Messaging Technology

1.1 Introduction

The IIR has continued to report on quantitative trends in spam and its content, and as described in IIR Vol. 39, from here out we will be focusing on explaining and chronicling the spread of messaging technologies, including technologies designed to combat spam.

In this issue, we go over the results of a survey on the spread of sender authentication technologies, particularly DMARC, and explain MTA-STS, a mechanism described in an RFC last year that relates to TLS encryption connection policies for email delivery channels, as well as SMTP TLS Reporting, a mechanism for reporting of TLS connection information. In relation to messaging, we report on the JPAAWG 1st General Meeting, held last year and co-hosted alongside the Anti-spam Conference, as well as on JPAAWG itself.

1.2 Spoofed Emails and Information Breaches

Emails spoofed to appear as though they were sent by someone else cause so many kinds of problems that they are given names like phishing emails and BECs (Business Email Compromises). The damage caused by such emails is both serious and wide ranging, including financial damages and breaches of confidential and personal information resulting from the capturing of IDs and passwords, malware infections, and the like.

A number of incidents have spurred these sorts of occurrences on. A spate of information breaches from a variety of Web services have occurred, with email addresses included in the information exposed in almost all cases, making it possible for spammers to direct spam with precision. News also came of a massive breach of personal information from a major hotel chain last year, and reams of spam have subsequently made their way into the inboxes of the exposed addresses. Some such spam messages even contain a login password. Services available via the Web can be convenient, but the service provider’s security cannot always be trusted. Users need to properly understand the strength of passwords and other information they set on Web services, as well as the types of services for which the same passwords are used.

1.3 Sender Authentication Rates

We have noted previously that sender authentication is an effective countermeasure against email spoofing. Settings

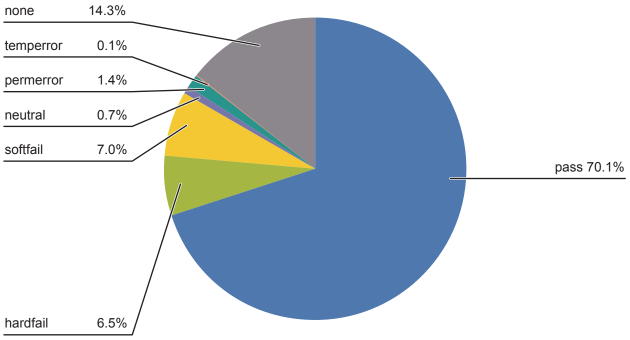

Figure 1: Breakdown of SPF Authentication Results

need to be configured on both the sending and receiving ends. Email recipients can use authentication to detect spoofed emails, while senders can configure their systems to ensure that their emails can be distinguished from spoofed ones.

If we are to promote the spread of sender authentication, we first need to understand how far it currently permeates the space. Here, we report on two sets of survey results, one looking at volume-based deployment rates among senders from a recipient perspective, and the other looking at the proportion of registered domain names on which sender authentication is implemented.

1.3.1 Volume-based Deployment rates

Here, we go over the spread of sender authentication among senders from an email recipient perspective based on email authentication data for emails received via IIJ's email services in April 2019.

Figure 1 shows a breakdown of SPF authentication results. Of all emails received, SPF authentication returned “none” for 14.3%. A value of “none” indicates that SPF authentication was not possible, so turning this around, it means that 85.7% of emails received were from senders that have implemented SPF. The year-earlier (April 2018) figure for “none” was 16.0%, around the same level, which indicates that sender authentication has spread to a point that the vast majority of received emails can be SPF authenticated.

Figure 2 shows a breakdown of DKIM authentication results. As a proportion of the total, the figure for “none” is 62.2%, which indicates that less than 40% of emails received were from senders that have implemented DKIM. The year-earlier figure for “none” was 62.4%, so DKIM deployment rates have not changed that much.

