

Split Alignment: Diffusing SPF Vulnerabilities with DMARC

Muhammad Hamza

muhamza@vt.edu

Virginia Tech

Blacksburg, VA, USA

Mattijs Jonker

m.jonker@utwente.nl

University of Twente

Enschede, Netherlands

Raffaele Sommese

r.sommese@utwente.nl

University of Twente

Enschede, Netherlands

Simon Fernandez

fernands@univ-grenoble-alpes.fr

Univ. Grenoble Alpes, CNRS,

Grenoble INP, LIG

Grenoble, France

Olivier Hureau

hureauo@univ-grenoble-alpes.fr

Univ. Grenoble Alpes, CNRS,

Grenoble INP, LIG

Grenoble, France

Eric Pauley

pauley@vt.edu

Virginia Tech & Terrace Networks

Blacksburg, VA, USA

Taejoong Chung

tijay@vt.edu

Virginia Tech

Blacksburg, VA, USA

Abstract

Email remains the backbone of online communication, yet its authentication mechanisms continue to lag behind evolving attacker capabilities. Email authentication protocols are layered, with their interactions propagating and amplifying underlying vulnerabilities. This is clearly evident in the case of SPF vulnerabilities, where pairing them with DMARC's organizational context of authentication can lead to organization-wide spoofing. In this paper, we perform a comprehensive empirical analysis of this phenomenon. Despite SPF vulnerabilities being well known, our large-scale dataset reveals 528 k vulnerable policies. When paired with DMARC, these policies enable spoofing of 79 k domain hierarchies. To mitigate this, we propose Split Alignment, a practical configurational technique that protects established email-sending domains from all SPF failures in the domain hierarchy. We then use our novel methodology to validate the security and deliverability of Split Alignment in a consistent and reproducible manner. In doing so, we identify and disclose critical bugs in the DMARC authentication flows of 4 major email services. In this way, we strengthen the security of the email ecosystem by mitigating long-standing SPF vulnerabilities.

CCS Concepts

• Security and privacy → Security protocols; • Networks → Network measurement.

Keywords

Email, Email Security, SPF, DKIM, DMARC, Internet Measurement

ACM Reference Format:

Muhammad Hamza, Mattijs Jonker, Raffaele Sommese, Simon Fernandez, Olivier Hureau, Eric Pauley, and Taejoong Chung. 2026. Split Alignment: Diffusing SPF Vulnerabilities with DMARC. In *Proceedings of the 2026 ACM*

Internet Measurement Conference (IMC '26), October 12–16, 2026, Karlsruhe, Germany. ACM, New York, NY, USA, 15 pages. <https://doi.org/10.1145/3777912.3809146>

1 Introduction

Email remains the most used form of asynchronous communication on the Internet, with over four billion active users [31]. Despite its ubiquity, the ecosystem provides no security by default, which has enabled widespread abuse, most notably through spoofing. In response to these weaknesses, the community standardized a series of layered protocols. While this approach eased adoption through backwards compatibility, it also introduced critical security issues: the chain of protocols underpinning email security is anchored by its weakest link, a fragility further exacerbated by protocol interactions.

At the center of email authentication lies DMARC, a protocol designed to authenticate the domain visible to the user. It does not define its own authentication mechanism but instead relies on Sender Policy Framework (SPF), an IP address-based scheme originally developed for spam mitigation. Recent studies [4, 6, 12, 26, 33] have demonstrated significant weaknesses in SPF, all of which are directly inherited by DMARC. Moreover, DMARC amplifies the impact of these weaknesses through its organizational authentication context, allowing SPF vulnerabilities in subdomains to propagate across the domain hierarchy, thereby enabling organization-wide spoofing. An adversary could exploit this to send spoofed emails on behalf of established domains trusted by recipients.

In this work, we present the first measurement study of SPF and DKIM vulnerabilities, in conjunction with DMARC alignment. Using our large-scale email-centric subdomain dataset, we discover 528 k vulnerable SPF policies, the largest reported to date. We further correlate these vulnerabilities with DMARC alignment tags and show that they enable the spoofing of 79 k domain hierarchies, affecting the integrity of email at a significant scale.

Our findings suggest that SPF is unsuitable as the basis for email authentication. Wholesale replacement is infeasible as a new protocol would face insurmountable barriers to adoption. To address



This work is licensed under a Creative Commons Attribution 4.0 International License. *IMC '26, Karlsruhe, Germany*

© 2026 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-2327-8/26/10

<https://doi.org/10.1145/3777912.3809146>

these challenges, we propose a novel policy configuration and email sending technique called Split Alignment. This ready-to-deploy technique leverages DMARC’s alignment mechanism and delivers organization-wide protection against all SPF vulnerabilities while preserving email deliverability.

Beyond conformance to the DMARC specification, we validate the security and deliverability guarantees of Split Alignment through empirical evaluation. To this end, we develop a novel methodology to test DMARC enforcement on the receiver side in a consistent and reproducible manner. Using this approach, we find that Split Alignment provides security and deliverability across all major email services. We also uncovered critical implementation flaws in the DMARC authentication flow of four major email services, which we responsibly disclose, with several already remediated and acknowledged through bug bounty awards. As a result, DMARC is now correctly enforced, and Split Alignment can be used today as the first widely-deployable solution for mitigating the risks and weaknesses of the SPF protocol.

Our work provides a comprehensive look at DMARC and its underlying protocols, and in doing so, we make the following novel contributions:

- We compile the largest email-centric subdomain dataset to date, comprising 34.6 M subdomains under 3.2 M organizational domains using Certificate Transparency (CT) Logs and Passive DNS.
- We perform a large-scale measurement of how SPF and DKIM vulnerabilities in subdomains, in conjunction with DMARC’s relaxed alignment, make the entire domain hierarchy spoofable.
- We find that 79 k out of 81 k domain hierarchies are spoofable through a vulnerable policy within their subdomains.
- We propose Split Alignment as a practical configuration technique to protect established email sending domains from all SPF failures across the domain hierarchy.
- We formulate a generalizable methodology for evaluating email deliverability across email receivers in a reproducible and consistent manner.
- We open-source our testing harness as an artifact for future research.
- We discover and responsibly disclose bugs in 4 major email services, which have been fixed and recognized through bounties.

2 Background

Modern email relies on a variety of protocols to ensure the confidentiality and integrity of messages. Our work focuses on email integrity. We provide an overview of email identities and supporting protocols below.

2.1 Email Sender Identities

Email security protocols were introduced incrementally, with each defining a new identity or adopting an existing one. As a result, modern emails carry multiple identities, each corresponding to a domain that claims responsibility by sending the email, signing it, or appearing in the sender address.

Envelope Domain. Emails are transmitted between Mail Transfer Agents (MTAs) using the SMTP protocol [20]. During SMTP communication, the sending MTA establishes a TCP connection and issues a sequence of commands to identify itself and transmit the email data. One of these commands is MAIL FROM, which has the following syntax:

```
MAIL FROM: <local-part@domain>
```

The MAIL FROM command specifies an email address composed of a local-part and a domain. The domain in this address identifies the entity responsible for sending the message and is also used to process bounced messages if delivery fails. Conceptually, it serves the same role as the sender information written on the envelope of physical mail. In this paper, we refer to the domain of the MAIL FROM address as the **Envelope Domain**.

Signer Domain. During transmission, an email may pass through multiple MTAs before reaching the recipient’s inbox. By default, email provides no protection against modifications in transit, which means any MTA along the path can alter its contents. To address this lack of integrity, DKIM (subsection 2.2) was introduced as a mechanism that allows a domain to bind its identity to the content of an email by signing it. The domain that signs the email is embedded in the email body within the DKIM-Signature header:

```
DKIM-Signature: ... d=example.com; s=default; ...
```

In this paper, we refer to the domain that signs the email as the **Signer Domain**.

Visible Domain. An email contains multiple headers, one of which is the From header with the following syntax:

```
From: display-name <local-part@domain>
```

The From header specifies the domain traditionally regarded as the true origin of an email. Its value is also what email applications display to users, which makes the domain in the From header the primary target of spoofing attacks. In this paper, we refer to this domain as the **Visible Domain**.

Each protocol authenticates its own identity. However, the layered nature of these protocols means that the identities are inherently interconnected. The next section introduces these protocols and explains their interconnections.

2.2 SPF & DKIM

Email authentication is built upon two primary protocols, SPF and DKIM, each validating different identities and employing different underlying mechanisms.

2.2.1 Sender Policy Framework (SPF). SPF is one of the earliest email security protocols, proposed to mitigate spam as its main purpose [19]. It provides an IP-based authentication mechanism that allows domain owners to specify, via a DNS TXT record, which IP addresses are authorized to send messages on behalf of their Envelope Domain. When an email is transmitted, the destination MTA uses the Envelope Domain to retrieve the SPF policy and verify the sending server’s IP.

