A First Look at Brand Indicators for Message Identification (BIMI)

Masanori Yajima[1(✉)], Daiki Chiba[2], Yoshiro Yoneya[3], and Tatsuya Mori[4]

[1] Waseda University, Tokyo, Japan
y-masa22@nsl.cs.waseda.ac.jp
[2] NTT Security (Japan) KK, Tokyo, Japan
daiki.chiba@ieee.org
[3] Japan Registry Services Co., Ltd., Tokyo, Japan
yoshiro.yoneya@jprs.co.jp
[4] Waseda University/NICT/RIKEN AIP, Tokyo, Japan
mori@nsl.cs.waseda.ac.jp

Abstract. As promising approaches to thwarting the damage caused by phishing emails, DNS-based email security mechanisms, such as the Sender Policy Framework (SPF), Domain-based Message Authentication, Reporting & Conformance (DMARC) and DNS-based Authentication of Named Entities (DANE), have been proposed and widely adopted. Nevertheless, the number of victims of phishing emails continues to increase, suggesting that there should be a mechanism for supporting end-users in correctly distinguishing such emails from legitimate emails. To address this problem, the standardization of Brand Indicators for Message Identification (BIMI) is underway. BIMI is a mechanism that helps an email recipient visually distinguish between legitimate and phishing emails. With Google officially supporting BIMI in July 2021, the approach shows signs of spreading worldwide. With these backgrounds, we conduct an extensive measurement of the adoption of BIMI and its configuration. The results of our measurement study revealed that, as of November 2022, 3,538 out of the one million most popular domain names have a set BIMI record, whereas only 396 (11%) of the BIMI-enabled domain names had valid logo images and verified mark certificates. The study also revealed the existence of several misconfigurations in such logo images and certificates.

Keywords: BIMI · Email · Measurement

1 Introduction

As promising countermeasure technologies against phishing emails, sender authentication techniques such as Sender Policy Framework (SPF) [38], Domain-based Message Authentication, Reporting & Conformance (DMARC) [26], and DNS-Based Authentication of Named Entities (DANE) [23] have been standardized and have become widespread. In addition to these technologies,

A. Brunstrom et al. (Eds.): PAM 2023, LNCS 13882, pp. 479–495, 2023.
https://doi.org/10.1007/978-3-031-28486-1_20

the standardization of Brand Indicators for Message Identification (BIMI) [16] is underway. The idea behind BIMI is to display the trademarked logo of a company or organization, along with information regarding its certification, in an email message. The recipient of the email can visually verify the legitimacy of the email sender by checking for the existence of a brand logo image with which they are familiar. BIMI technology has gained popularity since receiving official support from Google in July 2021.

For SPF, DMARC, and DANE, which are already widely used, many measurement studies have been conducted on the adoption, misuse, and misconfiguration of technologies. However, to the best of our knowledge, there have been no comprehensive measurement studies conducted on BIMI. Given this background, we set the following research questions to identify best practices and open research questions regarding the BIMI operation:

- *How widespread is BIMI currently?*
- *How do DNS administrators configure the BIMI records for their domain names?*
- *Is BIMI configured with other DNS-based email security mechanisms?*
- *What are the typical misconfigurations of BIMI?*
- *Are there any cyberattacks exploiting BIMI?*

To address these research questions, we conducted the first large-scale measurement study of BIMI in the wild. We examined the presence and configuration of BIMI records for a list of one million popular domain names. We collected logo images and Verified Mark Certificates (VMC) for BIMI records and verified the validity of each setting. In addition, we examined the domain names extracted from 114,915 phishing emails collected by our spam trap and the open database of phishing websites and investigated whether there are any attack cases that exploit BIMI.

The contributions and findings of this study are as follows:

- This is the first large-scale measurement study of the adoption and operation of BIMI in the wild.
- Of the one million popular domain names, 3,538 have BIMI records.
- Of the 3,538 domain names with a BIMI configuration, only 11% had a valid logo image and VMC.
- In domain names that had set up a VMC for BIMI, DMARC was set up in 99.5% of the domain names.
- We found 16 BIMI misconfigurations/violations in BIMI records, 1,224 in logos, 58 in VMCs, and 14 in the DMARC configuration.
- We found 45 domain names having differences between the images contained in the VMC and the images provided on the server.
- In this study, we found no cases of attacks exploiting BIMI.

2 Background

In this section, we first review the email security mechanisms. We then describe the specification of BIMI. For reference, we present the survey results of BIMI implementations for major mail user agents in Appendix.

