purpose of MDF, mail messages are categorized into the following types:

Category	Label	Description
communication	comm	A message intended to become a part of a bidirectional conversation.
transaction	trans	A message regarding money transaction/purchase confirmation.
notification	notif	An one-way message to report an event. No reply is usually expected.
promotion	promo	An advertisement message.
mailing-list	ml	A message delivered from a mailing-list server to the members of that list.
multi-hop	mh	A message is delivered through multi-hop path.
default	default	Fallback category when none of the above is applicable, or specified.

In MDF, the definition of a category SHOULD be agreed upon among senders, receivers, and RSPs so that the reputation feedback works well. DIVIDE records express the receiver's view what categories of message it is willing to receive by separate servers. Each category is advertised in the DIVIDE record with corresponding label.

The sender decides whether the mail message to be sent falls into any of the receiver-designated categories. If a category is found suitable to describe the message, it is used for subdomaining the recipient address.

"default" category is used as a fallback. When the message category is "communication" and the sender does not advertise "p=comm" in the DIVIDE record, the sender looks for "p=default". If an entry corresponding "default" is found, it is used. Otherwise, the message is sent without MDF.

Mail Category for a message MAY be decided by user-interaction, by MSA's context analysis, or by other means. For example, when an outgoing MTA is configured specifically for notification, it can use "notification" for all messages.

"mailing-list" and "multi-hop" do not describe the contents of a message. These instead correspond to delivery mechanisms. See sections Section 8 and Section 9 for details.

5. DIVIDE Records

Domain administrators declare DIVIDE specific DNS TXT records to specify DIVIDE configurations similar to SPF and DMARC. Henceforward, we call this TXT records as "DIVIDE records" in this document. We will show the details of the DIVIDE record in this section.

5.1. DNS Resource Records Syntax

A DIVIDE record is a DNS record that declares separated receiving servers for each Mail Categories, together with sender authentication policy and RSPs.

A DIVIDE record is declared to the "_divide" subdomain of target domains. The MSAs in the mail source domains query the TXT records for the mail destination domains to obtain the appropriate subdomains to deliver the mail messages. For example, if the destination domain of a mail message is "example.com", the MSA located inside the source domain of the message make a query to find the TXT record for "_divide.example.com".

The generic formats of DIVIDE records are:

```
_divide   IN  TXT "divide specific text"
_divide.example.com.   IN  TXT "divide specific text"
```

Multiple parts separated with semicolons compose the "divide specific text". These parts are called "Entry" in this document. Each Entry has several tags detailed in the following part of this section. Amongst each Entry in a DIVIDE record, the first Entry MUST be the one containing only a v (Version) tag. Currently, the only available value for the v tags is DIVIDE1.

The table below shows tag parameters of a DIVIDE Entry. Every tag in this table is mandatory for each DIVIDE Entry.

Tag	Format	Value	Notes
a	a=XXX	SPF, PRA, SenderID, DKIM, DMARC	to declare the authentication method
p	p=XXX:YYY	XXX=comm, trans, notif, promo, ml, mh, default	bind DIVIDE category and mail destination subdomains.
		YYY="subdomain name to be added", or "none"	"none" to specify no subdomaining.
rsp	rsp=XXX	FQDN or IP address of an RSP	specify an RSP for this DIVIDE entry.

Note that a DIVIDE record does not cover subdomains under the declared domain. For example, when an operator desires to add a DIVIDE record for "_divide.a.example.com." in addition to the one for "_divide.example.com.", the operator MUST add a new record for "_divide.a.example.com.".

5.2. Multiple DNS Records

Operators MUST NOT declare more than one DIVIDE record for each (sub) domain.

5.3. Record Size

As discussed in section 3.4 of [RFC7208], a DIVIDE record size SHOULD be small enough to fit in a single UDP packet of a DNS answer. When a DNS answer data size becomes greater than 512 octets, old DNS server implementations might fallback to TCP. The fallbacks may cause the performance degradations to the DNS answer procedures. In [RFC7208], it is recommended to adjust the length of the DNS name and the TXT record bound to it SHOULD be under 450 octets. The DIVIDE records SHOULD follow this guideline.

6. Reputation Service Providers

Reputation Service Provider keeps track of a white list of MDF-conforming senders. Receiving party can perform a query to see how a specific sender is conforming to MDF. Reputation reporting architecture [RFC7070] is adopted in MDF.

