



computer
emergency
response
team

CERT-EU
for the EU institutions, bodies
and agencies

CERT-EU Security Whitepaper 17-001

DMARC — Defeating E-Mail Abuse

Christos Koutroumpas

ver. **1.3**

February 9, 2017

TLP: WHITE

1 Preface

E-mail is one of the most valuable and broadly used means of communication and most organizations strongly depend on it. The *Simple Mail Transport Protocol* (SMTP) – the Internet’s underlying email protocol – was adopted in the eighties and is still in use after 35 years. When it was designed, the need for security was not so obvious, and therefore security was not incorporated in the design of this protocol.

As a result, the protocol is susceptible to a wide range of attacks. Spear-phishing campaigns in particular can be more successful by *spoofing* (altering) the originator e-mail address to impersonate a trusted or trustworthy organization or person. This can lead to luring the recipient into giving away credentials or infecting his/her computer by executing malware delivered through the e-mail.

While raising user awareness on how to avoid e-mail fraud is recommended, the Verizon Data Breach Investigations Report indicates that more needs to be done. The DBIR report reveals that **30% of all phishing e-mail messages were opened** by the recipients and with **12% clicked on the content and executed malicious code**. The median time for the first user of a phishing campaign to open the malicious email is **1 minute, 40 seconds**. The median time to the first click on the attachment was **3 minutes, 45 seconds**.

These statistics highlight the risk for an organization on the receiving end of spear-phishing e-mails. However – in addition – organizations should be aware of the risk of their brand being misused in e-mail campaigns. Nobody wants to be associated with fraud against their customers, or their potential customers, or the society at large. So, if you don’t protect your domain and your brand with email authentication, you’re making it easy for criminals to capitalize on your goods.

Mitigating the risk of email abuse is the objective of implementing Domain-based Message Authentication, Reporting & Conformance (DMARC) standard by helping email senders and receivers work together to better secure emails and to protect users and brands from painful and costly abuse.

Big mail providers such as Gmail and Yahoo are using DMARC. Facebook, Amazon, and LinkedIn, as well as the major financial institutions such as PayPal, VISA, Bank of America, American Express are also using it because DMARC improves the trustworthiness and reliability of Internet e-mail and protects their customers, brands, and trademarks.

Implementing DMARC in particular helps in:

- Mitigating the risk to **your organization** by stopping spear-phishing e-mails before they reach **your users**.
- Protect **other organizations** by decreasing the risk of them receiving spear-phishing e-mails, which misuse **your domains**.
- Be informed in real-time of new spear-phishing e-mail campaigns, which may put **your organization** or **your community** at risk.

2 Introduction

E-mail abuse is not exclusively a technical issue but it also depends on human factors. However, the effort that fourteen leading IT companies did in an alliance called DMARC, is significantly addressing it. The problem that DMARC addresses is the lack of assurance on the identity of the sender of a message. With the rise of the Internet services and the expansion of e-commerce, it becomes more difficult for users/recipients to establish whether to trust or distrust any particular domain/sender. There are a lot of examples where criminals impersonate a brand and then use the brand's reputation to steal personal information or credit card details from e-mail recipients or to infect their infrastructure with malware. Also, for service providers and system administrators of e-mail providers, it becomes excessively difficult to deal with complaints about spam that appears to have originated from their systems, but in reality did not, and to continuously install, configure, and fine tune e-mail anti-abuse mechanisms thus resulting in wasting money and time.

Another problem is that on the route to its final destination an e-mail often travels between multiple independent Mail Transport Agents (MTA) run by independent e-mail service providers. They usually (before DMARC) had no arrangements among one another and applied different rules and policies on message transmission. Without a specific mechanism it is very difficult to assign confidence and accountability to parties and to avoid or detect the injection of malicious e-mails into the Internet mail infrastructure.