2.1 DNS-Based Email Security Mechanisms

In the following, we present the overview of the major DNS-based email security mechanisms, except BIMI, which will be described in the next subsection.
The **Sender Policy Framework** (SPF) [38] is a mechanism used to verify the legitimacy of the sender of an email based on IP addresses. By registering SPF information in the DNS TXT record, mail server administrators can explicitly specify IP addresses that are allowed to send emails to the domain name in question.
DomainKeys Identified Mail (DKIM) [25] is a mechanism used to achieve authentication by adding a digital signature when sending email. To use DKIM, the domain name administrator must set up a public key for digital signatures on the DNS server. In addition, by setting a label called a selector, multiple public keys can be operated with a single domain name.
Domain-based Message Authentication, Reporting & Conformance (DMARC) [26] is a mechanism to verify the legitimacy of an email sender by referring to SPF and DKIM records. Like SPF, DMARC can be used by setting a TXT record on the authoritative DNS server of the domain name of the mail sender.
MTA-STS is a mechanism used to enforce STARTTLS on the sender of email, where STARTTLS [22] is a mechanism for encrypting the sending and receiving of email.
DNS-based Authentication of Named Entities (DANE) [15] is a mechanism used to guarantee the authenticity of mail destinations and the confidentiality of mail. DNSSEC [33–35] is used to determine the legitimacy, and STARTTLS is used to achieve confidentiality. To use DANE, a TLS public key must be set up on the email server.
TLS Reporting (TLSRPT) In MTA-STS and DANE, mail may not be delivered because of a failed authentication process. TLSRPT [29] is a function reporting such failures.

2.2 BIMI Specifications

BIMI presents an email to a user with an authenticated brand logo. This allows email recipients to visually distinguish the legitimacy of the email sender without having to look at the subject line or body of the email. The widespread use of BIMI is expected to reduce the success rate of phishing emails. However, for BIMI to be effective, the brand logo displayed by BIMI must be recognized by users [4,5,7]. As with DKIM, multiple logos can be set for a single domain name by setting the selector.

Table 5 in appendix summarizes the DNS records that must be set for each of the security mechanisms described above. "Configure" indicates who needs to configure the record.

BIMI Record: To enable BIMI for a domain name, the following data must be added to the TXT record of the domain name of the MX server:

`v=BIMI1;l=<logo link>;a=<vmc link>`,

where `logo link` describes the brand logo link and `vmc link` describes the link for the VMC. Among these links, only `https` is allowed as a schema.

Logo Image: The brand logo images used by BIMI must be provided in the SVG file format defined in RFC 6170 [36]. SVG Tiny P/S, currently proposed as an Internet Draft [17], sets the following restrictions:

- A title tag must be included (64 characters or less is recommended).
- The following attributes must be set in an svg tag:
  `xmlns="http://www.w3.org/2000/svg"`,
  `version="1.2"`,
  `baseProfile="tiny-ps"`.
- The inclusion of a desc tag is also recommended.
- The size of the logo is recommended to be less than 32 KB.

VMC: VMC is a digital certificate used to certify the ownership of a logo. Currently, DigiCert and Entrust are two CAs that can issue a VMC [14].

DMARC: In DMARC, the domain name owner can set a policy regarding what action should be taken by the email recipient when the source authentication by SPF or DKIM fails. The three policies are as follows:

- "none" indicates that no specific action will be taken.
- "quarantine" indicates that the email recipient will treat as suspicious email that fails the DMARC mechanism check. The email recipient must take action, such as placing the email in the spam folder or conducting further investigations.
- "reject" indicates that an email that fails the DMARC mechanism check is rejected.

DMARC allows one domain name and its subdomain names to be independently configured. A pct is a field that allows the domain name administrator to gradually implement the DMARC mechanism. By setting the pct, it is possible to apply a strong denial policy with a certain probability; otherwise, the next-strongest denial policy is applied. To use BIMI, domain name administrators must fully implement the DMARC mechanism. When using BIMI, "none" should not be applied.

Vetting Process. In order to use BIMI, it is necessary to obtain a valid VMC for the target logo as an email client will test both BIMI record and VMC. A user wishing to obtain a VMC for their logo submits the trademarked logo and information verifying the identity of the user to the VMC-issuing CA. The CA will review the submitted information and also conduct a video conference with the user. If no problems are found, a VMC associated with the logo will be issued. The two VMC-issuing CAs clearly describe in their Certification Practice Statement (CPS) that they meet the official security requirements for issuing

the VMC [6,8,9]. They are subject to an external audit in order to conduct the business of issuing VMCs. This audit is similar to the external audit that CAs issuing server certificates in Web PKI undergo.

3 Measurement Method

In this section, we present the list of domain names we target for our analysis and the data collection methodology.

3.1 Target Domain Names

In this study, we adopt the domain names used for popular websites, those of phishing email senders, and those of phishing websites as our research target domain names.

Tranco: We adopted the one million domain names published by Tranco [13] on February 20, 2022. To ensure that these Tranco domain names contained enough legitimate targets for phishing, we conducted a preliminary study. Specifically, we determined how many of the 382 brands targeted by phishing sites listed on OpenPhish [10] between January 22, 2022 and February 20, 2022 were included in these Tranco domain names. As a result, 96% (= 365/382) of the brands were included in them. This indicates that Tranco domain names are a reasonable target for our BIMI study.

