Thorn in the USA: On SPF, DMARC, and CAA Adoption and Misconfiguration by Institutions

Alexander Gebhard
Marquette University
alexander.gebhard@marquette.edu

David Plonka
WiscNet
plonka@wiscnet.net

Cameron Fronczak
Marquette University
cameron.fronczak@marquette.edu

Debbie Perouli
Marquette University
despoina.perouli@marquette.edu

***Abstract*—The security of email and the web depends in part on the integrity of the Domain Name System (DNS): both the authenticity of DNS records for web services and that of records for email domains, which are frequently impersonated to deliver malicious emails. This paper examines the adoption of three DNS-based protocols in the United States: Sender Policy Framework (SPF), Domain-based Message Authentication, Reporting, and Conformance (DMARC), and Certificate Authority Authorization (CAA). These protocols help establish legitimacy in an email sender's identity and in email content. From February 2023 to December 2025, SPF, DMARC, and CAA records from domains belonging to government organizations were collected, parsed, and semantically analyzed for security misconfigurations. In addition to that nationwide study, we perform a cluster-based analysis and report common misconfigurations for the hundreds of member institutions from one state's Research & Education Network (REN), comprising many municipal institutions and most of its educational institutions, statewide, in 2025. We found that adoption of all three protocols is slowly increasing, nationwide, however many configuration errors are repeated, common to very many domains, suggesting email and online content that purports to originate from these organizations and institutions often is not thoroughly authenticated, leaving recipients vulnerable to attacks by impersonation.**

I. Introduction and Motivation

The United States is home to more than 90,075 state and local governments, including 38,779 general-purpose local governments such as towns, counties, and municipalities [1]. These institutions play a trusted role in delivering essential public services—such as voting information, tax administration, utilities, and emergency medical services. Combined with anchor institutions, such as schools and hospitals, such organizations are trusted to manage sensitive personally identifiable information. Therefore, it is crucial for citizens to have confidence that email content that they receive originates from the legitimate entities rather than malicious parties masquerading as those trusted institutions.

To aid in the verification of email and service identities, there are three key protocols visible as types of Domain Name System (DNS) records: Sender Policy Framework (SPF), Domain-based Message Authentication, Reporting and Conformance (DMARC), and Certification Authority Authorization (CAA). This study specifically examines the scanning and parsing of these three DNS record types across governments in the U.S., which include local, county, state, interstate, and tribal governments. It also analyzes these in more detail, by clustering common protocol configurations, for similarly trusted entities in one state, which include municipalities, hospitals, libraries, public and private schools.

Our research has two primary objectives: first, to educate IT administrators about common implementation errors that suggest solutions for remediation; second, to inform protocol designers and policymakers of the real-world challenges faced in implementation. Acknowledging that many of the institutions we study operate with limited cybersecurity resources, we have developed free, open-source tools that assess their security posture.

II. Background

The increased reliance on digital infrastructure within public schools, municipalities, and public libraries implies that DNS security should be a critical concern. These institutions serve thousands of users daily, providing essential services such as public internet access, email communication, and online government resources funded by taxpayer dollars. However, the implementation of DNS security measures, such as SPF, DMARC, and CAA may be limited by resource constraints and the complexity of DNS management.

A. Sender Policy Framework (SPF)

SPF, first introduced circa 2003, is one of the standard methods for preventing or combating spam. Implemented using DNS TXT records, SPF is used to authenticate emails ostensibly sent from a domain and either permit or deny delivery, accordingly. Its original title was "Sender Permitted From." The Sender Policy Framework (SPF) is an email authentication mechanism that uses a DNS TXT record to define which mail servers are permitted to send email on behalf of a given domain. When a mail server receives a message, it checks the SPF record of the purported sender's domain to verify its authenticity. Here is an example of an SPF record: `v=spf1 ip4:192.0.2.25 include:_spf.example.com -all`. All SPF records start with text "`v=spf1`". This allows mail servers to determine if the TXT record is an SPF record. After the header, the SPF record contains a series of mechanisms or modifiers. These are used to describe the set of hosts which are permitted to send mail for the domain.

The eight defined mechanisms in SPF are as follows: `all` matches any sender and is often used at the end of a record with modifiers like fail (-) or soft fail (˜); `ip4` specifies

an IPv4 address or range as permitted senders; `ip6` works similarly for IPv6 addresses. The mechanism `a` matches the domain's A record, which reflects the domain's IP address, while `mx` matches the domain's MX records, which denote mail exchange servers. The `ptr` mechanism refers to the sender's PTR record, although it is deprecated and should not be used. The `include` mechanism allows the inclusion of another domain's SPF record, thereby permitting authorized senders from that domain. SPF records can use multiple of these mechanisms to specify valid hosts that are permitted to send email on behalf of a host. By defining which servers are allowed to send email for a domain, SPF helps protect against attackers who masquerade as users within that domain.

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

DMARC, officially introduced circa 2012 by PayPal, Google, Microsoft, and Yahoo, is a critical component of email security. It instructs receiving email servers on actions to take after checking a domain's SPF and DomainKeys Identified Mail (DKIM) records. By authenticating email senders, DMARC protects domain owners from unauthorized use, such as email spoofing and phishing. It verifies that emails claiming to come from a specific domain are legitimate, thus blocking attacks that exploit trust.

DMARC functions by requiring domain owners to publish a TXT record under the `_dmarc` subdomain. Organizations configure SPF and DKIM alongside DMARC, as they support its enforcement. A sample DMARC record looks like this: `v=DMARC1; p=quarantine; rua=mailto:dmarc-reports@example.com; pct=100`. Each DMARC record begins with "`v=DMARC1`", followed by the policy option designated by `p`. This policy can take one of three values: `reject` which blocks the email outright, `quarantine` which accepts the email but directs it to the spam folder, or `none` which permits the email without taking any unusual action. In this example, the policy is to quarantine any emails that fail SPF or DKIM checks. Additionally, domain administrators can configure options such as `rua` and `ruf` specifying email addresses that receive aggregate and forensic reports, respectively. The `pct` option allows them to specify the percentage of messages to which the DMARC policy should be applied.