Most of the common e-mail abuses, like spam and mass phishing¹, are distributed using spoofed e-mail addresses by forging the e-mail `Mail From` field in SMTP envelope sender or the `From` field in the e-mail header. In most spamming and phishing campaigns, the attacker will not use a real return address because they would have their account blocked very quickly. The result will be a *backscatter*, i.e. one would receive a bounce message with a Non Delivery Report (NDR) for an e-mail that one has never sent.

That is why e-mail abusers often use either faked addresses (currently this is in decline) or lists of valid e-mail addresses just for the purpose of diverting bounced messages. If a message bounces, spammers do not want it coming back to their mailbox.

Other types of e-mail abuse, where the sender address is forged or spoofed, are done by:

- spear-phishing e-mails, where the attackers want to impersonate well-known, trusted identities in order to steal passwords or other financial/personal data or download malicious files and exploits;
- fraudsters who want to cover their tracks and remain anonymous;
- computer worms;
- brand name impersonation.

Originators of Internet Mail need to be able to associate reliable and authenticated domain identifiers with messages, communicate policies about messages that use those identifiers, and report about mail using those identifiers and here is where DMARC comes into play.

DMARC is a **mechanism** to define a **coherent e-mail policy** that can effectively be **used by both the sender and the receiver** of the e-mail messages. The senders can list the authentication mechanisms they have put in place, and the receivers are informed what the sender suggests them to do, if the authentication fails on any message that claims to originate from them.

¹**phishing**, is an attempt to acquire sensitive, personal information such as usernames, passwords, bank account information, and credit card numbers by posing as a trustworthy source.

3 What is DMARC?

DMARC stands for Domain-based Message Authentication, Reporting & Conformance, and it is a technical specification initially created by:

- E-mail providers (receivers) such as: AOL, Comcast, Gmail, Hotmail, Yahoo! Mail
- Financial institutions and service providers (senders): Bank of America, Fidelity Investments, PayPal
- Social media properties: American Greetings, Facebook, LinkedIn
- E-mail security solutions providers: Agari, Cloudmark, eCert, Return Path, Trusted Domain Project

The DMARC technical specification defines the method/policy on how e-mail receivers would treat incoming mails using two pre-existing mechanisms, Sender Policy Framework (SPF) and Domain Keys Identified Mail (DKIM). DMARC in its current version is specified in RFC 7489.

A DMARC policy allows an e-mail sender to indicate that their e-mails are protected by SPF and/or DKIM, and tells their receiver what to do if neither of those authentication methods passes — such as junk or reject the message. DMARC also provides a way for the e-mail receiver to report back to the sender about messages that pass and/or fail DMARC evaluation. Both mechanisms are used for e-mail authentication in contrast with e-mail reputation systems widely used today.

So DMARC improves the trustworthiness or reliability of Internet e-mail by providing a DNS based policy publication, feedback, and enforcement mechanisms necessary to build secure e-mail channels, upon which trust can be established. DMARC can provide:

- E-mail senders with a mechanism on how to publish:
 - a policy that dictates how e-mail receivers treat unauthenticated e-mail and send reports back to e-mail senders,
 - policies on how to deploy existing e-mail authentication mechanisms such as SPF and DKIM.
- E-mail receivers with information on how to discover DMARC policies, how to process DMARC compliant e-mail, and how to generate feedback to senders.

3.1 DMARC Specifications

Sender Policy Framework (SPF) is an open standard specifying a technical method for e-mail validation. It is designed to prevent e-mail spam by detecting e-mail sender address spoofing by verifying sender IP addresses. In short, SPF allows domain owners to specify who can send e-mail on behalf of their domain.

Mail exchangers (MX) use the DNS to check if an e-mail from a given domain is being sent by a host who has the permission to do so. They can check if the sending MTA's IP address is allowed to send messages on behalf of the given domain by analyzing the message's envelope return address from the `Mail From` field defined in RFC 5321 and querying the domain's DNS record for permitted MTA IP addresses.

SPFv1 protects and verifies the envelope sender address or return-path that is used during the transport of the message from mail server to mail server. It does not use the header sender address contained in the `From:` header, thus does not prevent forgery of these addresses, if the spammer uses accounts in a domain with a sender policy, or abuse a compromised system in this domain.