Over time, SPF has been shown to suffer from myriad limitations, undermining its utility as an integrity mechanism:

- **IP Addresses as security principals.** SPF authenticates emails by their originating IP address, treating controlled IPs as *security principals*. This assumption has broken down over time as IP address leasing has been embraced by Internet service and public cloud providers [28]. Recent works have shown practical exploitation of IP address reuse [1, 23], particularly against SPF [33].
- **Brittle configuration.** SPF allows the use of CIDR ranges to configure whitelisted IPs. However, admins often misuse this feature by configuring excessively broad IP ranges [12]. The protocol also supports inheritance (include mechanism) which allows one domain to incorporate the SPF policy of another, leading to bloated records that lead to spoofing [4].
- **Permissive by default.** SPF defines qualifiers to specify the disposition of messages, with the default being + (pass). This syntax is error-prone and minor mistakes can have severe consequences. For example, if a domain owner mistakenly publishes a policy ending with +all or all, the result is effectively an authorization of all possible IP addresses [26].

Such design limitations, combined with susceptibility to misconfiguration, significantly weaken SPF's security guarantees. DMARC (subsection 2.3), being layered on top of SPF, inherits these weaknesses, thereby enabling organization-wide spoofing.

2.2.2 DomainKeys Identified Mail (DKIM). DKIM is the main protocol that provides integrity during email delivery by allowing senders to cryptographically sign the contents of an email [11]. The signature is reflected in the DKIM-Signature header, and the corresponding public key is published as a DNS TXT record, binding the identity of a domain to the email's content. To verify the signature, the receiver parses the DKIM-Signature header and extracts the d= tag, which specifies the signing domain, and the s= tag, which identifies the selector. Then it retrieves the public key from DNS under selector ._domainkey.domain and validates the signature.

The security of DKIM is based on asymmetric cryptography and has shown resilience to attacks over time. It also uses a more robust policy syntax, which is less prone to misconfigurations [34] compared to SPF. The usage of DKIM has increased over time, especially because major providers such as Gmail and Yahoo require the use of DKIM for bulk email delivery [15, 35]. Building on this trend, third-party email sending services, such as Mailgun and Amazon SES, enforce the use of DKIM by requiring users to set it up during the domain verification step [2, 25].

In practice, DKIM has faced limited operational challenges, such as insufficient key rotation or improper header signing, but these issues are well-documented and mitigated by implementing best practices [10, 29], using secure keys [5], or using automated tools [30].

Building on top of SPF and DKIM, we have DMARC, which is the primary protocol for sender authentication in modern email. The next section describes its functionality and how it integrates its underlying protocols.

2.3 Domain-based Message Authentication, Reporting, and Conformance (DMARC)

DMARC [21] is the primary protocol for the authentication of the Visible Domain, the domain shown to the user upon delivery, and the main target of spoofing. Beyond authentication, it also provides centralized policy enforcement and reporting. In effect, DMARC unifies all underlying identifiers to provide a single authenticated identity, the Visible Domain. DMARC also forms the basis of additional user-facing identity assurance through protocols such as BIMi [7], which allows organizations to display their verified brand logos alongside authenticated emails.

Unlike other email protocols, DMARC policies apply to the entire domain hierarchy, unless a subdomain defines its own policy, making the Organizational Domain (eTLD+1, e.g., example.com) the primary point of control. A DMARC policy is expressed with the following syntax:

```
v=DMARC1; p=reject; adkim=s; aspf=r;
```

This policy specifies how DMARC authentication should be applied (aspf and adkim tags) and what action to take if authentication fails (p tag).

Protocol layering and alignment. DMARC relies on the authenticated domains provided by SPF and DKIM, requiring an email to pass at least one of these mechanisms before comparing the authenticated domain with the visible domain in a process called *alignment*. If the two align, the email passes DMARC. Figure 1 provides an overview of the DMARC authentication flow. Alignment can be configured separately for the SPF and DKIM authenticated domain through the aspf and adkim tags. There are two primary modes of alignment:

- **Relaxed (default):** Only the Organizational Domain of the authenticated domain and the Visible Domain must match to successfully align. This places authentication in the organizational context, which means that any authenticated domain within a domain hierarchy can be used to pass DMARC on behalf of any other domain within the hierarchy. For example, subdomain.example.com will align with both mail.example.com and example.com and can be used to send DMARC-passing emails on behalf of both.
- **Strict:** The Authenticated Domain must exactly match the Visible Domain to successfully align and pass DMARC. This enforces authentication at the domain level rather than the organizational level and isolates each domain. For example, subdomain.example.com will only align successfully with itself and thus cannot be used to send DMARC-passing emails on behalf of another .example.com or example.com.

Based on DMARC's authentication flow, we observe that in its default state, it carries two major flaws that lead to organization-wide spoofing:

DMARC can be passed through either **SPF OR DKIM**: An attacker can compromise the weaker of the two protocols, i.e. SPF, to pass DMARC authentication.




Figure 1: DMARC authentication flow with relaxed alignment (default): SPF verifies the Envelope Domain’s IP, DKIM verifies the Signer Domain’s signature, and both are compared to the Visible Domain for alignment. At least one mechanism must align to pass DMARC.

DMARC **defaults** to **relaxed** alignment: An adversary only needs to exploit the weakest configured subdomain to pass DMARC for all domains in the hierarchy.

As explored in this work, these weaknesses in DMARC yield practically-realizable attacks. In studying the structure and prevalence of these, we propose a new configurational approach that prevents organization-wide spoofing.

3 The Threat of DMARC Spoofing

DMARC aims to provide email authentication at the organizational level. However, in its default **relaxed** configuration, DMARC leaves organizations vulnerable to attacks anywhere in the weakest link of the DNS hierarchy. The perils of propagating trust through the domain hierarchy have long been known in the service hosting space [9, 23, 28], but the same issues arise in an organization’s email presence when DMARC’s relaxed alignment is used.

Consider the following scenario: An organization sets up a DMARC policy on their Organizational Domain `example.com` that rejects unauthenticated messages. Their usual sending practices involve using a major email vendor, so they provision DKIM records at the apex, and SPF records under `spf.example.com`. They follow provider guidance for setting up their DMARC policy as follows:

```
v=DMARC1; p=reject; adkim=r; aspf=r
```

This policy is recommended by the email provider because, with SPF records deployed on a subdomain, relaxed alignment is required to send DMARC-passing emails on behalf of `@example.com` addresses, an established sending domain (an assumption we will show is invalid in section 6).

At the same time, the organization decommissions its own cloud servers previously used to send email. These servers used SPF records deployed at `spf.example.com`. However, the organization fails to remove these SPF records. Because the records refer to cloud IP address ranges, they are vulnerable to the BreakSPF [33] attack. This attack scenario is one of several weaknesses in the SPF that underpins DMARC.

Eventually, an attacker discovers this SPF configuration left by the organization, allocates a matching cloud IP address, and begins sending phishing emails from `@example.com` addresses. Because DMARC alignment is set to **relaxed**, weaknesses in subdomain configurations allow direct sending with an established Visible Domains, i.e., `example.com`. Notably, this ability is far stronger than the ability to send emails from only a subdomain address, as spoofed emails can match existing contacts in recipients’ email inboxes, interleave with conversations from the organization’s new production email workflows, and provide validated brand iconography through systems such as BIMi [7].

In effect, in the scenario above, DMARC’s relaxed alignment allows for complete organizational email takeover through compromising any subdomain of the Organization. Our work investigates the prevalence of this vulnerability, the DMARC deployment practices that lead to it, and the practicality of new sending strategies that mitigate systemic risk to organizational email authentication in practice.

4 DMARC Spoofing In the Wild

In this section, we perform the first large-scale study of SPF and DKIM vulnerabilities through the lens of DMARC authentication and alignment. Prior works (mentioned in section 8) have examined these protocols in isolation, whereas we are the first to study them in conjunction to understand how their interactions affect the security of the whole ecosystem. In doing so, we aim to answer the following questions:

- RQ1.** How prevalent are SPF and DKIM vulnerabilities?
- RQ2.** What is the distribution of deployed DMARC alignment configurations?
- RQ3.** How do deployed DMARC alignment configurations magnify or rectify the effect of underlying vulnerabilities?

4.1 Data Acquisition

Conducting a comprehensive study of DMARC is challenging, as it requires examining the SPF and DKIM policies, which are often deployed at the subdomain level. Analyzing these subdomains at

Table 1: Breakdown of our DNS dataset. (All Organizational Domains have a DMARC Record.)

	O	P	$O \cap P$	$O \cup P$
Org. Domains	2.5 M	1.4 M	747.6 k	3.2 M
Subdomains	22.6 M	15.6 M	3.6 M	34.6 M
→ w/ SPF	20.2 M	14.7 M	3.6 M	31.4 M
→ w/ DKIM	2.4 M	1.3 M	12.4 k	3.7 M
→ w/ DMARC	1.4 M	1.6 M	548.3 k	2.5 M

O = OpenINTEL CTLogs, P = Farsight Passive DNS.

scale is difficult because subdomain discovery itself is a nontrivial task.

To address this challenge, we compile the largest email-centric subdomain dataset to date. We use Certificate Transparency (CT) logs and Passive DNS as complementary data sources for subdomain discovery.

CTLogs. A practical approach to discovering subdomains in the wild is to mine the Subject Alternative Name (SAN) and Common Name (CN) fields of TLS certificates. This can be done at scale by leveraging Certificate Transparency (CT) logs, a comprehensive repository of all issued certificates. Building on this, we extract Organizational Domains with valid DMARC records from the OpenINTEL dataset and retrieve their associated certificates from CTLogs. From these, we parse SAN and CN fields to obtain approximately 434.5 M raw subdomains.

