On the Role of Forward-Confirmed reverse DNS in E-mail Authentication

Alfred Arouna
SimulaMet & OsloMet
alfred@simula.no

Rick Fontein
University of Twente
r.l.h.fontein@utwente.nl

Bernd Meijerink
University of Twente
bernd.meijerink@utwente.nl

Ioana Livadariu
SimulaMet
ioana@simula.no

Mattijs Jonker
University of Twente
m.jonker@utwente.nl

***Abstract*—Forward-confirmed reverse DNS (FCrDNS) (also known as *iprev*), considered as client authenticity validation, is an early anti-spam effort to thwart spammers. Over time, anti-forgery techniques evolved with the advent of standards such as the Sender Policy Framework (SPF). SPF includes a `ptr` mechanism that encompasses an FCrDNS check. In this paper, we characterize the deployment of FCrDNS on prospective senders in practice, study to which extent *iprev* and SPF `ptr` checks are performed by large email providers on the receiving side, and how (non)compliance with FCrDNS affects email delivery. Our results show that two-thirds of the prospective senders with `PTR` record are FCrDNS-configured and that not all large providers strictly require *iprev* authentication. Our email campaigns show that large email providers exhibit different treatments of the SPF `ptr` mechanism, possibly because the SPF standard lacks clarity.**

***Index Terms*—FCrDNS, email, anti-spam, DNS, PTR**

I. Introduction

The Simple Mail Transfer Protocol (SMTP) [1] is vital for communication. As of 2023, more than 343.7 B emails are exchanged daily, and this number is expected to increase to 408.2 B by 2027 [2]. However, as the initial design of SMTP did not account for security, it is vulnerable to attacks that target sender identity and message integrity. Due to extensive efforts to improve email security, spam and unsolicited email is steadily decreasing, from 80% (2011) to below 50% (2023) [3]. Spam has also become part of more sophisticated attacks such as Business Email Compromise (BEC), which caused losses of $2.4 billion in 2022 [4], and bulk phishing experienced in 2022 by 85% of companies worldwide [5]. Reportedly, 91% of cyber attacks begin with a phishing email to an unsuspecting victim [6], which can then lead to data breaches, identity theft, and ransomware [7].

The community has put great effort into addressing spam. The Anti-Spam Recommendations [8], which were published as BCP30 more than 25 years ago, suggested using host names rather than IP addresses to combat spam, noting that it needs to be verified that reverse (i.e., address to name) and forward (i.e., name to address) DNS records match. This is also known as *iprev* or *Forward-Confirmed reverse DNS (FCrDNS)*. The continuing recommendation of the Messaging, Malware and Mobile Anti-Abuse Working Group (M^{3}AAWG) is that sending Mail Transfer Agents (MTAs) meet FCrDNS [9]. Members of the group [10], which include large providers, thus require senders to meet *iprev* [11]–[17].

978-3-903176-74-4 ©2025 European Union

Nowadays, the *Sender Policy Framework (SPF)* [18] and *Domainkeys Identified Mail (DKIM) Signatures* [19] are *de facto* DNS-based email authentication protocols. SPF's `ptr` mechanism requires a FCrDNS-configured sender address, followed by a host name match. The current SPF standard, i.e., RFC7208 (2014) [18], states that the `ptr` mechanism is part of the SPF protocol and thus *must* be supported by SPF implementations. However, the same RFC also states that the `ptr` mechanism *should not* be published in SPF policies.

Several works have studied the deployment of email security techniques. Durumeric et al. [20] were among the first to study email delivery security, including SPF. Recently, Czybik et al. reported that 56.5% of 12 million Top-ranked domains specify an SPF policy, revealing an increase in SPF adoption [21].

To our knowledge, the role of FCrDNS in email authentication has not previously been studied. We aim to address this. We characterize FCrDNS state on prospective senders in the wild. We investigate to which extent *iprev* and SPF `ptr` checks are performed by selected providers upon receipt of mail and we examine *iprev* checks in popular mail software.

We make the following observations:

1) Two-thirds of prospective senders with a `PTR` record are FCrDNS-configured.
2) The two most-used Mail Transfer Agents (MTAs), *Postfix* and *Exim*, enable *iprev* authentication checks by default, but do not act on the result without additional configuration.
3) Our measurement campaign reveals that every provider considered performs *iprev* checks upon receiving messages, but that only some strictly act on the result. Of ten providers, eight delivered messages to the recipient even though our SMTP sender was not FCrDNS-configured.
4) We observe varying treatment of the SPF `ptr` mechanism by providers, in some cases depositing mail from explicitly forbidden senders to the inbox.

Our work expands the picture of email authentication by adding the perspective of `PTR` records. Moreover, our study sheds light on a widespread lack of awareness of the *iprev* procedure and SPF `ptr` mechanism by the research community. Considering our results, we posit that ambivalence regarding the SPF `ptr` mechanism can, in some cases, lead to circumvention of email authentication. Therefore, we recommend the community to make a collective decision on its fate.

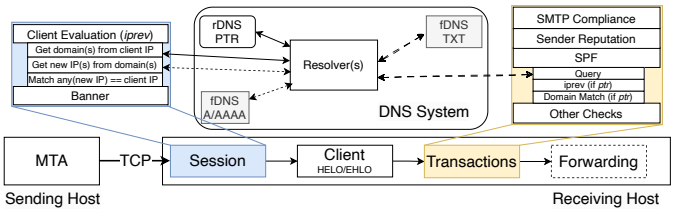

Fig. 1: A high-level breakdown of an SMTP exchange. Receiving MTAs start the SMTP session with *iprev* authentication checks, which is not part of the SMTP standard. SPF is evaluated during the *Transactions* step.

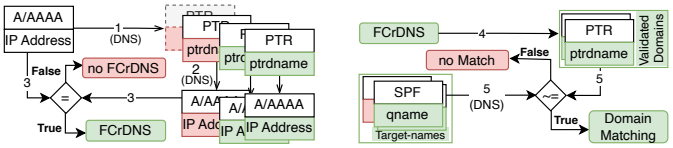

(a) *iprev* (b) SPF `ptr`

Fig. 2: *iprev* authentication and SPF `ptr` validation. SPF `ptr` extends *iprev* with an additional *domain matching* step on the *validated domains* (steps 4 and 5).

II. Background

A. Simple Mail Transfer Protocol (SMTP)

The initial design of the Simple Mail Transfer Protocol (SMTP) [1] focused on transport. Its latest specification is shared among several protocols, such as Extended SMTP (ESMTP) [22] and SMTP Service Extension for Authentication [23]. In a simplified view, after a mail has been submitted, a MTA will transfer the email, which can involve contacting other MTAs. As shown in Figure 1, the sending MTA contacts the receiving MTA and the following stages are activated, in order: 1) *Session Initiation*; 2) *Client Initiation*; 3) *Transactions (authentication, ACLs & restrictions)*; and 4) *Forwarding (routing & transport)* [22]. After accepting a connection from a client (i.e., the sending MTA), an SMTP server (i.e., the receiving MTA), will respond with an SMTP banner. Procedures that are not part of the SMTP standard can be implemented before this banner exchange, as shown by the *iprev* block in Figure 1 (more details will follow in §V). Next, the client sends the `HELO/EHLO` command to the server, containing identity information. Once *Client Initiation* is complete, the client can issue three consecutive commands to create a transaction: `MAIL FROM`, `RCPT TO` and `DATA`. If the message header and body are accepted, the server will forward the message for local delivery (subject to verification) or identify other MTAs to reach out to.[1]

Clients can be rejected during *Session Initiation*. Moreover, messages can be rejected during the *Transactions* and *Forwarding* steps as well. The steps implementations and interdependencies vary from one SMTP implementation server to another (see §V). Modern MTAs have a complex, modular architecture with features that are not strictly related to the SMTP specification but affect delivery and authentication.