6.1. White-list Management

MDF's effectiveness depends on whether the senders correctly label mail messages for the purpose of DIVIDE record lookup and selecting the receiving servers. If an abusive server sends SPAM messages to "c.beta.example.com", the advantage of the traffic separation is diluted. When a sender labels a message as "communication", the degree of how this labeling is correct is evaluated and accumulated as a reputation of this sender for the category "communication". An RSP maintain reputation for sending domains associated with a set of Mail Categories.

When a sending party is not known to the RSP that the recipient trusts, the sender SHOULD NOT be treated as MDF-conforming in the message handling. This is to prevent abusive senders from sending messages to MDF specific inboxes, by always using a new name and expect that a bad reputation would not be built in the RSP.

After looking up the DIVIDE record, the sending MTA SHOULD check whether it has already registered itself to the RSP specified by the recipient. If it has not, it SHOULD fall back to non-MDF mail delivery. In the meantime it registers itself to the specified RSP. Once it is recognized as MDF-conforming by the RSP, it can use MDF for the message delivery.

Methods to register a sender to an RSP are beyond the scope of this document.

Note that a receiver MAY specify itself as the RSP. In that case, MDF is applied only by an explicit consent between the sender and the receiver.

6.2. Reputation Query and Result Caching

Receiving MTA can make a reputation query for the sender domain for the category of the received message, to the RSP that it trust. The query can be performed as defined in [RFC7072].

[RFC7072] defines a URL template for a query as follows:

https://{service}/{application}/{subject}/{assertion}

For the purpose of MDF, the application context "email-divide" is used. Mail Category is used for assertion.

The query result can be cached according to "expires" field in the response, as described in Section 5 in [RFC7071].

An "https" URL with an HTTP over TLS transport SHOULD be used for privacy reasons. See Section 11.2.

6.3. Evaluation and Feedback

Abuse or improper categorization of received message SHOULD be reported to RSPs. ARF format [RFC6650] can be used for this purpose.

Methods of evaluating how the received message is correctly labeled for the Mail Category are beyond the scope of this document.

7. Result Handling

When the receiver MTA verified the sender is MDF-conforming, it generates an Authentication-Results header [RFC7410]. The header is added as the message is transmitted to the MDA.

MDA looks at Authentication-Results header of the mail message and see whether the message is delivered via MDF partitioned delivery network. The MDA puts the message into a separate inbox for the user.

The MUA identify the Authentication-Results header and make prominent sign on the display that the mail is delivered via MDF and verified its trust. For example, it MAY display a green icon to show that the mail message is verified in MDF.

8. Mailing-list

Mailing-list servers reformat the posted message and deliver it to list members. SPF can be used to authenticate the resending sender. Mail Category "ml" is reserved for this purpose, to accommodate a specifically configured authentication policy. Receiving server can advertise a separate RSP that is used for mailing-list senders than for communication.

For example the following DIVIDE Entry declares the mailing-list servers MUST authenticate itself with SPF and the trust is managed by "ml.repute.example.org".

a=SPF p=ml:ml rsp=ml.repute.example.org

9. Multi-hop Delivery

Mail Category "multi_hop" is reserved so that the recipient can express a policy for multi-hop messages.

For example,

- "a=spf p=multi_hop:mh" expresses that the receiving server rejects multi-hop messages.

- "a=dkim p=multi_hop:mh" expresses it accepts multi-hop messages only if DKIM authentication is used.

10. Security Considerations

10.1. DNS Spoofing

[TBD] Use DNSSEC if necessary.

11. Privacy Considerations

11.1. DNS queries

Sender MTA looks up a DIVIDE record under the subdomain "_divide" of the recipient domain. Watching for DNS queries can reveal that the sender is going to use MDF for the following outgoing mail. However, a "_divide" query does not reveal which category is in question.

After a successful DIVIDE lookup, the sender looks up the recipient subdomain's MX records. When MDF is in use, the domain depends on the category of the mail. This indicates that watching on MX queries can reveal the category of the mail that the sender is going to transmit. This is inevitable unless DNS queries are encrypted. A BoF on this topic was held in IETF-89, Encryption of DNS requests for confidentiality (dnse). Future works from that group can mitigate this risk.

11.2. Reputation queries

Queries for reputation server is performed according to [RFC7072]. [RFC7072] defines HTTP based query and optionally HTTPS can be used.

When a recipient MTA receives a mail for a category subdomain, it does a query to the corresponding reputation server. The query indicates the category of the mail in question in the "{assertion}" part of the URL. Thus there is a risk that the category can be observed by watching the traffic between sender and the receiver, combined with reputation queries.

To mitigate this risk, reputation query SHOULD be performed over HTTPS (HTTP over TLS), for the purpose of MDF.