Passive DNS. While CTLogs provide a strong foundation for subdomain discovery, they primarily capture certificates issued for web-facing services and may therefore miss subdomains used exclusively for email. To address this limitation, we augment our dataset with Passive DNS, ensuring broader coverage of subdomains that reflect real-world email use. This is particularly beneficial for identifying DKIM records, which reside under arbitrary subdomains (due to selectors) that are missed by active scanning methods.

For this study, we use the data from DomainTools Farsight Security Information Exchange (SIE) and extract TXT records that indicate email authentication mechanisms during the period from January 2024 to February 2025, resulting in 109.4 M observed subdomains.

Data Preparation and Characteristics. To account for differences in dataset collection dates and underlying shifts in the DNS, we perform a fresh scan of TXT records for SPF, DKIM, and DMARC policies for all subdomains and their corresponding Organizational Domains. The resulting dataset includes all subdomains that have either an SPF or DKIM record, along with an effective DMARC policy, whether explicitly defined or inherited from the Organizational Domain.

Our dataset, summarized in Table 1, represents the largest email-centric dataset of subdomains compiled to date. In total, the dataset contains 34.6 M Subdomains with deployed policies across 3.2 M Organizational Domains. Table 1 also compares our two data sources, showing minimal overlap and highlighting their complementary nature. This is especially true for DKIM, where OpenINTEL (O) only captures major DKIM selectors (e.g., `default` and `mail`), while

Table 2: Distribution of subdomains with vulnerable SPF and DKIM policies.

Vulnerability	Distinct Subdomains
SPF all/+all	35 K
BreakSPF	493 K
DKIM Compromised Keys	11
Total	528 K

Passive DNS (P) captures selectors queried during real-world email usage. This large-scale dataset uniquely enables our comprehensive measurement of the email ecosystem, which we discuss in the following section.

4.2 Measuring Vulnerable Subdomains

Because DMARC relies on SPF and DKIM for authentication, we examine our subdomain dataset (subsection 4.1) alongside known flaws and misconfigurations in these protocols to identify subdomains with vulnerable policies.

SPF. SPF-related vulnerabilities can be broadly classified into syntactic failures (i.e., due to its error-prone syntax) and abuse of IP control as a trust anchor.

We first focus on syntactic misconfigurations, particularly on policies that end with `+all` or `all`. Such policies effectively authorize all IP addresses to send emails, which means that any subdomain configured in this way can be leveraged for spoofing from virtually any IP. To identify such vulnerable subdomains, we retrieve the SPF policies from our dataset and search for instances of this misconfiguration.

In addition, inspired by BreakSPF [33], we analyze the IPs listed in SPF policies to determine whether they can be exploited through IPs obtainable from cloud providers. To that end, we retrieve the SPF policies from our dataset and recursively resolve their includes as specified in the RFC. From these expanded policies, we extract the listed IPs and cross-reference them against the IPs available through cloud providers. In doing so, we improve on the approach of BreakSPF by augmenting their dataset of 106 K IPs with additional IPs measured by the DScope Internet telescope [27], thereby increasing coverage by 11.5 M.

DKIM. As DKIM is a cryptographic scheme, its security depends on the keys used. To evaluate vulnerabilities in DKIM, we retrieve the DKIM policies from our dataset, extract the public keys, and evaluate them against known cryptographic weaknesses. Specifically, we test for vulnerabilities such as Fermat (CVE-2022-26320) and ROCS (CVE-2017-15361), as well as bugs including the Debian OpenSSL bug (CVE-2008-0166), the GitKraken bug (CVE-2021-41117), and weak keys exposed in the 2025 Fortinet/FortiGate leak (CVE-2022-40684).

Results. Addressing RQ1, our analysis reveals 528 k unique vulnerable subdomains, a large majority of which are due to SPF policies, while only 11 stem from compromised DKIM keys. Despite our study’s tight scope to domains with active DMARC deployments, this count represents the largest number reported in recent email studies, highlighting the breadth of our data sources and the

Table 3: Breakdown of aspf and adkim alignment modes across policies deployed on organizational and subdomains. The use of relaxed alignment remains common in the wild.

Domain	Alignment	aspf	adkim	Total
Org.	ℝ	70.07%	71.70%	3.2 M
	R	18.32%	17.99%	
	S	11.61%	10.31%	
Sub.	ℝ	39.24%	42.38%	2.5 M
	R	15.47%	15.13%	
	S	45.28%	42.49%	

ℝ = Implicit Relaxed, R = Explicit Relaxed, S = Explicit Strict.

rigor of our analysis. Table 2 provides an overview of the observed vulnerabilities.

Our results highlight SPF’s persistent structural weaknesses. Even with well-documented weaknesses and published best practices in SPF use, we continue to observe mass vulnerability among organizations. This persistence of vulnerabilities arises from the inherent design of SPF, which rests upon assumptions that do not hold for the modern internet. Updating SPF to address these issues is impractical given its age and the slow pace of adoption that is characteristic of the email ecosystem. Taken together, these factors show that SPF is not suitable as a basis of authentication.

This has significant implications for DMARC, which relies on either SPF or DKIM for authentication. As a result, its security is compromised by the weaker of the two protocols, most commonly SPF. DMARC integrates its underlying protocols through a configurable alignment mechanism, which we study in the next section.

4.3 Deployed DMARC Alignment

DMARC introduces the organizational authentication context through its relaxed alignment mechanism, which allows a vulnerable subdomain to compromise the security of the entire domain hierarchy. This means that the usage of relaxed alignment directly implies the potential for spoofing. To study the prevalence of this issue, we retrieve the DMARC policies from our dataset (subsection 4.1) and analyze the usage of alignment tags (aspf and adkim).

Results. Addressing RQ2, Table 3 shows the distribution of alignment configuration of deployed DMARC policies for both Organizational Domains and Subdomains. Around 90% of Organizational Domains make use of relaxed (ℝ and R) alignment, either defined explicitly and implicitly. Since DMARC policies deployed at the Organizational level are inherited by subdomains, a large number of domain hierarchies remain susceptible to spoofing.

A closer look shows that roughly 70% of the Organizational Domains do not explicitly set the optional DMARC alignment tags and fallback to the default relaxed (ℝ) mode. Because these tags are optional, administrators often have little incentive to investigate or configure them, leading to high reliance on the default configuration. Among roughly 18% of domains that explicitly configure alignment, still choose relaxed (R). The use of strict (S) alignment remains the lowest at around 10%.

Table 4: Spoofability of domain hierarchies with a vulnerable subdomain. The majority of hierarchies remain spoofable through their subdomains.

Subdomain	Organizational Domain	
	Spoofable	Not Spoofable
Spoofable	97.14%	0.01%
Not Spoofable	0.24%	2.61%
Total	81 K	

For subdomains, the use of DMARC policies expectedly remains low, with only 2.5 M out of 34.6 M subdomains publishing a DMARC policy. The presence of a DMARC policy at the subdomain level suggests an intention to improve security, as shown by the relatively high use of strict alignment (S), which appears in roughly 42% to 45% percent of cases across both alignment tags. However, relaxed (ℝ and R) alignment still remains common among the subdomains.

Our analysis reveals that a large portion of domain hierarchies use relaxed alignment configuration, leaving them susceptible to spoofing. The next section pairs discovered vulnerabilities in subdomains (from subsection 4.2) with DMARC policies to measure the prevalence of domain hierarchies that are practically spoofable.

4.4 Organizational Spoofing with DMARC

To measure the true scale of organization-wide spoofing, we need to study the underlying SPF and DKIM policies in conjunction with the deployed DMARC policies. To this end, we use the vulnerable subdomains identified in subsection 4.2 and retrieve the DMARC policies of their Organizational Domains and sibling domains. If a sibling domain does not have its own DMARC policy, we apply the Organizational Domain’s policy as specified in the RFC [21]. Next, we examine the relevant alignment tag (aspf or adkim) based on the type of deployed policy (SPF or DKIM). For example, if the vulnerable policy is for SPF, aspf will be the relevant alignment tag. Finally, if the effective alignment for a certain domain is set to relaxed, which is often the case, it is classified as susceptible to spoofing.

Results. Addressing RQ3, Table 4 shows the distribution of the spoofable domain hierarchies in the wild. In total, we identify 81 k Domain Hierarchies with at least one vulnerable subdomain. Among these, 97.14% (79 k) of the domain hierarchies are entirely spoofable through their vulnerable subdomains. This is due to the use of relaxed alignment, which allows local vulnerabilities to propagate through the rest of the domain hierarchy.

Next, we analyze the effective DMARC alignment being used by the domains and show how DMARC’s design, being permissible by default, leads to organization-wide spoofing. Table 5 details the organizational domains and subdomains along with their effective DMARC alignment tags, scoped to organizations with a vulnerable policy. The results show that 97.38% of Organizational Domains use the relaxed alignment configuration (ℝ or R), leaving them susceptible to spoofing. In particular, 83.54% of organizations do not specify these tags at all, falling back to the default relaxed (ℝ) mode.

Table 5: DMARC policies of domains with a vulnerable subdomain. Most domains remain spoofable due to their use of relaxed alignment.

