

Network Working Group
Internet-Draft
Intended status: Best Current Practice
Expires: 2 November 2026

A. Brotman
Comcast
T. Zink
Zink Magical Contraptions
J. Bradshaw
Fastmail
1 May 2026

General Guidance for Implementing Branded Indicators for Message
Identification (BIMI)

draft-brotman-ietf-bimi-guidance-14

Abstract

This document is meant to provide guidance to various entities so that they may implement Brand Indicators for Message Identification (BIMI). This document is a companion to various other BIMI drafts, which should first be consulted.

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 2 November 2026.

Copyright Notice

Copyright (c) 2026 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<https://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Revised BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Revised BSD License.

Table of Contents

1. Introduction	3
1.1. Terminology	3
2. Goals for BIMI	3
3. Who should implement BIMI?	4
3.1. Brands	4
3.2. Receiver	4
3.3. MUA Authors	5
3.4. MTA Authors	5
4. Terminology	5
5. Receivers	5
5.1. Site implementations	5
5.2. Validation of a BIMI message	6
5.2.1. BIMI processing requirements	6
5.2.2. Verified Mark Certificate (VMC) Validation	7
5.3. Communicating BIMI results between the MTA and the MUA	7
5.4. Leveraging ARC for MTA MUA communication	8
5.5. Image Retrieval	8
5.6. Limited use of HTTP Redirects	9
5.7. TTL of cached images	9
6. MUA Authors	9
6.1. Image Display	9
6.2. Security Concerns	10
6.3. Privacy Concerns	10
7. Brands	10
7.1. Logo Hosting Considerations	11
7.2. CDN Considerations	11
7.3. Domains listed in your evidence document	11
7.4. Deployment Guidance for LPS and AVP	11
8. Logo Designers	12
8.1. Known Issues	12
8.2. Adherence to SVG P/S	12
8.3. Tools	12
8.4. Caveats	12
9. Basic flow example	12
9.1. Message Classification	13
10. Domain Reputation	14
10.1. Rolling up based upon domain vs organizational domain	14

10.2. VMC Root of Trust	15
11. Security Concerns Relating to Message Authentication	15
11.1. SPF Concerns	15
11.2. DKIM Concerns	16
12. BIMI Playbook Checklist	16
13. Public documentation	17
13.1. Documentation For Brands:	17
13.2. Documentation For Users:	17
14. Appendix	17
14.1. Glossary	17
15. Contributors	19
16. References	19
17. Normative References	19
18. Informative References	19
Authors' Addresses	19

1. Introduction

The Brand Indicators for Message Identification (BIMI) specification introduces a method by which Mail User Agent (MUA, e.g., an email client) providers combine DMARC-based message authentication with cryptographic methods to ensure the identity of a sender. If the identity is ensured, the MUA can then retrieve sender-selected iconography to display within the MUA. This displayed iconography grants the sender brand impressions via the BIMI-capable MUA, and should be a driving factor for the adoption of authenticated email.

1.1. Terminology

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [BCP 14] [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

2. Goals for BIMI

As stated in other BIMI drafts, BIMI intends to advance email authentication by granting a sending party brand impressions as long as the message passes authentication mechanisms and meets other receiver qualifications (reputation, encryption, allow listing, et cetera). DMARC currently has wide adoption by some of the Internet's larger brands, but there is still a long tail of small-to-medium size brands (and many large ones) that do not have it. Furthermore, many domains are not employing DMARC enforcement via quarantine or reject policy, which may allow domain impersonation to continue. Because BIMI provides a visual presence in the inbox, and because visual impressions are desirable for brands, BIMI provides an incentive for

marketers to spur DMARC adoption, whereas a concern purely from security may not.

3. Who should implement BIMI?

3.1. Brands

Organizations take great care to create and promote the image associated with their brand. By implementing BIMI, and creating additional impressions, an organization can foster a stronger tie with customers. In exchange for positive authentication, and strong DMARC policies, the MBP and MUA may show the associated logos with those messages. It should be noted that the domain holder must implement those strong policy on not just a sub-domain, but also the Organizational Domain.