B. Email Authentication

SMTP is inherently insecure. RFC5321 states that only end-to-end methods can be considered real email security [22]. Such approaches include *Security Multiparts for MIME* [24]; *OpenPGP* [25]; and *S/MIME* [26]. The adoption of these approaches has been a challenge for the Internet community. Alternative anti-spam methods have been developed over the past two decades, with varying degrees of success. Some approaches have been standardized by the IETF, and others are best practices promoted by the email industry [9].

Examples of standards are the: *Sender Policy Framework (SPF)* [18]; the *DomainKeys Identified Mail (DKIM) signatures* [19]; the *Domain-based Message Authentication, Reporting, and Conformance (DMARC)* [27]; and the *Authenticated Received Chain (ARC) Protocol* [28]. Examples of best practices include: statistical content filtering; rule-based ACLs; rate limiting; graylisting; header and/or body checks; DNS Blocklists and Allowlist [29], [30]; and so-called *iprev* authentication.

C. The iprev Authentication Method

FCrDNS is a networking configuration in which the reverse DNS `PTR` record of a given IP address points to a domain name (i.e., host name) for which, in turn, exists a forward DNS address record that matches the given IP address. Figure 2a shows the process of matching DNS records. If a `PTR` record does not exist, the outcome is also negative.

The notion of matching forward and reverse records was first given in an informational RFC on *Common DNS Operational and Configuration Errors* (1996) [31] and later in *Anti-Spam Recommendations for SMTP MTAs* (BCP30) as prerequisite for hostname-based whitelisting (1999) [8]. The *Message Header Field for Indicating Message Authentication Status* (RFC5451) next proposed *iprev* authentication, which involves an FCrDNS check [32]. As anti-spam technique, *iprev* requires email senders to have matching forward and reverse records. A stricter *iprev* requires a single `PTR` record. This requirement is promoted by the Messaging, Malware and Mobile Anti-Abuse Working Group (M^{3}AAWG) [9].

Large email providers publish requirements for senders. At times, these requirements are open to interpretation and can be understood to be less stringent, requiring, e.g., a `PTR` record to be configured in the DNS, but not a forward match per se [11]–[17]. This prompts questions about how important FCrDNS configuration is. Also, as *iprev* is not part of SMTP, it is unclear at what stage in the SMTP exchange *iprev* authentication is performed in practice. To the best of our knowledge, there is no public document that describes the interplay between SMTP and *iprev*. For this reason, as part of our study, we investigate whether and when *iprev* authentication is performed. In this paper, we use *iprev* to refer either to the process of checking for matching records or to the result of such a check (i.e., the matching state).

[1]The number of MTAs depends on the recipient domains and use of relays.

D. SPF and ptr mechanism

SPF is the *de facto* email authentication standard [18], [33]. It allows receivers to check if a sender – identified by its IP address – is authorized to send mail on behalf of a sender-provided domain name.

SPF policies are published in DNS TXT records. Policies contain SPF mechanisms, modifiers, and qualifiers to specify authorized senders. The ip4 mechanism allows domain operators to specify authorized IP address ranges. Other mechanisms indirectly lead to authorized addresses (e.g., via additional DNS MX and A record lookups).

The include mechanism and the redirect modifier allow policies to be expanded or resolved via other domain names. An + (or omitted) qualifier in before a mechanism indicates SPF pass (i.e., the address range is authorized). Listing 1 shows an example in which clients with an IP address in the network 10.0.0.4/31 are authorized to send due to the ip4 mechanism, as are clients that share IP address with sender.org due to the a mechanism. The all mechanism matches any other client and explicitly forbids them from sending with the – qualifier.

```
sender.org TXT "v=spf1 ip4:10.0.0.4/31 a -all"
```

Listing 1: Example SPF policy in TXT record.

The SPF ptr mechanism allows for a FCrDNS check. Effectively, this mechanism is an extension of *iprev*. Its evaluation requires two steps of an SPF validator: 1) a FCrDNS check on the client IP address; and 2) comparison with a *target-name*. Figure 2b shows the additional step. Specifically, given a sender IP address and PTR record with PTRDNAME value, PTRDNAME is considered a *validated domain* if an A record for PTRDNAME exists that matches the sender IP address. Multiple PTR records can result in multiple *validated domains*. Next, any one of the *validated domains* need to match *or* be in the zone of the *target-name*. For instance, given a *validated domain* hostA.sender.org and *target-name* sender.org, there is a match. The latest version of the SPF standard (RFC7208 [18]) states that the ptr mechanism must be supported.[2] Therefore, a compliant SPF validator must support the mechanism. However, the standard discourages the use of this mechanism. In this study, we investigate whether select email providers perform ptr checks for SPF evaluation.

SPF evaluation results in one of the following seven outcomes: none, neutral, pass, softfail, fail, TempError or PermError, which can factor into, e.g., *Transaction* decisions by a receiving MTA (see Figure 1). Our test scenarios (more details will follow in §VIII) use expected outcomes to evaluate if an SPF ptr check is performed.

E. DKIM & DMARC

DomainKeys Identified Mail (DKIM) [19] is an SMTP security extension that relies on cryptography to detect message forgery and tampering. The Domain-based Message Authentication, Reporting, and Conformance (DMARC) [27] is a mechanism that relies on the SPF and DKIM to set a policy for authenticating received email. An email can be marked as spam, blocked or delivered. DMARC also offers a mechanism to report information to the sender domain, which we leverage as telemetry in our measurement campaign.

[2]According to RFC2119 (BCP14), the keyword MUST implies an absolute requirement, while SHOULD NOT means that there may be valid reasons why the particular behavior is acceptable or even useful.

III. Related Work

Email is one of the main forms of online communication. Consequently, many studies have focused on evaluating email authentication, deployment and configuration practices. Durumeric *et al.* were among the first to characterize the adoption of STARTTLS, SPF, DKIM, and DMARC, by studying mail exchanger configurations for Alexa Top Million domains, and more than a year's worth of SMTP connections to and from Gmail [20]. They reported that adoption is driven mainly by the top email providers that encrypt and authenticate messages, that almost half of the domains implemented SPF, and that 1.1% specified a DMARC authentication policy. Concurrently, Foster *et al.* [34] leveraged Top email hosts from the 2013 Adobe data breach [35] to evaluate if Top senders support TLS, DNSSEC, SPF and DKIM. They reported that popular providers usually support security extensions such as SPF and DKIM. Three years later Hu *et al.* [36] collected, over the span of a year, three DNS records snapshops for Alexa Top 1M domains, and reported high adoption rates for SPF for popular domains (i.e., 73% of the Top 1k domains use SPF). The authors also conducted end-to-end measurement experiments for 35 popular email providers to understand how email providers handle spoofed emails. The study reported that forged email still reaches user inbox despite email providers' spoofing detection protocol.

Several recent studies reported on SMTP security extensions deployment and configuration practices [21], [37]–[39]. The study published by Wang *et al.* [38] is the first large-scale study of DKIM deployment. Based on passive DNS data and DKIM signatures, the authors reported that 28.1% of the Alexa Top 1M domains have enabled DKIM, while also reporting cases of mismanagement. Maroofi *et al.* [37] relied on two large-scale measurement campaigns that covered 236 M domains (139 countries), to evaluate SPF and DMARC adoption as well as analyze how the lack of these mechanisms or misconfiguration enables spoofing. The study finds 65.9% (34.3%) SPF (DMARC) adoption rate for the Top 500 domains. The work publised by Czybik *et al.* analyzed 12 M domains for SPF records [21]. While the authors report an increase in the adoption (56.5%), they also find security issues in SPF implementations, such as errors and lax configurations.

Large email providers publish best practices for SMTP senders that aim to improve email security, such as the existence of a hostname [11], [12], [14]–[17]. Although the standing M^{3}AAWG recommendation is to configure FCrDNS, only Fastmail unequivocally recommends this [16].