Org. Domain with a Vulnerable Subdomain			
Org. Domain's Policy			
	ℝ	R	S
	83.54%	13.84%	2.62%
Total	81 K		
Subdomain with a Vulnerable Sibling			
Org. Domain's Policy			
Subdomain's Policy	ℝ	R	S
∅	79.02%	16.80%	0.41%
ℝ	0.75%	0.93%	0.06%
R	0.05%	1.94%	0.00%
S	0.02%	0.00%	0.02%
Total	1.7 M		

∅ = No DMARC Policy, ℝ = Implicit Relaxed, R = Explicit Relaxed, S = Explicit Strict, Shaded = Spoofable.

Furthermore, we observe the role of DMARC as a protocol for centralized control across the entire domain hierarchy. In practice, this can be beneficial since a strong DMARC policy at the Organizational Domain level extends protection to the entire hierarchy. However, this same mechanism backfires when the Organizational Domain defines a permissive policy, defining relaxed alignment either implicitly or explicitly. Our findings show that 95.82% of subdomains inherit this relaxed alignment configuration, leaving them vulnerable to spoofing through their sibling domains. Moreover, the subdomains that do specify a DMARC policy of their own tend to use relaxed alignment. Overall, 99.55% (1.7 M) of the analyzed subdomains are spoofable.

Takeaways. Our large-scale measurement of the DMARC protocol stack shows that SPF is unsuitable as a basis for authentication due to persistent structural weaknesses. Combined with DMARC's default relaxed alignment configuration, which enables authentication in the organizational context, these weaknesses propagate across the entire domain hierarchy, including established email-sending Visible Domains, the primary target for email spoofing.

5 Root Causes of DMARC Failure

Previous studies [22, 24, 33] have shown that the email ecosystem is highly centralized, with a small number of major providers managing a large fraction of domains. Given this central role, they have a marked impact on whether domain owners deploy effective DMARC policies that strengthen authentication and mitigate spoofing. To understand how well this role is fulfilled in practice, we conduct a qualitative analysis of email provider documentation to answer the following research questions:

RQ4. What role do Email Service Providers play in disseminating DMARC policies and practices?

Table 6: Documentation survey of major email service providers. Most email services either give no recommendation or promote relaxed alignment.

Email Service	DMARC Explained	Alignment Explained	Alignment Recommended
Google	✓	✓	Mixed
Mailgun	✓	✓	Relaxed
Amazon SES	✓	✓	Relaxed
Twilio	✓	✓	Relaxed
Mailjet	✓	✓	Relaxed
MailChannels	✓	✓	None
MailChimp	✓	✓	None
Outlook	✓	✗	None
Salesforce	✓	✗	None
SMTP2GO	✓	✗	None
Dotdigital	✓	✗	None

RQ5. What are the operational factors that shape these DMARC policy recommendations?

To this end, we perform a comprehensive survey of the documentation of common Email Service Providers (ESPs) identified from a previous study [33]. For each provider, we identify both core documentation (e.g., main provider support portal) and ancillary sources (e.g., blog posts) that would contribute to customer understanding of DMARC, alignment, and recommended configurations. In addition to recommendations, we also noted underlying reasons to understand provider sentiment on the practical impact of DMARC alignment mechanisms. We aim to be as thorough as possible and review all referenced materials to ensure completeness.

Policies and Practices. Table 6 summarizes the findings of our survey. Addressing RQ4, even though all providers describe the overall functionality of DMARC, only 7 of the 11 surveyed Email Services mention and explain the functionality of the optional alignment tags (aspf/adkim). Alignment and its configurable tags are a core aspect of DMARC policy design and should be discussed in detail.

In examining how providers promote best practices, 6 of the 11 providers do not recommend a specific DMARC alignment policy, which effectively results in the use of implicitly configured relaxed alignment. This aligns with the high prevalence of implicit relaxed DMARC policies observed in real-world deployments. Among the 5 that do provide recommendations, 4 advise using relaxed alignment, a choice that enables organization-wide spoofing. Out of these, the only outlier is Google, which notes that relaxed alignment provides sufficient protection in most cases but that strict alignment can be used in certain scenarios for enhanced security. This trend suggests that providers often prioritize deliverability and ease of deployment over security.

Operational Factors. Addressing RQ5, we examine the motivations behind these recommendations and find that operational considerations favor the use of relaxed alignment, as it is more permissive and therefore less prone to configurational errors. Many mail services deploy MX servers dedicated to handling bounced messages. By design, the SPF identifier (Envelope Domain) serves as the bounce address. To support this, providers use a dedicated

subdomain as the Envelope Domain, allowing the third-party MX servers to handle bounced messages without interfering with the domain’s regular email traffic. Under this configuration, the Envelope Domain would not align with the Visible Domain under the strict configuration (`aspf=s`). This technical incompatibility seemingly makes strict alignment impractical for many email services, a limitation we address in section 6.

Takeaways. The use of relaxed alignment is widespread across providers, a practice reinforced by provider documentation. Our findings suggest that this tendency stems from deployability concerns, but in turn leads to weak policies that lead to organization-wide spoofing. Practical defense against email spoofing demands an approach that can be seamlessly adopted by large-scale providers while enabling email senders to address the vulnerabilities inherent in SPF.

6 Split Alignment

Our proposed attack leverages faulty SPF records and relaxed alignment configuration in the DMARC policy. Introducing a new protocol to address these issues is impractical as it would cause disruptive change to the ecosystem.

An alternative solution would be to fix SPF policies and use strict alignment. However, both of these strategies are impractical and suboptimal. Practicing good policy hygiene for SPF requires constant monitoring and policy maintenance, making this impractical in operational settings, particularly in large organizations where email configurations are managed by myriad independent sub-units. Furthermore, simply setting alignment tags to strict does not fully secure visible domains (as discussed in subsection 6.2) and is not compatible with many third-party email-sending services (as discussed in section 5).

To address these practical limitations, and motivated by our empirical and analytical analysis of the DMARC ecosystem (in section 4), we propose Split Alignment, a policy configuration and email sending technique that achieves three key properties:

- **Security:** Protects established Visible Domain from all SPF failures.
- **Deployability:** Compatible with existing protocol implementations and third-party email sending services.
- **Deliverability:** Preserves normal email delivery behavior for benign emails.

6.1 Design

Split Alignment protects established Visible Domains, which can be an organizational domain or a subdomain, from SPF failures across the entire domain hierarchy while being deployable and ensuring deliverability. This is done by anchoring DMARC authentication on DKIM while preserving SPF authentication.

Split Alignment uses DMARC’s strict alignment mode and configures SPF and DKIM such that only the DKIM domain (Signer Domain) aligns with the Visible Domain and is used for DMARC authentication, while the SPF domain (Envelope Domain), which is deployed on a subdomain of the Visible Domain, is intentionally misaligned and therefore excluded from evaluation. This entirely removes the dependence of DMARC authentication on the outcome

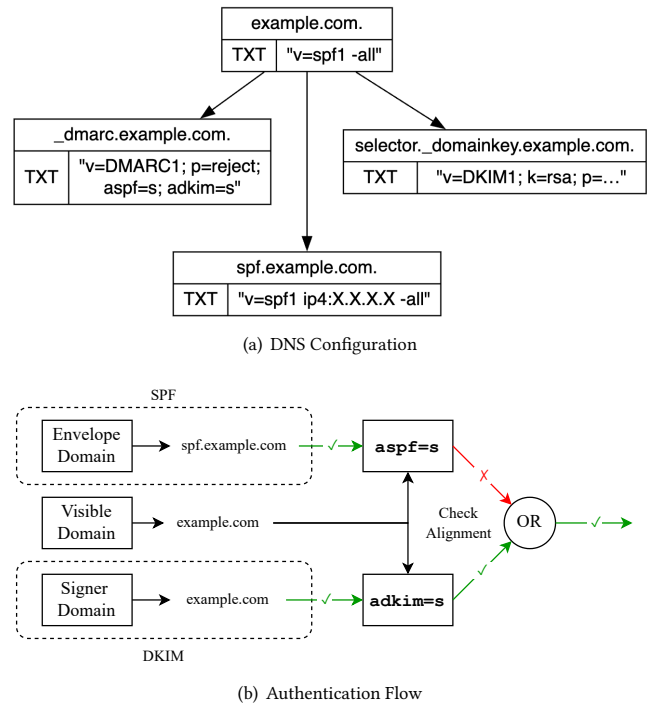


Figure 2: Split Alignment configuration and flow: SPF policies are relocated to `spf.example.com` and intentionally misaligned, excluding SPF from DMARC evaluation. DKIM (`d=example.com`) aligns under strict mode and passes DMARC, protecting the Visible Domain from SPF vulnerabilities.

of SPF, thus protecting the Visible Domain from all SPF failures and achieving **security**.

Since SPF authentication is still preserved, any receiver-side internal policy that relies on SPF results, such as spam detection, will continue to function, thereby achieving **deployability**.

Moreover, because the SPF policy is deployed on a subdomain, the configuration is compatible with email-sending services that require the use of a subdomain for SPF records (as discussed in section 5) and achieves **deliverability**.

The DNS configuration and authentication flow of Split Alignment is summarized in Figure 2. Split Alignment can be configured by following the steps described below.

DMARC Configuration. To prevent vulnerable subdomains from influencing authentication outcomes, configure the DMARC `aspf` and `adkim` tags to strict (`s`) at the organizational level.