SPFv1 allows the owner of a domain to specify their e-mail sending policy, e.g. which mail servers they use to send mail from their domain. This is achieved in two steps:

- The domain owner publishes this information in an SPF record (DNS TXT resource record) in the domain's DNS zone.
- When someone else's mail server receives a message claiming to come from that domain, the receiving server checks whether the message complies with the domain's stated policy. E.g., what to do with an e-mail coming from an unknown server.

DomainKeys Identified Mail (DKIM) is also an end-to-end authentication mechanism that allows a sender or a sender domain to take responsibility by associating a domain name to an e-mail message by affixing a digital signature to it. The difference with SPF is that it adds authentication capabilities for the existing e-mail infrastructure without the need to change the actual e-mail standards used, and allows a receiver to verify that the message was sent from the claimed source. The association is set up by means of a digital signature, which can be validated by recipients. The verifier recovers the signer's public key using a DNS TXT resource record and then verifies that the signature matches the actual message's content.

Unlike SPF, DKIM can deal with forwarder services and message paths spanning multiple hops because DKIM is independent of Simple Mail Transfer Protocol (SMTP) routing aspects in that it operates on the RFC 5322 message – the transported e-mail header and body – not the SMTP envelope defined in RFC 5321. Hence, the DKIM signature survives basic relaying across multiple MTAs.

The organization is a handler of the message, either as its originator or as an intermediary. Their reputation is the basis for evaluating whether to trust the message for further handling, such as delivery. Technically DKIM provides a method for validating a domain name identity that is associated with a message through cryptographic authentication and is thus independent of IP addresses.

The signing organization can be a direct handler of the message, such as the author, the originating sending site, or an intermediary along the transit path, or an indirect handler. For instance an independent service that is providing assistance to a direct handler. In most cases, the signing module acts on behalf of the author organization or the originating service provider, by inserting a DKIM-Signature to the header field. The verifying module typically acts on behalf of the receiver organization.

DKIM allows the signer to distinguish its legitimate mail stream. It does not directly prevent or disclose abusive behavior.

To publish a policy dictating what to do if a message fails authentication, a mechanism called Author Domain Signing Practices (ADSP) was adopted. It is an optional extension to the DKIM e-mail authentication scheme, whereby a domain can publish the signing practices it adopts when relaying mail on behalf of associated authors.

The previous mechanisms assure that only authorized owners of an e-mail address can send e-mails through their MTA by having them authenticate themselves to their Mail User Agent (MUA). Also domain authenticity is assured by letting receivers verify that the sending MTA is in charge of an e-mail address and that an e-mail is not originating from a potentially malicious source. Both parts together allow a receiver to be sure that an e-mail came from its alleged sender.

E-Mail receivers apply many different methods to analyze incoming messages, including SPF and DKIM, so in coordinating the above mechanisms DMARC comes in play. DMARC does not eliminate the need for additional forms of analysis, but it does provide a way for participating senders and receivers to streamline the process by coordinating their efforts.

4 How Does DMARC Work?

DMARC is designed to fit into an organization's existing inbound e-mail authentication process. The way it works, is to help e-mail receivers to determine if the purported message *aligns* with what the receiver knows about the sender. If not, DMARC includes guidance on how to handle the *non-aligned* messages.

DMARC alignment prevents spoofing of the e-mail header `From:` address by:

- matching the e-mail header `From:` domain name defined in RFC 5322 with the envelope `Mail From` domain name used during an SPF check, and
- matching the e-mail header `From:` domain name defined in RFC 5322 with the `d=domain` in the DKIM signature.

More in detail, DMARC requires that a message not only pass DKIM or SPF validation, but that its identifier domains are *in alignment*.

For SPF, the message must pass the SPF check and the domain in the (RFC 5322) `From:` header must match the domain used to validate SPF (must exactly match for strict alignment, or may be a sub-domain for relaxed alignment — this is the default).