As a Brand holder, you may need to satisfy these requirements:

- * Ability to alter DNS to host a new TXT record
- * A web server to host one or two files, depending on your implementation
- * If you choose to obtain an evidence document, you will need a person to act as a representative for your company
- * The desire to have DMARC enforcement (quarantine/reject) policies on both the organizational and sub-domains. (ex., example.com and sub.example.com)
- * In the DMARC record, pct must be absent or 100%

However, also note that BIMI may not be for every domain. For example, it seems unlikely that a domain would want to implement BIMI for person-to-person correspondence. Or if a domain is not meant to send email, the domain holder may want to explicitly ensure the domain is exempted from BIMI via the BIMI DNS record.

3.2. Receiver

If your site satisfies the requirements (#bimi-site-requirements), this is likely a "yes".

As email has evolved over the past three decades, it is no longer a medium of merely exchanging text, but of enabling people to build rich experiences on top of it. BIMI provides an incentive for brands to send email more securely because the desired behavior - a visual imprint in the inbox - first requires DMARC adoption.

3.3. MUA Authors

The Mail User Agent (MUA) is ultimately responsible for displaying BIMI logos. This could be an in-house/proprietary MUA, or something more generally available. While the MUA may enable the display of the logos, the responsibility for validating inbound messages lies with the Receiver/MBP. MUA Authors should also allow users the option to disable BIMI logo display.

3.4. MTA Authors

The receiving MTA at the destination is the system that is best suited to evaluate message authentication, as well as the DMARC and BIMI policies. The MTA would also be responsible for creating the additional headers that the MUA is meant to utilize. In an ideal world, all MTAs would support BIMI and allow the individual MBPs on deploying BIMI. The MTA would also ideally allow the MBP to alternately utilize a proxy instead of the direct URL retrieved from the BIMI record or evidence document.

4. Terminology

The following terms are used throughout this document.

- * MTA
- * MUA
- * DKIM
- * SPF
- * DMARC
- * MBP
- * Alignment
- * Verified Mark Certificate (VMC)
- * Recipient Domain
- * Sending Domain
- * MVA

For definitions of these terms, see the Appendix.

5. Receivers

5.1. Site implementations

In order for a site to correctly implement BIMI, the receiver must be able to perform the following:

- * Validate SPF
- * Validate DKIM signatures
- * Validate DMARC

- * Discover and fetch a BIMI assertion record using DNS
- * Fetch a SVG using HTTPS
- * Validate a SVG using a profile
- * Add Authentication-Results and BIMI-* Headers to a message

Optionally, for a site to correctly implement BIMI evidence document (VMC is one example) verification, the receiver must be able to perform the following:

- * Fetch the document using HTTPS
- * Validate the evidence document
- * Extract a SVG from the evidence document

A site may wish to implement URI alteration and image caching for hosted recipients. By implementing BIMI, a site agrees that through some combination of trust mechanisms, it will instruct a BIMI-capable MUA to display the image fetched from a URI within the message headers. This URI is created after the MTA authenticates a message, and is also (optionally) able to authenticate the evidence document associated with the sending domain. Discussion of these trust mechanisms is beyond the scope of this document.

5.2. Validation of a BIMI message

5.2.1. BIMI processing requirements

In the BIMI specification, a message **MUST** be authenticated via DMARC. As stated in the DMARC draft, this requires that only one of DKIM or SPF must successfully pass validation with alignment with the organizational domain in the From: address. However, for additional local security measures, a receiving site may choose to create additional requirements for senders in order to verify BIMI (that is, indicate to a downstream MUA that it is safe to load a BIMI logo in the email client)

This may include, but is not limited to:

- * Requiring both DKIM and SPF to validate and align with the organizational domain in the From: address (whereas DMARC only requires one of SPF or DKIM to align with the From: domain). See below for some Security Concerns.
- * SPF "strength" requirements (e.g., requiring "-all", disallowing usage of "?all" or not allowing inclusion of overly large address spaces)
- * SMTP delivery via TLS
- * Feedback Loop registration or other method of registration with the receiving site
- * Domain reputation via a DNS allow list or other reputation system

These localized requirements are at the discretion of the receiving site. In general, the stricter the criteria, the less chance there is of an MUA erroneously showing a logo and giving the wrong signal to a user.

Upon receipt of an email, a receiver that implements BIMI should remove or rename any previously existing BIMI-* headers other than BIMI-Selector, as they may have come from an attacker (as long as the BIMI-Selector is covered by the DKIM signature; if not, it should be removed, renamed, or ignored).