Difference to our study – Although SPF has been the subject of various studies on adoption, configuration and security

practices, SPF's `ptr` mechanism is either considered as *not recommended* [21] or simply ignored [37]–[39]. Moreover, while existing papers have studied various email authentication methods, we are unaware of any that studied *iprev* authentication. It is not part of the SMTP standard and while the SPF `ptr` mechanism can be used to perform *iprev* checks, the SPF standard discourages publishing this mechanism in SPF records. What has not previously been investigated is support for and evaluation of the `ptr` mechanism. Our work aims to shed light on this problem by analyzing if and when *iprev* checks are performed by large email providers. Moreover, our study expands the picture of email authentication by adding the perspective of `PTR` records. Over time, in addition to *iprev*, layers of security approaches have been added to SMTP software. Our work highlights that intended email authentication expressed through the SPF `ptr` mechanism can potentially be circumvented.

IV. Senders and *iprev* in the Wild

Given questions about the importance of *iprev*, we start by analyzing FCrDNS configuration on prospective senders. To this end, we use DNS data to infer a sizeable set of prospective senders and investigate if they are FCrDNS-configured. The data were provided to us by the OpenINTEL project [40], which performs active DNS measurements for research. Its *forward* DNS measurement is seeded, among others, by zone files and covers a significant part of the global DNS namespace. Its *reverse* DNS measurement covers the publicly routable IPv4 address space. We use data from February 26, 2025.

A. Prospective Senders

As SPF policies are used to authorize IP addresses to send mail on behalf of domains (see §II-D), our intuition was to infer prospective senders addresses from SPF records. We thus extracted SPF policies from OpenINTEL data. Out of the 298.1 M registered domains in the data, 125.9 M have an SPF policy. We extracted 12.1 M unique policies. Note that there can be overlap when mail is outsourced (we will look at often-used senders in a later section).

We extracted authorized IP address space from SPF data by processing positively qualified mechanism specifications (i.e., authorized senders). We expanded `include` and `redirect` chains up to a depth of 10, in consideration of the SPF DNS lookup limit (see RFC7208, section 4.6.4). Processing SPF `a` and `mx` specifications involved additional DNS lookups.[3]

We extracted IP networks of all sizes, from `/32` to `/0`. We calculated the distribution of the *widest* prefix per domain and found that *0.1%* of all domains authorize a prefix wider than `/8`. We shed the 1 k prefixes involved. This leaves us with 5.9 M unique IP networks which tally 2793 M prospective sender MTA IP addresses. While considering almost half of the allocated IPv4 space as prospective sender may seem disproportionate, previous studies have shown that lax SPF records are common [21]. Our approach to finding senders is best effort and we inherit imprecision from SPF policies.

B. Sender DNS Configuration

We use `PTR` and `A` records from the OpenINTEL data to check if prospective sender addresses are FCrDNS-configured.

Reverse DNS – 1027 M prospective senders have a `PTR` record. This means that, in terms of, e.g., the Microsoft guidelines [14], which suggest that senders merely need a `PTR` record rather than be FCrDNS-configured,[4] 36.7% of 2793 M prospective senders are sufficiently DNS configured.

We extracted 975 M unique `PTRDNAME` values from the `PTR` records. Multiple `PTR` records exist in the DNS for 4.13 M of these senders addresses (0.40% of 1027 M). Depending on the stringency of requirements (recall from §II-C that a stricter *iprev* requirement involves a single record), these sender IP addresses may not meet the *iprev* requirements imposed by receivers and potentially can not send email.

FCrDNS per Sender Address – For each of the 1027 M prospective sender IP addresses with a `PTR` record, we checked if the address is FCrDNS-configured, as outlined in §II-C. This holds for 668.8 M (65.1%), meaning two-thirds. As this is only 24% of 2793 M, a large number of prospective senders do not meet *iprev* requirements.

Matching Records – We found forward address record for 671.3 M `PTRDNAME` values and these `A` records contain 669.0 M unique addresses. For 99.9% `PTRDNAMEs` there exists exactly one `A` record. 99.6% of the 668.8 M FCrDNS-configured senders have a single `PTR` record and therefore can also meet more stringent *iprev* requirements. We also observe common `A` record values among 0.05% of the 671.3 M `PTRDNAMEs`. The Top-3 most commonly shared addresses appear for 1.3 k `PTRDNAMEs`. We traced these to organization that share "in-house" infrastructure between domains related to multiple products (e.g., mobile games) as well as a non-profit that offers shared services for higher education and research. We further consider often-used senders in §IV-C.

Nonmatching Records – We also investigated prospective senders with nonmatching records. In other words, senders with a `PTR` record where forward resolution of the `PTRDNAMEs` does not match the sender's address. This concerns 36.7 M of prospective senders. In 4% of these cases, the `A` records are covered by the IP network in SPF from which we inferred the prospective sender. We cannot exactly know the underlying reasons, but address adjacency suggests a minor DNS configuration misalignment.

FCrDNS per Sender Network – Not all prospective senders are actual senders. Some of the large IP networks that we extracted from SPF can be lax or misconfigured. For example, it is unlikely that all 2^{24} addresses of a `/8` are actually used to send mail on behalf of a given domain. Therefore, we analyzed FCrDNS state per IP network. In other words, for each of the 5.9 M networks we checked if at least one address

[3]For example, `MX` record resolution provides a mail exchanger name which, in turn, is resolved to one or more addresses.

[4]The guideline reads: "Email servers must have valid reverse DNS records."

is FCrDNS-configured. This criterion is met by 2.6 M (42.4%) of all 5.9 M networks. Focusing on the /32s, we find that 2.43 M are FCrDNS-configured which account for 58% of the overall /32 addresses. As the network size goes up, the ratio tends to go up as the number of addresses increases.

Figure 3 provides a different view and shows the average percentage of FCrDNS-configured addresses in sender networks of a given size, as well as the variance of this percentage. Here, the average tends to decrease, as larger networks contain a smaller fraction of FCrDNS-configured addresses. We see a brief upward trend for larger networks (e.g., /15). We investigated the "high-percentage" networks in this range and linked many to ISPs and cloud providers, which configure FCrDNS on all of their address space. Authorizing such large networks can be considered lax SPF configuration that increases the attack surface for phishing [21].

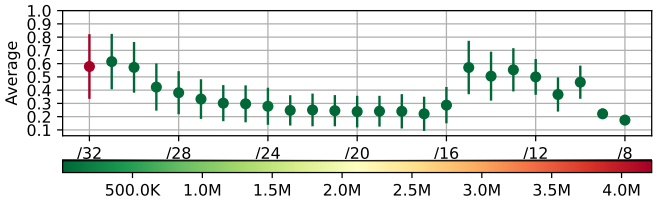

Fig. 3: Average percentage of FCrDNS-configured hosts per sender network of a given prefix length.

Sender Domains and FCrDNS – Recall from §II-C that an IP address can match on any host name (i.e., PTRDNAME value) to be FCrDNS-configured. The M^{3}AAWG sender BCP states that the *primary name* for an address should identify the party responsible for the host. Note that this is in line with the SPF ptr evaluation (see §II-D). Considering this, we checked for all FCrDNS-configured prospective sender addresses if the matching host name is part of the zone of the domain specifying the SPF policy. This is only the case for a marginal number of the FCrDNS-configured addresses, suggesting that the party responsible for the authorized sender is different from that of the domain. This aligns with our expectation that sending is largely outsourced to third parties [41] which comes with authentication-related challenges [42]–[44].

C. Frequent Senders and FCrDNS

SPF policies of different domains can overlap (as shown in §IV-A) and hence have authorized senders in common. We linked commonly authorized senders to providers such as *Google* and *Microsoft* (both M^{3}AAWG members). In light of existing works on email centralization trends, we expected these results [41]. Further analysis revealed that 96.8% of the addresses of frequent senders are FCrDNS-configured, while this applies to only 24% of all prospective sender addresses. Moreover, 93.9% involve a single PTR record, meaning most frequent senders can meet stricter *iprev* requirements, confirming adherence to M^{3}AAWG requirements by large providers. We believe that these figures are positive and show that outsourcing the sending of email to third-parties facilitates compliance with FCrDNS anti-spam requirements.