12. References

12.1. Normative References

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC4406] Lyon, J. and M. Wong, "Sender ID: Authenticating E-Mail", RFC 4406, DOI 10.17487/RFC4406, April 2006, <http://www.rfc-editor.org/info/rfc4406>.

[RFC4407] Lyon, J., "Purported Responsible Address in E-Mail Messages", RFC 4407, DOI 10.17487/RFC4407, April 2006, <http://www.rfc-editor.org/info/rfc4407>.

[RFC5234] Crocker, D., Ed. and P. Overell, "Augmented BNF for Syntax Specifications: ABNF", STD 68, RFC 5234, DOI 10.17487/RFC5234, January 2008, <http://www.rfc-editor.org/info/rfc5234>.

[RFC5321] Klensin, J., "Simple Mail Transfer Protocol", RFC 5321, DOI 10.17487/RFC5321, October 2008, <http://www.rfc-editor.org/info/rfc5321>.

[RFC5322] Resnick, P., Ed., "Internet Message Format", RFC 5322, DOI 10.17487/RFC5322, October 2008, <http://www.rfc-editor.org/info/rfc5322>.

[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, <http://www.rfc-editor.org/info/rfc6376>.

[RFC6409] Gellens, R. and J. Klensin, "Message Submission for Mail", STD 72, RFC 6409, DOI 10.17487/RFC6409, November 2011, <http://www.rfc-editor.org/info/rfc6409>.

[RFC7070] Borenstein, N. and M. Kucherawy, "An Architecture for Reputation Reporting", RFC 7070, DOI 10.17487/RFC7070, November 2013, <http://www.rfc-editor.org/info/rfc7070>.

[RFC7071] Borenstein, N. and M. Kucherawy, "A Media Type for Reputation Interchange", RFC 7071, DOI 10.17487/RFC7071, November 2013, <http://www.rfc-editor.org/info/rfc7071>.

[RFC7072] Borenstein, N. and M. Kucherawy, "A Reputation Query Protocol", RFC 7072, DOI 10.17487/RFC7072, November 2013, <http://www.rfc-editor.org/info/rfc7072>.

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

[RFC7410] Kucherawy, M., "A Property Types Registry for the Authentication-Results Header Field", RFC 7410, DOI 10.17487/RFC7410, December 2014, <http://www.rfc-editor.org/info/rfc7410>.

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

12.2. Informative References

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

[RFC6650] Falk, J. and M. Kucherawy, Ed., "Creation and Use of Email Feedback Reports: An Applicability Statement for the Abuse Reporting Format (ARF)", RFC 6650, DOI 10.17487/RFC6650, June 2012, <http://www.rfc-editor.org/info/rfc6650>.

[RFC7073] Borenstein, N. and M. Kucherawy, "A Reputation Response Set for Email Identifiers", RFC 7073, DOI 10.17487/RFC7073, November 2013, <http://www.rfc-editor.org/info/rfc7073>.

Appendix A. Collected ABNF

The following syntax specification of the DIVIDE record uses ABNF [RFC5234]. Terms not defined here are taken from [RFC5321].

```
divide-record         = divide-version
                        [divide-sep divide-authentication]
                        [divide-sep divide-policy]
                        [divide-sep divide-provider]
                        ; components other than divide-version
                        ; may appear in any order

divide-version        = "v" *WSP "=" *WSP
                        %x44 %x49 %x56 %x49 %x44 %x45 %x31

divide-sep            = *WSP %x3b *WSP

divide-authentication = "a" *WSP "=" *WSP
                        ( "SPF" / "PRA" / "SenderID" /
                          "DKIM" / "DMARC" )

divide-policy         = "P" *WSP "=" *WSP
                        ( "comm" / "trans" /
                          "notif"/ "promo" / "ml" / "mh" )
                        %x3a Domain

divide-provider       = "rsp" *WSP "=" *WSP ( Domain / address-literal )
```

Appendix B. Contributors and Acknowledgements

Appendix C. IANA Considerations

C.1. The DIVIDE DNS Resoruce Record Type

[TBD]

C.2. Email Authentication Methods

[TBD]

C.3. Email Authentication Property Types

[TBD]

C.4. Reputation Applications Registry

[TBD]

Authors' Addresses

Takehito Akagiri
Regumi, Inc.

Email: akagiri@regumi.net

Genki Yasutaka
Rakuten, Inc.

Email: genki.yasutaka@rakuten.com

Kouji Okada
Lepidum Co. Ltd.

Email: okd@lepidum.co.jp

Tatsuya Hayashi
Lepidum Co. Ltd.

Email: hayashi@lepidum.co.jp

Masaki Kase
Individual Contributor

Email: kase.masaki@softest.jp