# Security In Plain TXT
## Observing the Use of DNS TXT Records in the Wild

Adam Portier[1], Henry Carter[1], and Charles Lever[2]

[1] Villanova University {aporti01, henry.carter}@villanova.edu
[2] Georgia Institute of Technology chazlever@gatech.edu

**Abstract.** The Domain Name System is a critical piece of infrastructure that has expanded into use cases beyond its original intent. DNS TXT records are intentionally very permissive in what information can be stored there, and as a result are often used in broad and undocumented ways to support Internet security and networked applications. In this paper, we identified and categorized the patterns in TXT record use from a representative collection of resource record sets. We obtained the records from a data set containing 1.4 billion TXT records collected over a 2 year period and used pattern matching to identify record use cases present across multiple domains. We found that 92% of these records generally fall into 3 categories; protocol enhancement, domain verification, and resource location. While some of these records are required to remain public, we discovered many examples that unnecessarily reveal domain information or present other security threats (e.g., amplification attacks) in conflict with best practices in security.

**Keywords:** DNS · TXT Records · Security Protocols

## 1 Introduction

The Domain Name System (DNS) is a central piece of infrastructure that is relied upon by nearly every application on the Internet. Over its existence, it has grown beyond its motivating purpose as an IP address lookup directory, and is now used for applications including email routing, authentication, and cryptographic key repository, among others. By default, this information is made public, and may result in unintended information leakage.

As DNS has expanded its functionality, new record types have been developed and approved as parts of the DNS standard. Many of these applications embed information in TXT records, which were included in the original DNS specification to allow the storage of arbitrary text strings [32]. While some of these TXT-based applications are formally specified (such as SPF [41], DKIM [4], and DMARC [30]), many nonstandard applications have been developed without security vetting or technical review. With many applications using TXT records for information exchange, particularly in cloud computing, it is not clear how widely such records are used or if they are being used in a way that allows for malicious abuse.

In this paper, we take a first broad look at how TXT records are being used in production domains. Using a DNS record capture of 1.4 billion TXT records collected over the past two years, we first categorize and filter applications using known TXT record formats, after which we identify applications using the most common structured TXT records that did not match a known format. We found that 92% of these TXT records can be categorized into one of three application types: protocol enhancement, domain verification, or resource location records. We then show that only about 6% of these TXT records are deployed with DNSSEC verification and that the vast majority of these records reveal significant information about the infrastructure of a given domain. The public availability of this information makes developing a targeted intrusion or spear phishing attack significantly easier to mount.

Our work makes the following contributions:

- **Broad TXT categorization:** we collect and identify the most common applications for a representative set of TXT records, then develop a categorization that captures how TXT records are broadly used on the Internet.
- **Analysis of deployment:** Along with our categorization, we determine what fraction of these records are authenticated using DNSSEC, analyze the security level of deployed email verification policies, and measure TXT record size and entropy across our data set. Our measurements provide insight into how the most common TXT records are configured, and highlight outliers representing unusual or unsafe applications.
- **Security implications:** We conclude with a discussion of the security implications of our findings. Several observed applications are vulnerable to cache poisoning attacks without DNSSEC protection, leak information about a domain's infrastructure, or present an opportunity for other protocol-based exploits such as amplification attacks.

The remainder of our work is organized as follows: Section 2 outlines related research, Section 3 describes our data set and categorization methodology, Section 4 defines the TXT record applications and their observed usage, Section 5 discusses security implications, Section 6 provides concluding remarks.

## 2 Related Work

DNS was originally intended to map domain names to service locations, but the scope has expanded to include many use cases beyond this. As part of the original DNS specification [32], TXT records were specified as arbitrary text up to 255 characters in length. Subsequent protocol changes in EDNS(0) have extended the format to allow strings of any length. While the standard does not define a TXT record structure, RFC 1464 later proposed a "`key=value`" format, which is commonly (but not exclusively) used in practice. The permissive nature of the DNS TXT record has allowed for a wide variety of applications, including some that have gone through a formal RFC process and are used to combat spam email. These include Sender Policy Framework (SPF) [41], which provides

a list of servers allowed to send email on behalf of a domain. Related to this are DomainKeys Identified Mail (DKIM) [4] records, which use public key cryptography to validate email headers. Finally, Domain-based Message Authentication, Reporting and Conformance (DMARC) [30] records provide suggestions to mail transport agents (MTA) for what to do with email failing other checks. These protocols are well documented because of their origins in the RFC process, and security research has shown that they can provide strong security guarantees if they are configured correctly [20, 23]. However, many more informal and unverified use cases exist, which will be explored in this paper.