Takeaway – We presented an approach to infer prospective senders from DNS data at scale, characterized sender DNS configuration, and found that a large number will fail *iprev* checks, should receivers perform them. However, for frequently-used senders (i.e., those authorized to send mail on behalf of many domain names) the situation is much better, as nearly 97% are FCrDNS-configured.

V. *iprev* Support in MTA Software

We now turn to the receiver side and examine to which extent MTA software supports *iprev* authentication. We examine implementations that: 1) are open source, so we can perform code review; 2) are actively maintained; and 3) have a large market share. *SecuritySpace* crawls hundreds of thousands of mail servers monthly to provide a comprehensive report on MTA use. According to a recent report, 55% and 39% of mail servers worldwide use *Exim* or *Postfix* respectively [45] – both are actively being maintained and open source [46], [47].

Static Code Analysis – Our static code analysis revealed that both *Exim* and *Postfix* perform *iprev*. We provide details in Appendix B. Both implementations make *iprev* calls during the *Session Initiation* step, before banner exchange (see Figure 1). Thus, *iprev*-related DNS queries are potentially sent shortly after the connection with the sender is established.

Dynamic Lab Analysis – We setup two VMs with a recent Ubuntu release and installed packaged versions of *Exim* and *Postfix*. We ran the servers in debug mode to log DNS queries and left other settings as default. [5]We used telnet to emulate an SMTP client and, upon connection, both *Exim* and *Postfix* send reverse and forward DNS lookups to verify *iprev* on the sender's address. Although *iprev*-related DNS lookups are performed and evaluated by default, the result is not enforced under default settings. We configured enforcement (see Appendix D) and confirmed clients failing to pass *iprev* are disconnected before banner exchange.

Takeaway – We applied static code analysis to identify the support for *iprev* checks in widely used and open-source MTA software and ran dynamic analysis in a lab to further confirm our findings.

VI. Large Providers & Anti-Spam Mechanisms

We move beyond lab experiments and investigate *iprev* authentication by large email providers on the Internet. These providers cannot necessarily be expected to exhibit similar behaviors, as they may use complex, proprietary MTAs software [41]. Consequently, we instrumented a measurement campaign that involves sending mail from senders and sender domains for which we control forward and reverse DNS.

[5]We installed Exim 4.95 and Postfix 3.6.4 on Ubuntu 22.04.4 LTS.

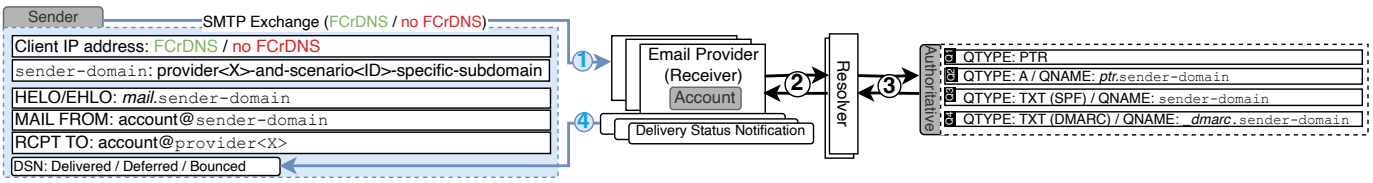

Fig. 4: Measurement scenarios. We use rDNS and fDNS configuration to steer the *iprev* and – where applicable – expected SPF outcomes. Gray components are under our control. Steps in blue and black, respectively, indicate SMTP and DNS exchanges.

A. Provider Selection

We target select large providers to study receiver-side behavior on the Internet. We scoped our investigation to ten because the experiments require provider-specific IP addresses and accounts. To make our selection, we took into account previous work that identifies mail providers that play an increasingly dominant role in email provisioning [36], [41]. Our selection is comprised of: *six* business e-mail providers (*Zoho*, *Mail.ru*, *AOL*, *Gmail*, *Outlook* and *Yahoo*);[6] *three* security and privacy-oriented providers (*Proton*, *Tutamail* and *Fastmail*); and *one* disposable mail provider (*Firefox Relay* {*FRelay*}). We created accounts with all providers, providing us with inboxes. *FRelay* works on the basis of mail forwarding, and we configured our account to forward mail to a private email address.[7] Although the providers may have different market shares, we treated them equally, consistent with other works [36], [43], [44].

B. Measurement Setup

Senders – We setup two VMs – for FCrDNS and non-FCrDNS – as senders and use *Postfix 3.7.11* for sending mail. Each VM has multiple, clean[8] IPv4 addresses that we use to initiate SMTP client connections. We control reverse DNS for all addresses. We followed *Postfix* configuration and security best practices to support virtual domains and multiple IP addresses. Moreover, we followed the robustness principle and added support for both TLS-encrypted and unsecured SMTP connections by using `STARTTLS`.

Sender Domains – We use a recognizable apex domain for measurement purposes. We coupled each sender with a child zone and deployed the authoritative name servers on a separate host. We use a unique subdomain per measurement scenario, consistent with previous studies [36], [43], [44].

Sender Designation – We designated one sender for scenarios that involve FCrDNS-configured addresses and the other for scenarios that don't. We configured FCrDNS for all IP addresses of the prior sender and did the opposite for the latter.

Measurement Scenarios – Each scenario involves sending a mail from one sender to each provider. We use DNS configuration to steer the *iprev* outcome of a scenario. Scenarios can also involve an SPF policy on the sender subdomain that is tailored towards an expected SPF outcome. We use a unique IP address and subdomain per provider scenario and embed a scenario identifier in the message body. To minimize the potential influence of subdomains on mail dispositioning to *Spam* [39], we allocated one day per scenario, with a 15-min interval [48] between messages to different providers. We ran out measurement in the period from March to May 2024.

Figure 4 visualizes our approach. Given scenario ID and provider X, the designated sender will use the scenario-specific subdomain, `sender-domain`, in the SMTP exchange and to send mail. The receiver may send FCrDNS-related DNS queries for `PTR` or `A` records, and may also perform SPF and DMARC lookups (`TXT`).

C. Telemetry

The gray components in Figure 4 are under our control and we use the following telemetry.

DNS Queries – We run authoritatives with *NSD* and use `dnstap` to capture queries. As we use scenario-specific subdomains and IP addresses, we can confidently link forward and reverse DNS queries to scenarios and providers. We set a low TTL on DNS records and wait 15 minutes between measurements to reduce (unexpected) caching effects and account for TTL violations [48]. This allows us to infer, for a given scenario, if and during which SMTP stage *iprev* and SPF related queries are sent.

Delivery Status Notifications (DSNs) – Our senders log SMTP activity, providing us with server responses received during the *Transactions* step, as well as message DSNs. According to RFC3464 [49], a DSN can indicate `Delivered`, `Failed`, `Delayed`, `Relayed` or `Expanded`. A DSN can include other details such as having failed anti-spam requirements. We use sender-dependent transport on *Postfix* such that each scenario uses a different `(IP,HELO/EHLO)` tuple.

DMARC Reports – We make use of *DMARC* Reporting to ask receivers for feedback on messages that we send. These reports provide, among others, the SPF result, which can complement information learned from DSNs and mail dispositioning. To enable reporting, we configure a `_dmarc` record on scenario-specific subdomains. We include `rua` and `ruf` tags, which specify email addresses to send feedback to. We also specify a DMARC policy of `none`, which means DMARC should not affect mail dispositioning. We can then compare the DMARC report to the expected SPF outcome.

User Account – As we send mail to our accounts with the select providers, we can observe if mail sent under a given scenario was delivered to the inbox or the spam folder. For

[6] *Yahoo* and *AOL* are owned by the same company. They have infrastructure overlap, but, as we will show, sometimes exhibit different behaviors.