Phishing Email Sender: We analyzed the phishing emails received by our spam trap, and extracted domain names from the email address of the email sender. We collected the domain names of email addresses in the From and Received headers of emails received on April 1 – April 28, 2022. Random sampling resulted in 84,730 unique domain names.

Phishing Website: We employed domain names published by OpenPhish [10] as the domain names of the phishing sites. We obtained this list of domain names on May 2, 2022. A total of 30,221 domain names were examined.

3.2 Data Collection Methodology

This section describes how to determine whether a domain name employs BIMI and other DNS-based email security mechanisms described in Sect. 2.1.

We first send a query to each domain name to look up the BIMI, SPF, DKIM, DMARC, MTA-STS, TLSRPT, or DANE records. Queries were sent using dnspython [32]. We recorded the response to each query and determined that each mechanism is operational if the responses matched the signatures listed in Table 5. In the following, we describe specific notes on collecting data for each security mechanism.

BIMI: In a BIMI study, we adopted `default` as the selector. We downloaded data from the URLs of the logo image and the VMC listed in the BIMI record. In this study, we defined three levels of operation in BIMI, as listed in Table 1.

SPF: In our SPF study, we covered both TXT and SPF records.

Table 1. The levels of BIMI configuration.

Level	Description
1	Has a valid BIMI record shown in Table 5
2	Has a valid logo available for download
3	Has valid logo and VMC available for download

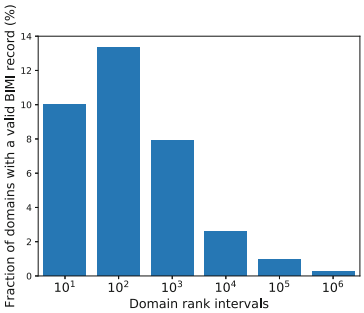

Fig. 1. Fractions (%) of the domain names with valid BIMI records. 10^n represents the logarithmic rank interval ranging from the $10^{n-1}+1$ th domain to the 10^n th domain.

DKIM: In the DKIM survey, we used `default` and `key1` as the selectors.
DANE: In the DANE study, the domain names listed in the MX records were targeted. If at least one of the domain names listed in the MX record supports DANE, the domain name is determined to have adopted DANE.

4 Understanding BIMI in the Wild

In this section, we report on our measurement study of the adoption of BIMI in the wild and its correlation with other DNS-based email security mechanisms described in Sect. 2. We also investigate cases where BIMI has been used in attacks.

4.1 Adoption of BIMI

Among the Tranco one million domain names examined, 3,581 domain names with BIMI records existed (Level 1). We obtained logos from 3,034 domain names (Level 2). However, surprisingly, only 396 of these domain names had a valid VMC available for download (Level 3). We believe that the reason why so few domain names today have had their VMC correctly set up is due to the high cost of obtaining a VMC. To obtain a valid VMC, a brand logo must be registered as a trademark, and a certificate must be issued by a third-party organization based on an examination. We expect that the fact that the cost of operating BIMI is not low will serve as a barrier to attacks that exploit BIMI using fake logos.

Figure 1 presents the number of BIMI-compatible domain names (Level 1 and above) in each rank interval expressed in logarithms, where the rank indicates the popularity of the website corresponding to the domain name on the Tranco

Table 2. Correlations of the email security mechanisms: BIMI vs. other mechanisms. The rows indicate other email security mechanisms and the columns indicate the BIMI setting level. The numerical values in the table indicate the number of domain names.

	Total	BIMI level 1	BIMI level 2	BIMI level 3
All	1,000,000	3,581	3,034	396
MX-enabled	745,746	3,552	3,012	392
SPF	600,672	3,529	2,993	392
DKIM	107,633	3,72	309	14
DMARC	194,123	3,450	2,929	394
MTA-STS	1,310	182	163	23
DANE	8,219	58	50	2
TLSRPT	2,187	249	218	35

list. As expected, the higher the ranking of a domain name, the higher the rate of BIMI adoption; for the top-100 domains, more than 10% of domain names have configured a valid BIMI record. On the other hand, we can see that a certain number of domain names with low rankings have also adopted BIMI, suggesting that the use of BIMI is spreading. For reference, we analyzed the breakdown of the domain names that have configured BIMI. The results are shown in Appendix.

4.2 Correlations Between BIMI and Other DNS-Based Email Security Mechanisms

We analyzed the correlation between BIMI and other DNS-based email security mechanisms, i.e., whether they are simultaneously employed. Table 2 presents the results. "MX-enabled" indicates that the results are restricted to only domain names for which MX records existed. As described in Sect. 2.1, if an email recipient retrieves BIMI data for a domain name, the domain name must pass the DMARC authentication, and the configured policy must be "quarantine" or "reject." Therefore, a high percentage of BIMI-enabled domain names have adopted SPF and DMARC.

We found that the number of domain names configuring BIMI is larger than those of MTA-STS and TLSRPT. This result suggests that BIMI is attracting the attention of more domain name administrators despite being a relatively new security mechanism. If a domain name operates BIMI with Level 3 and DANE, the domain name has an extremely high security level. We found that only two domain names meet these criteria. DANE requires DNSSEC [18,28,31,33–35] settings, which are difficult to configure.

