Neither Snow Nor Rain Nor MITM . . . An Empirical Analysis of Email Delivery Security

Zakir Durumeric[†] David Adrian[†] Ariana Mirian[†] James Kasten[†] Elie Bursztein[‡]
Nicolas Lidzborski[‡] Kurt Thomas[‡] Vijay Eranti[‡] Michael Bailey[§] J. Alex Halderman[†]

[†] University of Michigan [‡] Google, Inc. [§] University of Illinois, Urbana Champaign

{zakir, davadria, amirian, jdkasten, jhalderm}@umich.edu
{elieb, nlidz, kurtthomas, vijaye}@google.com
mdbailey@illinois.edu

ABSTRACT

The SMTP protocol is responsible for carrying some of users' most intimate communication, but like other Internet protocols, authentication and confidentiality were added only as an afterthought. In this work, we present the first report on global adoption rates of SMTP security extensions, including: STARTTLS, SPF, DKIM, and DMARC. We present data from two perspectives: SMTP server configurations for the Alexa Top Million domains, and over a year of SMTP connections to and from Gmail. We find that the top mail providers (e.g., Gmail, Yahoo, and Outlook) all proactively encrypt and authenticate messages. However, these best practices have yet to reach widespread adoption in a long tail of over 700,000 SMTP servers, of which only 35% successfully configure encryption, and 1.1% specify a DMARC authentication policy. This security patchwork — paired with SMTP policies that favor failing open to allow gradual deployment — exposes users to attackers who downgrade TLS connections in favor of cleartext and who falsify MX records to reroute messages. We present evidence of such attacks in the wild, highlighting seven countries where more than 20% of inbound Gmail messages arrive in cleartext due to network attackers.

Keywords

SMTP, Email, Mail, TLS, STARTTLS, DKIM, SPF, DMARC

1. INTRODUCTION

Electronic mail carries some of a user's most sensitive communication, including private correspondence, financial details, and password recovery confirmations that can be used to gain access to other critical resources. Users expect that messages are private and unforgeable. However, as originally conceived, SMTP — the protocol responsible for relaying messages between mail servers — does not authenticate senders or encrypt mail in transit. Instead, servers support these features through protocol extensions such as STARTTLS, SPF, DKIM, and DMARC. The impetus for mail servers to adopt these features is entirely voluntary. As a consequence, gradual rollout has led to a fractured landscape where mail servers must tolerate unprotected communication at the expense of user security. Equally problematic, users face a medium that fails to alert clients when messages traverse an insecure path and that lacks a mechanism to enforce strict transport security.

In this work, we measure the global adoption of SMTP security extensions and the resulting impact on end users. Our study draws from two unique perspectives: longitudinal SMTP connection logs spanning from January 2014 to April 2015 for Gmail, one of the world's largest mail providers; and a snapshot of SMTP server configurations from April 2015 for the Alexa Top Million domains. We use both perspectives to estimate the volume of messages and mail servers that support encryption and authentication, identify mail server configuration pitfalls that weaken security guarantees, and ultimately expose threats introduced by lax security policies that enable wide-scale surveillance and message forgery.

From Gmail's perspective, incoming messages protected by TLS have increased 82% over the last year, peaking at 60% of all inbound mail in April 2015. Outgoing messages similarly grew by 54%, with 80% of messages protected at the conclusion of our study in April. This improvement was largely fueled by a small number of popular web mail providers, including Yahoo and Outlook, enabling security features mid-year. However, such best practices continue to lag for the long tail of 700,000 SMTP servers associated with the Alexa Top Million: only 82% support TLS, of which a mere 35% are properly configured to allow server authentication. We argue that low adoption stems in part from two of the three most popular SMTP software platforms failing to protect messages with TLS by default.

A similar split-picture emerges for the adoption of technologies such as SPF, DKIM, and DMARC that authenticate senders and guard against message spoofing. In terms of sheer volume, during April 2015, Gmail was able to validate 94% of inbound messages using a combination of DKIM (83%) and SPF (92%). However, among the Alexa Top Million mail servers, only 47% deploy SPF policies and only 1% provide a DMARC policy, the absence of which leaves recipients unsure whether an unsigned message is invalid or expected. When mail servers specify SPF policies, 29% are overly broad (covering tens of thousands of addresses.)

This security patchwork — paired with opportunistic encryption that favors failing open and transmitting messages in cleartext, so as to allow incremental adoption — enables network attackers to intercept and surveil mail. In one such attack, network appliances corrupt STARTTLS connection attempts and downgrade messages to non-encrypted channels. We identify 41,405 SMTP servers in 4,714 ASes and 193 countries that cannot protect mail from passive eavesdroppers due to STARTTLS corruption on the network. We analyze the mail sent to Gmail from these hosts and find that in seven countries, more than 20% of all messages are actively prevented

IMC'15, October 28–30, 2015, Tokyo, Japan.
ACM 978-1-4503-3848-6/15/10.
DOI: http://dx.doi.org/10.1145/2815675.2815695.

from being encrypted. In the most severe case, 96% of messages sent from Tunisia to Gmail are downgraded to cleartext.

In the second attack, DNS servers provide fraudulent MX records for the SMTP servers of common mail providers. We searched for DNS servers that provide fraudulent addresses for Gmail's SMTP servers, and we find 14,600 publicly accessible DNS servers in 521 ASes and 69 countries. We investigate the messages that Gmail received from these hosts and find that in 193 countries more than 0.01% of messages from each country are transited through these impostor hosts. In the largest case, 0.08% of messages from Slovakia are relayed from a falsified IP, which can intercept or alter their contents.

Drawing on our measurements, we discuss various challenges and attacks, present current proposals for securing mail transport, and propose directions for future research. We hope that our findings can both motivate and inform further work to improve the state of mail security.

2. BACKGROUND

Simple Mail Transfer Protocol (SMTP) is the Internet standard for sending and relaying email [34, 40]. Figure 1 illustrates a simplified scenario: a client sends mail by transmitting it to its local outgoing SMTP server, which relays each message to the incoming SMTP server for the recipient's domain. In practice, mail forwarding, mailing lists, and other complications result in messages traversing multiple SMTP relays before arriving at their final destination.

As originally conceived in 1981, SMTP did not support protecting the confidentiality of messages in transit or authenticating messages upon receipt. Due to these shortcomings, passive observers can read message content on the wire, and active attackers can additionally alter or spoof messages. To address this gap in security, the mail community developed protocol extensions, such as STARTTLS, DKIM, DMARC, and SPF, to encrypt message content and authenticate senders.

2.1 Protecting Messages in Transit

STARTTLS is an SMTP extension introduced in 2002 that encapsulates SMTP within a TLS session [28]. In a typical STARTTLS session, a client first negotiates an SMTP connection with the server, after which the client sends the command `STARTTLS`, which initiates a standard TLS handshake. The client then transmits mail content, attachments, and any associated metadata over this cryptographically protected channel.

STARTTLS aims to protect the individual hops between SMTP servers, primarily protecting messages from passive eavesdroppers. As we will discuss in Section 4, STARTTLS is typically not used to authenticate destination mail servers, but rather provides opportunistic encryption. (This differs from the behavior of HTTPS clients, which strictly require TLS.) In almost all cases, mail servers do not validate presented certificates and will relay messages over cleartext if STARTTLS is not supported. Because STARTTLS only protects hops between individual relays, each relay still has access to messages and can freely read and modify message content.

The STARTTLS RFC [28] does not define how clients should validate presented certificates. While it suggests that the recipient's domain (e.g., `gmail.com`) should be present in the certificate, it also permits checking the fully qualified domain name (FQDN) of the MX server. This removes the need for third-party mail servers (e.g., shared hosting like Google Apps for Work) to present a trusted certificate for each hosted domain. However, it also enables network-level attackers to falsely report MX records that point to an attacker-controlled domain. Without additional security add-ons (e.g., DANE [14]), this attack remains a real threat.

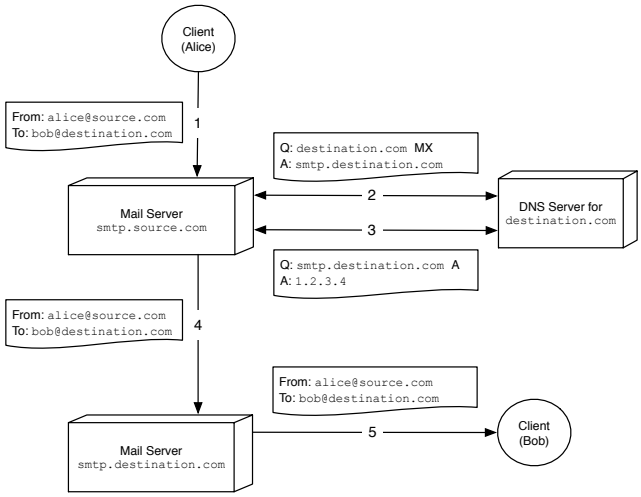

Figure 1: **SMTP Protocol**—A client sends outgoing mail by connecting to its organization's local SMTP server (❶). The local server performs a DNS lookup for the mail exchange (`MX`) record of the *destination.com* domain, which contains the hostname of the destination's SMTP server, in this case *smtp.destination.com* (❷). The sender's server then performs a second DNS lookup for the destination server's IP address (❸), establishes a connection, and relays the message (❹). The recipient can later retrieve the message using a secondary protocol such as POP3 or IMAP (❺).