[7] *Firefox Relay* offers content-based blocking, which we left disabled.

[8] We checked that none of the IP addresses were on common blocklists.

FRelay we observe if the mail was forwarded instead, for which it provides counters in the interface. Some providers also annotate messages in case anti-spam mechanisms were triggered. We extend prior work [36], [39], [43], [44] approaches to investigate email authentication and message delivery outcome based on controlled inputs (*iprev* and SPF `ptr`) and a combination of observed outputs (DSN, DMARC Report, final email dispositioning).

TABLE I: *iprev* measurement scenarios. Under S_{I} and S_{II} our senders are not FCrDNS-configured and under S_{III} they are.

S_{id}	***iprev* Configuration**	**Explanation**
I	$no\_iprev_{no\_ptr}$	Missing `PTR`.
II	$no\_iprev_{has\_ptr}$	Mismatching `PTR` with forward.
III	$is\_iprev_{has\_ptr}$	Matching `PTR` with forward.

VII. Receiver-Side *iprev* Authentication

We investigated *iprev* authentication by the select providers.

A. Measurement Scenarios

We used the setup described in §VI-B and tested three scenarios per provider. These scenarios are shown in Table I. In the first case (S_{I}) our sender does not have a `PTR` record and hence is not FCrDNS-configured. In the second case (S_{II}) our sender does have a `PTR` record, but we configured the forward to not match, so the sender is not FCrDNS-configured. In the third case (S_{III}) we configured the `PTR` and forward to match and hence the sender is FCrDNS-configured. With respect to Figure 4, none of these scenarios involve SPF or DMARC configuration.

B. Connection-Time iprev Authentication and Enforcement

Queries captured at the authoritatives confirm that all providers performed *iprev*-related lookups after our senders initiated a connection with the receiver MTA. We also observe that none of the providers disconnected senders during the *Client Initiation* step. Thus, the providers allow SMTP transactions to be started by clients that fail *iprev* authentication.[9]

C. iprev Auth and Message Delivery

We analyzed the SMTP logs and DSNs for *iprev*-related server responses during *Transactions*. Table II breaks this down per provider and scenario. For the scenario S_{I} (sender without `PTR`), three providers signalled an issue during the SMTP exchange. *Gmail* and *Zoho* rejected the message during the *Transactions* step. The prior explicitly stated in the DSN that FCrDNS is required, whereas the latter mentioned anti-spam more generically. *Tutamail Deferred* the message due to a missing client host name and an (automatic) later retry by our sender led to a DSN of *Delivered*. The remaining providers proceeded without impediment and indicated *Delivered*. We recorded similar results for scenario S_{II} (forward does not match), except that *Tutamail* did not *Defer* the message this time. As such, our messages were *Delivered* with eight providers under the scenarios where our SMTP senders were not FCrDNS-configured. Under scenario S_{III} (FCrDNS-configured sender), *Zoho* now *Delivered* the message, leaving only *Gmail* to still reject outright with a DSN stating that senders are required to authenticate via SPF or DKIM.

[9]We assume uniform configuration and anti-spam policies within these large providers, but cannot exclude that receiving MTAs could exhibit different treatments. We consider revealing inconsistencies future work.

TABLE II: *iprev* is not always enforced, except by *Zoho* and *Gmail*. Some messages were initially `Deferred` and later `Delivered` (*).

Provider	S_{I}	S_{II}	S_{III}
AOL	`Delivered`	`Delivered`	`Delivered*`
Fastmail	`Delivered`	`Delivered`	`Delivered`
Gmail	`Failed`	`Failed`	`Failed`
Mail.ru	`Delivered`	`Delivered`	`Delivered`
FRelay	`Delivered`	`Delivered`	`Delivered`
Outlook	`Delivered`	`Delivered`	`Delivered`
Proton	`Delivered`	`Delivered`	`Delivered`
Tutamail	`Delivered*`	`Delivered`	`Delivered`
Zoho	`Failed`	`Failed`	`Delivered`
Yahoo	`Delivered`	`Delivered`	`Delivered`

TABLE III: From the recipient perspective and half of the providers, the messages were present but consistently in *Spam* regardless of the *iprev* status of senders. *Absence* for *Zoho* and *Gmail* aligns with the `Failed` DSN.

Provider	S_{I}	S_{II}	S_{III}
AOL	Spam	Spam	Spam
Fastmail	Inbox	Inbox	Inbox
Gmail	Absent	Absent	Absent
Mail.ru	Spam	Spam	Spam
FRelay	Inbox	Inbox	Inbox
Outlook	Spam	Spam	Spam
Proton	Spam	Inbox	Inbox
Tutamail	Absent	Inbox	Inbox
Zoho	Absent	Absent	Inbox
Yahoo	Spam	Spam	Spam

D. The Recipient's Perspective

Table III shows the perspective from our user accounts with the providers. From this perspective, a message can be present in the *Inbox* or *Spam* folder, or can be altogether *Absent*.

The *absence* for *Zoho* and *Gmail* aligns with the failure DSNs. For *AOL*, *Mail.ru*, *Outlook* and *Yahoo*, messages are present but consistently in *Spam*, regardless of *iprev* status. We attribute this to the "new sender domain effect" (i.e., domain not previously seen by provider [43]). *Fastmail* always delivered to *Inbox* and *FRelay* forwarded all messages. *Proton* is special in that it appears to treat scenarios S_{I} (sender without `PTR`) and S_{II} (forward mismatch) differently and delivers messages for only the prior to *Spam*.

Our observations for *Zoho*, *Gmail*, *AOL*, *Mail.ru*, *Outlook* and *Yahoo* align with restrictive policies, which one could arguably expect from business email providers (see §VI-A). *Fastmail* and *Proton* are among the three security and privacy-oriented providers that we selected, and these two providers did not behave as we had initially expected. However, we learned from our provider outreach efforts (see §VIII-E)

that providers have well-considered reasons behind such mail dispositioning decisions.

Takeaway – We are the first to provide empirical evidence that *iprev* checks are performed by large email providers, although the result is not always enforced. As a consequence, even though spammers cannot control the *iprev* outcome on compromised hosts that are unfavorable in this respect, spammers can still use such hosts to deliver forged mail.

VIII. Receiver-Side SPF ptr Evaluation

A. Measurement Scenarios

We expect an RFC-compliant SPF validator to evaluate the `ptr` mechanism. We devised four measurement scenarios to test to which extent providers do this. Each scenario uses a distinct policy that we tailored to expected SPF outcomes, as shown in Table IV. All four scenarios involve a FCrDNS-configured sender. We devised scenarios with and without a `target-name` (see §II-D), which is always validatable, meaning the `ptr` mechanism should match in each scenario. The positively qualified `ptr` mechanisms in scenarios S_A and S_B (+ is the default) lead to an expected SPF outcome `pass`. The negative qualifiers in scenarios S_C and S_D lead to an expected SPF `fail`. We always use a reverse qualifier on the `all` mechanism in attempt to solicit the opposite SPF outcome if the `ptr` mechanism is not supported. However, we do not assume that SPF validators will necessarily "jump over" the `ptr` mechanism in case they do not support it.

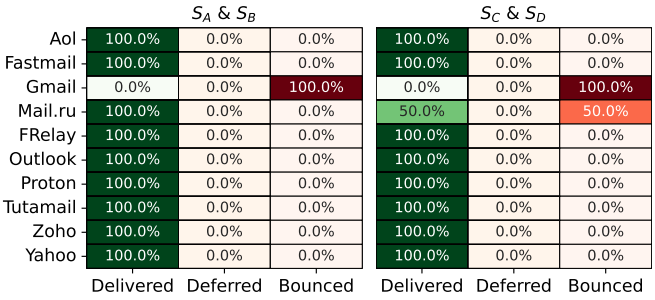

Fig. 5: Delivery status for SPF `ptr` mechanism measurement scenarios S_A and S_B (SPF `pass` expected) and S_C and S_D (SPF `fail` expected).