There have been several vulnerabilities discovered in the DNS protocol itself. The most well known is Dan Kaminsky's 2008 work on cache poisoning attacks, and the push toward DNSSEC that followed [24]. The introduction of DNSSEC records has increased the size of many DNS responses, and has in turn enabled other protocol abuses such as amplification attacks [1, 11]. Other common vulnerabilities rely on human error, such as various forms of "squatting" (typosquatting [46, 2], bitsquatting [17, 35], soundsquatting [34] and combosquatting [25]) as well as domain "parking" for serving advertisements [48]. The abuse of normal recursion paths for purposes of censorship [37, 49] and advertising [47] have also been observed. While these studies focus on protocol vulnerabilities, there has been much less attention given to how DNS TXT records have been used, or in some cases, misused. For example, large TXT records have been observed in past amplification attacks [3]. While techniques such as response rate limiting (RRL) and TCP fallback [8] have been developed to reduce the risk of DNS amplification attacks, the inconsistency of implementing these techniques across resolvers leaves TXT-based amplification attacks a possibility [31, 43]. Unfortunately, no broad measurement of the presence of large TXT records has been performed to determine whether these attacks are still viable in practice.

Several studies have been performed to try and measure how DNS is being used. A significant number of these studies have looked at DNSSEC adoption rates and configuration [36, 22, 38, 12, 13, 45, 27]. Other studies have quantified the deployment of formally specified DNS records and protocols (DANE, CAA, and CT) used for verifying TLS certificates issued to a domain [6, 43, 40]. Finally, several studies have measured DNS deployment and misuse by examining open resolvers [15], modified recursion paths [47] and typo-squatting [46]. We identified three prior studies specifically looking at how TXT records are being used in the context of email security. A 2007 study by Stefan Görling did an investigation into SPF adoption rates in the Sweden country code TLD (.se) [21]. In 2015, Durumeric et al. [19] performed a study examining how email delivery security is implemented in SMTP servers in the Alexa top 1 million domains, which included measurements of SPF, DKIM, and DMARC TXT record deployment. Most recently, Szalachowski and Perrig performed a deployment study of DNS-based security protocols in the Alexa top 100k, including both TXT record based email security protocols as well as protocols which have their own DNS record types specified [43]. Our work seeks to provide a broader view of TXT record applications and deployment including both formally specified and otherwise.

**Table 1.** ActiveDNS Dataset June 2016 - May 2018

| RR Type | RR Count |
|---|---|
| TXT | 1,410,219,403 |
| MX | 1,784,771,811 |
| RRSIG | 338,693,718 |
| Total | 3,533,684,932 |

## 3 Methodology

Given the versatility of the TXT record, we planned to answer three questions related to their use. First, we expected to gain some insight into how much the formally defined uses of TXT records were used in practice. Second, we expected to discover many informal uses of TXT records and to tie these back to a service that consumes this record. Finally, we expected that the permissive nature of the TXT record format to be misused in some way, and introduce new vulnerabilities that did not exist previously for the domain operator. The objective of this work was to not only catalogue how these records are being used, but also to observe potential drawbacks to their use and make some recommendations accordingly.

The study was conducted using the publicly available Active DNS project dataset [26] run by the Georgia Institute of Technology. Their collection infrastructure performs an active DNS scrape once per day of every domain and record type they are able to resolve based off of a growing list of approximately 400 million seed domains. These seed domains are compiled from a combination of sources, including the Alexa top 1 million, the TLD zone files for COM, NAME, NET, ORG, and BIZ, sites captured by the Common Crawl project, multiple public domain blacklists, private security vendor lists, and other popular domain lists. Unfortunately, since their domain list does not contain all subdomains for each seed domain, we cannot observe records in highly specific subdomains (such as ACME records [7]). However, the breadth of coverage for domains and common subdomains still yields significant insight into the broader application of TXT records. From this dataset, we scraped unique resource records (RRs) observed every month from an average of 48 million responding domains per month. The aggregate record counts collected for this work are shown in Table 1.

We then developed a regular expression-based pattern matching filter by investigating occurrences common across domains in the top 10,000 domains as ranked by the OpenDNS platform top 1 million list [14]. The filter was implemented using regular expressions in Python, applied in a series against the string value of the record data until a match was found. Records that had standard documented formats were classified as such. Records that did not match this format were then matched on two common attribute formats: "key=value" and "key: value". Any patterns occurring on at least 3 records were manually inspected and assigned a corresponding application label using publicly available documentation to identify the service that uses the record and its purpose.