Once a DMARC policy is established, receiving servers validate senders using SPF and/or DKIM before delivering the messages. DMARC also enables domain owners to receive reports on authentication failures providing insights into threats.

The policy enforcement allows domain owners to dictate the handling of emails that fail authentication checks, often leading to quarantining such messages temporarily. In summary, DMARC strengthens email security by preventing spoofing and phishing while improving trust in email communication and offering transparency into domain usage, aiding in the identification of security weaknesses or malicious behavior.

C. Certificate Authority Authorization (CAA)

CAA, the newest of the three records, was introduced in RFC 6844 [2] (2013) and is currently defined by RFC [3] (2019.) Having broader applications than SPF and DMARC, the CAA mechanism lets domain owners specify which certificate authorities (CAs) are allowed to issue digital certificates for their domain. CAA-compliant certificate authorities consult the DNS CAA resource record before issuing any certificates [4], so that a website identified by domain name can be determined to be genuine. CAA therefore helps reduce the risk of unauthorized certificate issuance, but it does not itself establish that a website is genuine. For example, when a user visits `www.example.com` over HTTPS, the browser validates the website's certificate including its chain of trust, validity period, and association with the requested domain. If validation fails, the browser applies its security policy, typically blocking access or displaying a warning that the connection is not trusted. Some browsers may nevertheless allow users to proceed after an additional warning [5]. By making unauthorized certificate issuance more difficult, CAA can help limit certain forms of website impersonation and phishing. This capability is particularly relevant for government and public-sector domains, where it helps ensure that citizens can trust the authenticity of the information and the web services provided in email content from such agencies.

III. Related Work

Previous researchers have scanned SPF adoption in the wild. Bennett *et al.* [6] scanned mail servers to determine patching related to two critical software vulnerabilities discovered in libSPF2. The libSPF2 software is a popular, open-source library to evaluate SPF records. Those authors found 18,660 mail servers vulnerable to the libSPF2 vulnerabilities. Even after public disclosure, they find that 80% of the mail servers still remained vulnerable. *Tatang et al.* [7] uses the Alexa top 1 million, Majestics, and Tranco lists of popular domains to scan over 2 million domains. Their work briefly discusses findings for federal ".gov" domains: 92% of federal domains implement SPF and 88% of domains implement DMARC.

Deccio *et al.* [8] looked to examine the number of mail servers that use SPF, DMARC, and DKIM to perform validation on incoming email messages. The authors used a dataset of 26,695 domains collected for an unrelated study focusing on vulnerability notification. Of those domains to which email was sent, 22,703 (85%) were found to be SPF-validating, 21,814 (82%) domains were found to be DKIM-validating, and 14,436 (54%) were found to be DMARC-validating. This demonstrates widespread server usage of email validation technologies such as SPF, DKIM, and DMARC.

Ashiq*et al.* [9] scanned the Internet to determine the adoption of Mail Transfer Agent Strict Transport Security (MTA-STS). MTA-STS is an email security protocol by which domains can require TLS encryption and valid X509 certificates for SMTP delivery. Mail server administrators publish their MTA-STS record as a file at the path:

`mta-sts.<domain>/.well-known/mta-sts.txt`. In their study, the authors scan over 86 million domains across four different Top Level Domains (TLDs): .com, .org, .net and the .se country-code TLD (ccTLD). The authors found adoption with MTA-STS, between 3,800 (0.07%) and 7,355 (0.12%) for the .com and .org TLDs respectively. Additionally, the authors find that 29.6% of domains with MTA-STS records in their latest scan have faulty MTA-STS setup.

There is previous work on understanding the adoption of cyber security technologies in governments. Norris *et al.* [10] conducted a survey of 409 local governments from across the United States on a wide range of topics related to cyber security. The authors found that over half of the respondents reported increases in their cyber security budget over the past five years. However, this increase did not apply to other areas such as funding for higher staff compensation, additional staff, training for staff, or policies and procedures. The authors also found many local governments were unaware of and/or had not implemented applicable and highly recommended cybersecurity standards. Previous work examined DNSSEC deployment in federal government domains [11].

Our work differs from prior works in three key ways. First, we include government domains that do not use the .gov TLD. This accounts for 9,094 more government domains. Second, we focus not only on federal but also state, county, and local domains. These often have limited cybersecurity resources, but similarly provide essential services. Finally, we have developed and publicly released a parser and analyzer to verify the syntax of SPF and DMARC records.

IV. Methodology

We present DNSScanner [12], a comprehensive DNS record validation tool that systematically evaluates the security and configuration of SPF, CAA, and DMARC records. The implementation leverages Go (Golang) due to its high-performance networking capabilities and efficient concurrency model.

DNSScanner is extensible to adopt new security resource record types and employs multiple design patterns for flexibility in processing. DNSScanner uses the Chain-Of-Responsibility design pattern for collecting DNS records and analyzing associated records for security. Chain of Responsibility allows passing requests along a chain of handlers. Upon receiving a request, each handler decides either to process the request or to pass it to the next handler in the chain. Each DNS record we collect is collected by a handler specific to that record type. Once a record is collected and analyzed by its handler, it gets added to a list of results which is returned in JSON (JavaScript Object Notation).

For easy use online, a scan can be triggered by visiting the following url, specifying the target *domain*:

https://securegovernment.us/api/*domain*

To parse DMARC records, DNSScanner leverages the *participle* parsing library [13]. It is a lightweight yet powerful parsing tool that facilitates the import of extended Backus–Naur form (EBNF) grammars for text parsing. We have translated the augmented Backus–Naur form (ABNF) [14] specified in section 6.4 of the DMARC RFC into the EBNF notation required by *participle*. This grammatical transformation enables a tight validation of the DMARC specification's syntactic structure as defined in the original RFC.