Additionally:

- * It may be useful to have messages exiting a site to have those BIMI-* headers removed as well.
- * It is useful for a site that has not implemented BIMI to remove those headers so that an MUA that does make use of those headers would not accidentally display a BIMI image when the message has not been properly authenticated by the email receiver (even though an MUA should not make use of BIMI headers and instead rely upon settings from the mail store, it is possible that some MUAs will nevertheless use headers without taking appropriate precautions).

5.2.2. Verified Mark Certificate (VMC) Validation

(Currently, see document in Reference below)

5.3. Communicating BIMI results between the MTA and the MUA

In order for a receiver that has implemented BIMI to notify an MUA that it should display the images:

- * An MTA must verify BIMI, and if it passes, add additional headers containing the logo to be displayed.

The MUA must check to see if a message passed BIMI before loading the BIMI image.

While the MTA MAY stamp BIMI-related information in the message headers, they should not be relied upon by an MUA without additional checks to make sure they were added by a trusted source, for example, making sure the MTA strips existing headers on ingress, or by checking for a bimi pass in a trusted Authentication-Results header.

5.4. Leveraging ARC for MTA MUA communication

If both the MTA and MUA support ARC then this MAY be used by the MUA to check that the BIMI evaluation was undertaken by a trusted MTA. In this case the MTA MUST add bimi entries to the Authentication-Results and ARC-Authentication-Results headers. The MUA MUST evaluate ARC, and only use ARC sets which have passed and were added by known good servers. This evaluation MUST stop at the first ARC fail.

If the MUA is configured to require an evidence document then it MAY check for a policy.authority=pass in the bimi Authentication-Results and decline to show a logo if that is not present.

If the MTA has added a BIMI-Indicator header containing the encoded SVG then the MTA SHOULD also add a short hashed checksum of this SVG into the Authentication-Results set. The MUA MUST disregard any BIMI-Indicator headers which do not have a matching hash in the Authentication-Results headers. The hashed checksum for the Indicator MAY be added in the policy.indicator-hash entry. If no BIMI-Indicator is present, or if the hash does not match then the MUA MAY retrieve the indicator from the evidence document at URL specified in the policy.authority-uri if present, from the SVG at the URL specified in the policy.indicator-uri if present and if the MUA does not require a verified evidence document, or by evaluating BIMI directly using the domain and selector from the bimi Authentication-Results entry. The BIMI-Location header is not protected from forgery in the ARC set, and MUST NOT be used.

NOTE: This needs to be added to draft-brand-indicators-for-message-identification and relevant entries registered with IANA A hashed Indicator will need to be added to the AR set policy.indicator-uri to be added to draft

5.5. Image Retrieval

A core part of the BIMI specification is that the MUA will retrieve an image file to display for each BIMI-validated message. There are multiple ways to accomplish this, for example:

- * In its most basic setup, a BIMI-capable MUA could retrieve the image file directly from the site specified in the BIMI-Location header.
- * A BIMI capable MTA will add a header containing the Base64 encoded SVG of the image file. The MUA can use this header to retrieve the already validated image file for display. This is the recommended method of image retrieval as the work of retrieval and

validation has already been done by the MTA. A consideration for this method may be the additional storage requirements for adding a base64-encoded version of the SVG, where the original file could be between 1 and 30 kilobytes, and encoding may add 35% to that size.

- * Other providers may choose to cache the associated images in a local store which could be used as the BIMI resource address in the headers of a BIMI-approved message in a sort of proxy configuration.

5.6. Limited use of HTTP Redirects

- * Receivers may choose not to follow HTTP redirects when retrieving images or evidence documents, or may choose to follow only a limited number of redirects.
- * When setting up BIMI, senders should eliminate, or limit the use of HTTP redirects to avoid images being unretrievable by receivers who either do not support the use of HTTP redirection, or have limited its use.

5.7. TTL of cached images

In some circumstances it is necessary to cache the images that an MUA would want to load. For example, if a domain owner has a short TTL time, it would force the MUA to look it up in an unreasonably short period of time. In this case, a receiver may want to set its own TTL.

One option is to set it to several hours, or a day; another option is to set the TTL to the same as the expiration period in the evidence document that contains the BIMI image. The downside is that the caching mechanism might need to check for certificate revocation, and then re-fetch images.