B. SPF ptr and Message Delivery

Figure 5 shows the delivery status that we recorded for the four measurement scenarios for each provider. *Gmail bounced* all four cases and the DSNs indicate that the senders had failed to authenticate via SPF or DKIM, which indicates that scenarios S_A and S_B did not lead to the expected `pass` outcome. *Mail.ru bounced* scenario S_D but not S_C. The other providers all *Delivered* all four messages (100%). With respect to scenarios S_C and S_D, we note that an SPF `fail` outcome does not mean that a receiver will choose to *bounce* the message. It may also discard it or place it in a recipient's *Spam* folder. We further discuss this when we examine DMARC telemetry and the recipient's perspective.

TABLE IV: Measurement scenarios to test SPF `ptr` mechanism evaluation by providers. The scenario-specific SPF policies are tailored towards expected `pass` and `fail` outcomes if the mechanism is evaluated.

S_{id}	SPF Policy	Expected
A	`v=spf1 ptr -all`	`pass`
B	`v=spf1 ptr:domain-spec -all`	
C	`v=spf1 -ptr +all`	`fail`
D	`v=spf1 -ptr:domain-spec +all`	

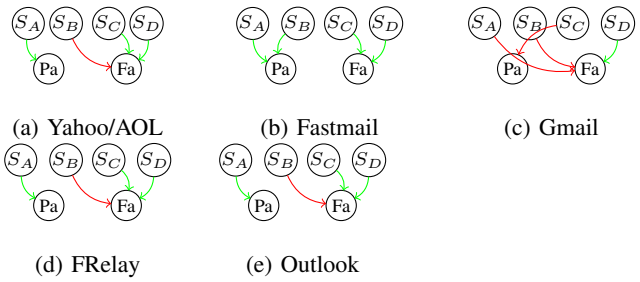

(a) Yahoo/AOL (b) Fastmail (c) Gmail

(d) FRelay (e) Outlook

Fig. 6: Reported SPF outcome. The starting nodes are for the measurement scenarios S_A — S_D, the end nodes for reported Pa (`pass`) and Fa (`fail`) SPF outcomes, and red- and green-coded transitions to indicate alignment with expectation.

C. Expected versus Reported SPF Outcome

We examined DMARC reports to learn how providers processed the four scenario-specific SPF policies. Of particular interest to us is if the expected and reported outcomes align. That is, did providers evaluate and match the `ptr` mechanism, regardless of whether they chose to *bounce* messages from unauthenticated senders outright or not, which *Gmail* did.

Although we defined both `ruf` and `rua` addresses in our DMARC policies, we received aggregate reports only. Given that we used unique subdomains per provider scenario, we can reliably tie reports specific measurements. Figure 6 shows the six providers from which we received reports. We merged *Yahoo* and *AOL* as a result of their infrastructure overlap. The figure contains scenario nodes, `pass` and `fail` outcome nodes, and green arrows for expected and red otherwise.

Fastmail is the only provider that reported all outcomes as expected. *Fastmail* have their own SPF implementation [50] which is officially listed SPF project software [51]. We tested the four policies directly against their implementation and got the same results. The *Yahoo/AOL*, *FRelay* and *Outlook* reports mostly align with our expectations, except for scenario S_B. *Gmail* is a special case. The reported results for scenarios S_A and S_B are `fail`, which aligns with the DSNs and suggests that *Gmail* did not evaluate the `ptr` mechanism. However, scenario S_C was reported as `pass`, even though their MTA *bounced* the messages of all four scenarios.

D. The Recipient's Perspective

Figure 7 shows the recipient's perspective for the four measurement scenarios. *Gmail bounced* all messages, which explains their *Absent* state. *Fastmail* and *Proton* deposited all

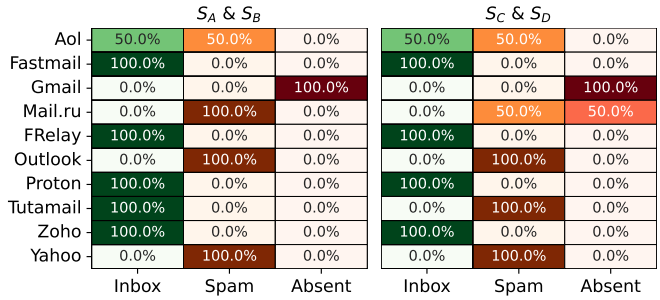

Fig. 7: Recipient perspective for SPF `ptr` mechanism measurement scenarios S_A and S_B (SPF `pass` expected) and S_C and S_D (SPF `fail` expected).

messages to *Inbox*. *Tutamail* annotated messages for which we expected a `fail` result with an SPF failure warning. *Proton* annotated all messages except scenario S_A's with a SPF failure warning, which notably also includes S_B, for which we expected a `pass` outcome. Counter-intuitively, *AOL* delivered the S_A message to *Spam*, despite the reported `pass` result. We attribute this to the "new sender domain effect" (mentioned in §VII-D), which we confirmed by seeing the messages delivered to *Inbox* after repeating the measurement. Despite the unexpectedly reported `fail` result for S_B, *AOL* deposited the message to *Inbox*. As we could not find any SPF warnings in the user interface of *AOL*, we suspect the message was not treated as a regular `fail` case. For measurement scenarios S_C and S_D in which the expected outcome is `fail`, we observe that providers such as *Aol*, *Fastmail*, *FRelay* and *Outlook* delivered messages to *Spam* or *Inbox* while also reporting a `fail` result.

A lesson that runs contrary to what we expected when designing our SPF measurement scenarios is that care must be taken when interpreting DMARC reports. The reported result is not always aligned with the actual treatment of the message.

E. Provider Outreach

We notified the providers that deposited either or both messages for scenarios S_C and S_D in the user's *inbox*. Our rationale here is that, under these scenarios, our senders are explicitly forbidden from sending by SPF. We would argue that messages should then be deposited to *Spam*. This concerns *AOL*, *Fastmail*, *Proton*, *Zoho*. Conversations with *Fastmail* and *Proton* have made clear that they rely on internal authentication and delivery policies. For example, while SPF is expected to `fail`, both providers confirmed that mail is dispositioned to user inboxes in accordance with internal policy.

Takeaway – We unveiled that, with the exception of *Gmail*, most large providers show signs of some degree of support for the SPF `ptr` mechanism. However, this support is only partial, as specifications with an explicit `target-name` lead to unexpected SPF outcomes for several providers. Moreover, our results show that, in cases where domain name operators intend to explicitly forbid senders using the `ptr` mechanism, messages may still be deposited to user inboxes as the result of provider policies.

IX. Discussions & Conclusions

We investigated the role of FCrDNS in e-mail authentication. We inferred prospective sender IP addresses from SPF policies in DNS data to characterize FCrDNS deployment in the wild. We found a quarter of all addresses to be FCrDNS-configured and almost 97% of addresses that often appear in SPF policies, showing that these senders can meet *iprev* authentication requirements if imposed by the receiving side.

Our e-mail campaign involving ten select providers revealed that even when our sender was not FCrDNS-configured, messages were often delivered. Only *Gmail* and *Zoho* outright rejected messages from a sender that fails *iprev* authentication. Our analysis of SPF `ptr` treatment by providers offers similar insights. We explicitly forbade some senders from sending mail on behalf of a domain, but in some cases these messages were deposited in the user inbox, which we learned from our provider outreach efforts was a matter of policy.

We posit that diverse treatment of the SPF `ptr` mechanism relates to different interpretations of the standard and that this may lead to misalignment between intended authorization (as expressed by the domain operator) and treatment by receivers.