For SPF records, we developed a dedicated parser that functions as its own library within DNSScanner. This not only performs syntactic analysis, as outlined in Section 12 of RFC 7489, but also conducts semantic analysis of the SPF records' contents. Unlike DMARC, SPF records involve semantic rules which require evaluation of the entire record to verify its conformance to RFC 7489. For example, RFC 7489 specifies that the number of DNS lookups required to assess an SPF record must not exceed ten. This evaluation cannot be accomplished through syntax alone; hence, we implemented a syntactic and semantic analyzer. After reviewing RFC 7489 for SPF requirements, we developed a list of thirteen different validation and security rules. These are listed in Table I.

A. Results Clustering Analyses, statewide

To explore how the security situation might be improved, e.g., by protocol improvements or operator training, we perform additional analyses and share results focusing on common errors and how they might be corrected en masse. Results in Section VI-B focus on a set of domains in the Research & Education Network (REN) community, where there is an available channel of communication amongst the members for sharing solutions and a cooperative trust-based relationship that is characteristic of RENs.

To process and interpret the JSON results pertaining to the REN members, Python scripts were used. These scripts performed several key functions to enable grouping, parsing, visualization, and reporting of DNS and email security configurations.

The *grouping script* performed the following tasks:

- Grouped JSON outputs by configuration similarities across member domains (e.g., similar SPF mechanisms and identical DMARC policies).
- Generated a human-readable summary of the different groups and their associated errors.

The *parsing script* focused on identifying configuration errors and deviations from best practices:

- Parsed each JSON object to detect specific configuration errors or misalignments with best practices.
- Generated human-readable summaries and error reports for each domain.
- Consolidated results into a format suitable for both technical and non-technical readers.

The *graphing scripts* for visualization and presentation of results:

- Parsed each JSON object to detect configuration errors or misalignments with best practices.
- Generated results in tables, charts, and graphs.

The final output included both aggregate findings and detailed error reports per domain, highlighting common mis-

Error	Description
MULTIPLE_HEADERS	Multiple TXT records with the string `v=spf1` are present
HEADER_NOT_FIRST	A TXT record includes the string `v=spf1`, but it is not positioned as the first entry.
UNKNOWN_MECH	An unrecognized string appeared in the SPF record that is not included in the specification
MECH_AFTER_ALL	If an `all` mechanism is present, it must appear after all other mechanisms
PASS_ALL	There is a `+all` which allows all messages to pass SPF
DEPRECATED_PTR	The record contains the `ptr` mechanism which is deprecated and should no longer be used
ALL_WITH_REDIRECT	The `all` mechanism cannot be present with `redirect` modifier
MORE_THAN_10_LOOKUPS	The SPF specification limits DNS queries for evaluation to 10 (including recursive lookups)
UNRESOLVABLE_DOMAIN	A domain in the an `include`, `a`, `mx`, or `redirect` did not resolve when queried.
CIRCULAR_REFERENCE	A domain specified in an `include` or `redirect` has already been resolved.
MECH_AFTER_MODIFIER	A mechanism comes after a `redirect` or `exp` modifier
TOTAL_FAILED_MORE_THAN_2	The SPF specification limits DNS failures for evaluation to 2 (including recursive lookups)
DUPLICATE_MODIFIER	Only one of `redirect` and `exp` modifier can appear in a SPF record

TABLE I: Misconfiguration issues that DNSScanner can detect

configurations to identifying member organizations that may benefit from support or training.

B. SPF Policy Considerations

There is no single SPF configuration suitable for all organizations. However, the version tag `v=spf1` must appear at the beginning of every SPF record, and a default mechanism (`all`, `-all`, or `?all`) must be specified:

- `all`: Soft fail; accepts mail but marks it as suspicious (commonly used during testing).
- `-all`: Fail; rejects unauthorized senders (strongest security).
- `?all`: Neutral; no stated policy (not recommended outside of testing).
- `+all`: Pass; accepts all senders and is strongly discouraged due to spoofing risk.

C. Evaluation Criteria

The following criteria were used to evaluate configurations:

SPF

- Check whether SPF is enabled.
- Identify missing or misconfigured mechanisms.
- Ensure DNS lookups do not exceed 10.
- Detect usage of `+all`.
- Detect use of private IP addresses (unreachable remotely).

DMARC

- Check whether DMARC is enabled.
- Validate the policy value (`none`, `quarantine`, or `reject`).
- If the policy is `none`, verify the presence of a `rua` or `ruf` contact email.
- Check the `rua` tag for a valid reporting address.

CAA

- Check whether CAA is enabled.
- Verify presence of the `issue` tag.
- Verify presence of the `iodef` tag.

After performing these syntactic validations, we clustered similar results according to the criteria outlined above. The resulting clusters were represented for all possible combinations. The clustering methodology and results are described in Section VI-B3.

V. Experiments

Our time-series measurement experiments begin with an initial 8-month scanning period in 2023. We conducted follow-up scans in April 2025 and December 2025 to assess long-term trends and measure whether adoption rates and misconfiguration patterns have evolved across years. To identify the domains for these scans, we utilized two distinct and reputable data sources. The first source is the publicly available repository for the .gov TLD from the Cyber Security and Infrastructure Security Agency (CISA). CISA is the TLD manager for the .gov TLD. The .gov repository provides a comprehensive list of government domains that leverage the .gov TLD. The second source is managed by the U.S. General Services Administration (GSA), which maintains a repository of government domains that do not utilize the .gov TLD. This repository serves as a crucial resource for GSA's all-government search engine hosted on usa.gov. In total, 23,897 domains were scanned throughout our scanning campaigns. The breakdown of domains includes 14,212 Local, 4,063 County, 2,255 State, 1,984 Federal, 301 Independent Intrastate, 373 Special District, 335 Native Sovereign Nation, 69 Quasigovernmental, 92 Interstate, and 213 Unknown domains. Figure 1 presents a heat map of the self-reported states across all domains.

At the start of each month, we refreshed our list of domains to scan using data from these two sources. This ensures we had the current government domain registrations. Once the list was finalized, we initiated the monthly scans to retrieve, store, parse, and analyze SPF, DMARC, and CAA records.