6. MUA Authors

6.1. Image Display

Although BIMI does not define an aspect ratio for Brand Indicators it is expected that the majority of receivers will display them in a square or circular space. Is it recommended to brands that their Indicators should be constructed to display in a 1:1 aspect ratio, receivers should design the user interface display for BIMI Indicators with this in mind.

6.2. Security Concerns

Receivers should consider the impact of XML bomb or "billion laughs" Denial of Service attacks when handling XML documents such as when validating SVG documents. CVE-2003-1564 (<https://cve.mitre.org/cgi-bin/cvename.cgi?name=CVE-2003-1564>) is an example of this attack.

When validating XML documents, receivers should consider the security and privacy implications of retrieving external entries referenced in those documents.

6.3. Privacy Concerns

There is some concern that the retrieval of the iconography could result in a privacy leak.

As the images are retrieved, it's possible that the image provider could track the retrieving system in some way. This has implications whether it be the sender or provider that is hosting the image. For example, a sender could include a singular selector for a single recipient, or a provider could append a tracking string to the image URI in the header.

A receiver may choose to track the number of selectors an organizational domain is permitted to use and deny processing if this exceeds a defined limit. Similarly, a receiver may choose to track and limit distinct Indicator URLs.

MTAs are encouraged to cache BIMI Records, evidence documents, and Indicators to limit tracking.

MUAs are encouraged to extract Indicators from the BIMI-Indicator header rather than retrieving them directly from the source, as doing so will limit any data exposure to the MTA processing the message. The BIMI approved SVG profile prohibits an SVG from loading external elements, this removes the risk of tracking when an Indicator is shown in the client.

An in-depth discussion of all the potential privacy leaks with respect to loading or embedding images is outside the scope of this document.

7. Brands

7.1. Logo Hosting Considerations

The logo you wish to associate with your brand can be hosted anywhere, not necessarily within the domain that will be used to send the messages. Doing so may make it easier to associate during inspection, though it is understood that not all entities have a web server at the domain associated with their email messages.

7.2. CDN Considerations

If the logo is behind a CDN (Content Delivery Network) this may prevent automated systems from reaching the resource. The automated systems may not appear to be a proper browser experience, and would not be able to correctly respond to a challenge that the CDN may use to protect a site, and therefore unable to retrieve the logo file. If possible, those BIMI logos/resources should be marked as unprotected, allowing any who request the resource to do so without possibility of a challenge.

7.3. Domains listed in your evidence document

While obtaining an evidence document, an entity is expected to provide at least one domain name. There exists the opportunity to list additional domains in the "SAN" field of the certificate. These domains may or may not match the 5322.From domain, but must match the domain being used in the BIMI assertion record. When using the organizational domain, other third-level domains can take advantage of the evidence document as well. Within the core specification, it is discussed how the evaluator should look at the original domain being used, as well as the Organizational Domain.

7.4. Deployment Guidance for LPS and AVP

There exist two methods by which the owner of a domain can directly influence when BIMI logos for email messages might be displayed at a mailbox provider, provided the mailbox provider has implemented support for either/both of these options. In the domain's BIMI declaration, there are options for 'lps' and 'avp'. LPS allows for the domain owner to specify a list of sender addresses which will explicitly receive the BIMI treatment. AVP allows the domain owner to state whether personal avatars should have preference over a BIMI declaration.

LPS would typically be used in a situation where a domain owner uses the same domain for some marketing, but also for employee messages. The domain owner wants marketing messages from 'marketing@example.com' to display a logo, while employee messages

from 'alice@example.com' should display any underlying profile pictures instead. The LPS option does require publication of an additional DNS BIMI assertion record per logo, and potentially an additional DNS lookups for that record.

AVP can also be used to exert influence over display, however, they would be using a messaging platform that allows for users to supply their own avatar, or contain an address book where the receiver can add an image for the sender. In these cases, the domain owner can opt to supersede those personalized images in place of the BIMI-specified logo. This could also be useful in a situation where a company would like all their employees to recognize when an email comes from another employee (without scrutinizing a small photo of a face), or perhaps a specific partner company.