`v=DMARC1; p=reject; adkim=s; aspf=s;`

Setting alignment to strict ensures that only the SPF and DKIM records on the Visible Domain are considered for authentication. This isolates the Visible Domains from all influence of subdomains or sibling domains.

This configuration substantially strengthens the security of the Visible Domains. However, they may still be exposed through exploitable SPF policies configured on the Visible Domains themselves.

We address this limitation in the next steps by refining the SPF and DKIM policies.

SPF Update and Relocation. To keep DMARC authentication anchored on DKIM while preserving SPF for operational needs (i.e., email deliverability and compatibility with third-party email services), relocate each SPF policy away from the established Visible Domains to dedicated subdomain(s). Each Visible Domain can have a dedicated subdomain, or there can be a singular subdomain with an SPF policy that applies to all Visible Domains. This dedicated subdomain will be used as the Envelope Domain (MAIL FROM) during email sending.

Next, for every established Visible Domain, update the existing SPF policy to be maximally restrictive, e.g.: `v=spf1 -all`, which explicitly disallows all IPs. This insulates Visible Domains from any failure in their own SPF policies.

Since each SPF policy is deployed on a dedicated subdomain, the Envelope Domain no longer aligns with the Visible Domain under strict alignment. Consequently, SPF authentication results are excluded from DMARC evaluation. Nevertheless, emails remain SPF-authenticated for receivers that rely on SPF results for ancillary functions such as spam detection.

DKIM Configuration. Our measurement results (in section 4) show that vulnerabilities in DKIM are largely absent, especially when contrasted with those in SPF. Based on this, under Split Alignment DMARC uses DKIM as the only basis of authentication. To achieve this, sign outgoing emails with DKIM and deploy the corresponding DKIM records on each email-sending Visible Domain, such that the Signer Domain and the Visible Domain align under a strict policy. Specifically, publish the record at:

```
<selector>._domainkey.<visible-domain>
```

This configuration enables DMARC to pass via DKIM.

Note that, under Split Alignment, at a minimum, the email needs to be DKIM signed.

6.2 Split Alignment vs. Strict Alignment

In this section, we present a qualitative comparison between Split Alignment and Strict Alignment. They look similar in terms of structure and naming, but they are entirely distinct concepts.

Strict Alignment is a configuration setting for the `aspf` and `adkim` tags in a DMARC policy that changes how DMARC authentication is applied. On the other hand, Split Alignment uses DMARC's strict alignment configuration, along with carefully structured DNS records, to provide a configuration that is both secure and compatible with diverse email sending scenarios.

Table 7 provides a comparison of the required changes needed, security guarantees, and compatibility of Split Alignment and Strict Alignment against the baseline of Relaxed Alignment. The table assumes that the domain being configured represents the worst-case scenario and requires the most extensive configurational changes. In this scenario, the domain deploys an SPF policy but lacks any DKIM record.

Relaxed Alignment. The default configuration of DMARC sets both `aspf` and `adkim` to relaxed and is used by the majority of domains today (as shown in subsection 4.3). Being the most permissive configuration, it is the easiest to deploy and is compatible with

third-party email-sending services. However, it does not provide any security against vulnerable policies across the domain hierarchy and is one of the core mechanisms enabling our proposed attack.

Strict Alignment. An improvement over relaxed is to set the alignment tags (`aspf` and `adkim`) to strict. Under this configuration, if the SPF and DKIM authenticated domains differ from the Visible Domain, they will be misaligned, and DMARC will fail. Without any configurational change, Strict Alignment is incompatible with third-party email-sending services, which require a dedicated subdomain as the SPF identifier (Envelope Domain) to handle bounce messages (as discussed in section 5). Furthermore, Strict Alignment protects against attacks originating from other domains within the hierarchy, but does not completely isolate the Visible Domain from its own SPF policy, since that can still be used to pass DMARC. Therefore, simply setting alignment tags to strict is neither optimal nor practical.

Split Alignment. Applying careful configurational changes to SPF and DKIM records and setting alignment to strict brings us to the Split Alignment Configuration. Split Alignment completely isolates the established Visible Domains from all SPF vulnerabilities across the domain hierarchy while remaining compatible with third-party email-sending services and ensuring deliverability. Split Alignment can be implemented by following the steps described in subsection 6.1.

6.3 Split Alignment and DMARC Design

DMARC was originally designed as a fault-tolerant protocol, with redundancy provided by either relying on SPF or DKIM for authentication. However, this design assumes that both authentication mechanisms provide integrity of comparable strength. This assumption simply does not hold in modern networks, where SPF's IP-based mechanisms degrade the trust model of the DMARC authentication chain instead of serving as a safety net.

Furthermore, critics may argue that Split Alignment introduces an apparent mismatch between the network layer, where SPF authentication succeeds, and the policy layer, where DMARC alignment fails. We contend that this behavior is intentional as it preserves the existing infrastructure for message routing and basic reputation checks through an SPF pass, while at the same time informing modern policy engines that an IP address should no longer be treated as a reliable proxy for organizational identity. By separating message transport authorization from identity verification, Split Alignment formalizes the ongoing transition in email security from trust based on IP addresses (SPF) to trust grounded in cryptographic keys (DKIM).

Takeaways. By deploying Split Alignment, all established Visible Domains are isolated from SPF failures throughout the domain hierarchy while being compliant with diverse email sending setups. In doing so, we practically address a core security issue in email without introducing disruptive changes. We next perform a detailed evaluation of our technique to validate our claimed guarantees.

7 Evaluating Split Alignment

Previous studies [3, 4, 6, 8, 10, 13, 16, 29] have shown that receiver-side protocol behaviors often deviate from specifications, either due

Table 7: Comparison of different alignment configurations. (Assumption: The Required Changes column assumes that the Visible Domain has a deployed SPF policy but no deployed DKIM policy.)

Name	Alignment Policy		Required Change		Secure Against			3rd Party Compatible
	aspf	adkim	SPF Policy	DKIM Policy	Own SPF Vuln.	Hierarchy's SPF Vuln.	Hierarchy's DKIM Vuln.	
Relaxed Alignment	r	r	N/A	N/A	✗	✗	✗	✓
Strict Alignment	s	s	○	○	✗	✓	✓	✗
Split Alignment	s	s	●	●	✓	✓	✓	✓

○ = Configurational Change Not Applied, ● = Configurational Change Applied, ✗ = Not Provided, ✓ = Provided.

to software bugs or, more commonly, the application of internal policies. This implies that even minor changes in sending configurations, intentional misalignment of the SPF identifier in the case of Split Alignment, can have a meaningful and measurable effect on deliverability in practical settings. Therefore, while Split Alignment is theoretically sound, it must undergo empirical evaluation to determine whether its security and deliverability claims hold across the diverse protocol implementations observed in the wild.

We evaluate the utility and practicality of Split Alignment by asking the following research questions:

- RQ6.** Does Split Alignment negatively affect the *deliverability* of emails?
- RQ7.** Does Split Alignment provide the claimed *security* against our proposed attack?
- RQ8.** Do email receivers correctly implement the DMARC authentication flow (*deployability*)?

Guided by these questions, we design a reproducible methodology to consistently evaluate Split Alignment and DMARC authentication flows of major email services.

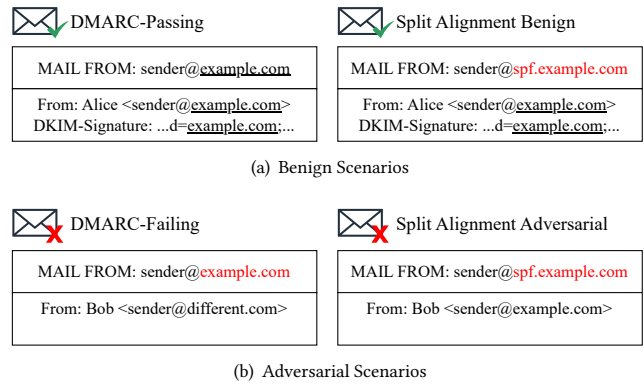
7.1 Methodology

Deliverability and security are closely associated with DMARC outcomes, which form the central focus of our evaluation. There are several ways to infer DMARC outcomes, such as by examining the Authentication-Results header or inspecting SMTP error codes. However, these indicators are not consistently available across all email services, as some omit headers, while others authenticate asynchronously and do not return errors for DMARC failures. As a result, reliably determining DMARC outcomes consistently across email receivers remains a challenging task.

Moreover, email services often enforce internal policies that can override DMARC decisions. Consequently, an email that fails DMARC may still be delivered even when the policy explicitly requires rejection, and the opposite can also occur. These internal policies operate on numerous variables, such as domain reputation, IP reputation, and historical email volume, many of which are difficult for researchers to control. This shows that deliverability outcomes alone do not provide a reliable basis for inferring DMARC results, and can vary across different email receivers.

To address this fundamental challenge, we develop a methodology based on differential analysis of email sending configurations to systematically compare how various configurations affect DMARC outcome and deliverability in conjunction. The core intuition behind this methodology is that emails with the same theoretical

Figure 3: Email sending configuration for evaluated scenarios. (Successful alignments are underlined; alignment failures under strict appear in red.)