To detect any regional differences in the DNS query responses, we deployed three geographically diverse servers to perform the scans simultaneously. These servers were diversely located in Amazon Web Services' (AWS) data centers, specifically in Virginia, Oregon, and Ohio. All servers utilized Google's Public DNS for querying. Through our analysis, we did not observe any significant variations in the data collected from different regions. The findings presented here are sourced from the AWS server located in Ohio.

For the statewide analysis, we conducted a single comprehensive scan of all 501 member domains of a state-level REN in March 2025. This REN's member institutions are a

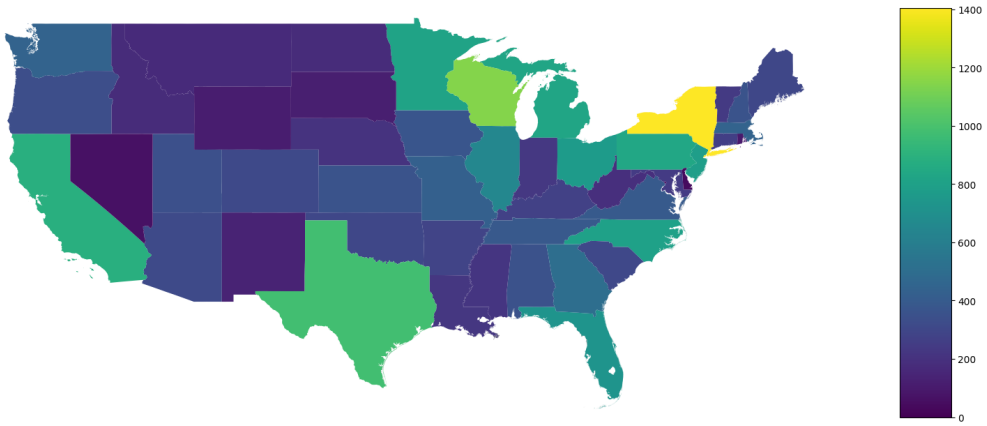

Fig. 1: Geographical heatmap of government domains self-reported locations

mix of educational institutions, municipal governments, and public service organizations. The REN also has established communication channels that enable coordinated remediation efforts among members.

Data collected in this project was limited to publicly accessible DNS records. No sensitive or private data was accessed, stored, or processed.

VI. Results

We provide a comprehensive analysis of the results for the three key security-related DNS records we collected: SPF, DMARC, and CAA. First, we present our time-series findings for government entities at the national level, followed by the analysis of 501 REN member institutions across a state.

A. Government Domains, nationwide

In December 2025, we scanned 23,652 domains, a notable increase from 17,870 domains in February 2023.

1) SPF Adoption: Over the course of our study, the adoption of SPF records has increased, rising from 72.56% (n=12,966) in February 2023 to 76.23% (n=18,032) in December 2025. Notably, SPF adoption across all government types has shown uninterrupted growth, with no month-over-month declines. Moreover, the data reveals a marked shift in the awareness and implementation of SPF records, particularly in sectors that had previously lagged behind in adopting email security standards. Government entities, which have historically faced challenges in embracing new technologies, appear to be making significant strides, with most government types reporting adoption rates around 75%.

Even though total adoption increased, the adoption of SPF varied across different government types. Figure 2 presents adoption by government types from February 2023 to December 2025. We find that SPF adoption has increased between February 2023 to December 2025. Interestingly, the highest increase in SPF adoption among government types is the local, state, and county governments. The observed adoption of SPF among certain government types is a direct result of federal government domains already implementing SPF prior to our analysis period. The Cybersecurity and Infrastructure Security Agency (CISA) issued BOD 18-01 in October 2017, mandating that all federal Executive branch agency domains have valid SPF/DMARC records, with at least a DMARC policy of "p=none." Additionally, BOD 18-01 requires federal Executive branch agencies to secure mail servers by requiring STARTTLS and removing support for SSLv2 and SSLv3. Federal Executive branch agencies were required to comply by January 2018 to implement SPF. However, we did find many executive branch agencies covered by BOD 18-01 that have not implemented SPF. Out of the 287 .gov federal domains that do not have SPF in December 2025, 110 of those domains are federal executive branch domains. This points to lack of compliance with BOD 18-01. Worryingly, many of these domains provide critical services. In addition, many of the domains related to new initiatives also lack SPF. There is still work to do to secure federal civilian executive branch domains that fall under BOD 18-01.

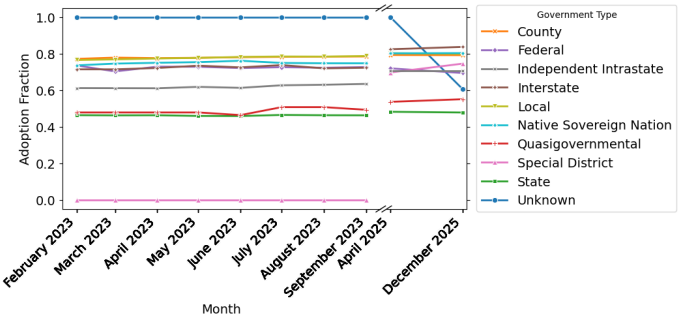

Fig. 2: SPF Adoption By Government Types

2) SPF Misconfiguration: In addition to evaluating SPF adoption, we also tracked the number of SPF issues found in each record. Table I presents possible SPF issues that can be detected during the scan. These misconfigurations create concrete security risks. For example, domains with +all (`PASS_ALL`) policies allow any server to send mail claiming to originate from government addresses, which could allow

Error	Number of Occurrences
MORE_THAN_10_LOOKUPS	4279
DEPRECATED_PTR	2414
UNRESOLVEABLE_DOMAIN	2388
UNKNOWN_MECH	1646
MECH_AFTER_ALL	1014
NO_HEADER	183
PASS_ALL	147
MULTIPLE_HEADERS	107
MECH_AFTER_MODIFIER	75
TOTAL_FAILED_MORE_THAN_2	19
ALL_WITH_REDIRECT	16
CIRCULAR_REFERENCE	0
HEADER_NOT_FIRST	0
Total	12306

TABLE II: Total count of misconfigurations, all scans

for convincing phishing campaigns. Additionally, exceeded lookup limits (MORE_THAN_10_LOOKUPS) cause SPF validation to fail open, negating protection. Unresolved domains produce indeterminate results that many mail servers treat permissively.