Figure 3 breaks down DMARC authentication results. And by the same measure, the figure for “none” here is 76.9%, indicating that around 20% of emails received were from domains where the sender has implemented DMARC. DMARC authentication uses the results of SPF and DKIM, so it is predictable that authentication rates here will be lower than for SPF and DKIM. That said, and although the resending of emails is an issue, emails can be DMARC authenticated via

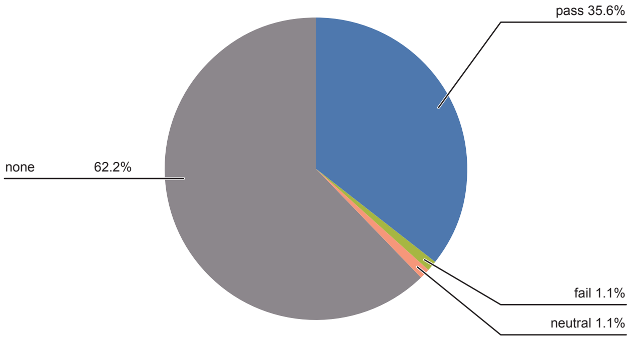

Figure 2: Breakdown of DKIM Authentication Results

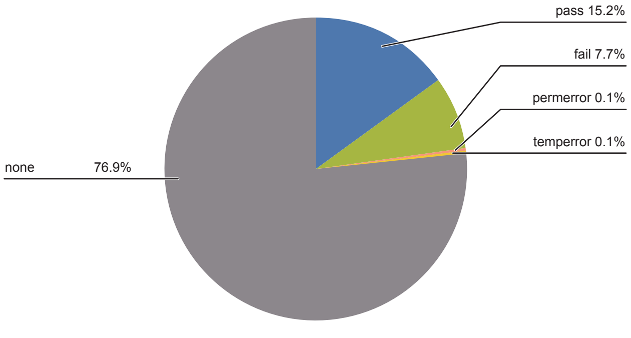

Figure 3: Breakdown of DMARC Authentication Results

SPF alone, so the figure is really quite low considering that the SPF deployment rate is above 80%.

Figure 4 shows the results for DMARC since January 2016. Initially, the figure for "none" was 87.5%, so on a volume basis, the deployment rate has risen by around 10 percentage points over almost three years. The proportion has roughly doubled. Although changes in the proportion of "fail" results have fluctuated over time, the data show that the proportion of all emails that can be authenticated, including those for which authentication fails, is gradually rising.

1.3.2 Deployment Rates Based on Registered Domain Names

Next, for registered jp domain names, we look at whether SPF or DMARC is implemented. As noted previously in Vol. 39, we have a joint research agreement with the Japan Registry Services (JPRS)—which manages jp domain names—and the Japan Data Communications Association for the purpose of gauging the spread of sender authentication technology. I am taking part in the studies as a visiting researcher for the Japan Data Communications Association.

Figure 5 shows the results for March 2018 onward. For domain names configured with an MX record, which indicates the domain name is used for email, the graph shows what proportion had a DMARC record configured, broken down by type of jp domain. According to the latest data from May, this was 0.95% of jp domains overall. By type, ad.jp tops the list, but still only with a figure of 3.4%. Next down the list with 2.1% is go.jp, which has a step-function look to it on the graph because the number of such domains registered is small.

Materials[*1] disclosed by NISC (the Cabinet Office's cybersecurity center) indicate that the use of SPF, DKIM, and DMARC on the sender and receiver sides is listed as a measure for preventing email spoofing within the information

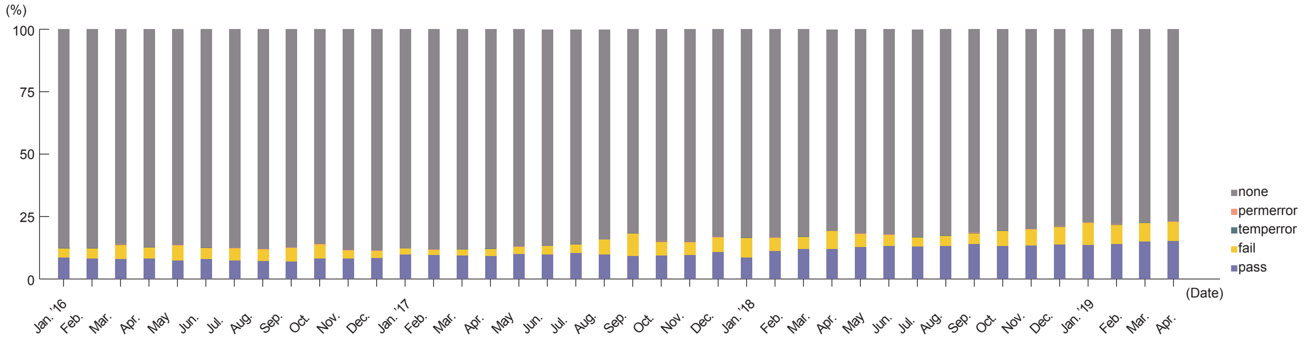

Figure 4: Breakdown of DMARC Authentication Results Over Time

[*1] "Guidelines on the Formulation of Information Security Measures for Government Agencies and Related Bodies (2018 Edition)" (in Japanese at https://www.nisc.go.jp/active/general/pdf/guide30.pdf).

security strategy for government bodies. This means we can expect the proportion of go.jp domains with a DMARC record configured to increase ahead. Note also that registered go.jp domains top the list for the proportion with an SPF record configured (Figure 6).