DMARC outcome should exhibit consistent deliverability for a given receiver, regardless of differences in the sending configuration:

$$D(c_1, r, o) \equiv D(c_2, r, o)$$

where D denotes the deliverability outcome, c_1 a baseline and c_2 a different sending configuration, r is the email receiver, and o is the theoretical DMARC outcome. Any deviation from this expectation indicates either an implementation bug or an artifact of an applied internal policy.

This methodology enables consistent and reproducible evaluation of sending configurations across email receivers. Next, as summarized in Figure 3, we apply this methodology to test different scenarios under Split Alignment.

Benign Scenario. This scenario is used to evaluate **deliverability** and represents the standard email sending configuration, where the email passes DMARC and is successfully delivered. Using our methodology, we test whether Split Alignment achieves the same level of deliverability as a perfectly configured DMARC-passing email.

To establish a baseline, we first send a correctly configured DMARC-passing email that is authenticated through both DKIM and SPF, with all identifiers aligned with the Visible Domain, and record its delivery outcome. We then send an email configured under Split Alignment and compare its delivery outcome to that of the baseline. This comparison reveals whether partial alignment of

identifiers has any negative impact on the email due to the application of internal policies.

Adversarial Scenarios. This scenario is used to evaluate the security and represents an adversarial email sending configuration, where the email should fail DMARC and be rejected. Using our methodology, we test whether Split Alignment provides security against the adversary by achieving the same lack of deliverability as a DMARC-failing email.

To establish a baseline, we first send a DMARC-failing email that is not DKIM signed and has a SPF passing identifier that is completely different from that of the Visible Domain and record its delivery outcome. We ensure that SPF passes because SPF authentication occurs before DMARC evaluation, and we want the deliverability outcome to result specifically from a DMARC failure rather than a SPF failure. We then send an email crafted under the adversarial Split Alignment configuration, where an adversary uses a subdomain as the SPF identifier to send an email on behalf of an established Visible Domain. In this scenario, the adversary will not have the ability to sign the email with DKIM. Provided that DMARC is implemented correctly, the SPF authenticated domain (Envelope Domain) will not align with the Visible Domain under strict configuration, and the email will fail DMARC. We record the delivery outcome and compare it with the baseline to evaluate whether Split Alignment correctly triggers DMARC failure and prevents the delivery of adversarial emails. This comparison reveals whether the DMARC authentication flow, and in particular the alignment logic, is correctly implemented by email receivers.

We apply this methodology to widely used email services that support DMARC, as identified in prior studies [8, 16]. Since an email’s Visible Domain can be an Organizational Domain or a Subdomain, we repeat the experiment under both configurations to capture all scenarios. All experiments are performed with the DMARC policy set to reject (`p=reject`) and alignment tags set to strict (`aspf=s` and `adkim=s`).

7.2 Results

Our evaluation shows that Split Alignment offers both deliverability and security (excluding confirmed bugs under remediation) across all major email services. Table 8 summarizes the results of our analyses with detailed deliverability outcomes mentioned in Appendix B.

Benign Scenario. Addressing RQ6, our evaluation shows that Split Alignment achieves the same level of deliverability as a perfectly configured DMARC-passing email, across all email services we evaluated, for both the Organizational Domain and Subdomain as the Visible Domain. The results also confirm that partial alignment of identifiers has no measurable impact on deliverability. This implies that Split Alignment has no measurable effect on deliverability, and email senders can deploy it today to achieve improved security.

Adversarial Scenario. Addressing RQ7, evaluation shows that in the adversarial scenario, Split Alignment provides security against all SPF vulnerabilities by achieving the same lack of deliverability as a DMARC-failing email across all email services we evaluated. Furthermore, our differential analysis reveals cases where Split

Table 8: Evaluation of Security and Deliverability of Split Alignment. Split Alignment provides security and deliverability across the tested email services.

Visible Domain	Deliverability		Security	
	Org.	Sub.	Org.	Sub.
gmail.com	✓ _i	✓ _i	✓ _r	✓ _r
outlook.com	✓ _s	✓ _s	✓ _r	✓ _r
proton.me	✓ _i	✓ _i	▲ _i →✓ _s	✓ _s
yahoo.com	✓ _i	✓ _i	✓ _r	✓ _r
aol.com	✓ _i	✓ _i	✓ _r	✓ _r
zoho.com	✓ _i	✓ _i	✓ _r	✓ _r
naver.com	✓ _i	✓ _i	▲ _i →✓ _b	✓ _b
tuta.com	✓ _i	✓ _i	✓ _s	▲ _i →✓ _s
onet.pl	✓ _i	✓ _i	✓ _r	✓ _r
yandex.com	✓ _s	✓ _s	✓ _w	✓ _w
mailo.com	✓ _i	✓ _i	✓ _r	✓ _r
mail.com	✓ _i	✓ _i	✓ _s	✓ _s
inbox.lv	✓ _i	✓ _i	✓ _s	✓ _s
sapo.pt	✓ _i	✓ _i	✓ _r	✓ _r
seznam.cz	✓ _i	✓ _i	▲ _s →*	✓ _b

✓ = Expected Behavior, ▲ = Confirmed Bug, * = Pending Fix, _i = Inbox, _s = Spam, _w = Spam w/ Warning, _r = Rejected, _b = Blackholed.

Table 9: Status of disclosures.

Email Service	Confirmed	Status	Bounty
proton.me	✓	Fixed	Offered
naver.com	✓	Fixed	Offered
tuta.com	✓	Fixed	-
seznam.cz	✓	Intended	-

Alignment’s security outcomes deviate from expectations due to incorrect DMARC implementation by email services. We disclose our findings to the affected services, who confirm them as bugs, and several have since been remediated (see subsection 7.3).

Outliers. A notable outlier in our evaluation is Yandex Mail. In the adversarial scenario, we observed a slight difference in the deliverability between Split Alignment and DMARC-failing emails. When an email fails DMARC because the authenticated identifier and Visible Domain are entirely different, Yandex rejects it outright. However, when an email fails DMARC due to a lack of alignment under the strict configuration, Yandex delivers it to the Spam folder with a warning. Although the deliverability outcomes of both cases differ, the displayed warning clearly states that “Email failed DMARC verification,” and we therefore consider both functionally equivalent. This case is a result of Yandex’s application of internal policies.

7.3 Disclosures

Our differential analysis reveals bugs in DMARC authentication in four major email services (RQ8). The affected services have confirmed the bugs, deployed patches, and in some cases offered bug bounty rewards. Table 9 summarizes the status of these disclosures.

As a result of our work, Split Alignment is broadly compatible and deployable today.

Proton Mail & Naver. For both Proton Mail and Naver, emails sent using Organizational Domain as the Visible Domain in the adversarial scenario pass DMARC and are delivered to the Inbox. Instead, they should fail DMARC due to a lack of alignment under strict configuration and fail delivery. Both services confirm this as a bug in their DMARC authentication flow and have already deployed fixes, with bug bounties issued in acknowledgment. Following the patch, emails in the adversarial scenario now fail DMARC and exhibit the same lack of deliverability as regular DMARC-failing emails.

Tuta Mail. For Tuta Mail, emails sent using a Subdomain as the Visible Domain are delivered to the Inbox. This occurs in both the DMARC-failing and the adversarial Split Alignment scenario. In the DMARC-failing case, emails correctly fail DMARC, as indicated by the Authentication-Results header, but are incorrectly delivered to the Inbox. Under the adversarial Split Alignment scenario, emails incorrectly pass DMARC and are delivered to the Inbox rather than failing DMARC due to a lack of alignment under the strict configuration. Tuta Mail has confirmed both cases as bugs in its DMARC authentication implementation and has deployed fixes. Following the patch, emails in the adversarial scenario now fail DMARC and are sent to Spam.

Seznam. For Seznam, similar to Proton Mail and Naver, emails with the Organizational Domain as the Visible Domain sent under the adversarial scenario pass DMARC and are delivered to Spam rather than failing DMARC and being blackholed. We disclose this behavior to Seznam, who confirm the issue but describe it as an expected outcome resulting from their internal email-handling policies. Specifically, they state that this is a “conscious choice aimed at improving overall deliverability,” even though the behavior does not comply with the DMARC specification.

Takeaways. Using our methodology, we confirm that Split Alignment provides security against organization-wide SPF failure without compromising deliverability. Furthermore, we uncover critical bugs in implementations of DMARC authentication for four major email services, which we responsibly disclose. Our disclosures suggest that the email ecosystem lacks a consistent way to ensure the validity of authentication implementations, an issue that our methodology directly addresses. To this end, we release our testing harness as an open-source artifact,¹ both to support future academic research and to provide email providers with a practical tool to audit their DMARC implementations.

8 Related Works

Email security broadly consists of approaches that provide integrity (SPF, DKIM, DMARC, BIMI) and confidentiality (STARTTLS, DANE, MTA-STS, OpenPGP, S/MIME), which have been extensively studied from sender and receiver perspectives. Our paper focuses on email integrity, particularly DMARC authentication and its dependencies.

¹<https://doi.org/10.5281/zenodo.19562165>

Work on DMARC can be taxonomized into studies of the protocol’s deployment, reporting mechanism, and authentication functionality. Studies focusing on deployment have mainly investigated usage and support trends. Durumeric et al. [14] and Tatang et al. [32] perform an analysis of the usage of DMARC, finding high prevalence of permissive policies ($p=none$), while Deccio et al. [13] examine DMARC support on the receiver side, finding broad provider support. Complementing these, Hureau et al. [17] present a comprehensive measurement study of all DMARC policy tags, finding that the usage of alignment tags remains low, something that aligns with our analysis.