Across all the domains analyzed, a total of 12,306 misconfiguration issues were identified. Table II shows the frequency of encountered issues. The most frequent issue was the occurrence of more than 10 lookups, which is likely due to domain administrators being unaware of this SPF requirement. Our scans revealed that many domain administrators included excessive "include" mechanisms, often adding SPF records with their own "includes." Since the lookup limit applies to both the base domain and any included domains, administrators must carefully manage the includes within their SPF records to avoid exceeding the limit.

The second most common misconfiguration found in the SPF records was the inclusion of the deprecated `ptr` mechanism. The `ptr` (Pointer) record maps an IP address to a domain name via reverse DNS lookups. However, its use is discouraged due to inefficiencies, the potential for errors, and the added load it places on DNS servers. Domains that currently use the `ptr` mechanism should instead specify IP addresses directly using the `ip4` or `ip6` mechanisms.

The third most frequent misconfiguration was the inclusion of an unresolved domain in the SPF record. This occurs when a domain in an `include`, `a`, `mx`, or `redirect` statement fails to resolve to an IP address upon query. According to the SPF specification, if more than two unresolved domains are encountered, a `permerror` must be returned. It is then up to the mail server to determine how to handle this `permerror`.

In addition to examining SPF misconfiguration issues, we also analyzed these misconfigurations by government type. The results, broken down by government category, are shown in Figure 3. Interestingly, the prevalence of misconfigurations has fluctuated over the course of the scan. In 2023, issues in local and county government domains were nearly twice as frequent as those found in federal domains. However, in the most recent scan, the error rates in local (3.53%) and county (3.34%) domains have narrowed significantly, now

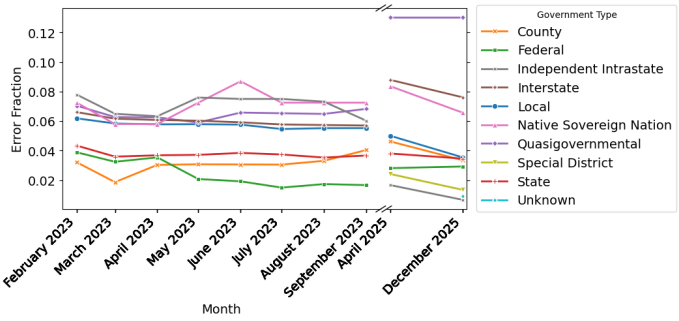

Fig. 3: SPF Misconfiguration by Month and Gov't Type

closely aligning with the error rate in federal domains (2.92%). This shift can primarily be attributed to local and county domains fixing misconfiguration issues in the SPF record. When comparing local and county domains that existed in both the February 2023 dataset and the December 2025 dataset, 703 errors were fixed, while 351 errors were added. This results in a net reduction of 352 errors in local and county domains.

While misconfiguration issues in county, local, and independent interstate domains have decreased over the course of our scan, we have observed an increase in misconfigurations within Native Sovereign Nations, Quasi-governmental, and Interstate domains. These might not indicate trends, though, given the relatively small number of domains in these categories.

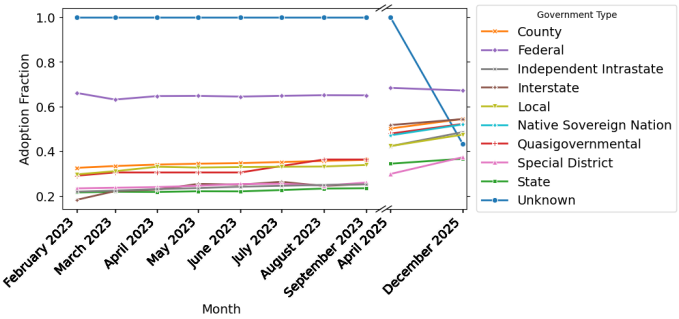

Fig. 4: DMARC Adoption by Month and Gov't Type

3) DMARC Adoption: Throughout our data collection, we observed a consistent increase in DMARC adoption across all domains scanned. This trend holds even when analyzing DMARC adoption by government sector, as shown in Figure 4. The exception to this pattern is seen within federal government domains. Similar to SPF, DMARC adoption was mandated for federal executive domains under BOD-18-01. This mandate likely explains the uniformity in DMARC adoption among federal government domains. Additionally, there is a drop in DMARC adoption for domains which are "Unknown". The drop in DMARC adoption between April 2025 and December 2025 was due to addition of 212 domains which were uncategorized in December 2025.

In addition to verifying the presence of DMARC records, we also assessed the policies defined within them. It's essential

to highlight that the `p` tag in DMARC records enables domain owners to specify the actions that email recipients should take when a message fails validation. There are three policy options available: `none` (no action taken), `quarantine` (the message is directed to Spam or Junk), and `reject` (the message is completely discarded). Of these, `reject` provides the strongest security, ensuring that any email that fails validation is entirely blocked from a recipient's inbox.

In December 2025, among the 11,373 domains with valid DMARC records, 39.91% adopted the `none` policy, 30.31% utilized `reject`, and 29.33% implemented the `quarantine` policy. A closer examination of the `none` policy across different levels of government reveals significant disparities. For instance, of the 1,306 federal government domains employing DMARC, only 6.45% (n=84) opted for the `none` policy. In stark contrast, more than 40% of domains at the state, local, and interstate levels adopted this same policy. Whereas a majority of federal domains follow standards from BOY-18-01, this does not apply to local and state domains. This lack of standardization leaves local and state domain administrators to choose more insecure settings than their federal counterparts.