2.2 Authenticating Mail

Mail servers deploy several complementary mechanisms for authenticating and verifying the integrity of received mail. While STARTTLS protects individual hops between servers, these additional protocols allow recipients to verify that messages have not been spoofed or modified, and they provide a mechanism to report forged messages. A more detailed discussion of each protocol and its limitations is available from MAAWG [12]. We depict the interplay between these mechanisms in Figure 2.

DKIM DomainKeys Identified Mail (DKIM) lets SMTP servers detect whether a received message has been spoofed or modified during transit (RFC 6376 [11]). To utilize DKIM, a sender appends the `DKIM-Signature` field to the message header. This header contains a digital signature of the message signed with the private key tied to the sender's domain. Upon delivery, the recipient can retrieve the sender's public key through a DNS request and verify the message's signature. DKIM does not specify what action the recipient should take if it receives a message with an invalid or missing cryptographic signature. Instead, the organization must have a predetermined agreement with the sender.

SPF Sender Policy Framework (SPF) allows an organization to publish a range of hosts that are authorized to send mail for its domain (RFC 7208 [33]). To deploy SPF, the organization publishes a DNS record that specifies which hosts or CIDR blocks belong to the organization. Upon receiving mail, the recipient performs a DNS query to check for an SPF policy and can choose to reject messages that do not originate from the specified servers. SPF further allows organizations to delegate a portion or the entirety of their SPF policy to another organization, and they commonly delegate SPF settings to a cloud provider (e.g., Google Apps for Work.)

DMARC Domain-based Message Authentication, Reporting, and Conformance (DMARC) builds upon DKIM and SPF and allows senders to suggest a policy for authenticating received mail

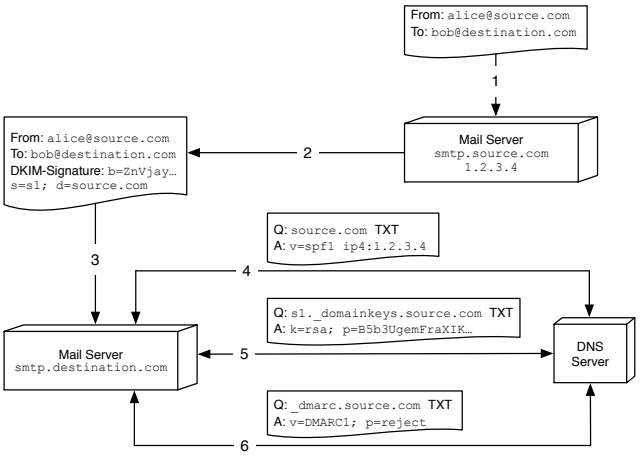

Figure 2: **Mail Authentication**—SPF, DKIM, and DMARC are used to provide source authentication. The outgoing server digitally signs the message (❷). The receiving mail server performs an SPF lookup (❹) to check if the outgoing server is whitelisted, a DKIM lookup (❺) to determine the public key used in the signature, and a DMARC lookup (❻) to determine the correct action should SPF or DKIM validation fail.

(RFC 7489 [35]). Senders publish a DNS TXT record (named `_dmarc.domain.com`) that indicates whether the sender supports mail authentication (i.e., DKIM and/or SPF), and what action recipients should take if authentication fails (e.g., the DKIM signature is missing or invalid). DMARC further allows organizations to request daily reports on spoofed messages that other servers receive.

3. DATASET

Our study is based on two unique datasets: logs of the SMTP handshakes negotiated for mail sent to and from Gmail from January 2014 to April 2015, and a snapshot of SMTP server configurations from April 2015 for the Alexa Top Million domains.

Gmail Inbound and Outbound Messages Google publicly reports statistics about encrypted inbound and outbound messages via its Transparency Report.[1] This dataset explicitly excludes spam messages as not to conflate user security with bulk automated messages. We obtain a companion dataset via a collaboration with Google that contains the set of all ciphers negotiated with external SMTP servers during the same period, as well as any authentication performed on behalf of the sending party. We rely on this dataset for making observations about the volume of mail protected by encryption and authentication. We note that this dataset is noticeably skewed towards a handful of large web mail providers, including Yahoo and Outlook, as well as personal mail accounts provided by local ISPs that relay the bulk of the mail.

Alexa Top Million Mail Servers For a second perspective on organizational support for mail security, we examine the SMTP security features enabled by mail servers belonging to the Alexa Top Million ranked websites [1]. On April 26, 2015, we performed `MX` record lookups for the Alexa Top Million domains. For domains with mail servers, we followed up with a DNS query to identify supported SMTP security extensions (i.e., SPF and DMARC) and attempted an SMTP and STARTTLS handshake using ZMap [16, 19] to identify whether the mail servers support encryption. Our previous work [17, 19] describes our scanning methodology and ethical considerations involved and the practices we use to minimize potential harms.

[1]Data is available at http://www.google.com/transparencyreport/safermail.

Status	Top Million Domains	
No MX records	152,944	(15.29%)
No resolvable MX hostnames	5,447	(0.55%)
No responding SMTP servers	49,125	(4.91%)
SMTP Server	792,494	(79.25%)

Table 1: **Organizational SMTP Deployment**—We investigate how domains in the Alexa Top Million have deployed SMTP.

In total, 792,494 domains (79.2% of the Alexa Top Million) have operational mail servers, as detailed in Table 1. The remaining domains include sites such as `t.co`, `googleusercontent.com`, and `blogspot.com` that do not need incoming mail servers.

4. CONFIDENTIALITY IN PRACTICE

We measure the state of mail confidentiality based on the volume of messages protected by STARTTLS that are sent and received by Gmail, and the fraction of mail servers that support and correctly configure encryption.

4.1 Gmail

As of April 26, 2015, Gmail successfully initiated STARTTLS connections for 80% of outgoing messages, while 60% of incoming connections initiated a STARTTLS session. This represents a 54% increase (52% to 80%) in outbound and an 82% increase (33% to 60%) in inbound connections that utilize STARTTLS since January 2014. While this growth is encouraging, overall gains arrived in bursts rather than consistent growth, as shown in Figure 3. We point out two periods of immediate interest. Between May 10 and 30, 2014, outbound encryption jumped from 47% to 71%. This was likely due to Yahoo and Outlook deploying STARTTLS. Second, between October 8 and 17, outbound STARTTLS dropped from 73% to a low of 50%. The lowest point occurred on October 14, which corresponds with the public disclosure of the POODLE vulnerability [15]. We suspect the correlated drop was a result of mail server misconfigurations introduced by administrators attempting to disable SSLv3.

Influence of Major Mail Providers Major mail providers, such as Gmail, Yahoo, and Outlook, heavily skew the apparent adoption of STARTTLS in contrast to the long tail of organizations that run their own mail servers. Of the 877 most common domains that Gmail transited mail to on April 26, 2015, only 58% accepted 100% of messages over TLS. Similarly, only 29% of 26,406 inbound mail domains encrypted 100% of messages. We explore this skew further in Section 4.2 from the perspective of the Top Million domains.

Along these same lines, we argue that the periodicity present in Figure 3 stems from users sending less business mail on weekends and instead relying on personal accounts provided by major providers. In particular, on weekdays between April 1 and 26, 2015, Gmail encrypted 79.8% of outbound messages, while mail servers encrypted 53.7% of incoming connections. Weekends during this same period saw an average 7.2% increase in the number of secured messages. As we discuss in Section 4.4, major mail providers support encryption by default.

Negotiated Cipher Suites We analyzed the cipher suites chosen by incoming Gmail connections on April 30, 2015, and found that 84.2% of TLS connections (45.2% of all incoming connections) chose a perfect forward secret cipher suite. However, similar to HTTPS, 45.63% of clients continue to prefer RC4 despite its known

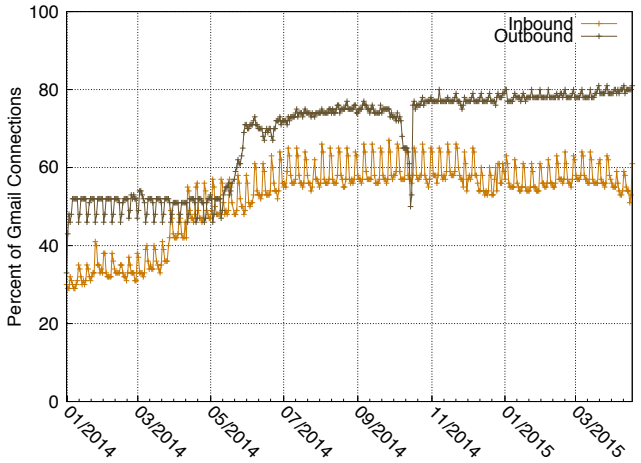

Figure 3: **Historical Gmail STARTTLS Support**—Inbound connections that utilize STARTTLS increased from 33% to 60% for weekdays between January 2014 and April 2015. Weekends consistently have close to 10% more connections that support STARTTLS than weekdays. Support for outgoing STARTTLS increased from 52% to 80% during this period.

weaknesses [2, 45]. For the remaining connections, 95% utilized AES-128-GCM and 5% used AES-128 (Table 2). We note that while this perspective does not show any known-broken ciphers (e.g., EXPORT suites), these may still occur between other mail servers. These do not appear in our dataset because Google does not support these ciphers, and if a client does not support any modern ciphers the TLS handshake will fail.