We used DMARC Reporting to ask feedback on the SPF validation outcome. We received reports from some and observed discrepancies between the outcome stated in the report, the expected SPF outcome, and message placement at the recipient's side. Our work sheds light on a widespread lack of awareness and/or comprehension of the *iprev* procedure and SPF `ptr` mechanism by the research community. Our work reveals that ambivalence regarding the `ptr` mechanism in the SPF standard can lead to circumvention of email authentication. We thus recommend that the Internet community, including the anti-spam industry, reach a collective decision on the fate of this `ptr` mechanism. We also suggest that email providers that do not yet unequivocally state and implement specific FCrDNS requirements for senders to start doing so. As future work, we plan to further investigate the complex correlation between the observed unseen domain effect, internal policies and mail dispositioning.

Acknowledgments

We thank our shepherd and the anonymous reviewers for their insightful feedback and help in improving our paper. This research is a collaboration between Simula Metropolitan Center for Digital Engineering (SimulaMet) and the University of Twente, and was additionally made possible by the Oslo Metropolitan University (OsloMet). The results in this paper were enabled by OpenINTEL, a joint project of the University of Twente, SURF, SIDN, and NLnet Labs; and operational support from the Twente University Centre for Cybersecurity Research (TUCCR). This research was partially funded by the European Commission under the GN5-2 project (grant agreement ID 101194278).

References

[1] J. Postel, "Rfc0821: Simple mail transfer protocol," 1982.

[2] Statista, "Number of e-mails per day worldwide 2017-2026," 2024. [Online]. Available: https://www.statista.com/statistics/456500/daily-number-of-e-mails-worldwide/

[3] ——, "Spam: share of global email traffic 2011-2023," 2024. [Online]. Available: https://www.statista.com/statistics/420400/spam-email-traffic-share-annual/

[4] FBI, "Business Email Compromise and Real Estate Wire Fraud," FBI 2022 Congressional Report on BEC and Real Estate Wire Fraud, 2022. [Online]. Available: https://www.fbi.gov/file-repository/fy-2022-fbi-congressional-report-business-email-compromise-and-real-estate-wire-fraud-111422.pdf/view

[5] Statista, "Phishing - Statistics & Facts," 2024. [Online]. Available: https://www.statista.com/topics/8385/phishing/

[6] Deloitte, "8 simple practices towards cyber-resilience," 2020. [Online]. Available: https://www2.deloitte.com/my/en/pages/risk/articles/91-percent-of-all-cyber-attacks-begin-with-a-phishing-email-to-an-unexpected-victim.html

[7] FBI, "Internet Crime Report," 2023. [Online]. Available: https://www.ic3.gov/Media/PDF/AnnualReport/2023_IC3Report.pdf

[8] G. Lindberg, "Anti-Spam Recommendations for SMTP MTAs," RFC 2505, Feb. 1999.

[9] M3AAWG, "M3AAWG Sender Best Common Practices," 2015. [Online]. Available: https://www.m3aawg.org/sites/maawg/files/news/M3AAWG_Senders_BCP_Ver3-2015-02.pdf

[10] ——, "M3AAWG Members," 2024. [Online]. Available: https://www.m3aawg.org/about/roster

[11] Yahoo!, "Sender Requirements & Recommendations," 2023. [Online]. Available: 20240323072302/https://senders.yahooinc.com/best-practices/

[12] Google, "Email sender guidelines," 2024. [Online]. Available: https://support.google.com/a/answer/81126?hl=en

[13] AOL, "Anti-Spam Technical Alliance," 2004. [Online]. Available: http://postmaster.aol.com/asta/proposal_html.html

[14] Microsoft, "Policies, Practices, and Guidelines," 2024. [Online]. Available: https://sendersupport.olc.protection.outlook.com/pm/policies.aspx

[15] Zoho Mail, "Reverse IP Lookup," 2023. [Online]. Available: https://www.zoho.com/toolkit/reverse-lookup.html

[16] Fastmail, "Best practices for SMTP servers," 2024. [Online]. Available: https://www.fastmail.help/hc/en-us/articles/1500000278362-Best-practices-for-SMTP-servers

[17] Mail.ru, "Technical and administrative requirements for sending email to Mail.Ru," 2023. [Online]. Available: https://help.mail.ru/enmail-help/rules/technical

[18] S. Kitterman, "Sender Policy Framework (SPF) for Authorizing Use of Domains in Email, Version 1," RFC 7208, Apr. 2014.

[19] M. Kucherawy, D. Crocker, and T. Hansen, "DomainKeys Identified Mail (DKIM) Signatures," RFC 6376, Sep. 2011.

[20] Z. Durumeric, D. Adrian, A. Mirian, J. Kasten, E. Bursztein, N. Lidzborski, K. Thomas, V. Eranti, M. Bailey, and J. A. Halderman, "Neither snow nor rain nor mitm... an empirical analysis of email delivery security," in *Proceedings of the 2015 Internet Measurement Conference*, 2015, pp. 27–39.

[21] S. Czybik, M. Horlboge, and K. Rieck, "Lazy gatekeepers: A large-scale study on spf configuration in the wild," in *Proceedings of the 2023 ACM on Internet Measurement Conference*, 2023, pp. 344–355.

[22] D. J. C. Klensin, "Simple Mail Transfer Protocol," RFC 5321, Oct. 2008.

[23] R. Siemborski and A. Melnikov, "SMTP Service Extension for Authentication," RFC 4954, Jul. 2007.

[24] S. Murphy, J. Galvin, S. Crocker, and N. Freed, "Security Multiparts for MIME: Multipart/Signed and Multipart/Encrypted," RFC 1847, Oct. 1995.

[25] H. Finney, L. Donnerhacke, J. Callas, R. L. Thayer, and D. Shaw, "OpenPGP Message Format," RFC 4880, Nov. 2007.

[26] B. C. Ramsdell, "Secure/Multipurpose Internet Mail Extensions (S/MIME) Version 3.1 Message Specification," RFC 3851, Jul. 2004.

[27] M. Kucherawy and E. Zwicky, "Domain-based Message Authentication, Reporting, and Conformance (DMARC)," RFC 7489, Mar. 2015.

[28] K. Andersen, B. Long, S. Blank, and M. Kucherawy, "The Authenticated Received Chain (ARC) Protocol," RFC 8617, Jul. 2019.

[29] J. R. Levine, "DNS Blacklists and Whitelists," RFC 5782, Feb. 2010.

[30] A. Church, "DNS Blacklists Considered Harmful," Internet Engineering Task Force, Internet-Draft draft-church-dnsbl-harmful-01, Aug. 2005, work in Progress.

[31] D. Barr, "Common DNS Operational and Configuration Errors," RFC 1912, Feb. 1996.

[32] M. Kucherawy and S. Inc., "Message Header Field for Indicating Message Authentication Status," RFC 5451, Apr. 2009.

[33] W. Schlitt and M. W. Wong, "Sender Policy Framework (SPF) for Authorizing Use of Domains in E-Mail, Version 1," RFC 4408, Apr. 2006. [Online]. Available: https://www.rfc-editor.org/info/rfc4408

[34] I. D. Foster, J. Larson, M. Masich, A. C. Snoeren, S. Savage, and K. Levchenko, "Security by any other name: On the effectiveness of provider based email security," in *Proceedings of the 22nd ACM SIGSAC conference on computer and communications security*, 2015, pp. 450–464.

[35] B. Arkin, "Important Customer Security Announcement," 2013. [Online]. Available: https://blog.adobe.com/en/publish/2013/10/03/important-customer-security-announcement

[36] H. Hu and G. Wang, "End-to-End measurements of email spoofing attacks," in *27th USENIX Security Symposium (USENIX Security 18)*. USENIX Association, 2018.

[37] S. Maroofi, M. Korczyński, A. Hölzel, and A. Duda, "Adoption of email anti-spoofing schemes: a large scale analysis," *IEEE Transactions on Network and Service Management*, vol. 18, no. 3, pp. 3184–3196, 2021.

[38] C. Wang, K. Shen, M. Guo, Y. Zhao, M. Zhang, J. Chen, B. Liu, X. Zheng, H. Duan, Y. Lin *et al.*, "A large-scale and longitudinal measurement study of {DKIM} deployment," in *31st USENIX Security Symposium*, 2022.