For DKIM, the message must pass the DKIM check and the `d=domain` of the valid signature must align with the domain in the (RFC 5322) `From:` header (must exactly match for strict alignment, or must be a sub-domain for relaxed alignment). It is important to mention again that even if SPF and DKIM pass authentication DMARC will still fail if the identifiers are not aligned.

Identifier alignment forces the domains authenticated by SPF (typically the `Mail From` domain but can be the `HELO` domain if the `Mail From` was empty) and DKIM (the DKIM signing domain as shown in the DKIM-Signature header's `a` field) to have a relationship to the header `From:` domain which is more typically visible to a user in e-mail clients. In *strict* alignment mode the domains must be an exact match. In *relaxed* alignment mode the domains can be different sub-domains of the same organizational domain.

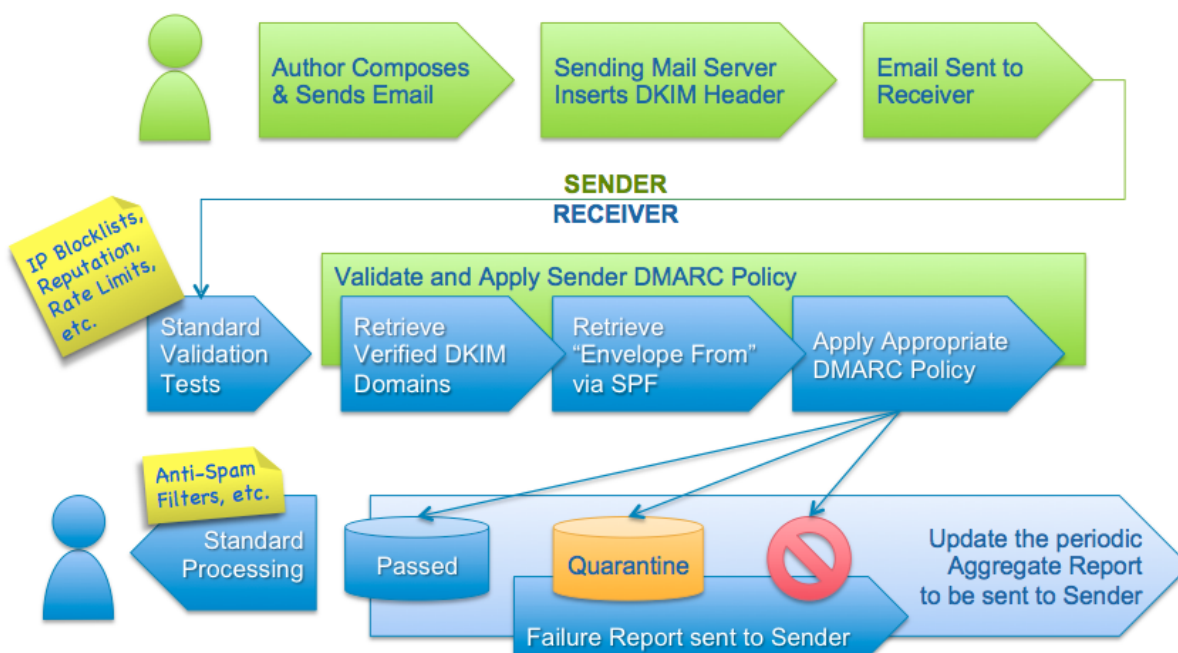


Figure 1: Example flow for DMARC.

To implement DMARC, all SPF or ADSP policies dictating how to treat a message that fails authentication must be discarded and only DMARC policies must be advertised. For example, assuming that a receiver deploys SPF and DKIM plus its own spam filters, the flow may look something like in Figure 1.

As we can see, the testing for alignment according to DMARC is applied at the same point where ADSP would be applied in the flow.

An important aspect of DMARC is the reporting. There are two kinds of reports:

The first and more typical seen are called aggregate reports, where each e-mail account provider includes summary information about all the messages they saw using your domain. It includes the number of messages, which IP addresses they came from, whether they authenticated with DKIM or SPF, whether there was a DMARC *pass*, what was done with the message, and so on. The report is generated using the XML format. The reports can be tens or hundreds of kilobytes. It all depends on the size of the e-mail account provider and the amount of traffic sources.