4) DMARC Errors: While DMARC records are syntactically simpler than SPF records, we developed a dedicated parser to assess the syntactic validity of DMARC configurations. We categorize syntactic DMARC errors into three distinct types: misunderstandings related to publishing an SPF record at the DMARC domain, typographical errors within the DMARC record itself, and issues concerning domain validation at the DMARC domain.

In our December 2025 scan, we found that 167 out of 355 domains with parsing failures were due to an SPF record being incorrectly published at the `_dmarc` subdomain. This misconfiguration likely stems from domain administrators misunderstanding both specifications. This demonstrates the need for clarity regarding the placement of TXT records for DMARC and SPF to prevent future errors.

The second most common parsing error involved DMARC records that correctly included a `v=DMARC1` header but contained typos. This issue accounted for 116 of the 355 parsing errors. Many of these typos were simple errors, such as writing "p-none" instead of "p=none." Additionally, some DMARC records with the `v=DMARC1` header failed to position the policy tag immediately after the header (for example, "`v=DMARC1; p=reject; pct=100`" is valid, whereas "`v=DMARC1; pct=100; p=reject`" is not). These errors further reflect a gap in understanding among domain administrators.

The third category of errors accounted for the remaining 72 of the 355 DMARC parsing failures. Various services require domain administrators to publish a TXT record containing a specific random value to verify domain ownership. For instance, Microsoft requests that administrators publish a TXT record at the domain root in the format "MS=ms*ID*" (where *ID* represents a unique ID from the Microsoft admin center). We observed that these 72 errors stemmed from domains incorrectly placing such TXT verification strings for Microsoft, Google, Apple, DocuSign, and others under the `_dmarc` subdomain. The `_dmarc` subdomain is intended solely for publishing DMARC records; thus, these verification records should not be present. The reasons behind this misplacement by some administrators remain unclear.

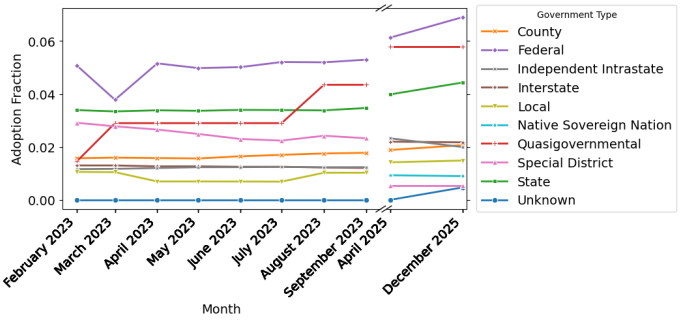

Fig. 5: CAA Adoption by Month and Gov't Type

5) CAA Adoption: Among SPF, DMARC, and CAA adoption, CAA adoption remains the lowest. Figure 5 illustrates the adoption of CAA records across different government types. While BOD 18-01 does not mandate CAA records for executive federal government domains, federal domains show the highest rate of adoption. In contrast, other government types exhibit minimal adoption of the CAA record. We hypothesize this is largely due to a general lack of awareness about the CAA record. Furthermore, email impersonation is more easily recognized by IT administrators, whereas certificate misissuance is a rare occurrence, typically caused by misconfigurations on the part of a Certificate Authority (CA).

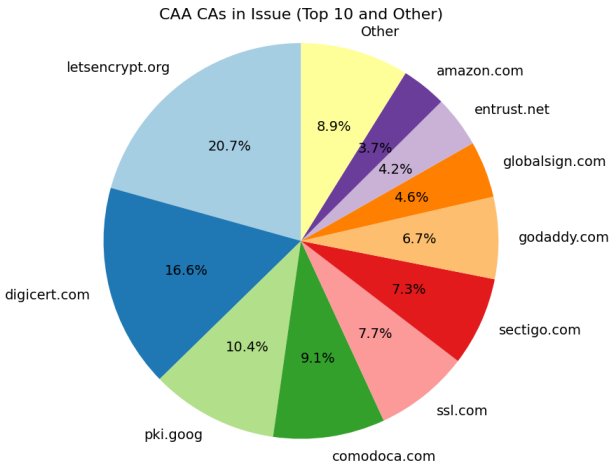

Fig. 6: Certificate Authorities in CAA Issue Field

Additionally, we analyzed the latest scan of domains that implemented CAA, breaking it down by CA. The most prominent CAs are displayed in Figure 6. Notably, Let's Encrypt stands out as the leading CA, accounting for 20.7% of CAs in the CAA issue field. Let's Encrypt is a free, non-profit certificate authority managed by the Internet Security

TABLE III: Adoption and Configuration Errors, statewide

Protocol	% Implemented	% Correctly Configured	Types of Errors, e.g., ...
SPF	83% (418)	61% (305)	Exceed lookup limit, `+all`, syntax issues
DMARC	65% (327)	31% (158)	Missing `p`, `rua`, or having invalid tags
CAA	30% (151)	23% (117)	Missing `issue` or `iodef` fields

Research Group. Let's Encrypt's dominance is expected, given its market share of 64.1% among the top 1,000 websites [15].

B. A REN Members' Domains, statewide

We consider results and more detailed analyses for member institution's domains in one state's non-profit Research and Education Network (REN), scanned in March 2025. Statewide, a total of 501 such domains were scanned using the same tool employed in our nationwide results. Each domain's results, in JSON, were evaluated by a Python script for the presence and proper configuration of SPF, DMARC, and CAA protocols.

1) Adoption: Statewide protocol adoption results are summarized in Table III and in Figure 7(a) showing the percentage of member domains having some form of DMARC, SPF, and/or CAA implemented. Figure 7 further breaks these adoption results into subsets having each Top Level Domain (TLD). In Figure 7, note that "None" indicates that none of the protocols' records (SPF, DMARC, or CAA) are valid.