4.3 Attacks Exploiting BIMI

We applied BIMI record lookups on the domain names of phishing emails and websites, which we describe in Sect. 3.1. We found no BIMI records for 114,915 domain names in the two datasets combined; that is, as of today, we have not observed any phishing attempts that exploit BIMI records. We expect that this

observation is due to the fact that the trademark registration process contributes to raising barriers to BIMI record operations. However, there is no assurance that BIMI-abusing domain names will not appear in the future, and it is therefore necessary to keep a close watch on this aspect.

5 Incorrect BIMI Configurations

In this section, we present a measurement study focused on the typical incorrect configurations of BIMI records, logo images, and VMC.

5.1 BIMI Record

We first study the inherent configuration errors we found with respect to the format of the BIMI records collected. It is meaningful to summarize such information and share explicit knowledge of the mistakes that administrators are prone to make.

Logo Setting: Two of the domain names did not have a field to set the logo. In one of these two cases, only a link to the certificate existed. In addition, although 11 domain names had a field for setting a logo, the content was empty, where the empty content in the logo setting field indicates that the domain name in question explicitly refuses to participate in BIMI.

Use of HTTP: There are five domain names whose logo URLs used `http` instead of `https`. None of the five domain names has a URL for the certificate. Similarly, one domain name was used `http` in the URL pointing to the certificate. The URL pointed to the Let's Encrypt server and not the certificate.

Typos: Six domain names were incorrectly used `I=` instead of `l=` as the field for setting the logo. The certificate link did not exist for any of the six domain names.

Unnecessary Parentheses: One domain name existed in which the domain name was described as `l=[<logo link>]` when setting the logo. The domain name in question does not contain a certificate link set.

Invalid String: Two domain names existed, in which invalid character strings were set in records that should describe the URLs.

These misconfigurations were found in domain names that had set only a logo or had not set a logo at all.

5.2 Brand Logo Image

We analyzed logo images in SVG format retrieved from the URLs listed in the BIMI records. A total of 3,034 logo images were analyzed. In the following, we show the cases that violated the mandatory and recommended conditions described in the Internet Draft [17] of SVG shown in Sect. 2.2. Of the domain names with VMC configured, only five domain names failed to configure SVG in the correct format.

Title Tag—*mandatory:* There were 1,008 (33%) logo images without title tags. Two images with empty title tags are found.

Table 3. Frequencies of issuers.

Issuer	Count
Entrust, Inc	166
Digicert, Inc	225
Sectigo Limited	2
Let's Encrypt	3

Table 4. BIMI configuration policies for the target domain (rows) vs. subdomains (columns).

	Reject	Quarantine	None
Reject	258	6	**4**
Quarantine	2	114	**2**
None	**0**	**0**	**6**

SVG Tag—*mandatory:* There were 1,224 (40.3%) logo images that did not conform to the svg tag format.
Desc Tag—*recommended:* A total of 2,905 (95.7%) logo images did not contain a desc tag.
Image Size—*recommended:* In total, 241 (7.9%) logo images exceeded the recommended 32 KB.
Aspect Ratio—*recommended:* Logos displayed on email clients are often circles or squares. It is therefore recommended that the aspect ratio of the logo be 1:1 [1], and 496 (16.3%) of the logo images do not have this aspect ratio.

5.3 VMC

We analyzed VMCs obtained from the URLs listed in the BIMI records. The analysis covered 396 certificates collected from domain names with Level 3 BIMI settings, as shown in Table 1.
Certificate Issuer: Table 3 shows a breakdown of the issuers of the collected certificates. Currently, certificates issued by parties other than Entrust and Digicert are invalid for BIMI, among which there are five such cases. These certificates did not contain logo images, whereas all certificates issued by Entrust and Digicert contained image data.
Certificate Validity Period: We analyzed the validity period of the collected certificates. As a result, 13 certificates had expired. One of these is the domain name `entrustdatacard.com`, which was used by Entrust. The domain name redirects https://www.entrust.com/. However, BIMI records, logos, and certificate links are still accessible.
Legitimacy of Images Extracted from the VMC: We verified whether the 391 logo images extracted from the collected VMCs matched the logos collected from the URLs listed in the BIMI records. We found 45 domain names for which there was a difference between the two logo images. The differences included the use of completely different images, the presence of line breaks in the files, differences in the image size, and differences in the SVG titles.

5.4 Violation of DMARC Policy

We analyzed DMARC policies for 396 domain names using Level 3 BIMI settings. Of the 396 cases, four domain names did not have DMARC configurations. Table 4 presents the results of the analysis of the DMARC policy settings. The

rows represent the configuration policies for the target domain names, and the columns represent configuration policies for the subdomain names. In the table, bold numbers indicate the number of policy violations, 12 of which were present.

6 Discussion

6.1 Current Status of BIMI

Here, we discuss three perspectives on the current status of BIMI as revealed by our results.