Similarly, the proportion of all jp domain names with an SPF record configured was 59.7%. This is a 2.8-percentage-point increase vs. the previously reported figure of 56.9% (Vol. 39). The fact that this SPF adoption rate is still rising seems to indicate that awareness of SPF is quite high. Unfortunately, the rate of increase for DMARC is quite low compared with that for SPF, so we will need to boost awareness of DMARC somehow.

1.3.3 Deployment rates Overseas

According to a survey[*2] by the US-based Valimail, 80% of federal government domains in the United States have DMARC records. This was the highest rate among the industries surveyed. As I reported last time, this increase likely traces to a legally binding order[*3] issued by the United States Department of Homeland Security. And according to DMARC.org, a group that advocates for the use of DMARC technology, the number of domains in the DNS with DMARC records increased by over 2.5-fold in 2018[*4].

1.4 Encryption of Email Delivery Channels

Email is used not only to exchange simple messages but also as a means of transferring various types of data via attachment capabilities (MIME). Meanwhile, it does seem that users do not give much consideration to what route an email that contains data will take when being delivered nor to what level of data leakage risk exists. The SMTP email delivery protocol can use TLS (STARTTLS) as an extension. Here, we discuss issues with this conventional STARTTLS protocol and standards established to address those issues, namely MTA-STS and SMTP TLS Reporting.

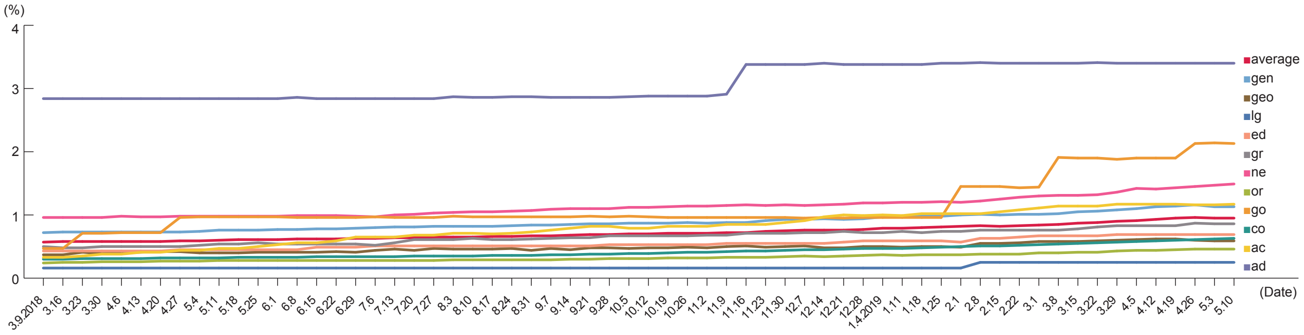

Figure 5: Proportion of jp Domains with a DMARC Record

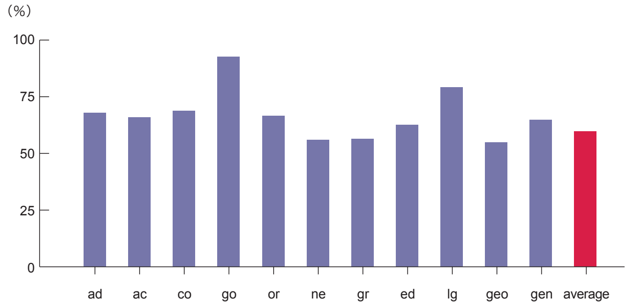

Figure 6: Proportion of Domains with an SPF Record Declaration

[*2] Email Fraud Landscape, Q4 2018 (https://www.valimail.com/resources/email-fraud-landscape-q4-2018/).

[*3] DHS, "Binding Operational Directive 18-01" (https://cyber.dhs.gov/bod/18-01/blank).

[*4] DMARC Policies Up 250% In 2018 (https://dmarc.org/2019/02/dmarc-policies-up-250-in-2018/).