In the area of DMARC reporting, Hureau et al. [18] find that most email services do not fully comply with the reporting specification. Consistent with this, Ashiq et al. [3] exploit this to demonstrate a practical denial-of-service attack.

Prior research on DMARC authentication is limited, largely because its functionality depends on SPF and DKIM, making it necessary to consider the full protocol stack to understand its behavior. Maroofi et al. [26] provide an early analysis that considers SPF and DMARC policies together to study email spoofing. From the receiver perspective, Blechschmidt et al. [8] investigate various DMARC and SPF sending configurations, finding that the deliverability outcomes vary widely across major email services.

Our work complements these studies by examining both SPF and DKIM through the lens of DMARC, considering both sender and receiver perspectives. These protocols are interconnected through DMARC’s alignment mechanism, a gap in prior study that forms the core focus of our work. Building on our analysis, we further propose a practical mitigation to enhance the overall security of the email space.

In the broader context of email integrity protocols, many works [4, 6, 12, 33] focus on SPF, examining its adoption, misconfigurations, and vulnerabilities. Work on DKIM follows a similar direction, with Wang et al. [34] performing a large-scale study of DKIM records, finding widespread use of shared and weak keys. Yajima et al. [36] extend this analysis to BIMI, providing an examination of its adoption and use in practice. In addition, Chen et al. [10] and Shen et al. [29] shift the focus to the receiver side, investigating parsing and UI-related attacks that enable email spoofing.

9 Discussion and Conclusion

The security of the email ecosystem has been consistently undermined by a combination of protocol weaknesses and misconfigurations. We find that SPF vulnerabilities are widespread, and are further compounded by DMARC’s relaxed alignment mode, which is enabled by default. Both of these issues can be traced back to the assumptions under which these protocols were developed, which no longer hold in the modern internet. As demonstrated in our large-scale measurement study, these vulnerabilities can be readily exploited to enable organization-wide spoofing.

SPF’s reliance on IP ownership as a security principle does not hold in the context of IP leasing, making it unsuitable for use. In the case of DMARC, relaxed alignment being the default eased deployment. But when paired with SPF, DMARC exacerbates its weaknesses. Even though Split Alignment provides a practical mitigation, the email ecosystem needs to move towards a new protocol

that addresses these weaknesses, unifies multiple protocols into one, and allows future extensions.

A key finding that permeates our study is the ecosystem's emphasis on deliverability. This was visible across specification (relaxed alignment), provider behavior (treatment of disposition in policy tags), and provider responses to disclosure (e.g., seznam.cz explicitly deviating from the DMARC specification). This prioritization of email deliverability is insufficiently addressed in current protocols, to the detriment of ecosystem integrity. A future solution that addresses deliverability and security in tandem could leverage this exact prioritization for rapid adoption while improving on the oversights of previous protocols.

Until then, Split Alignment offers a practical, domain-wide security solution that preserves deliverability without requiring disruptive changes in the form of a new protocol or changes to existing ones.

Acknowledgments

We thank the anonymous reviewers for their helpful feedback. This material is based upon work supported by the National Science Foundation under Grant No. CNS-2339378 and the Commonwealth Cyber Initiative.

References

- [1] Eihal Alowaisheq, Siyuan Tang, Zhihao Wang, Fatemah Alharbi, Xiaojing Liao, and XiaoFeng Wang. 2020. Zombie Awakening: Stealthy Hijacking of Active Domains through DNS Hosting Referral. In *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security (CCS '20)*. Association for Computing Machinery, New York, NY, USA, 1307–1322. <https://doi.org/10.1145/3372297.3417864>
- [2] Amazon. 2025. Creating and Verifying Identities in Amazon SES - Amazon Simple Email Service. <https://docs.aws.amazon.com/ses/latest/dg/creating-identities.html#verify-domain-procedure>.
- [3] Md. Ishtiaq Ashiq, Weitong Li, Tobias Fiebig, and Taejoong Chung. 2023. You've Got Report: Measurement and Security Implications of DMARC Reporting. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, Anaheim, CA, 4123–4137. <https://www.usenix.org/conference/usenixsecurity23/presentation/ashiq>
- [4] Md. Ishtiaq Ashiq, Weitong Li, Tobias Fiebig, and Taejoong Chung. 2024. SPF Beyond the Standard: Management and Operational Challenges in Practice and Practical Recommendations. In *33rd USENIX Security Symposium (USENIX Security 24)*. USENIX Association, Philadelphia, PA, 3081–3098. <https://www.usenix.org/conference/usenixsecurity24/presentation/ashiq>
- [5] badkeys. 2025. badkeys. <https://github.com/badkeys/badkeys>.
- [6] Nathaniel Bennett, Rebekah Sowards, and Casey Deccio. 2022. SPFail: discovering, measuring, and remediating vulnerabilities in email sender validation. In *Proceedings of the 22nd ACM Internet Measurement Conference*. ACM, Nice France, 633–646. <https://doi.org/10.1145/3517745.3561468>
- [7] Seth Blank, Peter Goldstein, Thede Loder, Terry Zink, Marc Bradshaw, and Alex Brotman. 2025. *Brand Indicators for Message Identification (BIMI)*. Internet Draft draft-brand-indicators-for-message-identification-10. Internet Engineering Task Force. <https://datatracker.ietf.org/doc/draft-brand-indicators-for-message-identification/12/>
- [8] Birk Blechschmidt and Ben Stock. 2023. Extended Hell(o): A Comprehensive Large-Scale Study on Email Confidentiality and Integrity Mechanisms in the Wild. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, Anaheim, CA, 4895–4912. <https://www.usenix.org/conference/usenixsecurity23/presentation/blechschmidt>
- [9] Kevin Borgolte, Tobias Fiebig, Shuang Hao, Christopher Kruegel, and Giovanni Vigna. 2018. Cloud Strife: Mitigating the Security Risks of Domain-Validated Certificates. In *Proceedings of the 25th Network and Distributed System Security Symposium (NDSS)* (25 ed.) (San Diego, CA, USA, 2018-02), Patrick Traynor and Alina Oprea (Eds.). Internet Society (ISOC). <https://doi.org/10.14722/ndss.2018.23327>
- [10] Jianjun Chen, Vern Paxson, and Jian Jiang. 2020. Composition Kills: A Case Study of Email Sender Authentication. In *29th USENIX Security Symposium (USENIX Security 20)*. USENIX Association, 2183–2199. <https://www.usenix.org/conference/usenixsecurity20/presentation/chen-jianjun>
- [11] D. Crocker, T. Hansen, and M. Kucherawy. 2011. *DomainKeys Identified Mail (DKIM) Signatures*. Technical Report RFC6376. RFC Editor. RFC6376 pages. <https://doi.org/10.17487/rfc6376>
- [12] Stefan Czybik, Micha Horlboge, and Konrad Rieck. 2023. Lazy Gatekeepers: A Large-Scale Study on SPF Configuration in the Wild. In *Proceedings of the 2023 ACM on Internet Measurement Conference (IMC '23)*. Association for Computing Machinery, New York, NY, USA, 344–355. <https://doi.org/10.1145/3618257.3624827>
- [13] Casey Deccio, Tarun Yadav, Nathaniel Bennett, Alden Hilton, Michael Howe, Tanner Norton, Jacob Rohde, Eunice Tan, and Bradley Taylor. 2021. Measuring email sender validation in the wild. In *Proceedings of the 17th International Conference on Emerging Networking EXperiments and Technologies (CoNEXT '21)*. Association for Computing Machinery, New York, NY, USA, 230–242. <https://doi.org/10.1145/3485983.3494868>
- [14] Zakir Durumeric, David Adrian, Ariana Mirian, James Kasten, Elie Bursztein, Nicolas Lidzorski, Kurt Thomas, Vijay Eranti, Michael Bailey, and J. Alex Halderman. 2015. Neither Snow Nor Rain Nor MITM...: An Empirical Analysis of Email Delivery Security. In *Proceedings of the 2015 Internet Measurement Conference*. ACM, Tokyo Japan, 27–39. <https://doi.org/10.1145/2815675.2815695>
- [15] Google. 2025. Email Sender Guidelines - Google Workspace Admin Help. <https://support.google.com/a/answer/81126?hl=en>.
- [16] Hang Hu and Gang Wang. 2018. End-to-End Measurements of Email Spoofing Attacks. In *27th USENIX Security Symposium (USENIX Security 18)*. USENIX Association, Baltimore, MD, 1095–1112. <https://www.usenix.org/conference/usenixsecurity18/presentation/hu>
- [17] Olivier Hureau, Jan Bayer, Andrzej Duda, and Maciej Korczyński. 2024. Spoofed Emails: An Analysis of the Issues Hindering a Larger Deployment of DMARC. In *Passive and Active Measurement*. Vol. 14537. Springer Nature Switzerland, Cham, 232–261. https://doi.org/10.1007/978-3-031-56249-5_10 Series Title: Lecture Notes in Computer Science.
- [18] Olivier Hureau, Andrzej Duda, and Maciej Korczyński. 2024. Stress Testing the DMARC Reporting System: Compliance with Standards and Ways of Improvement. In *Proceedings of the 20th International Conference on emerging Networking EXperiments and Technologies*. ACM, Los Angeles CA USA, 1–9. <https://doi.org/10.1145/3680121.3697809>
- [19] S. Kitterman. 2014. *Sender Policy Framework (SPF) for Authorizing Use of Domains in Email, Version 1*. Technical Report RFC7208. RFC Editor. RFC7208 pages. <https://doi.org/10.17487/rfc7208>
- [20] J. Klensin. 2008. *Simple Mail Transfer Protocol*. Technical Report RFC5321. RFC Editor. RFC5321 pages. <https://doi.org/10.17487/rfc5321>
- [21] M. Kucherawy and E. Zwicky. 2015. *Domain-based Message Authentication, Reporting, and Conformance (DMARC)*. Technical Report RFC7489. RFC Editor. RFC7489 pages. <https://doi.org/10.17487/rfc7489>
- [22] Ruixuan Li, Chaoyi Lu, Baojun Liu, Yanzhong Lin, Haixin Duan, Qingfeng Pan, and Jun Shao. 2025. Understanding and Characterizing Intermediate Paths of Email Delivery: The Hidden Dependencies. In *Proceedings of the 2025 ACM Internet Measurement Conference (IMC '25)*. ACM, Madison, WI, USA, 1–14. <https://doi.org/10.1145/3730567.3764488>
- [23] Daiping Liu, Shuai Hao, and Haining Wang. 2016. All Your DNS Records Point to Us: Understanding the Security Threats of Dangling DNS Records. In *Proceedings of the 2016 ACM SIGSAC Conference on Computer and Communications Security (CCS '16)*. Association for Computing Machinery, New York, NY, USA, 1414–1425. <https://doi.org/10.1145/2976749.2978387>
- [24] Enze Liu, Gautam Akiwate, Mattijs Jonker, Ariana Mirian, Stefan Savage, and Geoffrey M. Voelker. 2021. Who's Got Your Mail? Characterizing Mail Service Provider Usage. In *Proceedings of the 21st ACM Internet Measurement Conference (IMC '21)*. Association for Computing Machinery, New York, NY, USA, 122–136. <https://doi.org/10.1145/3487552.3487820>
- [25] Mailgun. 2025. Domain Verification. <https://documentation.mailgun.com/docs/mailgun/user-manual/domains/domains-verify>.
- [26] Sourena Maroofi, Maciej Korczyński, Arnold Hölzel, and Andrzej Duda. 2021. Adoption of Email Anti-Spoofing Schemes: A Large Scale Analysis. *IEEE Transactions on Network and Service Management* 18, 3 (Sept. 2021), 3184–3196. <https://doi.org/10.1109/TNSM.2021.3065422>
- [27] Eric Pauley, Paul Barford, and Patrick McDaniel. 2023. DScope: A Cloud-Native Internet Telescope. In *32nd USENIX Security Symposium (USENIX Security 23)*. USENIX Association, Anaheim, CA, 5989–6006. <https://www.usenix.org/conference/usenixsecurity23/presentation/pauley>
- [28] Eric Pauley, Ryan Sheatsley, Blaine Hoak, Quinn Burke, Yohan Beugin, and Patrick McDaniel. 2022. Measuring and Mitigating the Risk of IP Reuse on Public Clouds. In *2022 IEEE Symposium on Security and Privacy (SP)*. IEEE, San Francisco, CA, USA, 558–575. <https://doi.org/10.1109/SP46214.2022.9833784>
- [29] Kaiwen Shen, Chuhan Wang, Minglei Guo, Xiaofeng Zheng, Chaoyi Lu, Baojun Liu, Yuxuan Zhao, Shuang Hao, Haixin Duan, Qingfeng Pan, and Min Yang. 2021. Weak Links in Authentication Chains: A Large-scale Analysis of Email Sender Spoofing Attacks. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 3201–3217. <https://www.usenix.org/conference/usenixsecurity21/presentation/shen-kaiwen>