Prevalence of BIMI: Compared with other security mechanisms, BIMI has a relatively high adoption rate despite its novelty (see Sect. 4.2). This is because BIMI is relatively easy to set up, and includes setting up the BIMI records and registering the SVGs. However, our results show that only a small fraction of domain names are correctly configured up to VMC. This is because setting up a VMC increases the difficulty of setting up BIMI and incurs certain financial costs.

Misconfiguration of BIMI: Currently, many documents on the Web introduce BIMI settings, and we assume that domain name administrators refer to these documents to set up BIMI. However, it is highly likely that the SVG conversion tool [2] and the BIMI configuration check tool [4,5,7] are not correctly introduced in such documents since misconfiguration of BIMI exists. In the future, further dissemination of these tools is essential to reduce BIMI misconfiguration by domain name administrators and to enable them to self-check whether the correct settings have been made.

Abuse of BIMI: The results of our study show that there is still no evidence of BIMI abuse in phishing emails or in the domain names of phishing sites. BIMI is not yet fully deployed, even for well-known services, and end users are not yet familiar with BIMI. Thus, there is no advantage for attackers in configuring BIMI. However, there is no guarantee that attackers will not continue to implement BIMI abuse in the future. It is therefore necessary to continuously monitor the existence of BIMI abuses.

Challenges for BIMI to Scale: Our measurement study revealed that the adoption of BIMI is not high at the present time. In the following, we discuss approaches that may be effective in increasing the adoption of BIMI. We examined information about MUAs that have implemented BIMI and the categories of domain names that have registered BIMI (see Appendix for detials.) First, we found that there are MUAs that do not currently support BIMI. Although we surveyed major MUAs, there were cases where MUAs did not support BIMI. We hope that MUA vendors will understand the effectiveness of BIMI for protecting their users and implement it in the near future. It is also important for MUAs to provide a usable interface for displaying BIMI so that end-users can recognize and utilize BIMI correctly. In addition, in order to increase the number of BIMI compliant domain names, it would be effective to reduce the cost of setting up BIMI [42]. We expect that the availability of open tools and knowledge of the procedures required to register BIMI will increase its popularity.

6.2 Limitations

Our study has the following three limitations. First, in our study, we sent only a minimum number of queries (up to three) to avoid overloading the target. This means that if the target server was offline during our study, the data might not have been correctly retrieved. Second, our study only investigated the specific selectors for BIMI and DKIM. Therefore, if the target of our survey is to use individual selectors for each sending destination, it may be judged as unsupported in our study. Finally, our study did not clarify the current status of BIMI from the viewpoint of administrators and email recipients. To investigate the current issues in setting up BIMI and the effectiveness of BIMI from the viewpoint of the recipients, it is necessary to conduct an interview study.

6.3 Possibility of Registering Fake Logos

To register a brand logo with BIMI and obtain a legitimate certificate, it must be registered as a trademark. This is expected to make the registration of fake logos more difficult. By contrast, approximately 90% of the domain names that currently have BIMI records operate BIMI without valid certificates. It has also been pointed out that some email clients display BIMI brand logo images without certificate validation [3]. Based on this background, we investigated whether there were any cases of fake logos registered with BIMI. We employed a perceptual hash (pHash) [11], which calculates the similarity between two images. In addition, pHash is widely used to detect a copyright infringement. The analysis revealed several cases in which the same logo was used for multiple domain names. Most of these cases involve the use of several different TLDs for the same service, such as amazon.com and amazon.co.uk. By contrast, there was one domain name using the digicert logo for a completely different service, which we concluded was a misconfiguration. At this point, no obviously fake logos have been found, although we plan to monitor this situation closely.

6.4 Ethical Considerations

Our measurement study discovered several domain names with incorrect BIMI settings. As an ethical consideration, we decided to notify the administrators of those domain names to prevent their misuse. In particular, we are in the process of making a responsible disclosure to the administrators of domain names with VMC configured but with some misconfiguration. We also plan to notify the administrators of domain names that have only SVG configured.

7 Related Work

Several measurement studies have been conducted on DNS-based email security mechanisms. This section divides such studies into two broad categories: those that focus on SPF, DKIM, and DMARC, and those focusing on other areas.

SPF, DKIM, and DMARC: In 2011, Mori et al. conducted an early study on SPF implementation by investigating the existence of SPF and the errors found in SPF policies [30]. In 2015, Durumetric et al. measured email servers supporting SPF, DKIM, and DMARC by analyzing SMTP connections on Google's email servers [20]. In 2015, Foster et al. investigated the prevalence of SPF and DMARC from the perspective of email providers [21]. Hu et al. studied the states of support for SPF, DKIM, and DMARC in 35 email providers in 2018, and conducted a phishing email measurement with end-users [24]. Deccio et al. measured the latest status of SPF, DKIM, and DMARC on several email servers in 2021 [19]. Tatang et al. continuously investigated the status of SPF, DKIM, and DMARC support for domain names listed in multiple top lists in 2021 for a period of 1.5 years [40]. Wang et al. conducted measurements of DKIM deployments using a 5-year Chinese Passive DNS dataset from 2015 to 2020 and server logs of an Chinese email provider in 2020 [41].