The second kind of reports are failure reports, which are not sent by all e-mail account providers. Microsoft's Hotmail is probably the best-known e-mail account provider that sends them. In this case, if the domain owners requested them through their DMARC policy, the e-mail account provider sends one failure report for every message they receive that uses the domain, but which fails to authenticate.

DMARC process flow:

- An e-mail is sent from the **sender** to the **receiver**.
- The **receiver** requests the **sender's** Resource Records (RR).
- The **sender's** DNS server responds with SPF/TXT, DKIM, and DMARC RR to the **receiver's** request.
- The **receiver** checks DKIM and/or SPF and then applies DMARC policy.
- If the e-mail does not pass DMARC policy, a failure report is sent to the **sender** domain (as defined in the message).
- Optionally a periodic aggregate report is sent to the **sender's** domain.

The *not pass* DMARC policy consists of the following three choices:

- Quarantine messages that fail DMARC (e.g., move to the spam folder), OR
- Reject messages that fail DMARC (e.g., don't deliver the mail at all), OR
- None (e.g., do nothing).

5 How Can DMARC be Technically Implemented?

These steps should be implemented in sequential order²:

1. It does not matter if you have SPF or DKIM deployed, just publish a DMARC record with `p=none` and a `rua=` pointing to a special mailbox to receive aggregate reports. **NOTE: Do not put a `ruf=` at this point as it may overwhelm your server and it is not yet needed.**
2. Read the aggregate reports, understand your e-mail infrastructure, and get at least one of SPF or DKIM to work correctly for all mail you care about.
3. If you use third party providers to send mail on your behalf, get them to be DMARC compliant (see the DMARC FAQ in the Reference section).

²A step-by-step *DMARC Setup Guide* created by The Global Cyber Alliance to make it easy for organizations of all size to implement DMARC is available in [12].

4. Once you know that your infrastructure is mostly correct, and you can predict the approximate number of failure reports you will receive, you can decide whether you would like a `ruf=` record, and if so add one pointing to a different address from the `rua=` record.
5. Add DMARC filtering on your incoming e-mail infrastructure, and re-check all the aggregate and failure reports.
6. Fix some more of your e-mail infrastructure, because the DMARC filtering on incoming e-mails is likely to show you more forgotten areas as important e-mails are likely to be between third party senders your employees use, and your employees' organizational mail system.
7. Whitelist in your DMARC filter some well-known forwarders (mainly some third-party senders you are using).
8. Whitelist in your DMARC filter all the mailing lists your employees are using.
9. Do you have a phishing problem? If not, then you may have fixed enough your infrastructure. The next steps are harder to put in place and only warranted if the benefits of fighting phishing outweighs the complexity that will arise from a very restricted infrastructure.
10. Be mindful there are at least two cases where DMARC is likely to reject e-mails:
 - e-mail forwarding,
 - mailing lists.
11. Consider moving transactional mail to a separate domain from employee mail, which you can do more aggressive enforcement on.
12. Due to compliance issues related to mailing lists and DMARC, unless strictly necessary, advise employees to refrain from subscribing to mailing lists using their work e-mail.
13. Get all third party providers to be DMARC compliant (see the DMARC FAQ in the Reference section)
14. You are ready to move to `p=quarantine` and/or `p=reject`

Last, you have to resolve all the cases that DMARC does not cover, such as cousin domains, friendly display, and receivers with no DMARC filtering.