1.4.1 Issues with STARTTLS

The SMTP extension STARTTLS (TLS) is used to encrypt the channel when emails are delivered. The procedure is as follows: if the Recipient mail server supports STARTTLS (determined by the response when connecting), the sender sends a STARTTLS command to start a TLS session. So the channel cannot be encrypted under the following conditions.

- **Recipient mail server does not have STARTTLS (does not return a response to STARTTLS)**
- **The STARTTLS command is sent in order to start a TLS session but the available TLS version and cipher suites do not match**

Cipher suites are combinations of encryption algorithms, key length, and so on. Encrypted communications are not possible unless both sending and receiving ends are able to use the same cipher suite. If the STARTTLS command cannot be executed, many sending email servers will switch to conventional unencrypted plaintext email transmission. This setup exposes email to a sort of man-in-the-middle attack because by intercepting the SMTP session and deleting the intended recipient server's STARTTLS response, an attacker can force plaintext transmission and snoop the contents of email. This sort of technique is also called a downgrade attack.

1.4.2 MTA-STS and TLSRPT

MTA-STS[*5] is a mechanism in which recipient domains use a combination of DNS and HTTPS to publish their receiving policies. This mechanism allows you to determine whether TLS authentication is supported before sending an email and what action the sender should take if a TLS connection cannot be established.

Recipient domains should make the following settings.

(1) Configure an MTA-STS record
(2) Set a "well-known" path so that the MTA-STS policy can be fetched

The MTA-STS record is usually a TXT record that is named by adding "_mts-sts" to the destination domain and that starts with the string "v=STSv1". So if the mail destination domain is "example.com", the record is configured as follows.

```
_mta-sts.example.com.  IN TXT "v=STSv1; id=20160831085700Z;"
```

The id parameter is a string that can be used to determine when the policy has been updated. By first referring to this MTA-STS record, the sender can check whether the recipient domain supports MTA-STS.

[*5] SMTP MTA Strict Transport Security, RFC 8461

To fetch the MTA-STS policy, the sender refers to the “well-known” path on the policy domain prepended with “mta-sts”. The “well-known” path is described in RFC 5785. In the case of MTA-STS, it is fetched via an HTTPS GET request for the following path.

```
https://mta-sts.example.com/.well-known/mta-sts.txt
```

The MTA-STS policy contains line feed-separated (CRLF-separated) key/value pairs. The currently allowable parameters are shown in Table 1.

“max_age” specifies how long the policy should be cached. “mx” specifies patterns matching hostnames given in the MX record. Multiple hosts and patterns can be set. Table 2 shows the allowable values for operation mode (“mode”). The sending MTA determines whether to continue sending emails based on the value of this “mode” field.

An example of an MTA-STS policy appears below.

```
version: STSv1
mode: enforce
mx: mail.example.com
mx: *.example.net
mx: backupmx.example.com
max_age: 604800
```

The TLSRPT*6 specification is used to report to the sender if the policy verification succeeds or fails under MTA-STS or other mechanisms such as DANE*7. Senders use the DNS to publish a TLSRPT policy for receiving reports. Email recipients that support TLSRPT first determine whether this TLSRPT policy has been specified by the sender domain, and if fetchable, a report is sent to the report recipient, if specified, in that policy. The TLSRPT policy settings can be retrieved by prepending “_smtp._tls” to the target domain. The parameters are quite similar to those for DMARC*8 but differ in that “v=TLSRPTv1” specifies version 1 of TLSRPT and the “rua=” field, which specifies where the report is to be submitted, can specify the mailto schema (“rua=mailto:”) as well as HTTPS (“rua=HTTPS:”). An example of a TLSRPT policy record appears below.

```
_smtp._tls.example.com. IN TXT "v=TLSRPTv1;rua=mailto:reports@example.com"
```

When emailing reports to destinations specified using “rua=mailto:”, the report must contain a DKIM signature by the sender domain. The DKIM record of the sender providing the DKIM signature SHOULD contain the “s=tlsrpt” service type declaration.