Comparison With Prior Estimates We compare Gmail's perspective with prior estimates published by Facebook [23]. During May 2014, Facebook—which sends mail notifications for friend requests and new user activity—successfully encrypted 58% of outbound messages. Of the mail servers contacted, 76% supported STARTTLS. By August 2014, Facebook successfully encrypted 95% of outbound notifications after several large webmail providers, notably Microsoft and Yahoo, deployed STARTTLS.

During this same time period, outbound encrypted messages for Gmail increased from 47% to 74%. Despite the same opportunistic STARTTLS policy, we find Gmail generally has a lower percentage of outgoing mail protected by STARTTLS than Facebook. We believe this stems from Facebook primarily communicating with personal mail accounts provided by major providers (e.g., Gmail, Yahoo, and Outlook) as opposed to businesses. We note that our Gmail measurements may have similar but less pronounced biases towards large providers, which we investigate further in the next section.

4.2 Organizational Deployment

Given the skew present in Gmail's message volume towards major mail providers, we provide an alternate perspective by analyzing STARTTLS support for the 792,494 Alexa Top Million domains that advertise mail servers. In total, 648,030 (81.8%) of mail-enabled domains supported STARTTLS, as shown in Table 3. Only 5 domains within the Alexa Top 50 did not: `wikipedia.org`, `vk.com`, `weibo.com`, `yahoo.co.jp`, and `360.cn`. Support for encrypted mail was bolstered in part by 25% of domains outsourcing their mail servers to common third-party providers, such as Gmail, GoDaddy, Yandex, QQ, and OVH, all of which support STARTTLS. We give more details about these providers and their popularity in Table 4.

TLS Version	Key Exchange	Symmetric Cipher	HMAC	Inbound Traffic
TLSv1.2	ECDHE	AES-128-GCM	SHA-256	51.500%
TLSv1	ECDHE	RC4	SHA-1	29.225%
TLSv1	RSA	RC4	SHA-1	14.403%
TLSv1.2	ECDHE	AES-128	SHA-1	1.586%
TLSv1.2	RSA	RC4	SHA-1	1.147%
TLSv1	ECDHE	AES-128	SHA-1	0.999%
TLSv1.1	ECDHE	RC4	SHA-1	0.723%
TLSv1.2	RSA	AES-128-GCM	SHA-256	0.203%
SSLv3	RSA	RC4	SHA-1	0.060%
TLSv1.2	ECDHE	RC4	SHA-1	0.060%
TLSv1	RSA	AES-128	SHA-1	0.050%
TLSv1.1	RSA	RC4	SHA-1	0.024%
TLSv1.1	ECDHE	AES-128	SHA-1	0.011%
TLSv1.1	ECDHE	AES-256	SHA-1	0.004%
TLSv1.2	RSA	AES-256	SHA-1	0.003%
TLSv1.2	RSA	AES-128	SHA-1	0.001%
TLSv1	RSA	RC4	MD5	0.001%

Table 2: **Cipher Suites for Inbound Gmail Traffic**—80% of inbound Gmail connections are protected by TLS. Here, we present the selected cipher suites for April 30, 2015.

Key and Cipher Suites With the exception of two mail servers, all present certificates with RSA keys: 10.0% of domains use 1024-bit keys, 86.4% use 2048-bit keys, and 3% use 4096-bit or larger keys. Only 316 domains present 512-bit RSA certificates (which provide little to no security in 2015 [26]). We found 25.3% of domains support perfect forward secrecy and completed an ephemeral Diffie-Hellman key exchange. In addition, 59.2% of domains use RC4 and 40.8% use AES; only 25 domains select 3DES. In summary, most domains that deploy STARTTLS also deploy secure certificates. However, as in the HTTPS ecosystem, domains are slow in deploying modern, secure cipher suites.

Certificate Validity As part of the STARTTLS handshake, each mail server presents an X.509 certificate. While RFC 3207 [28] suggests that certificates match the mail domain (e.g., `gmail.com`), it also permits certificates that only match the name of the server in the domain's MX record (e.g. `aspmx.l.google.com`). However, certificates that match the MX server do not provide true authentication unless the MX records for the domain are cryptographically signed. Otherwise, an active attacker can return the names of alternate, attacker-controlled MX servers in the initial MX query. In practice, DNSSEC has not been widely deployed—recent studies have found that less than 0.6% of `.com` and `.net` domains have deployed DNSSEC [46]—and so operators cannot rely on this protection.

In our scan, 414,374 Top Million domains (52% of domains with valid SMTP servers, and 64% of domains that support STARTTLS) present certificates that validate against the Mozilla NSS root store [38], as detailed in Table 5. However, only 0.6% of domains present trusted certificates that match their domain name, while 34.2% present trusted certificates that match their MX server.

Surprisingly, 18.1% of domains present trusted certificates that match neither. These are primarily due to several mail hosting providers, including `psmtp.com` and `pphosted.com`, that incorrectly deployed wildcard certificates. In the remaining cases, certificates were simply for different domains. Another 33,281 domains present expired certificates, 60 certificates are signed by unknown CAs, and 55 certificates are invalidly signed by a parent certificate whose type mismatched the child certificate.

In summary, the certificates used by mail servers are in disarray. Less than 35% of mail servers with STARTTLS can be authenticated in any form, and a sender can only confirm that their connection had

Status	Top Million Domains	
SMTP Server — No STARTTLS support	144,464	(18.2%)
SMTP Server — STARTTLS support	648,030	(81.8%)

Table 3: **STARTTLS Deployment by Top Million Domains** — Our scan results show that 79% of Alex Top Million domains have incoming SMTP servers, of which 81.8% support STARTTLS.

Mail Provider	Domains	STARTTLS	Trusted Certificate	Certificate Matches
Gmail	126,419 (15.9%)	Yes	Yes	server
GoDaddy	36,229 (4.6%)	Yes	Yes	server
Yandex	12,326 (1.6%)	Yes	Yes	server
QQ	11,295 (1.4%)	Yes	Yes	server
OVH	8,508 (1.1%)	Yes	Yes	mismatch
Other	597,717 (75.4%)	–	–	–

Table 4: **Top Mail Providers for Alexa Top Million Domains** — Five providers are used for mail transport by 25% of the Top Million domains. All five support STARTTLS for incoming mail.

	Matches Domain	Matches Server	Matches Neither
Trusted	4,602 (0.6%)	270,723 (34.2%)	143,113 (18.1%)
Untrusted	4,345 (0.6%)	21,057 (2.7%)	181,242 (22.9%)
Total	8,947 (1.1%)	291,780 (36.8%)	324,355 (41.0%)

Table 5: **Certificates for Top Million Domains** — While 52% of domains' SMTP servers present trusted certificates, only 34.2% of trusted certificates match the MX server, and only 0.6% are valid for the recipient domain.

not been intercepted by an active attacker for 0.6% of domains. This has likely occurred because, as we discuss in the next two sections, common SMTP implementations and popular mail providers do not validate certificates, so server operators have little incentive to purchase and maintain a certificate.

4.3 Common Software Implementations

In order to understand why such a large number of organizations have not deployed STARTTLS and why only half of inbound connections to Gmail initiate a STARTTLS connection, we investigated the five most popular SMTP implementations, which account for 97% of identifiable mail servers for the Top Million domains. We tested whether each implementation initiated STARTTLS connections, whether it supported STARTTLS for incoming connections, and how it validated certificates. We installed the latest version of each SMTP server on an Ubuntu 14.04.1 LTS system, except for Microsoft Exchange, which was readily documented online [37]. The results are summarized in Table 6.

By default, Microsoft Exchange, Exim, and Sendmail initiate STARTTLS connections when delivering messages. Postfix and qmail — which together account for nearly 35% of all identifiable mail servers on the public IPv4 address space — send all messages over cleartext unless explicitly configured to use STARTTLS. All of the servers we tested fail open and send mail in cleartext if STARTTLS is not available.

Postfix and Microsoft Exchange Server support inbound STARTTLS connections by default by generating a self-signed certificate upon installation. This provides immediate protection against passive attacks without user configuration. The remaining servers do not accept TLS connections without manual configuration. Postfix and Exchange — the two servers that have confidentiality protection enabled by default — account for 22% of servers associated with Top Million domains.

Postfix was the only server capable of performing both server-based and domain-based certificate validation, although its documentation specifically recommends *against* enabling validation when interacting with the greater Internet [41]. Exim, qmail, Sendmail, and Microsoft Exchange do not support validating the destination domain when relaying mail.

4.4 Popular Mail Providers

Since a large fraction of mail is transited through a small number of popular providers, a single change can have a large impact on the entire ecosystem, as previously demonstrated in Figure 3. We measured inbound and outbound STARTTLS support for 19 common webmail providers and Internet service providers. We created an account on each provider and then sent mail to a Postfix server that we configured to support STARTTLS with a self-signed certificate. To test incoming STARTTLS support, we connected to the mail servers listed in each domain's MX record and initiated a STARTTLS handshake.