Example of DMARC TXT resource record of `ec.europa.eu` domain:

```
v=DMARC1; p=none; rua=mailto:dmarcreports@ec.europa.eu; fo=s; adkim=s; aspf=s; sp=none
```

- `v=dmARC1` – **version** – the DMARC record version.
- `p=none` – **policy** – the policy to apply to an e-mail that fails the DMARC test. The value can be `none`, `quarantine`, OR `reject`.
- `rua=mailto:dmarcreports@ec.europa.eu` – **receivers** – a list of URIs for receivers to send XML feedback to. URIs are required to be added in the format of `mailto:address@example.com`.
- `fo=s` – **forensic reporting** – forensic reporting options. Possible values: `0` to generate reports if all authentication mechanisms fail to produce a DMARC *pass* result, `1` – reports if **any** mechanisms fail, `a` – report if DKIM signature failed to verify, `s` if SPF failed.
- `adkim=s` – **alignment mode DKIM** – *alignment mode* for DKIM signatures. `r` is for *relaxed* – allows authenticated DKIM `d=domains` that share a common organizational domain with an e-mail's header `From: domain` to pass the DMARC test. `s` is for *strict*, which requires an exact match between the DKIM `d=domain` and an e-mail's header `From: domain`.
- `aspf=s` – **alignment mode SPF** – *alignment mode* for SPF signatures. `r` is for *relaxed*, which allows SPF authenticated domains that share a common organizational domain with an e-mail's header `From: domain` to pass the DMARC test. `s` is for *strict*, which requires an exact match between the SPF domain and an e-mail's header `From: domain`.
- `sp=none` – **sub-domain policy** – the policy to apply to e-mail from a sub-domain of this DMARC record that fails the DMARC test. The value can be `none`, `quarantine`, OR `reject`.

6 Conclusion

DMARC is gaining momentum as we can see in Figure 2.