2) Configuration Errors: In Table III, we also show some types of errors observed. For instance, SPF misconfiguration issues include: (a) exceeded 10 DNS-lookup limit, (b) use of `+all`, which permits spoofing and is strictly advised against, and (c) improper `include:` or `redirect=` chains. DMARC misconfiguration issues were: (a) policy (`p`) tag missing or set to `none` without a monitoring email address, (b) missing `rua` (email recipient for aggregate reports) or `ruf` (email recipient for forensic reports), and (c) syntax errors in record tags. CAA misconfiguration issues were: (a) missing `issue` tag (no CA authorized to issue certificates), (b) missing `iodef` reporting contact email address, and (c) conflicting or redundant flags. To study prevalence of configurations and associated errors, next we turn to clustering our observations.

3) Configuration Clustering: The results were processed using two more Python scripts: the first to group domains by configuration likeness, and the second to analyze and report configuration errors. We grouped domains having identical or near-identical configurations, separating those common configuration pitfalls from outliers among the REN members. The 501 statewide domains were collected into 18 clusters, enumerated in Table IV, based on the similarity of their SPF, DMARC, and CAA configurations, representing all combinations sorted from most to least common. Notice that Clusters 1 and 2 account for more than 77% of all analyzed domains where their misconfigurations could be fixed in two common ways. The top four clusters represent 94% of the domains. This presents an opportunity to convey many common fixes in outreach or training by the REN and its members.

4) Configuration Maturity: Based on the resulting cluster sizes and knowledge of the protocols, Cluster 1 seems to represent a basic configuration having DMARC and SPF properly configured but not CAA. Intermediate configurations represented in Clusters 5, 16, and 18, exhibit only a single, easily correctable error. The most mature configurations – or *most secure* ones – were found in Cluster 8, where domains exhibited zero errors across DMARC, SPF, and CAA. These domains demonstrate best practices in DNS and email security configuration yet, disappointingly, represent only 2% of the domains.

TABLE IV: Cluster Results, 501 Domains statewide

Cluster	Domains	DMARC Status	SPF Status	CAA Status
1	280	OK	OK	Missing `iodef`, `issue`
2	110	Invalid `p`; missing `rua/ruf`	OK	Missing `iodef`, `issue`
3	44	Missing `rua/ruf`	OK	Missing `iodef`, `issue`
4	39	Invalid `p`; missing `rua/ruf`	Missing	Missing `iodef`, `issue`
5	5	OK	OK	Missing `iodef`
6	4	OK	Missing	Missing `iodef`, `issue`
7	4	Invalid `p`; missing `rua/ruf`	OK	Missing `iodef`, `issue`
8	3	OK	OK	OK
9	2	OK	Private IP in SPF	Missing `iodef`, `issue`
10	2	Missing `rua/ruf`	OK	Missing `iodef`
11	1	Invalid `p`; missing `rua/ruf`	Missing	Missing `iodef`, `issue`
12	1	Missing `rua/ruf`	Missing	Missing `iodef`, `issue`
13	1	Invalid `p`; missing `rua/ruf`	OK	OK
14	1	OK	Exceeds 10 lookups	Missing `iodef`, `issue`
15	1	Invalid `p`; missing `rua/ruf`	OK	Missing `iodef`
16	1	OK	OK	Missing `issue`
17	1	Invalid `p`; missing `rua/ruf`	Private IP in SPF	Missing `iodef`, `issue`
18	1	Invalid `p`; missing `rua/ruf`	OK	OK

C. Results Discussion

While adoption of SPF, DMARC, and CAA is significant, both nationwide and in the state studied, many institutions have only a shallow adoption: configurations not yet matured to significantly increase trustworthiness. Many institutions have a configuration that is essentially a "no-op" (no operation) for one or more of the three protocols. For example, a DMARC configuration with "Invalid `p`" (policy) means there is no rejection nor notification when untrustworthy email is received, ostensibly leaving users vulnerable to spoofing. Our subjective observation is that many institutions stop short of effective configurations when those configurations would likely result in increased problem notifications, both false positives and true positives. It may seem like less work to simply pause at only a partial configuration and, unfortunately, remain vulnerable. We recommend that widespread security assessments be performed, perhaps orchestrated by governments and/or non-profit RENs, and that misconfigurations be reported to the affected institutions. Judging from our independently initiated assessment, they will often show that more attention and resources should be directed to improving institutions' trustworthiness online.

VII. Conclusion

In this work, we study the uptake, configuration, and misconfiguration of three authentication protocols that leverage the DNS, namely: SPF, DMARC, and CAA. Specifically, we study these in (a) tens of thousands of government organizations, nationwide across the United States and (b) in hundreds of similarly-trusted entities in one state comprising government organizations; anchor institutions such as libraries; as well as public and private schools. We chose to study government organizations because public trust in them is