We summarize our results in Table 7. Only one provider, Lycos, did not support inbound STARTTLS. Two providers — facebookmail and OVH — presented certificates that matched neither their domain nor the hostname of their MX server. *None* of the providers presented a certificate that matched their domain, and thus none could have provided strong authentication in the presence of an active attacker who falsified the service's MX records.

Less than half of the providers negotiated a perfect-forward-secret cipher suites. When sending mail, three of the providers — Lycos, GoDaddy, and OVH — did not initiate STARTTLS connections. The remaining providers initiated STARTTLS connections but did not validate certificates; in effect, this provides opportunistic encryption but no authentication.

4.5 Takeaways

Our results show that there has been significant growth in STARTTLS adoption over the past year. However, much of this growth can be attributed to a handful of large providers. In contrast, as seen in our scans, smaller organizations continue to lag in deploying STARTTLS, and as of March 2015 nearly half of inbound weekday connections remain unencrypted. This may be due, in part, to several popular implementations failing to initiate STARTTLS connections by default.

In the cases where encryption is present, messages are protected opportunistically. Connections fail open to cleartext if any issues arise during the handshake or if STARTTLS is not supported. None of the popular providers and implementations we tested use TLS for authentication, and only one common implementation supports validating a certificate against the destination domain.

Unfortunately, in the protocol's current form, mail providers cannot fail closed in the absence of STARTTLS until there is near total deployment of the extension, and until organizations deploy valid certificates, relays will be unable to automatically authenticate destination servers.

5. THREATS TO CONFIDENTIALITY

As deployed in practice, STARTTLS protects connections against passive eavesdroppers but does not protect against active man-in-the-middle attacks. We examine two types of network attacks that this enables — downgrading STARTTLS sessions to insecure channels and falsifying MX records to re-route messages — and measure the prevalence of both methods in the wild.

Mail Software	Top Million Market Share	Public IPv4 Market Share	STARTTLS Incoming	STARTTLS Outgoing	Server Validation	Domain Validation	Reject Invalid Certificates	TLS Version
exim 4.82	34%	24%	◐	●	○	○	○	1.2
Postfix 2.11.0	18%	21%	●	◐	◐	◐	◐	1.2
qmail 1.06	6%	1%	◐	◐	○	○	○	1.2
sendmail 8.14.4	5%	4%	◐	●	○	○	○	1.2
Exchange 2013	4%	12%	●	●	◐	○	◐	1.0
Other	3%	<1%						
Unknown	30%	38%	● default behavior | ◐ supported but not default | ○ no support					

Table 6: **Popular Mail Transfer Agents (MTA)**—We investigated the default behavior for five popular MTAs. By default, Postfix and qmail do not initiate STARTTLS connections. All five MTAs we tested fail open to cleartext if the STARTTLS connection fails.

Provider	Incoming TLS Version	Incoming Key Exchange	Incoming Cipher	Certificate Matches	Outgoing TLS Version	Outgoing Key Exchange	Outgoing Cipher
Gmail	1.2	ECDHE	AES-128-GCM	server	1.2	ECDHE	AES-128-GCM
Yahoo	1.2	ECDHE	AES-128-GCM	server	1.0	ECDHE	RC4-128
Outlook	1.2	ECDHE	AES-256-CBC	server	1.2	ECHDE	AES-256
iCloud	1.2	ECDHE	AES-128-GCM	server	1.2	DHE	AES-128-GCM
Hushmail	1.2	RSA	RC4-128	server	1.2	ECDHE	AES-256-GCM
Lycos	–	–	–	–	–	–	–
Mail.com	1.2	ECHDE	AES-256-CBC	server	1.2	DHE	AES-256-GCM
Zoho	1.0	RSA	RC4-128	server	1.0	RSA	RC4-128
Mail.ru	1.2	RSA	RC4-128	server	1.2	ECDHE	AES-256-GCM
AOL	1.0	RSA	RC4-128	server	1.0	DHE	AES-256-CBC
QQ	1.1	RSA	RC4-128	server	1.0	DHE	AES-256-CBC
Me.com	1.2	ECHDE	AES-128-GCM	server	1.2	DHE	AES-128-GCM
facebookmail	1.0	RSA	AES-128-CBC	mismatch	1.0	ECDHE	AES-128
GoDaddy	1.2	RSA	RC4-128	server	–	–	–
Yandex	1.2	RSA	AES-128-GCM	server	1.2	ECDHE	AES-256-CBC
OVH	1.2	RSA	AES-128-GCM	mismatch	–	–	–
Comcast	1.2	RSA	RC4-128	server	1.2	DHE	AES-128-CBC
AT&T	1.2	ECDHE	AES-128-GCM	server	1.0	ECDHE	RC4-128
Verizon	1.2	RSA	AES-128-GCM	server	1.0	DHE	AES-128-CBC

Table 7: **Encryption Behavior of Mail Providers**—We measured support for incoming and outgoing STARTTLS among various popular mail providers. While most providers supported STARTTLS, *none of them* validated our certificate, which was self-signed.

Provider	Servers Providing Invalid MX Answers	Servers Providing Invalid IP Answers	Unique Invalid MX Servers	Unique Invalid IPs	Responsive Invalid Mail Servers
Gmail	30,931	23,134	146	1,150	144
Yahoo	31,219	55,459	130	1,117	114
Outlook.com	29,618	23,145	117	1,059	110
Mail.ru	31,214	25,796	97	1,053	110
QQ	30,091	55,467	122	1,171	111

Table 8: **Fraudulent DNS Responses**—We scanned the public IPv4 address space for DNS servers that returned falsified MX records or SMTP server IP addresses for five popular mail providers. This data excludes loopback addresses and obvious configuration errors.

	Nov. 2013	Apr. 2015	Change
Overall failure rate	10.65%	6.14%	−4.42%
Crypto failures:			
Weak crypto key (<1024 bits)	21.00%	15.08%	−5.92%
Key is revoked	0.02%	0.01%	−0.01%
Signature algorithm not supported	0.27%	0.26%	−0.02%
Key is expired		0.06%	
Body hash doesn't match signature		18.66%	
Protocol version incorrect	0.59%	3.32%	+2.73%
Some DKIM tags are duplicated		0.05%	
Other error	77.91%	62.55%	−15.36%

Table 9: **Gmail DKIM Errors**—We present the breakdown of Gmail DKIM validation failures for November 2013 and April 2015.

Scan Result	IPv4 Hosts
TCP port 25 open	14,131,936
Responsive SMTP server	8,850,664
Successful STARTTLS handshake	4,620,561

Table 10: **IPv4 SMTP Scan Results** — We could perform a STARTTLS handshake with 52% of the SMTP servers that our IPv4 scans identified.

Category	IPv4 Hosts	
Command not echoed	3,606,468	(85.26%)
`STARTTLS` echoed correctly	617,093	(14.59%)
`STARTTLS` replaced	5,756	(0.14%)
Command truncated to four characters	786	(0.02%)

Table 11: **Detecting STARTTLS Manipulation** — We could extract an echoed command from 14.75% of servers that sent errors in response to our `STARTTLS` command. 0.14% of these responses indicate that the command was tampered with before reaching the server.

Type	ASes	
Corporation	182	(43.0%)
ISP	74	(17.5%)
Financial	57	(13.5%)
Academic	35	(8.3%)
Government	30	(7.1%)
Healthcare	14	(3.3%)
Unknown	12	(2.8%)
Airport	9	(2.1%)
Hosting	7	(1.7%)
NGO	3	(0.7%)

Table 12: **ASes Stripping STARTTLS** — We categorize the 423 ASes for which 100% of SMTP servers showed behavior consistent with STARTTLS stripping.

	Top Million Domains	IPv4 Hosts
Cisco-style tampering	2,563	41,405
`BLUF` tampering	0	6

Table 13: **Styles of STARTTLS Stripping** — The most prominent style of manipulation matches the advertised behavior of Cisco security devices and affects 41K SMTP servers.

Category	IPv4 Hosts
DNS servers	13,766,099
Responsive DNS servers	8,860,639
Any invalid MX responses	234,756
Class of invalid behavior:	
Identical response regardless of request	131,898
Returns loopback address	16,015
Returns private network address	7,680
Flipped bits in response	56,317
Falsified DNS record	178,439

Table 14: **Invalid or Falsified MX Records** — We scanned the IPv4 address space for DNS servers that provided incorrect entries for the MX servers for five popular mail providers.

5.1 STARTTLS Corruption

An active attacker — or a legitimate organization with a vested interest in snooping mail — can prevent mail encryption by tampering with the establishment of a TLS session. In this attack, a network actor takes advantage of the fail-open design of STARTTLS — where SMTP servers fall back to cleartext if any errors occur during the STARTTLS handshake — to launch a downgrade attack. A network actor can manipulate packets containing the `STARTTLS` command to prevent mail servers from establishing a secure channel, or alter a mail server's `EHLO` response to remove `STARTTLS` from the list of server capabilities. To measure whether STARTTLS sessions are being downgraded in the wild, we attempted to initiate STARTTLS connections with SMTP servers throughout the public IPv4 space and looked for evidence of tampering.

Scanning Methodology To find servers where STARTTLS is blocked, we build on the fact that SMTP servers frequently report back invalid commands they receive — which would include any corrupted `STARTTLS` command. We performed a TCP SYN scan of the public IPv4 address space on port 25 and attempted to perform an SMTP and STARTTLS handshake with responsive hosts. We performed this scan on April 20, 2015, from the University of Michigan campus using ZMap [19].