8. Logo Designers

8.1. Known Issues

8.2. Adherence to SVG P/S

There may be a few issues that designers may experience when trying to adhere to SVG P/S.

- * SVG P/S is based on SVG Tiny 1.2, which does not allow for certain types of gradients. When trying to convert/save as SVG Tiny 1.2, it will typically result in an embedded raster file. This is not compliant with SVG P/S, and could result in display issues.
- * When exporting to SVG Tiny 1.2 with Adobe Illustrator, the application will insert x and y attributes within the svg element. These need to be removed to comply with SVG P/S.

8.3. Tools

8.4. Caveats

9. Basic flow example

One sample implementation of BIMI by a receiver, who does everything on-the-fly, is as following:

- * Upon receipt of a message, the receiver checks to see if the message passes aligned-SPF or DKIM, and DMARC, and ensures that the sending domain has a DMARC policy of quarantine or reject per local receiver policy, while properly applying the appropriate DMARC policy to the message.

- * If the message passes prior checks, the receiver will then check to see if the domain in the From: address has a BIMI record (or, if the message has a BIMI-Selector header that is covered by the DKIM-Signature, uses that to do the BIMI query in DNS).
- * If a BIMI record is found, the receiver then retrieves the VMC from the location that the BIMI record points to, and attempts to verify the VMC using a trusted root certificate. .
- * Upon successful verification of the VMC, the receiver extracts the verified image from the VMC. If the SVG also passes the SVG validation steps then this is a successful BIMI verification.
- * If the BIMI verification fails then the MTA must not indicate to the MUA to show a BIMI image. The MUA MAY show a default image such as a set of initials, or unidentified sender.
- * The email receiver then does the rest of its anti-spam, anti-malware, and anti-phishing checks as discussed in Message Classification (#message-classification) below.
- * The email receiver then adds the relevant Authentication-Results and BIMI-* headers to the message to signal to the downstream email client that the message passed BIMI and that is safe to load the logo.
- * Eventually, the MUA checks the BIMI-* headers, decodes the image in the BIMI-Indicator header, and displays it as the sender photo (or however else it chooses to render the BIMI logo in conjunction with the message).

9.1. Message Classification

The successful validation of BIMI does NOT indicate that a message is not spam, malware, or phishing.

It is expected that receivers undertake their usual message filtering and classification steps, and take the results of these checks into consideration when deciding if a BIMI Indicator should be shown to the user.

If classification is performed before BIMI is evaluated then a receiver MAY CHOOSE to skip BIMI processing for that message, in this case they SHOULD add a bimi=skipped entry to the Authentication-Results header for that message, and SHOULD add a comment stating the reasons for skipping BIMI processing.

If a message is classified as phishing or malware then the MUA SHOULD NOT display the logo.

If a message is classified as spam (meaning that the message comes from a known brand, but contains spammy content), then the email receiver MAY choose not to display the logo.

10. Domain Reputation

Receivers are advised to consider incorporating local sources of domain trust intelligence into the processes which ultimately determine whether or not BIMI logos are displayed. Simply because a sending domain passes BIMI requirements does not mean the images should automatically be displayed in the MUA; a site may impose further restrictions based on domain reputation.

One source of additional reputation intelligence could be a platform that the email provider has created to calculate domain trust based on historical traffic; another is an explicit list of trusted domains that has been curated by an individual provider; a third is a list that is purchased from a vendor that might be a pass/fail or a scored list; another option is some mix of any of the previous three.

10.1. Rolling up based upon domain vs organizational domain

BIMI is designed to be able to work on selectors, and so in theory a brand/domain could specify multiple BIMI logos and differentiate them on a per-domain (per-selector) basis. The advantage for the brand is that they can choose the image they want the user to see depending upon various conditions (e.g., seasonal images, regional images, etc.).

However, for an email receiver, it may be easier to roll up BIMI logos on an organizational domain basis. One reason may be for the purposes of reputation, another may be for simplifying management of images. In this case, it would need to be made clear to brands that this is how the loading of BIMI images works. This documentation could live on a postmaster site, under technical documentation, or other official page maintained by the receiver. It could then be referred to when sending organizations ask about how to on-board to BIMI at the receiver, and provide official guidance about the way it works at the site.