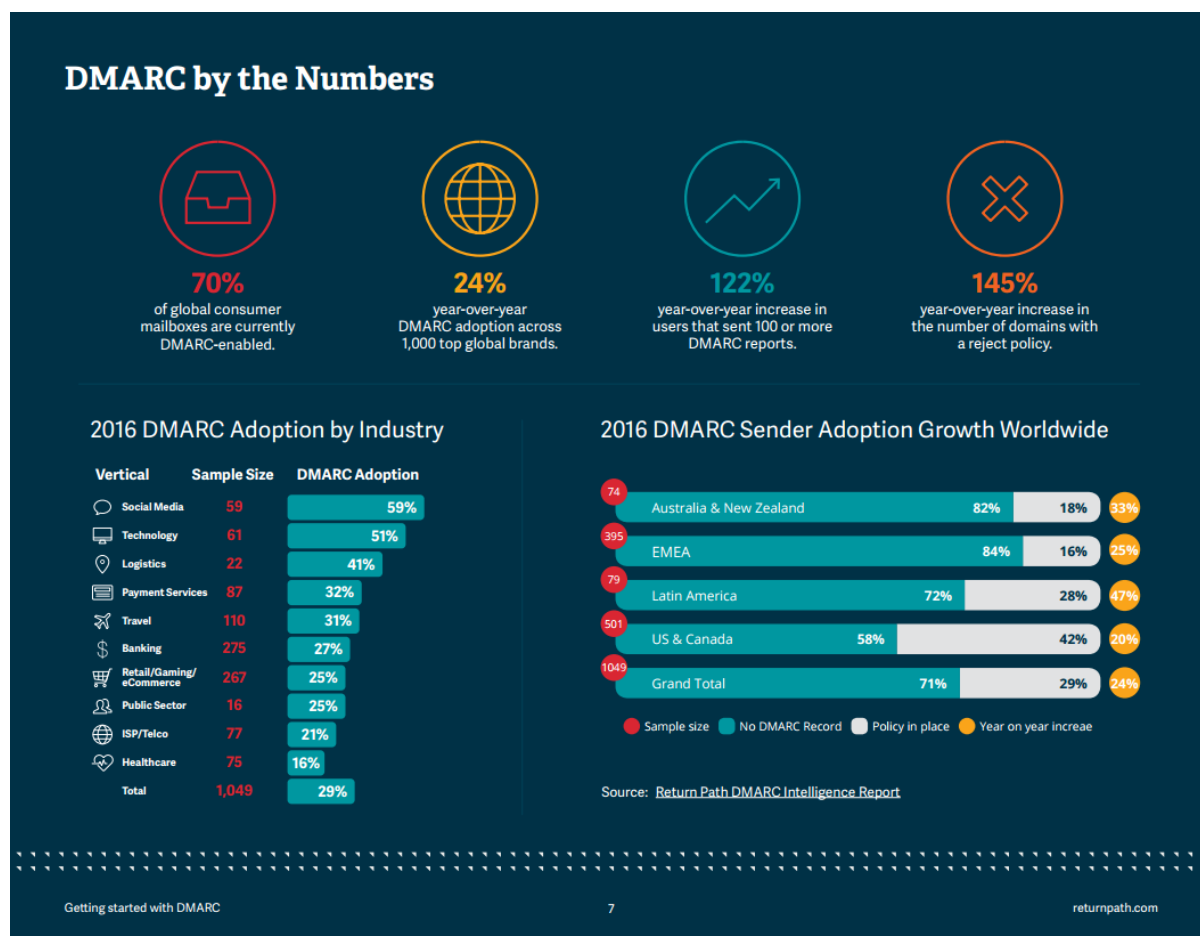


Figure 2: DMARC use statistics.

Until today, all e-mail anti-abuse mechanism have been proven ineffective in a large scale attack, especially in the case of spear-phishing. This is because, when criminals send targeted e-mails (aka spear-phishing) in low enough volumes, statistics based anti-spam engines can effectively be by-passed. DMARC introduces – even using pre-existing authentication mechanisms – a new way of considering e-mail security by creating secure e-mail channels between senders and receivers. It also generates more feedback with aggregate and forensic real-time reports than ever before. Currently vendor support is limited, but it is expected to increase as vendors will start using DMARC functionality in their products in the future. This specification has all it takes to become a *de facto* standard in the IT security market by changing e-mail's insecure nature.

7 References

- [1] RFC 7489 (DMARC) Internet Engineering Task Force <https://tools.ietf.org/html/rfc7489>
- [2] DMARC Organization FAQ <https://dmarc.org/wiki/FAQ#>
- [3] DMARC specification <https://dmarc.org/resources/specification/>

- [4] RFC 4871 (DKIM) Signatures — Internet Engineering Task Force <http://www.ietf.org/rfc/rfc4871.txt>
- [5] DKIM organization <http://www.dkim.org>
- [6] Wikipedia on SPF http://en.wikipedia.org/wiki/Sender_Policy_Framework
- [7] Sender Policy Framework OPEN SPF <http://openspf.org>
- [8] Creation of a DMARC record <http://www.agari.com/dmarc>
- [9] What is Identifier Alignment <https://agari.zendesk.com/hc/en-us/articles/202952519-What-is-Identifier-Alignment-> What is Identifier Alignment
- [10] Return Path DMARC FAQ <https://blog.returnpath.com/how-to-explain-dmarc-in-plain-english/>
- [11] Litmus DMARC FAQ <https://litmus.com/blog/dmarc-what-it-is-how-it-helps-protect-your-brand-against-e-mail-fraud#five>
- [12] DMARC (DMARC, SPF, DKIM) wizard of Global Cyber Alliance <https://dmarcguide.globalcyberalliance.org>

8 Appendix 1 — Example DMARC Record

This is a sample DMARC record for the `example.com` domain:

Tag	Purpose	Example
<code>v</code>	Protocol version	<code>v=DMARC1</code>
<code>pct</code>	% of messages subjected to filtering	<code>pct=20</code>
<code>ruf</code>	Reporting URI for forensic reports	<code>ruf=authfail@example.com</code>
<code>rua</code>	Reporting URI of aggregate reports	<code>rua=aggrep@example.com</code>
<code>p</code>	Policy for organizational domain	<code>p=quarantine</code>
<code>sp</code>	Policy for subdomains of the OD	<code>sp=reject</code>
<code>adkim</code>	Alignment mode for DKIM	<code>adkim=s</code> (strict)
<code>aspf</code>	Alignment mode for SPF	<code>aspf=r</code> (relaxed)