We found 14.1M hosts with port 25 open, 8.9M SMTP servers, and 4.6M SMTP servers that supported STARTTLS (see Table 10). Of the 4.2M hosts that failed to complete a TLS handshake, 623,635 (14%) echoed back the command they received. We classify these responses in Table 11. 617,093 (98.95%) of the responding hosts returned `STARTTLS` (and indeed do not to support it), 5,750 (0.92%) returned `XXXXXXXX`, 786 (0.14%) responded with `STAR` or `TTLS`, and 6 responded with `BLUF`.

The `STAR` and `TTLS` commands are four-character command truncations and are likely not due to an attack. Prior to ESMTP, SMTP commands were all four characters, and we were able to confirm that all commands were truncated to four characters on these servers. However, the `XXXXXXXX` and `BLUF` commands appear due to the `STARTTLS` command being altered to prevent the establishment of a TLS session.

Affected Servers Excluding four-character truncations, our scan found 5,756 servers that display evidence of a corrupted `STARTTLS` command. However, given that only 14% of servers reported back the received command, this is likely an underestimate. We extended our search for servers where the `STARTTLS` command was corrupted in the server's list of advertised features, which is returned in response to the `EHLO` command. This identified an additional 35,649 servers. When combined with the initial set, this yields a total of 41,405 servers that apparently have `STARTTLS` messages corrupted. These 41K servers are located in 4,714 ASes (15% of all ASes with an SMTP server) and 191 countries (86% of countries with SMTP servers). They transit mail for 2,563 domains in the Top Million.

In 423 ASes (736 hosts), 100% of SMTP servers are affected by STARTTLS stripping. The AS performing stripping on 100% of the inbound and outbound mail with the most SMTP servers (21) belongs to Starwood Hotels and Resorts (AS 13401). We show the classification of ASes with 100% stripping in Table 12. Overall, no single demographic stands out; the distribution is spread over networks owned by governments, Internet service providers, corporations, and financial, academic, and health care institutions. We note that several airports and airlines appear on the list, including an AS belonging to a subsidiary of Boingo (AS 10245), a common provider of in-flight and airport WiFi.

Our scanning methodology does not comprehensively find all servers where STARTTLS is blocked. Local SMTP servers may not

be accessible from the University of Michigan, STARTTLS might only be stripped for outgoing messages, or the command might be removed altogether instead of being corrupted in place. However, it appears that the practice is widespread.

Possible Causes The XXXXXXXX replacements are likely caused by security products that intercept and strip the command. In one prominent example, Cisco Adaptive Security Appliances (ASA) [7] and Cisco IOS Firewalls [8] both advertise replacing the STARTTLS command with Xs to facilitate mail inspection as part of their *inspect smtp* and *inspect esmtp* configurations. By inspecting messages, Cisco advertises that their products are capable of searching for and dropping messages with invalid characters in mail addresses, invalid SMTP commands, and long commands that may be attempting to exploit buffer overflows [9]. Table 13 shows the prominence of this style of tampering. While Cisco advertises this functionality, we cannot necessarily attribute every instance seen in the wild to Cisco devices, since others could implement stripping the same way.

We are unable to attribute the BLUF replacement to any commonly known security software. The six hosts affected by this replacement also had the PIPELINING and CHUNKING capabilities in the EHLO response masked to HIPELINING and PHUNKING, respectively. Only those six hosts displayed this behavior; all were located in Ukraine.

Impact on Transited Mail To understand the volume of mail affected by STARTTLS corruption, we measure the number of messages transited to/from these devices from Gmail's perspective. The overall fraction of mail affected is small, but a handful of countries have a high local stripping rate (see Table 15). In the most extreme example, 96.13% of mail transited from Tunisia to Gmail is affected by STARTTLS stripping. Another 8 countries experience over 10% stripping, and 16 experience more than 5% stripping.

It is important to note that the devices that are stripping TLS from SMTP connections are not inherently malicious, and many of these devices may be deployed to facilitate legitimate filtering. Regardless of the intent, this technique results in messages being sent in cleartext over the public Internet, enabling passive eavesdropping and other attacks. Furthermore, the Cisco documentation does not discuss the downsides of enabling this functionality, and administrators may not be aware that the setting puts users at risk. Instead of stripping TLS, manufacturers should consider deploying in-line devices that accept and initiate STARTTLS connections, allowing them to inspect messages before securely forwarding them.

5.2 DNS Hijacking

Mail security, like that of many other protocols, is intrinsically tangled with the security of DNS resolution. Rather than target the SMTP protocol, an active network attacker can spoof the DNS records of a destination mail server to redirect SMTP connections to a server under the attacker's control. We investigated the prevalence of DNS servers that provide false MX records or SMTP server IP addresses for: gmail.com, yahoo.com, outlook.com, qq.com (a popular Chinese webmail provider), and mail.ru (a popular Russian webmail provider). We find evidence that 178,439 out of 8,860,639 (2.01%) publicly accessible DNS servers provided invalid IPs or MX records for one or more of these domains (see Table 8).

Scanning Methodology We identified servers responding with falsified DNS records by scanning the IPv4 address space ten times with ZMap in search of open resolvers and subsequently requesting the MX and A record of gmail.com, yahoo.com, outlook.com, qq.com, and mail.ru. We performed these scans on April 25, 2015, from the University of Michigan. In total, we identify 13.8 million DNS servers, of which 8.9 million resolved at least one query and 235K provide an invalid or falsified MX record (see Table 14).

Tunisia	96.13%	Reunion	9.28%
Iraq	25.61%	Belize	7.65%
Papua New Guinea	25.00%	Uzbekistan	6.93%
Nepal	24.29%	Bosnia and Herzegovina	6.50%
Kenya	24.13%	Togo	5.45%
Uganda	23.28%	Barbados	5.28%
Lesotho	20.25%	Swaziland	4.62%
Sierra Leone	13.41%	Denmark	3.69%
New Caledonia	10.13%	Nigeria	3.64%
Zambia	9.98%	Serbia	3.11%

Table 15: **Countries Affected by STARTTLS Stripping** — We measure the fraction of incoming Gmail messages that originate from the IPs that we found were stripping TLS from SMTP connections. Here, we show the countries with the most mail affected by STARTTLS stripping and the affected percentage of each country's incoming mail between April 20 and 27, 2015.

Slovakia	0.08%
Romania	0.04%
Bulgaria	0.03%
India	0.02%
Israel	0.01%
Switzerland	0.01%
Poland	0.01%
Ukraine	0.01%

Table 16: **Countries Affected by Falsified DNS Records** — We measure the fraction of mail received by Gmail on May 21, 2015 from IP addresses pointed to by false Gmail DNS entries. Here, we show the breakdown of mail from each country that originates from one of these addresses for the countries with the most affected mail.

We then performed a secondary scan, in which we repeated the same queries as well as performed A record lookups for a domain unrelated to mail transit (umich.edu) and a nonexistent domain. Of the 235K servers that provided invalid responses in our first scan, 56K supplied correct results in the secondary scan and were incorrectly flagged due to erroneous bit flips. Excluding those hosts, 132K provide the same publicly accessible address for every DNS query regardless of the domain, 7.7K provide reserved or private addresses, 16K respond with a loopback address, and 17.2K did not appear to spoof answers but were missing at least one of the MX servers. The devices that provided identical responses to every query or were missing an MX server appeared to be improperly configured embedded devices rather than malicious. After removing these hosts, we were left with 14.6K hosts that provided invalid responses for mail servers. These hosts pointed to 1,150 unique falsified mail servers, of which 144 (12.5%) completed an SMTP handshake.

Our scans do not provide an exhaustive list of hosts that might be intercepting mail. Since open resolvers are frequently used to launch DDoS attacks, most recursive DNS resolvers are not externally accessible and will not appear in our scans. Similarly, many of the addresses we recorded are private, non-routable addresses, so we are unable to test whether mail transits through these hosts. However, our scan still finds upwards of 15K open resolvers that provide fraudulent responses when queried about mail providers and 1.2K false mail servers, which allows us to determine whether mail server DNS hijacking occurs in the wild.

Responsible Networks The DNS servers that provide fraudulent responses are located in 521 ASes. 83.6% of the hosts were located in five ASes: 62% Unified Layer (American Hosting Provider), 11.7% ChinaNet, 5.3% Telecom Italia (Italian ISP), 2.4% SoftLayer Technologies, and 2.0% eNom. In the case of Unified Layer, 9K

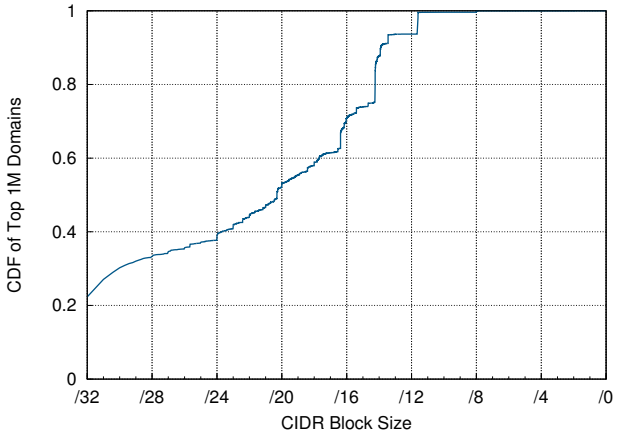

Figure 4: **Size of SPF Permitted Networks** — We show the CDF of the number of addresses whitelisted in a recursive resolution of the SPF records for Top Million domains.

hosts point back to seven servers, of which two completed SMTP handshakes. The hosts in the ChinaNet AS point to a range of loopback and private addresses and to 42 publicly routable servers, of which one completed an SMTP handshake. The devices in the Telecom Italia AS returned monotonically increasing IP addresses within 198.18.1.0/24 for all queries and do not appear to be specifically intercepting mail queries. The SoftLayer hosts respond with one of 18 addresses, of which 10 were SMTP servers. All eNom hosts point to a single IP, which did not accept SMTP connections. The remaining 2,386 DNS servers are located in 533 ASes and in 69 countries.