If rolling up by organizational domain, then it may make sense to use a "lowest common denominator" approach. That is, an organizational domain must meet all the requirements for BIMI, rather than only a sub-domain. The reason for this is that if sub.brand.com gets an image due to having strong authentication policies, but brand.com

does not, then this may cause confusion because a user may learn to associate sub.brand.com and its image with brand.com; and if brand.com can be spoofed even though sub.brand.com cannot, that can lead to users becoming more susceptible to phishing from brand.com.

To alleviate this, receivers may wish to show logos only for domains that have organizational domains with strong DMARC policies. Or, if an organizational domain does not have a strong DMARC policy but a sub-domain does, then it may treat the organizational domain as if it does have a strong DMARC policy so as to prevent a phisher or spammer from impersonating the brand or any of its sub-domains.

A strong DMARC policy may be defined as one which has some level of enforcement. For example, a p=quarantine policy with an effective pct=100, or a p=reject policy.

10.2. VMC Root of Trust

VMCs are verified back to their issuing Mark Verifying Authority (MVA). Receivers may wish to maintain their own list of trusted CAs for BIMI rather than relying on a generally available bundle of trusted Root Certificates such as those distributed with browsers or operating systems. The AuthIndicators Working Group will maintain a list of known VMC Root CA Certificates to help bootstrap such a list.

11. Security Concerns Relating to Message Authentication

BIMI relies upon the foundations of existing message authentication mechanisms. As of the writing of this document, those are DMARC, SPF, and DKIM. Each of these were created several years ago, and with time, some issues have been found, most specifically with SPF and DKIM. The items below are not specific to BIMI, and the referenced documents have more information.

11.1. SPF Concerns

SPF [RFC7208] is used to denote from where a message should be arriving, typically based on IP. SPF will only provide authentication for the first hop when sending from the originator to another internet mail site. There are some mechanisms within the SPF that could be misused in a number of ways:

- * SPF allows for a +all mechanism. This effectively allows all hosts on the internet to be authenticated as this domain
- * Some domains publish domains with includes that result in a large number of IPs that can be used by any number of other senders
- * A sender could typo a CIDR from an ip4/ip6 statement

- * A platform may not properly validate users are attached to a specified domain when sending
- * A platform may allow for some loose rules relating to forwards, which could permit an attacker to misuse a domain

It's suggested that a receiver should be extremely careful when allowing a message to be authenticated solely on SPF. Similarly, a sender should do as much as they can to utilize both SPF and DKIM, properly aligned.

11.2. DKIM Concerns

DKIM [RFC6376] is a cryptographic signature meant to protect against tampering with a message. This method is the most likely to survive forwarding to another internet site.

- * DKIM Replay is a method by which an attacker attempts to subvert a previously sent message, and use the signature to send something different. See [draft-chuang-dkim-replay-problem] for additional information.
- * Poor selection for signature algorithm, the key length, or the length which a key has been in use
- * Oversigning of headers (RFC6376, Section 8.15) is considered a partial protection against DKIM Replay, and should be considered by senders implementing BIMI
- * Expiration of signatures utilizing the -x option while signing

12. BIMI Playbook Checklist

There are several factors to consider for email receivers on things that can go wrong; below are a handful of considerations:

- * Failing to verify a VMC
- * Failing to extract an Indicator from a validated VMC
- * Failing to validate a SVG against the recommended profile
- * Failing to parse a gzipped SVG Indicator
- * Failing to load a logo in the email client
- * Failing to access the logo (e.g., permissions errors)
- * Connectivity problems to the logo
- * Failing to display a correct logo in the email client
- * Having the wrong logo stored for a brand (i.e., uploading it to a local store but associating it with the wrong brand)
- * Caching a logo for too long after it has updated

There are many reasons why a logo may fail to load; having tools to investigate (logs, headers in messages, internal documentation that is clearly written, having the knowledge pushed out to multiple escalation channels) is important for investigation.

13. Public documentation

13.1. Documentation For Brands:

It is ideal to publish the criteria that is used by your site to determine when BIMI will be displayed. It is fine to say that you use some internal domain reputation metrics as additional criteria to determine whether or not a logo should be displayed, and it isn't necessary to give away the exact nature of the algorithm other than to say "You must maintain good sending practices."

If you use an explicit allow list, a site may want to list the minimum requirements, and the method of applying to be listed. Similarly, a provider may wish to state what type of activity will revoke the decision to display logos previously approved.