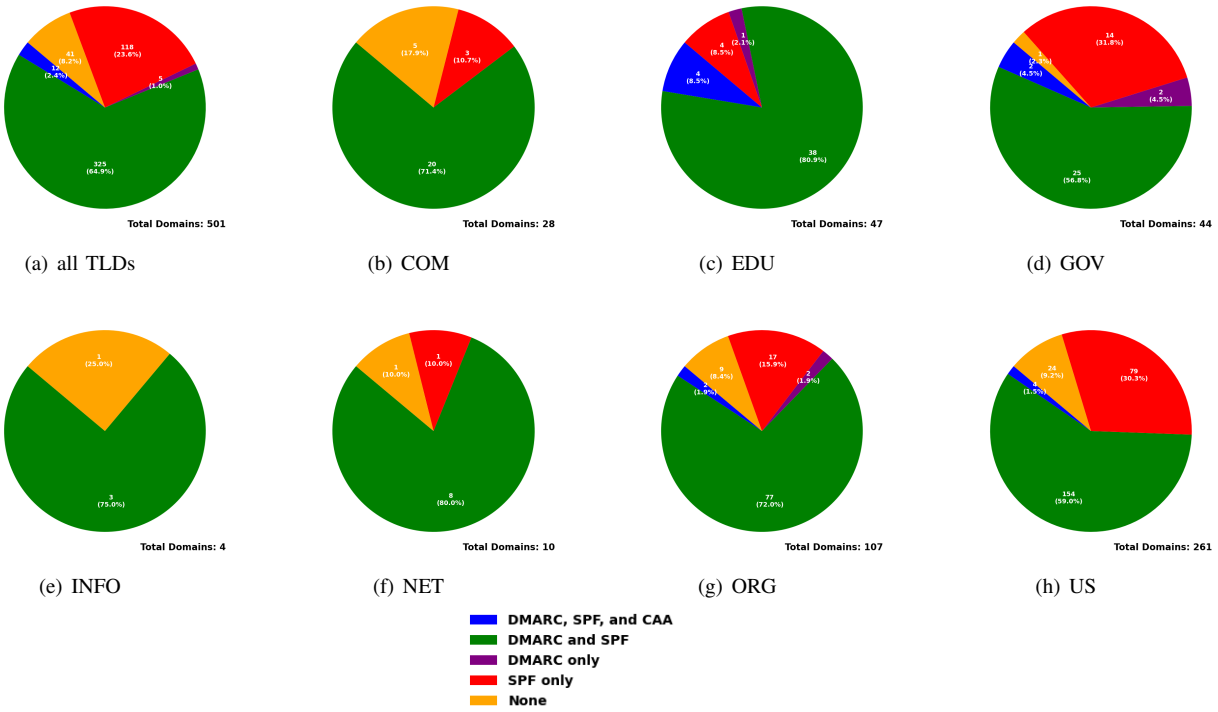

(a) all TLDs (b) COM (c) EDU (d) GOV

(e) INFO (f) NET (g) ORG (h) US

Fig. 7: Adoption by Top Level Domain, statewide

essential, given their responsibility for delivering election, public safety, and health information, and education. We find, nationwide, that government agencies of all sorts collectively show slowly increasing adoption of these security measures. We also find, studying one state, that configuration problems exist in a broad set of institutions that citizens also hope to trust. Our results show that email forgery remains a significant risk and, therefore, remains a vector by which users can be duped into accessing malicious and inauthentic web sites via links in email from unauthenticated entities. However, there is hope: (a) increased adoption and correct configuration coincide with federal mandates and (b) many institutions share the same configuration problems and, thus, the same remedies.

References

[1] U.S. Census Bureau, "2017 Census of Governments," U.S. Department of Commerce, 2017. [Online]. Available: https://www.census.gov/data/tables/2017/econ/gus/2017-governments.html

[2] P. Hallam-Baker and R. Stradling, "DNS Certification Authority Authorization (CAA) Resource Record," RFC 6844, January 2013. [Online]. Available: https://www.rfc-editor.org/info/rfc6844

[3] P. Hallam-Baker, R. Stradling, and J. Hoffman-Andrews, "DNS Certification Authority Authorization (CAA) Resource Record," RFC 8659, November 2019. [Online]. Available: https://www.rfc-editor.org/info/rfc8659

[4] Let's Encrypt. Certificate authority authorization (caa). Let's Encrypt. [Online]. Available: https://letsencrypt.org/docs/caa/

[5] Wikipedia contributors. (2025) Certificate authority. Wikimedia Foundation. [Online]. Available: https://en.wikipedia.org/wiki/Certificate_authority

[6] N. Bennett, R. Sowards, and C. Deccio, "Spfail: discovering, measuring, and remediating vulnerabilities in email sender validation," in *Proceedings of the 22nd ACM Internet Measurement Conference*, ser. IMC '22. New York, NY, USA: Association for Computing Machinery, 2022, p. 633–646. [Online]. Available: https://doi.org/10.1145/3517745.3561468

[7] D. Tatang, F. Zettl, and T. Holz, "The evolution of dns-based email authentication: Measuring adoption and finding flaws," in *Proceedings of the 24th International Symposium on Research in Attacks, Intrusions and Defenses*, ser. RAID '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 354–369. [Online]. Available: https://doi.org/10.1145/3471621.3471842

[8] C. Deccio, T. Yadav, N. Bennett, A. Hilton, M. Howe, T. Norton, J. Rohde, E. Tan, and B. Taylor, "Measuring email sender validation in the wild," in *Proceedings of the 17th International Conference on Emerging Networking EXperiments and Technologies*, ser. CoNEXT '21. New York, NY, USA: Association for Computing Machinery, 2021, p. 230–242. [Online]. Available: https://doi.org/10.1145/3485983.3494868

[9] M. I. Ashiq, T. Fiebig, and T. Chung, "Unraveling the complexities of mta-sts deployment and management in securing email," in *Proceedings of the 2025 ACM Internet Measurement Conference*, ser. IMC '25. New York, NY, USA: Association for Computing Machinery, 2025, p. 1–16. [Online]. Available: https://doi.org/10.1145/3730567.3732916

[10] Donald Norris, Laura Mateczun, Anupam Joshi, and Tim Finin, "Managing Cybersecurity at the Grassroots: Evidence from the First Nationwide Survey of Local Government Cybersecurity," *Journal of Urban Affairs*, vol. 43, no. 8, pp. 1173–1195, April 2020. [Online]. Available: https://doi.org/10.1145/3485983.3494868

[11] S. Rose, "Progress of dns security deployment in the federal government," in *Proceedings of the 26th International Conference on Large Installation System Administration: Strategies, Tools, and Techniques*, ser. lisa'12. USA: USENIX Association, 2012, p. 223–228.

[12] Alex Gebhard and contributors, "DNSScanner," 2025. [Online]. Available: https://github.com/a1g3/dns-scanner

[13] Alec Thomas and contributors, "A parser library for Go," 2025. [Online]. Available: https://github.com/alecthomas/participle

[14] D. Crocker and P. Overell, "Augmented BNF for Syntax Specifications: ABNF," RFC 5234, Jan. 2008. [Online]. Available: https://www.rfc-editor.org/info/rfc5234

[15] W3Techs, "Usage statistics and market shares of ssl certificate authorities for websites," 2025. [Online]. Available: https://w3techs.com/technologies/overview/ssl_certificate