Impact on Transited Mail While a number of servers appear to be intercepting mail, the impact on transited messages is unclear. We estimate the amount of mail affected by measuring the number of inbound messages that Gmail received from each of the servers. The vast majority of mail that transits from these hosts is spam, but a small number of non-spam messages are sent through these servers. As shown in Table 16, in the most extreme case, upwards of 0.08% of mail transited from Slovakia came from falsified servers.

Whether malicious or well-intentioned, STARTTLS stripping and falsified DNS records highlight the weakness inherent in the fail-open nature and lack of authentication of the STARTTLS protocol. These attacks are both readily found in the wild and pose a real threat to users, with more than 20% of mail being sent in cleartext within seven countries.

6. AUTHENTICATION IN PRACTICE

While STARTTLS protects messages against passive eavesdropping, it does not provide authentication. Mail can be modified or spoofed altogether, even in the presence of STARTTLS. As described in Section 2, SPF, DKIM, and DMARC have been developed to authenticate incoming mail. In this section, we measure how these protocols have been deployed in practice. We summarize the deployment of all three protocols in Table 18.

6.1 SPF

When supported, SPF allows recipients to check that incoming messages from a domain (e.g., gmail.com) originate from an IP range authorized by that organization. From Gmail's perspective, Google successfully authenticated 92% of inbound messages during April 2015 using SPF, as detailed in Table 19. Of the unauthenticated

Provider	SPF Policy	DMARC Policy
Gmail	soft fail	none
Yahoo	neutral	reject
Outlook	soft fail	none
iCloud	soft fail	none
Hushmail	soft fail	–
Lycos	soft fail	–
Mail.com	fail	–
Zoho	soft fail	–
Mail.ru	soft fail	none
AOL	soft fail	reject
QQ	soft fail	none
Me.com	soft fail	none
Facebook	fail	reject
GoDaddy	fail	none
Yandex	soft fail	–
OVH	neutral	–
Comcast	neutral	none
AT&T	–	–
Verizon	neutral	–

Table 17: **SPF and DMARC Policies** — The majority of popular mail providers we tested posted an SPF record, but only three used the "strict fail" policy. Even fewer providers posted a DMARC policy, of which only three used "strict reject."

messages, Gmail fails to validate 0.42% due to failures fetching the domain's SPF record; all other messages come from domains without an SPF policy.

Similar to STARTTLS deployment, the servers associated with the Top Million domains show significantly slower adoption with only 401,356 domains — 47% of the Top Million domains with MX records — publishing an SPF policy (see Table 20). Of those, 255,867 domains allow mail to be sent if the server has an MX record on the domain, and 104 domains allow mail from hosts with reverse DNS names that match the domain.

Record Delegation Of the Top Million domains, 10,432 redirect (or fully delegate their SPF policy) to another provider, and 213,464 (53.2%) include records from one or more other domains' SPF policies. While this could potentially open an attack vector if multiple organizations specified the same IP blocks, we find that this was not the case. Instead, domains commonly include records from or redirect to a few well-known cloud mail providers. In the case of redirection, 35.7% delegate to Yandex (a Russian mail provider), 16.7% to mailhostbox.com, 8.0% to nicmail.ru, 3.9% to serveriai.lt, and 3.5% to mail.ru. 3,813 (36.6%) of all redirects point to a Russian mail service. For includes, 136,473 domains (64% of domains with includes) include one of five large mail providers: Gmail (59,660), Outlook (44,216), websitewelcome.com (20,291), mandrillapp.com (16,606), and SendGrid (10,700).

SPF Policy Coarseness For the domains that specify SPF policies, we find evidence that some report overly broad IP ranges that potentially enable an attacker to spoof mail origins, as shown in Figure 4. We find that 133,490 domains (60.9%) specify SPF CIDR ranges larger than a /24, 99,698 (29.2%) specify CIDR ranges larger than or equal to a /16, and 1,333 (0.4%) specify more than a /8. The vast majority of the domains that include a /8 mistakenly allowed messages from 10.0.0.0/8. In rarer cases, we find evidence of blatant misconfiguration: 62 domains specify network ranges akin to 255.255.255.255/8, and 20 domains specify ranges larger than a /8.

Policy Types Of the domains that deploy SPF, 21.7% adopt hard-failure policies, where recipients should reject mail from outside of the specified network. Another 58.0% adopt soft-failure policies,

where recipients should accept the mail but consider it suspect (e.g., mark as spam), and 20.3% set no policy. If we restrict our analysis to the top mail providers, we find that most publish SPF records with a soft-fail policy (see Table 17). Exceptions include Facebook Mail, Mail.com, and GoDaddy, all of which have hard-fail policies. AT&T was the only provider we checked that does not have a valid SPF policy. We note that soft-fail policies allow more leeway for the destination domain to decide how to process a message, such as considering other spam indicators instead of downright rejecting messages.

6.2 DKIM

As a complement or alternative to SPF, DKIM allows a recipient to confirm the integrity and authenticity of inbound messages, even in the presence of a man-in-the-middle attack. In April 2015, 83.0% of the messages that Gmail received contained a DKIM signature. Of the signed messages, 6.14% failed to validate due to weak cryptographic keys, revoked keys, or protocol errors (see Table 9). This represents a 4.42% decrease in invalid signatures when compared to two years prior. We specifically call attention to the fact that 18.7% of failures in April 2015 arose due to DKIM signatures not matching a message's content. While we cannot distinguish between malice and misconfiguration, such failures reflect the importance of authentication to alert mail servers to potential tampering, and they emphasize the imperative for the remaining 17% of unauthenticated inbound Gmail messages to adopt DKIM signatures. Unfortunately, due to the nature of the DKIM protocol, we cannot directly measure how many domains in the Alexa Top Million have deployed DKIM.

6.3 DMARC

DMARC policies allow sending mail servers to alert recipients that they support DKIM and SPF and then inform recipients how to handle incoming messages that fail to validate or that lack a DKIM signature. In contrast to the 90% of the messages that Gmail can validate with SPF or DKIM, only 26.1% of all inbound Gmail messages come from domains with a published DMARC policy. This discrepancy limits the effectiveness of DKIM as, absent a publicized policy, recipients cannot determine whether the lack of a signature is intended or is an indication of spoofing. For those messages with a policy, we provide a detailed breakdown in Table 22. We found that the majority of policies favored rejection (13%), though a significant fraction did not specify any action (11%).

A similarly pessimistic picture emerges for the Alexa Top Million, for which only 1.1% of domains with MX records published DMARC policies. We provide a breakdown of all the policies in Table 22 and the policies of top mail providers in Table 17. Even when DMARC is present, the majority of Alexa domains and even the top mail providers specified an empty policy. Only Yahoo, AOL, and Facebook advertised DMARC reject policies.

We suspect that many organizations have yet to deploy DMARC due its relatively young age—RFC 7489 was first introduced in March 2013 and was last updated in March 2015, one month prior to our measurements. However, we note its necessity in enforcing SPF and DKIM policies, and we hope it will see increased deployment moving forward.

7. DISCUSSION

The mail community has retroactively applied several security measures to SMTP. Nearly 60% of incoming connections to Gmail are encrypted and 94% of messages are authenticated with DKIM or SPF. In many ways, this is a feat, given that SMTP did not originally provide any support for transport security. However, our two perspectives paint drastically different pictures of how mail security

Authentication Method	Nov. 2013	Apr. 2015	Change
DKIM & SPF	74.66%	81.01%	+6.31%
DKIM only	2.25%	1.98%	−0.27%
SPF only	14.44%	11.41%	−2.99%
No authentication	8.65%	5.60%	−3.00%

Table 18: **Gmail Incoming Mail Authentication**—During April 2015, 94.40% of incoming Gmail messages were authenticated with DKIM, SPF, or both.

SPF Policy	Gmail Messages
DNS timeout	<0.001%
Temporary error	0.184%
Permanent error	0.141%
Invalid record	0.098%

Table 19: **SPF Errors for Incoming Gmail Traffic**—We show the breakdown of errors fetching SPF records for incoming mail. Temporary errors can be fixed by retrying later; permanent errors mean the record was unable to be fetched.