Others: In addition, measurement studies were conducted to elucidate other individual protocols (see Sect. 2.1). Scheitle et al. were the first to examine the number of CAAs deployed in 2018 [37]. In 2020, Lee et al. conducted an extensive study to determine how widely DANEs are spread and managed at both the server and client sides [27]. Tatang et al. conducted the first large-scale measurement study of MTA-STS adoption in 2021 [39]. Yajima et al. measured the adoption rates of DNSSEC, DNS Cookies, CAA, SPF, DMARC, MTA-STS, DANE, and TLSRPT, which are security mechanisms that can be implemented in 2021 [42].

None of the studies above mentioned any quantitative results for BIMI, which is just beginning to spread, and our study is the first BIMI measurement approach as of November 2022.

8 Conclusion

In this study, we conducted the first large-scale measurement of BIMI in the wild. We investigated the prevalence of BIMI in one million domain names and found that 3,538 already had BIMI records, despite the BIMI mechanism having only recently begun to be used. We also found that there are intrinsic misconfiguration patterns and specification violations in BIMI records, logos, VMCs, and DMARCs. In addition, no evidence of BIMI abuse was found during our investigation. For the coming widespread use of BIMI, future work includes development of a tool that enables domain name administrators to configure BIMI settings easily and properly, conducting interviews with both domain name administrators and email users on the incentives of adopting/leveraging BIMI, and continuously measure the adoption status of BIMI. We hope that the findings we derived through our measurement study of the BIMI will contribute to its further spread and help thwart the damages caused by phishing attacks.

A BIMI Implementations of Major Mail User Agents

Table 5. DNS records used for configuring mail security mechanisms.

	Configure	Target domain name	RR	Signature
BIMI	sender	<selector>._bimi.<domain name>	TXT	v=BIMI1...
SPF	sender	<domain name>	TXT	v=spf1...
DKIM	sender	<selector>._domainkey.<domain name>	TXT	v=DKIM1...
DMARC	sender	_dmarc.<domain name>	TXT	v=DMARC1...
MTA-STS	receiver	_mta-sts.<domain name>	TXT	v=STSv1...
DANE	receiver	_25._tcp.<mail server domain name>	TLSA	n/a
TLSRPT	receiver	_smtp._tls.<domain name>	TXT	v=TLSRPTv1...

Table 6. BIMI adoption status of major MUAs. ✓ indicates that the valid BIMI logo was correctly displayed on the corresponding MUA.

MUAs (Webmail + browser)	Website 1 (Perfect)	Website 2 (Presence of logo)
Gmail (Chrome 107.0.5304.87)	✓	–
Fastmail (Chrome 107.0.5304.87)	✓	✓
Yahoo Mail (Chrome 107.0.5304.87)	✓	–
MUAs (Email apps)	**Website 1 (perfect)**	**Website 2 (presence of logo)**
Apple Mail (iOS 16)	✓	–
Gmail 6.0.221016 (iOS 16)	✓	–
Gmail 2022.09.18.479203120 (Xperia Z4, Android 6)	✓	–
Gmail 2022.09.18.479203120 (Galaxy S6 edge, Android 7)	✓	–
Microsoft Outlook (Windows 10, version 2202)	–	–
Thunderbird (Windows 10, version 102.3.3)	–	–

In the following, we summarize the current support status of BIMI by the major Mail User Agents (MUAs) – both webmail services and application-based email clients. As webmail services, we adopted Gmail, Fastmail, and Yahoo Mail. We used Google Chrome to study the BIMI adoption status of these webmail services. As email client apps, we adopt Apple Mail, Microsoft Outlook, and Thunderbird. For Gmail in particular, we checked Gmail apps that work on iOS and Android.

We picked up the two popular websites operated with the following BIMI-compatible domain names.

- Website 1 (perfect): Both VMC and logo are correctly registered.
- Website 2 (valid logo): Only logo is correctly registered. VMC is not properly configured.

Note that since our goal is not to expose the level of BIMI operation for specific institutions, and since the BIMI configuration status is likely to be updated in the future and is not invariant, we decided to refrain from naming the respective websites. In addition, since the purpose of this study is to evaluate the BIMI compatibility of MUAs, the type of website does not matter as long as the BIMI setting on the domain name side is consistent.

We registered email accounts on the two websites, where we used different email accounts for each MUA. Emails sent from each website were received by the MUAs used in the experiment to study the adoption of BIMI by MUAs.