- [30] Michael A. Specter, Sunoo Park, and Matthew Green. 2021. KeyForge: Non-Attributable Email from Forward-Forgeable Signatures. In *30th USENIX Security Symposium (USENIX Security 21)*. USENIX Association, 1755–1773. <https://www.usenix.org/conference/usenixsecurity21/presentation/specter-keyforge>
- [31] Shopify Staff. 2025. You Should Know These Email Marketing Stats in 2025. <https://www.shopify.com/blog/email-marketing-statistics>
- [32] Dennis Tatang, Florian Zettl, and Thorsten Holz. 2021. The Evolution of DNS-based Email Authentication: Measuring Adoption and Finding Flaws. In *24th International Symposium on Research in Attacks, Intrusions and Defenses (San Sebastian Spain, 2021-10-06)*. ACM, New York, NY, USA, 354–369. <https://doi.org/10.1145/3471621.3471842>
- [33] Chuhan Wang, Yasuhiro Kuranaga, Yihang Wang, Mingming Zhang, Linkai Zheng, Xiang Li, Jianjun Chen, Haixin Duan, Yanzhong Lin, and Qingfeng Pan. 2024. BreakSPF: How Shared Infrastructures Magnify SPF Vulnerabilities Across the Internet. In *Proceedings 2024 Network and Distributed System Security Symposium*. Internet Society, San Diego, CA, USA. <https://doi.org/10.14722/ndss.2024.23113>
- [34] Chuhan Wang, Kaiwen Shen, Minglei Guo, Yuxuan Zhao, Mingming Zhang, Jianjun Chen, Baojun Liu, Xiaofeng Zheng, Haixin Duan, Yanzhong Lin, and Qingfeng Pan. 2022. A Large-scale and Longitudinal Measurement Study of DKIM Deployment. In *31st USENIX Security Symposium (USENIX Security 22)*. USENIX Association, Boston, MA, 1185–1201. <https://www.usenix.org/conference/usenixsecurity22/presentation/wang-chuhan>
- [35] Yahoo. 2025. Sender Best Practices. <https://senders.yahooinc.com/best-practices/>.
- [36] Masanori Yajima, Daiki Chiba, Yoshiro Yoneya, and Tatsuya Mori. 2023. A First Look at Brand Indicators for Message Identification (BIMI). In *Passive and Active Measurement*. Springer Nature Switzerland, Cham, 479–495. https://doi.org/10.1007/978-3-031-28486-1_20

A Ethics

The DNS data used in our measurement study (section 4) was compiled in accordance with institutional policies and data-sharing agreements to protect sensitive data. Through our analysis, we identified vulnerable policies in approximately 81 k organizations (subsection 4.2). Although broad disclosure of discovered vulnerabilities is impractical at this scale, the mitigation technique proposed in this work enables administrators to ensure that they are secure against the discussed attack.

Regarding the impact of the experiments themselves, the only experiment that requires interaction with third parties is during our evaluation of Split Alignment (see section 7). To ensure ethical experimentation, we only send emails to our own controlled accounts from our dedicated servers and ensure that the total volume of sent emails remains minimal so as not to overwhelm the email receivers.

B Detailed Evaluation Results For Split Alignment

Table 10 mentions the detailed deliverability results for the evaluation of Split Alignment.

Table 10: Deliverability outcomes of emails under different sending configurations. (Box indicates confirmed bugs)

Visible Domain Sending Config.	Deliverability				Security			
	Org.		Sub.		Org.		Sub.	
	DMARC Pass	Split Alignment	DMARC Pass	Split Alignment	DMARC Fail	Adversarial Scenario	DMARC Fail	Adversarial Scenario
gmail.com	Inbox	Inbox	Inbox	Inbox	Rejected	Rejected	Rejected	Rejected
outlook.com	Spam	Spam	Spam	Spam	Rejected	Rejected	Rejected	Rejected
proton.me	Inbox	Inbox	Inbox	Inbox	Spam	Inbox	Spam	Spam
yahoo.com	Inbox	Inbox	Inbox	Inbox	Rejected	Rejected	Rejected	Rejected
aol.com	Inbox	Inbox	Inbox	Inbox	Rejected	Rejected	Rejected	Rejected
zoho.com	Inbox	Inbox	Inbox	Inbox	Rejected	Rejected	Rejected	Rejected
naver.com	Inbox	Inbox	Inbox	Inbox	Blackholed	Inbox	Blackholed	Blackholed
tuta.com	Inbox	Inbox	Inbox	Inbox	Spam	Spam	Inbox	Inbox
onet.pl	Inbox	Inbox	Inbox	Inbox	Rejected	Rejected	Rejected	Rejected
yandex.com	Spam	Spam	Spam	Spam	Rejected	Warning	Rejected	Warning
mailo.com	Inbox	Inbox	Inbox	Inbox	Rejected	Rejected	Rejected	Rejected
mail.com	Inbox	Inbox	Inbox	Inbox	Spam	Spam	Spam	Spam
inbox.lv	Inbox	Inbox	Inbox	Inbox	Spam	Spam	Spam	Spam
sapo.pt	Inbox	Inbox	Inbox	Inbox	Rejected	Rejected	Rejected	Rejected
seznam.cz	Inbox	Inbox	Inbox	Inbox	Blackholed	Spam	Blackholed	Blackholed