Policy	Top Million Domains	Recursive Top Million
SPF policy	401,356	401,356
Hard fail	84,801 (21.13%)	86,919 (21.65%)
Soft fail	226,117 (56.34%)	232,736 (57.99%)
Neutral	80,394 (20.03%)	81,701 (20.36%)
Redirect	10,045 (2.50%)	0 (0.00%)

Table 20: **SPF Policies for Top Million Domains**—We queried the SPF policies for the Top Million domains for both the top-level record and for full recursive resolution.

Record Type	Top Million Domains	Recursive Top Million
IPv4	200,976 (33.08%)	344,844 (40.22%)
IPv6	6,862 (1.13%)	108,086 (12.61%)
A	139,979 (23.04%)	148,688 (17.34%)
MX	249,345 (41.04%)	255,867 (29.84%)
REDIRECT	10,432 (1.72%)	0 (0.00%)

Table 21: **SPF Record Types for Top Million Domains**—We show how hosts are whitelisted within an SPF record for both the top-level SPF record and for full recursive resolution.

Published Policy	Gmail Messages	Top Million Domains
Quarantine	1.34%	709 (0.09%)
Empty	11.66%	6,461 (0.82%)
Reject	13.08%	1,720 (0.22%)
Not published	73.92%	783,851 (98.9%)

Table 22: **DMARC Policies**—We categorize DMARC policies for incoming Gmail messages from April 2015 and for Top Million domains with MX records on April 26, 2015.

has been deployed. As can be seen by the 51% jump in encrypted inbound messages when Microsoft and Yahoo deployed STARTTLS, much of this success can be attributed to large mail providers that are pushing security forward. Unfortunately, as our scans demonstrate, smaller organizations lag in deploying security mechanisms correctly. While mail delivery security is rapidly improving, there are several structural problems that the mail community must address to guarantee the confidentiality and integrity of mail.

7.1 Challenges for Confidentiality

There are several challenges to guaranteeing the confidentiality of mail in transit. First, unlike HTTPS, there is no mechanism in SMTP for servers to indicate that mail should be protected by TLS. Further, users cannot indicate that mail should only be transited securely nor can they detect whether a message traversed a secure path.

In HTTPS, users can choose not to communicate over an insecure channel, and HSTS allows sites to indicate that future connections must use HTTPS. However, in SMTP, messages are relayed in cleartext if TLS cannot be negotiated. As we showed in Section 5.1, this has led to organizations corrupting the STARTTLS negotiation to force mail to be sent in the clear. Whether this is being done for legitimate or nefarious purposes, it illustrates that STARTTLS provides no protection against frequently occurring man-in-the-middle (MITM) attacks.

Second, even when TLS is used, there is no robust way for a sender to verify the authenticity of the recipient mail server. Common MTAs can validate that a server's certificate matches the destination domain's MX record, but not the destination domain name itself. Unfortunately, this still leaves the server open to impersonation unless the DNS responses are separately authenticated. As we showed in Section 5.2, certain entities are using this weakness to redirect the flow of messages.

One potential option for preventing MITM attacks is to create a mechanism similar to HTTP Public Key Pinning [21] for SMTP. This would allow a mail server to indicate whether future connections should require TLS and specify a public key. Other protections being adopted for the HTTPS PKI might also be considered for STARTTLS, such as the use of Certificate Transparency [6] to guard against dishonest or compromised certificate authorities.

Finally, we note that end-to-end mail encryption, as provided by PGP [4] and S/MIME [42], does not address many of the challenges we discuss in this work. While these solutions do safeguard message content, they leave metadata, such as the subject, sender, and recipient, visible everywhere along the message's path. This information is potentially exposed to network-based attackers due to the lack of robust confidentiality protections for SMTP message transport. Although greater adoption of end-to-end encryption would undoubtedly be beneficial, for now, the overwhelming majority of messages depend solely on SMTP and its extensions for protection.

7.2 Challenges for Integrity

A major open question surrounding mail integrity is how to authenticate mail sent through mailing lists. Mailing lists frequently modify messages in transit, and DKIM signatures are invalidated by these modifications, which prevents large mail providers from publishing a DMARC reject policy. When Yahoo deployed a reject policy in 2014, it resulted in a heavy number of complaints and service malfunctions [10].

A second challenge is ensuring strong integrity as organizations move to cloud providers, where mail infrastructure and IP address blocks may be shared with other organizations. This infrastructure sharing is challenging in two respects. First, SPF has become less relevant, since, as explained in Section 6.1, SPF records tend to be overly broad. Second, DKIM becomes threatened by massive key compromises, as was the case for the SendGrid leak [5]. Overall, these two issues are part of a larger open question: How do we reliably establish the legitimacy of senders — whether for spam prevention or for integrity purposes — when many senders, good and bad, share common infrastructure?

The issue of shared infrastructure also affects mail confidentiality, as third-party providers would need to have certificates containing their clients' domains to allow strict certificate verification. This is problematic, as it opens the door to attacks where the third-party mail provider — or an attacker who breaches their systems — uses these certificates to impersonate the clients' domains, either for mail delivery or for HTTPS connections. This threat might be mitigated with a scope-reducing X.509 extension or through some other mechanism not yet devised.

8. RELATED WORK

There has been little formal measurement of the public key infrastructure that supports mail transport until recently. There are three works similar to ours.

The first is a set of Facebook blog posts [22, 23] that describe the STARTTLS configurations seen from the perspective of Facebook notifications. In May 2014, Facebook found that 28.6% of notification emails are transported over a STARTTLS connection with strict certificate validation, 28.1% are protected with opportunistic encryption (indicating a misconfigured STARTTLS server), and 41.0% of notifications are sent in cleartext. In August 2014, Facebook posted follow-up statistics in which they note that 95% of notification emails are sent over STARTTLS with strict certificate validation. Facebook further notes that this rise is primarily due to two major mail providers, Yahoo and Microsoft, deploying STARTTLS. The jump of encrypted messages from 28.6% to 95% is incredibly exciting. However, as noted by Facebook, their notification emails are skewed towards personal addresses and large hosting providers, such as Gmail and Yahoo Mail.

Concurrently with our work, Foster et al. [24] performed a similar study on the deployment of SMTP extensions for the Top Million domains present in the 2013 Adobe data breach [3] (ranked by number of accounts) and investigated how different types of senders (e.g., popular banks and e-commerce sites) protect mail. They found that 89% of popular mail providers deploy STARTTLS, 85% have SPF records, and 68% have DMARC policies. In comparison, at the termination of our study in April, 54% of incoming messages to Gmail were protected by STARTTLS, and 82% of the domains to which Gmail transited mail supported inbound STARTTLS. While protocol deployment appears higher in their work, this falls in line with the trends we see: popular providers have deployed security extensions more comprehensively than smaller organizations.

In June 2014, Sean Rijs [43] published a measurement study on the use of STARTTLS among 116 Dutch organizations which found that: 55% of their domains used STARTTLS, 34% did not support STARTTLS, and 11% could not be tested. Our results provide another perspective, including how incoming messages are protected, mail is authenticated, and organizations deploy STARTTLS .

Mail Redirection There is a large corpus of work on DNS servers that provide false responses in order to facilitate content filtering [13, 36, 39, 47]. However, our study is the first to measure the extent to which DNS servers are falsifying MX records for mail providers and the amount of mail sent through these servers. In 2014, the Electronic Frontier Foundation (EFF), Golden Frog, and Telecom Asia posted anecdotal evidence of several ISPs blocking STARTTLS sessions in the United States and Thailand [25, 29, 44]. The EFF proposed STARTTLS Everywhere, an open source project

that contains a public list of domains that support STARTTLS and scripts for generating configuration files for common MTAS that require STARTTLS for those domains [30]. Our work provides an additional perspective and estimates the amount of mail affected by STARTTLS stripping.

Internet-wide Scanning While Internet-wide scanning has not previously been used to measure the mail security ecosystem, it has become a standard practice for measuring the HTTPS ecosystem. In 2010, the EFF performed a distributed scan [20] of the IPv4 address space to identify certificate authorities. Later, in 2011, Holz et al. [31] scanned the Alexa Top Million in order to measure HTTPS deployment and commonly used CAs. In 2012, Heninger et al. [27] performed comprehensive scans of HTTPS servers and detected wide use of weak cryptographic keys. Again in 2013, Durumeric et al. [18] completed daily scans in order to identify weaknesses in the HTTPS CA ecosystem. In 2014, Huang et al. [32] scanned the Top Million to measure the deployment of Forward Secrecy.

9. CONCLUSION

While electronic mail carries some of users' most sensitive correspondence, SMTP did not originally include support for message confidentiality or integrity. Over the past fifteen years, the mail community has retrofitted SMTP with several security mechanisms, including STARTTLS, SPF, DKIM, and DMARC. In this work, we analyzed the global adoption of these technologies using data from two perspectives: Internet-wide scans and logs of SMTP connections to and from one of the world's largest mail providers over a sixteen month period. Our measurements show that the use of these secure mail technologies has surged over the past year. However, much of this growth can be attributed to a handful of large providers, and many smaller organizations continue to lag in both deployment and proper configuration. The fail-open nature of STARTTLS and the lack of strict certificate validation reflect the need for interoperability amidst the gradual rollout of secure mail transport, and they embody the old adage that "the mail must go through." Unfortunately, they also expose users to the potential for man-in-the-middle attacks, which we find to be so widespread that they affect more than 20% of messages delivered to Gmail from several countries. We hope that by drawing attention to these attacks and shedding light on the real-world challenges facing secure mail, our findings will motivate and inform future research.