Table 6 presents the results of studying whether or not each MUA displays the BIMI logo for emails sent from website 1 and website 2. The behavior of a correctly developed BIMI implementation is to display the logo for website 1, which has perfectly configured BIMI, and not for website 2, which has registered BIMI records but has incomplete VMC. The study revealed that, for webmail-based MUAs, Gmail and Yahoo Mail, accessed with Chrome, correctly implemented BIMI. Fastmail displays the BIMI logo for correctly configured domain names, but does not validate the VMC. Considering the risk of the above-mentioned fact being exploited in a phishing attack, we are currently in the process of making a responsible disclosure to Fastmail. In the email apps, Apple Mail and the all the versions of the Gmail apps correctly implemented BIMI. As of November 2022, Outlook and Thunderbird do not support BIMI.

B Categorization of Domain Names Adopting BIMI

We have categorized domain names that have adopted BIMI. To this end, we leveraged SimilarWeb [12], which is a commercial database that collects web traffic statistics and compiles website information collected from million-order devices deployed around the world. We made use of SimilarWeb to identify categories of domain names, both for those with BIMI records present, and for those with VMC set in addition to BIMI records. Table 7 presents the aggregated

Table 7. Top-10 categories of domain names with BIMI configuration. Level 1 (left) and Level 3 (right).

	Level 1		Level 3	
Rank	Category	Count	Category	Count
1	Computers and Electronics	697	Finance	66
2	Unknown	406	Computers and Electronics	59
3	Finance	373	Lifestyle	29
4	Business and Consumer Services	269	Business and Consumer Services	28
5	Science and Education	197	E-commerce and Shopping	26
6	Health	158	Arts and Entertainment	26
7	Lifestyle	155	Health	25
8	E-commerce and Shopping	140	News and Media	20
9	Travel and Tourism	133	Food and Drink	17
10	Food and Drink	127	Travel and Tourism	15

results for the top-10 categories for domain names with BIMI configuration of Level 1 and Level 3. Majority of Level-1 websites were dominated by Computers and Electronics, Finance, and Business uses. Note that "Unknown" indicates that the category of the website with that domain name was not identified in SimilarWeb. For the Level-3, the breakdown of the websites was different from the above, with Finance topping the list. This observation suggests that since financial websites are often the target of phishing attacks, there is an incentive for them to eagerly take measures using BIMI.

References

1. Creating BIMI SVG Logo Files (2020). https://bimigroup.org/creating-bimi-svg-logo-files/
2. SVG Conversion Tools Released (2020). https://bimigroup.org/svg-conversion-tools-released/
3. Fastmail now supports BIMI (2021). https://www.spamresource.com/2021/04/fastmail-now-supports-bimi.html
4. BIMI Inspector (2022). https://bimigroup.org/bimi-generator/
5. BIMI Record Checker - BIMI Record | EasyDMARC (2022). https://easydmarc.com/tools/bimi-lookup
6. DigiCert Certificate Policy/ Certification Practices Statement for Private PKI Services (2022). https://www.digicert.com/content/dam/digicert/pdfs/legal/digicert-private-pki-cp-cps-v3-9.pdf
7. Email Sender Identity Verification, Authentication & Security Solutions | Valimail (2022). https://domain-checker.valimail.com/bimi/
8. ENTRUST CERTIFICATE SERVICES Certification Practice Statement (2022). https://www.entrust.com/-/media/documentation/licensingandagreements/entrust-certificate-services-cps-3-11.pdf?la=en&hash=EA7E3B4CDEB02433939E7F7AB2762E60
9. Minimum Security Requirements for Issuance of Verified Mark Certificates Version 1.4 (2022). https://bimigroup.org/resources/VMC_Requirements_latest.pdf
10. OpenPhish (2022). https://openphish.com/
11. phash (2022). https://www.phash.org/
12. SimilarWeb (2022). https://www.similarweb.com/
13. Tranco (2022). https://tranco-list.eu/
14. VMC Issuer Information (2022). https://bimigroup.org/vmc-issuers/
15. Barnes, R.: Use Cases and Requirements for DNS-Based Authentication of Named Entities (DANE). RFC 6394, October 2011. https://doi.org/10.17487/RFC6394
16. Blank, S., Goldstein, P., Loder, T., Zink, T., Bradshaw, M., Brotman, A.: Brand Indicators for Message Identification (BIMI). Internet-Draft draft-brand-indicators-for-message-identification-01, Internet Engineering Task Force, April 2022. https://datatracker.ietf.org/doc/html/draft-brand-indicators-for-message-identification-01, work in Progress
17. Brotman, A., Adams, J.T.: SVG Tiny Portable/Secure. Internet-Draft draft-svg-tiny-ps-abrotman-03, Internet Engineering Task Force, April 2022. https://datatracker.ietf.org/doc/html/draft-svg-tiny-ps-abrotman-03, work in Progress
18. Chung, T., et al.: A longitudinal, end-to-end view of the DNSSEC ecosystem. In: Proceedings of USENIX Security Symposium (2017)