Table 1: MTA-STS Policy Parameters

Parameter	Meaning
version	The version (currently only “STSv1”)
mode	Expected behavior of sender if policy validation fails
max_age	Max lifetime of the policy (in seconds)
mx	Allowed MX record patterns

Table 2: MTA-STS Policy Modes

Operation mode	Meaning
enforce	Messages are not delivered to hosts that fail policy validation or TLS
testing	Report sent if sending MTA implements TLSRPT*6; messages continue to be delivered
none	Indicates that no explicit MTA-STS policy is applied

*6 SMTP TLS Reporting, RFC8460

*7 The DNS-Based Authentication of Named Entities (DANE) Transport Layer Security (TLS) Protocol: TLSA, RFC 6698

*8 Domain-based Message Authentication, Reporting, and Conformance (DMARC), RFC 7489

```
selector._domainkey.example.com  IN TXT
      "v=DKIM1; k=rsa; s=tlsrpt; p=Mlf4qwSZfase4fa=="
```

Reports sent via email are sent as attachments (MIME) in the same manner as DMARC reports. An example of a TLSRPT policy record for sending reports over HTTPS also appears below.

```
_smtp._tls.example.com. IN TXT "v=TLSRPTv1; rua=https://reporting.exam-
ple.com/v1/tlsrpt"
```

Report data should be compressed for both email and HTTPS transport. Whether applying compression or not, the media type should be consistent with the format ("application/tlsrpt+gzip" or "application/tlsrpt+json"). Reports are sent in JSON format, unlike DMARC reports. We do not go over the parameters given in the report data here, but details can be found in RFC 8460.

Based on this, Figure 16 shows plots of total traffic of the past 10 years. The data series are stacked. The outbound data are observations made at entry points, and the inbound data are observations made at exit points. Some traffic is eliminated within the backbone, such as that involved in attacks, but generally all traffic that comes into the backbone also exits at some point, so the totals are almost the same.

1.5 About JPAAWG

The IIR has mentioned the international antispam organization M^{3}AAWG[*9] several times in the past. Recently, it has also become a forum for a range of discussion on highly relevant security issues beyond that of email. It has also been supporting the establishment of regional organizations beyond North America and Europe, where many M^{3}AAWG members reside. A recently formed group is LAC-AAWG for Latin America and the Carribean. The organization is also working toward and supporting AFR-AAWG for Africa. This leaves only the issue of Asia and what to do there.

[*9] Messaging, Malware and Mobile Anti-Abuse Working Group

IIJ has long been an active member in M³AAWG since it was established, but the number of participants from Japan has not really risen as much the number from the US and Europe. To increase the number of participants, we have been publicizing the M³AAWG's activities in Japan and sounding out the prospects of holding an M³AAWG General Meeting in Japan or Asia from time to time. Against that backdrop, M³AAWG has been making efforts to support M³AAWG-linked activities in other regions. And out of that process emerged efforts among M³AAWG and participants from Japan to set up JPAAWG.

As an organization, JPAAWG (Japan Anti-Abuse Working Group) is entirely independent of, but receives considerable support from, M³AAWG. The JPAAWG 1st General Meeting on November 8, 2018, was held in conjunction with the Internet Association Japan's Anti-Spam Conference, an event that has been running for over a decade, and attracted many speakers and participants. Speakers included the chair and key members of M³AAWG. With the event's success, we made preparations for ongoing JPAAWG activities, culminating in JPAAWG being formally established on May 30, 2019. We hope JPAAWG's future activities will be of interest to you.

1.6 Conclusion

In this issue, we described MTA-STS, a technical specification for reliably ensuring encryption of email deliveries, and TLSRPT as a means of ascertaining what operations have taken place. So far, the IIR has looked at sender authentication technologies including DMARC, ARC, and DANE, but email-related technical specifications continue to evolve along with new specifications such as BIMI (Brand Indicators for Message Identification) and JMAP (JSON Meta Application Protocol). Going forward, the IIR will continue to discuss new technical specifications and the background to their development.

Shuji Sakuraba

Senior Manager, Application Service Department, Network Division, IIJ. Mr. Sakuraba is engaged in the research and development of communication systems. He is also involved in various activities in collaboration with related external organizations aimed at bringing about safe and secure messaging environments. He has been a member of M³AAWG since its establishment. He is the chair of the Japan Anti-Abuse Working Group (JPAAWG). He is acting chairperson of the Anti-Spam mail Promotion Council (ASPC) and a member of its administrative group, as well as chief examiner for the Technology Workgroup. He is chairman of Internet Association Japan's Anti-Spam Measures Committee. He is a member of the Email Security Conference program. He is a visiting researcher for the Japan Data Communications Association. And he is a visiting researcher at JIPDEC.