Acknowledgments

The authors thank Vern Paxson, Paul Pearce, Niels Provos, Eric Wustrow, and our shepherd, Alan Mislove, for their help and feedback. We thank the exceptional sysadmins at the University of Michigan for their help and support, including Chris Brenner, Kevin Cheek, Laura Fink, Dan Maletta, Jeff Richardson, Donald Welch, Don Winsor, and others from ITS, CAEN, and DCO. This material is based upon work supported by the National Science Foundation under grants CNS-1111699, CNS-1255153, CNS-1345254, CNS-1409505, CNS-1409758, and CNS-1518741, by the Google Ph.D. Fellowship in Computer Security, by the Morris Wellman Faculty Development Assistant Professorship, and by an Alfred P. Sloan Foundation Research Fellowship.

10. REFERENCES

[1] Alexa Internet, Inc. Alexa Top 1,000,000 Sites. http://s3.amazonaws.com/alexa-static/top-1m.csv.zip.

[2] N. J. AlFardan, D. J. Bernstein, K. G. Paterson, B. Poettering, and J. C. Schuldt. On the security of RC4 in TLS. In *22nd USENIX Security Symposium*, pages 305–320, Aug. 2013.

[3] B. Arkin. Adobe important customer security announcement, Oct. 2013. http://blogs.adobe.com/conversations/2013/10/important-customer-security-announcement.html.

[4] J. Callas, L. Donnerhacke, H. Finney, D. Shaw, and R. Thayerj. OpenPGP message format. RFC 4880, 2007. https://www.ietf.org/rfc/rfc4880.txt.

[5] D. Campbell. Update on security incident and additional security measures, 2015. https://sendgrid.com/blog/update-on-security-incident-and-additional-security-measures/.

[6] Certificate Transparency, 2015. http://www.certificate-transparency.org/.

[7] Cisco. Cisco ASA 5500-X series next-generation firewalls, 2015. http://www.cisco.com/c/en/us/products/security/asa-5500-series-next-generation-firewalls/index.html.

[8] Cisco. Cisco IOS Firewall, 2015. http://www.cisco.com/c/en/us/products/security/ios-firewall/index.html.

[9] Cisco. SMTP and ESMTP inspection overview, 2015. http://www.cisco.com/c/en/us/td/docs/security/asa/asa93/configuration/firewall/asa-firewall-cli/inspect-basic.html#pgfId-2490137.

[10] L. Constantin. Yahoo email anti-spoofing policy breaks mailing lists, 2014. http://www.pcworld.com/article/2141120/yahoo-email-antispoofing-policy-breaks-mailing-lists.html.

[11] D. Crocker, T. Hansen, and M. Kucherawy. DomainKeys Identified Mail (DKIM) signatures. RFC 6379, Sept. 2011. https://tools.ietf.org/html/rfc6376.

[12] D. Crocker and T. Zink. M3AAWG trust in email begins with authentication, 2015. https://www.m3aawg.org/sites/maawg/files/news/M3AAWG_Email_Authentication_Update-2015.pdf.

[13] H. Duan, N. Weaver, Z. Zhao, M. Hu, J. Liang, J. Jiang, K. Li, and V. Paxson. Hold-on: Protecting against on-path DNS poisoning. In *Workshop on Securing and Trusting Internet Names*, 2012.

[14] V. Dukhovni and W. Hardaker. SMTP security via opportunistic DANE TLS, July 2013. http://tools.ietf.org/html/draft-ietf-dane-smtp-with-dane-12.

[15] Z. Durumeric, D. Adrian, J. Kasten, D. Springall, M. Bailey, and J. A. Halderman. POODLE attack and SSLv3 deployment, 2014. https://poodle.io.

[16] Z. Durumeric, D. Adrian, A. Mirian, M. Bailey, and J. A. Halderman. A search engine backed by Internet-wide scanning. In *22nd ACM Conference on Computer and Communications Security*, Oct. 2015.

[17] Z. Durumeric, M. Bailey, and J. A. Halderman. An Internet-wide view of Internet-wide scanning. In *23rd USENIX Security Symposium*, Aug. 2014.

[18] Z. Durumeric, J. Kasten, M. Bailey, and J. A. Halderman. Analysis of the HTTPS certificate ecosystem. In *13th ACM Internet Measurement Conference*, Oct. 2013.

[19] Z. Durumeric, E. Wustrow, and J. A. Halderman. ZMap: Fast Internet-wide scanning and its security applications. In *22nd USENIX Security Symposium*, Aug. 2013.

[20] P. Eckersley and J. Burns. An observatory for the SSLiverse. Talk at Defcon 18 (2010). https://www.eff.org/files/DefconSSLiverse.pdf.

[21] C. Evans, C. Palmer, and R. Sleevi. Public key pinning extension for HTTP. RFC 7469, 2015. http://tools.ietf.org/html/rfc7469.

[22] Facebook. The current state of SMTP STARTTLS deployment, May 2014. https://www.facebook.com/notes/protect-the-graph/the-current-state-of-smtp-starttls-deployment/1453015901605223/.

[23] Facebook. Massive growth in SMTP STARTTLS deployment, Aug. 2014. https://www.facebook.com/notes/protect-the-graph/massive-growth-in-smtp-starttls-deployment/1491049534468526.

[24] I. Foster, J. Larson, M. Masich, A. Snoeren, S. Savage, and K. Levchenko. Security by any other name: On the effectiveness of provider based email security. In *22nd ACM Conference on Computer and Communications Security*, Oct. 2015.

[25] Golden Frog. The FCC must prevent ISPs from blocking encryption. http://www.goldenfrog.com/blog/fcc-must-prevent-isps-blocking-encryption.

[26] N. Heninger. Factoring as a service. Rump session talk, *Crypto* 2013.

[27] N. Heninger, Z. Durumeric, E. Wustrow, and J. A. Halderman. Mining your Ps and Qs: Detection of widespread weak keys in network devices. In *21st USENIX Security Symposium*, Aug. 2012.

[28] P. Hoffman. SMTP service extension for secure SMTP over transport layer security. RFC 3207, Feb. 2002. http://www.ietf.org/rfc/rfc3207.txt.

[29] J. Hoffman-Andrews. Isps removing their customers' email encryption. https://www.eff.org/deeplinks/2014/11/starttls-downgrade-attacks.

[30] J. Hoffman-Andrews and P. Eckersley. STARTTLS everywhere, June 2014. https://github.com/EFForg/starttls-everywhere.

[31] R. Holz, L. Braun, N. Kammenhuber, and G. Carle. The SSL landscape: A thorough analysis of the X.509 PKI using active and passive measurements. In *11th ACM Internet Measurement Conference*, 2011.

[32] L.-S. Huang, S. Adhikarla, D. Boneh, and C. Jackson. An experimental study of TLS forward secrecy deployments. In *Web 2.0 Security and Privacy (W2SP)*, 2014.

[33] S. Kitterman. Sender policy framework (SPF) for authorizing use of domains in email. RFC 7208, Apr. 2014. http://tools.ietf.org/html/rfc7208.

[34] J. Klensin. Simple mail transfer protocol. RFC 5321, Oct. 2008. http://tools.ietf.org/html/rfc5321.

[35] M. Kucherawy and E. Zwicky. Domain-based message authentication, reporting, and conformance (DMARC). RFC 7489, Mar. 2015. https://tools.ietf.org/html/rfc7489.

[36] G. Lowe, P. Winters, and M. L. Marcus. The great DNS wall of China. Technical report, New York University, Dec. 2007. http://cs.nyu.edu/~pcw216/work/nds/final.pdf.

[37] Microsoft. TLS functionality and related terminology, June 2014. http://technet.microsoft.com/en-us/library/bb430753%28v=exchg.150%29.aspx.

[38] Mozilla Developer Network. Mozilla Network Security Services (NSS). http://www.mozilla.org/projects/security/pki/nss/.

[39] Z. Nabi. The anatomy of web censorship in Pakistan. *arXiv preprint arXiv:1307.1144*, 2013.

[40] J. B. Postel. Simple mail transfer protocol. RFC 821, Aug. 1982.

[41] http://www.postfix.org/postconf.5.html#smtp_tls_security_level.

[42] B. Ramsdell and S. Turner. Secure/multipurpose Internet mail extensions (S/MIME) version 3.2 message specification. RFC 5751, 2010. https://tools.ietf.org/html/rfc5751.

[43] S. Rijs and M. van der Meer. The state of StartTLS, June 2014. https://caldav.os3.nl/_media/2013-2014/courses/ot/magiel_sean2.pdf.

[44] Telecom Asia. Google, Yahoo SMTP email severs hit in thailand. http://www.telecomasia.net/content/google-yahoo-smtp-email-severs-hit-thailand.

[45] M. Vanhoef and F. Piessens. All your biases belong to us: Breaking RC4 in WPA-TKIP and TLS. In *24th USENIX Security Symposium*, Aug. 2015.

[46] Verisign Labs. DNSSEC scoreboard, 2015. http://scoreboard.verisignlabs.com/.

[47] J.-P. Verkamp and M. Gupta. Inferring mechanics of web censorship around the world. *3rd USENIX Workshop on Free and Open Communications on the Internet (FOCI)*, Aug. 2012.