19. Deccio, C.T., et al.: Measuring email sender validation in the wild. In: Proceedings of the International Conference on emerging Networking EXperiments and Technologies (CoNEXT) (2021). https://doi.org/10.1145/3485983.3494868
20. Durumeric, Z., et al.: Neither snow nor rain nor MITM...: an empirical analysis of email delivery security. In: Proceedings of the ACM Internet Measurement Conference (IMC) (2015). https://doi.org/10.1145/2815675.2815695
21. Foster, I.D., Larson, J., Masich, M., Snoeren, A.C., Savage, S., Levchenko, K.: Security by any other name: On the effectiveness of provider based email security. In: Proceedings of the ACM Conference on Computer and Communications Security (CCS) (2015). https://doi.org/10.1145/2810103.2813607
22. Hoffman, P.E.: SMTP Service Extension for Secure SMTP over Transport Layer Security. RFC 3207, February 2002. https://doi.org/10.17487/RFC3207
23. Hoffman, P.E., Schlyter, J.: The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA. RFC 6698, August 2012. https://doi.org/10.17487/RFC6698
24. Hu, H., Wang, G.: End-to-end measurements of email spoofing attacks. In: Proceedings of the USENIX Security Symposium (2018). https://www.usenix.org/conference/usenixsecurity18/presentation/hu
25. Kucherawy, M., Crocker, D., Hansen, T.: DomainKeys Identified Mail (DKIM) Signatures. RFC 6376, September 2011. https://doi.org/10.17487/RFC6376. https://www.rfc-editor.org/info/rfc6376
26. Kucherawy, M., Zwicky, E.: Domain-based Message Authentication, Reporting, and Conformance (DMARC). RFC 7489, March 2015. https://doi.org/10.17487/RFC7489
27. Lee, H., Gireesh, A., van Rijswijk-Deij, R., Kwon, T., Chung, T.: A longitudinal and comprehensive study of the DANE ecosystem in email. In: Proceedings of the USENIX Security Symposium (2020). https://www.usenix.org/conference/usenixsecurity20/presentation/lee-hyeonmin
28. Lian, W., Rescorla, E., Shacham, H., Savage, S.: Measuring the practical impact of DNSSEC deployment. In: Proceedings of the USENIX Security Symposium (2013)
29. Margolis, D., Brotman, A., Ramakrishnan, B., Jones, J., Risher, M.: SMTP TLS Reporting. RFC 8460, September 2018. https://doi.org/10.17487/RFC8460
30. Mori, T., Sato, K., Takahashi, Y., Ishibashi, K.: How is e-mail sender authentication used and misused? In: Proceedings of the Collaboration, Electronic messaging, Anti-Abuse and Spam Conference (CEAS) (2011). https://doi.org/10.1145/2030376.2030380
31. Müller, M., Chung, T., Mislove, A., van Rijswijk-Deij, R.: Rolling with confidence: managing the complexity of dnssec operations. IEEE Trans. Netw. Serv. Manage. (2019). https://doi.org/10.1109/TNSM.2019.2916176
32. Nominum: dnspython (2022). https://github.com/rthalley/dnspython
33. Rose, S., Larson, M., Massey, D., Austein, R., Arends, R.: DNS Security Introduction and Requirements. RFC 4033, March 2005. https://doi.org/10.17487/RFC4033
34. Rose, S., Larson, M., Massey, D., Austein, R., Arends, R.: Resource Records for the DNS Security Extensions. RFC 4034, March 2005. https://doi.org/10.17487/RFC4034
35. Rose, S., et al.: Protocol Modifications for the DNS Security Extensions. RFC 4035, March 2005. https://doi.org/10.17487/RFC4035
36. Santesson, S., Housley, R., Rosenthol, L., Bajaj, S.: Internet X.509 Public Key Infrastructure - Certificate Image. RFC 6170, May 2011. https://doi.org/10.17487/RFC6170. https://www.rfc-editor.org/info/rfc6170

37. Scheitle, Q., et al.: A first look at certification authority authorization (CAA). Comput. Commun. Rev. (2018). https://doi.org/10.1145/3213232.3213235
38. Schlitt, W., Wong, M.W.: Sender Policy Framework (SPF) for Authorizing Use of Domains in E-Mail, Version 1. RFC 4408, April 2006. https://doi.org/10.17487/RFC4408
39. Tatang, D., Flume, R., Holz, T.: Extended abstract: A first large-scale analysis on usage of MTA-STS. In: Proceedings of the Detection of Intrusions and Malware, and Vulnerability Assessment (DIMVA) (2021). https://doi.org/10.1007/978-3-030-80825-9_18
40. Tatang, D., Zettl, F., Holz, T.: The evolution of dns-based email authentication: measuring adoption and finding flaws. In: Proceedings of the International Symposium on Research in Attacks, Intrusions and Defenses (RAID) (2021). https://doi.org/10.1145/3471621.3471842
41. Wang, C., et al.: A large-scale and longitudinal measurement study of DKIM deployment. In: Proceedings of the USENIX Security Symposium (2022)
42. Yajima, M., Chiba, D., Yoneya, Y., Mori, T.: Measuring adoption of DNS security mechanisms with cross-sectional approach. In: Proceedings of the IEEE Global Communications Conference (GLOBECOM) (2021). https://doi.org/10.1109/GLOBECOM46510.2021.9685960