[39] B. Blechschmidt and B. Stock, "Extended hell(o): A comprehensive Large-Scale study on email confidentiality and integrity mechanisms in the wild," in *32nd USENIX Security Symposium*, 2023.

[40] R. van Rijswijk-Deij, M. Jonker, A. Sperotto, and A. Pras, "A high-performance, scalable infrastructure for large-scale active dns measurements," *IEEE journal on selected areas in communications*, vol. 34, no. 6, pp. 1877–1888, 2016.

[41] E. Liu, G. Akiwate, M. Jonker, A. Mirian, S. Savage, and G. M. Voelker, "Who's got your mail? characterizing mail service provider usage," in *Proceedings of the 21st ACM Internet Measurement Conference*, 2021, pp. 122–136.

[42] J. Chen, V. Paxson, and J. Jiang, "Composition kills: A case study of email sender authentication," in *29th USENIX Security Symposium*, 2020.

[43] E. Liu, G. Akiwate, M. Jonker, A. Mirian, G. Ho, G. M. Voelker, and S. Savage, "Forward pass: On the security implications of email forwarding mechanism and policy," in *2023 IEEE 8th European Symposium on Security and Privacy (EuroS&P)*. IEEE, 2023, pp. 373–391.

[44] F. Holzbauer, J. Ullrich, M. Lindorfer, and T. Fiebig, "Not that simple: Email delivery in the 21st century," in *2022 USENIX Annual Technical Conference (USENIX ATC 22)*, 2022, pp. 295–308.

[45] Securityspace, "Mail (MX) Server Survey," 2025. [Online]. Available: https://www.securityspace.com/s_survey/data/man.202502/mxsurvey.html

[46] Exim, "EXperimental Internet Mailer," 2024. [Online]. Available: https://github.com/Exim/exim

[47] W. Venema, V. Dukhovni, "Postfix mail system," 2024. [Online]. Available: https://github.com/vdukhovni/postfix

[48] P. Bhowmick, M. I. Ashiq, C. Deccio, and T. Chung, "Ttl violation of dns resolvers in the wild," in *International Conference on Passive and Active Network Measurement*. Springer, 2023, pp. 550–563.

[49] G. Vaudreuil and K. Moore, "An Extensible Message Format for Delivery Status Notifications," RFC 3464, Jan. 2003.

[50] FastMail, "Mail::SPF," 2024. [Online]. Available: https://github.com/fastmail/mail-spf

[51] S. Kitterman, "SPF - Implementations," 2024. [Online]. Available: http://www.open-spf.org/Implementations/

[52] Exim, "DKIM, SPF, SRS and DMARC." [Online]. Available: https://www.exim.org/exim-html-current/doc/html/spec_html/ch-dkim_spf_srs_and_dmarc.html

[53] Postfix, "Add-on Software." [Online]. Available: http://www.postfix.org/addon.html

APPENDIX

A. Ethical considerations

Our email campaign involves sending one message per scenario per provider for four scenarios. We send messages to our own accounts only and wait at least 15 minutes between sending messages. Although we knowingly trigger *iprev* and SPF authentication failures, we believe such events to be common to receivers and not to amount to harm. Our prefixes have detailed contact information in the RIR database.

B. iprev Implementations

Exim and Postfix implemented *iprev* in various routines.

Exim Routines – *iprev* is evaluated in *Exim* with the routine `host_name_lookup` [10] called from the `main`[11] function through a cascade of routines. From a simplified point of view, `host_name_lookup` is called by `smtp_start_session` which is called by `main`.

Postfix Routines – *iprev* is implemented in *Postfix* in the function `smtpd_peer_init`.[12] which is called by a chain or routines from the SMTP deamon. From a very simplified point of view, `smtpd_peer_init` is called by `smtpd_state_init` which is called by `smtpd_service`.

C. iprev and SPF

In addition to other restrictions, Exim rely on the parameter `reverse_host_lookup` [13] while Postfix need in its configuration file `reject_unknown_client_hostname` [14] to enforce iprev consistency. Exim delegate SPF evaluation to an external *lipspf2* library [52]. According to Exim sender verification approach, *the combination of two sins is enough* to get an email flagged as spam. Similarly, Postfix rely on a plug-in policy daemon such as the *libspf2* [53] patch.

D. Hardening iprev

Exim and *Postfix* under default configuration send *iprev* authentication related queries for connecting clients (i.e., they check *iprev*) but do not act on the result. Both MTAs provide hardening parameters on client restrictions which are enforced during the SMTP *Transactions* step, mostly after receiving the SMTP client's `RCP TO` command. *Postfix* [15] can reject a client while *Exim* ACL [16] add a warning [17] in the header of the message. According to RFC5321, a client can be rejected during the *Session Initiation* and *Transactions* steps.

Reject Sender on *Exim* – Several configuration parameters have to be enabled in *Exim* configuration to reject senders that fail *iprev* authentication. First, the parameter `host_lookup` has to have its default value `host_lookup = *`. Second, some additional configuration parameters are required. Listing 2 provides an example for *Exim4*.

Reject Sender on *Postfix* – Assuming that parameter `smtpd_peername_lookup` is set to its default value `yes`, Listing 3 provides an example configuration for *Postfix v3*.

```
1  # Add an ACL for the session initiation stage (during
       the TCP connection). This ACL can be in file /etc/
       exim4/conf.d/main/02_exim4-config_options
2  [REDACTED]
3  .ifndef MAIN_ACL_CHECK_CONNECT
4  MAIN_ACL_CHECK_CONNECT = acl_check_connect
5  .endif
6  acl_smtp_connect = MAIN_ACL_CHECK_CONNECT
7  [REDACTED]
8
9  # Then we can create the ACL rules. This rules can be in
        file /etc/exim4/conf.d/acl/20_exim4-
       config_local_deny_exceptions
10 [REDACTED]
11 acl_check_connect:
12   deny
13     condition = ${if and{{def:sender_host_address}{!def:
       sender_host_name}}{yes}{no}}
14
15     message = X-Host-Lookup-Failed: Reverse DNS lookup
       failed for $sender_host_address (${if eq{
       $host_lookup_failed}{1}{failed}{deferred}})
16     log_message = X-Host-Lookup-Failed: Reverse DNS
       lookup failed for $sender_host_address (${if eq{
       $host_lookup_failed}{1}{failed}{deferred}})
17
18   # Otherwise, the sender address is OK.
19   accept
20 [REDACTED]
21
22 # Regenerate Exim configuration
23 $ sudo update-exim4.conf
```

Listing 2: Example of Exim4 configuration to reject clients that fail *iprev* authentication.

```
1 [REDACTED]
2 #By default, Postfix waits until the RCPT TO command
      before evaluating $smtpd_client_restrictions,
      $smtpd_helo_restrictions and
      $smtpd_sender_restrictions, or wait until the ETRN
      command before evaluating
      $smtpd_client_restrictions and
      $smtpd_helo_restrictions.
3
4 smtpd_delay_reject = no
5 smtpd_client_restrictions =
      reject_unknown_client_hostname
6 [REDACTED]
```

Listing 3: Example of Postfix configuration to reject clients that fail *iprev* authentication.

[10] https://github.com/Exim/exim/blob/master/src/src/host.c

[11] https://github.com/Exim/exim/blob/master/src/src/exim.c#L1764

[12] https://github.com/vdukhovni/postfix/blob/master/postfix/src/smtpd/smtpd_peer.c

[13] https://www.exim.org/exim-html-current/doc/html/spec_html/ch-access_control_lists.html

[14] https://www.postfix.org/postconf.5.html#reject_unknown_client_hostname

[15] https://www.postfix.org/postconf.5.html#smtpd_delay_reject

[16] https://www.exim.org/exim-html-current/doc/html/spec_html/ch-access_control_lists.html

[17] https://www.exim.org/exim-html-current/doc/html/spec_html/ch-the_default_configuration_file.html