13.2. Documentation For Users:

BIMI is not meant to instill additional trust in messages, and it is important to make this known to your users. All messages, even those with logos, should still be treated with (mild) skepticism, and any action regarding the message should still be individually evaluated. It's possible for a site that has a high trust value to become compromised and send fraudulent messages that could compromise a user's system. Ensure your customers have a place that documents BIMI and demonstrates how to check messages for fraud.

14. Appendix

14.1. Glossary

- * MUA - Mail User Agent - The application used to read messages by the end user. This could be a thick client or a web-based application.
- * MTA - Mail Transfer Agent - Software used to transfer messages between two systems, typically between two sites, using SMTP as the protocol.
- * MBP - Mailbox Provider - An organization who provides access to a user's mailbox via some method such as IMAP/POP/MAPI/Webmail.
- * SPF - Sender Policy Framework (<https://tools.ietf.org/html/rfc7208>) - SPF is a framework that designates which systems should be sending for a given domain. This can be a list of IPs, CIDRs, or references to DNS records. As the sender should be controlling their DNS, they should understand which IPs should be sending as their domain.

- * DKIM - DomainKeys Identified Mail (<https://tools.ietf.org/html/rfc6376>) - DKIM is a system by which a chosen set of headers, combined with the message contents, are cryptographically signed, and then validated by the receiving system. Using DNS, the receiving system can retrieve a public key, and then validate the signature within the headers of a message. When implemented properly, the systems responsible for sending the messages for a given domain name should be the only ones capable of creating messages that correctly validates.
- * DMARC - Domain-based Message Authentication, Reporting, and Conformance (<https://tools.ietf.org/html/rfc7489>) - DMARC is a message authentication mechanism that works with SPF and DKIM. The BIMI specification requires that a message passes DMARC. In order for a message to pass DMARC, one of SPF or DKIM must successfully validate, and the domain in the From: address must align with the domain that passed SPF or DKIM.
- * Alignment - Alignment refers to the organizational domain, as defined by DMARC, of the domain in the From: address being the same as the organizational domain that passed SPF or DKIM. For example, baz.example.com has an organizational domain of example.com; bar.foo.example.com also has an organizational domain of example.com. It aligns with org.example.com, because both have the same organizational domain. A definition of organizational domain and methods of discovery may be found in the DMARC (<https://tools.ietf.org/html/rfc7489>) RFC.
- * MVA - Mark Verifying Authority - An entity that a receiver uses to certify that the iconography that they intend to use with BIMI is properly/legally licensed for their use.
- * DRA - Dispute Resolution Authority - This organization will moderate between two entities that believe they are both entitled to use a logo. Receivers should then abide by the decision of the DRA as it pertains to logo usage in the MUA.
- * VMC - Verified Mark Certificate - An Extended Validation Certificate is used in conjunction with BIMI to create a place where information pertaining to iconography for a sending domain can be securely verified. In the case of BIMI, hashes for an MVA-approved set of iconography will be stored in a field within the certificate. This should allow a receiver site to validate the retrieved imagery before putting the BIMI image URI into the message headers.

15. Contributors

TBD

16. References

The full BIMI verification spec can be found at:
<https://github.com/authindicators/rfc-brand-indicators-for-message-identification> (<https://github.com/authindicators/rfc-brand-indicators-for-message-identification>)

Verified Mark Certificates Usage:

http://bimigroup.org/resources/VMC_Guidelines_latest.pdf
(http://bimigroup.org/resources/VMC_Guidelines_latest.pdf)

17. 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, <<https://www.rfc-editor.org/info/rfc2119>>.
- [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>>.

18. Informative References

- [RFC6376] Crocker, D., Ed., Hansen, T., Ed., and M. Kucherawy, Ed., "DomainKeys Identified Mail (DKIM) Signatures", STD 76, RFC 6376, DOI 10.17487/RFC6376, September 2011, <<https://www.rfc-editor.org/info/rfc6376>>.
- [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>>.

Authors' Addresses

Alex Brotman
Comcast
Email: alex_brotman@comcast.com

Terry Zink
Zink Magical Contraptions

Email: tzink@terryzink.com

Jemma Bradshaw
Fastmail
Email: jemma@fastmailteam.com