**Table 2.** Taxonomy Counts

| Category | RR Count | Percent | Apps |
|---|---|---|---|
| Protocol Enhancement | 1,080,278,464 | 76.60% | 5 |
| Domain Verification | 220,168,210 | 15.61% | 43 |
| Resource Location | 9,961 | 0.00% | 4 |
| Unknown | 109,762,768 | 7.78% | |
| Total | 1,410,219,403 | | 52 |

Finally, any remaining records that did not have common format, but rather common structure, were collected. These include records that were arbitrary integer strings, hexadecimal strings, Base64 encoded data, and so on. Records that did not mach any other classifier were classified as "unknown". This procedure was repeated until no further examples of 3 or more matching records could be found. The order in which regular expressions were tested was rearranged every time a new pattern was introduced so as to avoid incorrect matches. This filter was then used to label the full data set of 1.4 billion records using a combination of Hadoop and Python, using the same filter logic developed on the domain subset. We calculated record usage frequency across the entire set of records and performed a configuration analysis for the RFC-formalized record types.

While previous work has shown that top domain lists can skew statistical results related to domains [39], we found that the regular expression patterns identified in the top 10,000 domains matched 92% of the records in the larger dataset, showing that these applications are by far the most commonly found overall. A similar process used in developing the filter could have been applied to the larger dataset to identify use cases not present in the 10,000 domains; however, each new identified pattern applied to an ever smaller number of records, and documentation was generally not available. We instead used size and entropy analyze to identify trends in these long-tail records.

## 4 Taxonomy

We examined 1.4 billion records and identified 52 unique applications of TXT records. The most diverse category of these applications is Domain Verification, which are TXT records used to prove the ownership of a domain namespace for use in a Software as a Service (SaaS) solution. This is followed by the Protocol Enhancement records, which are used to enhance the security of another protocol by verifying information in DNS (and account for the largest raw count of records). Finally, we found 4 applications of Resource Location records, which are used to pass the location of another resource using DNS to a client. Of the records analyzed, roughly 8% either did not have a discernible pattern that met our criteria, were misconfigured examples of other patterns, or did not have enough unique occurrences for consideration (See Table 2). However, to determine how these unknown records may differ from more popular applications and

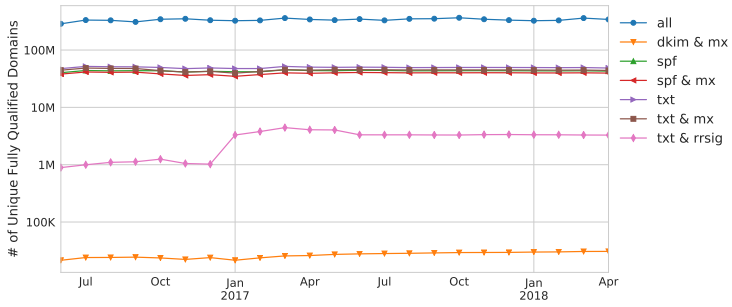


**Fig. 1.** The count over time of domains in our data set posting protocol enhancement TXT record types. We observed a sharp increase in the number of domains posting RRSIG records in January 2017, the vast majority of which were in the Western Samoa (.ws) TLD. Note the log-scale on the y-axis.

their potential for use in amplification attacks, we categorized the entropy and record size for all unknown records and compared against the distributions for known applications. To determine how records change over time, we examined unique record counts by month for the most common protocol enhancement and domain verification records. Because the monthly counts for these records were very consistent, we quantify the rest of our results in aggregate and note anomalous trends where present.

### 4.1 Protocol Enhancement

Records that fall into this category provide some form of security enhancement for another protocol or application. The general use case for these records is roughly the same: when a server receives a message from a client using another protocol, the server makes a DNS query to obtain information to prove the validity of that message. The most common use of this method is to verify email messages from a domain. All of the formally documented applications for TXT records fall into this category. The most common Protocol Enhancement record is SPF, which is also the most common TXT record found overall. Other applications in this category include SenderID, DKIM and DMARC, as well as a Base64 key signature used by Active Directory federations, for a total of 5 applications. In addition, related work in TXT usage has been almost exclusively focused on these record types. For each of the RFC-specified record types (SPF, DKIM, and DMARC), we compare our results to three other measurement studies performed over the past 12 years [21, 19, 43].

**SPF Records** The most commonly observed application of Protocol Enhancement records is SPF. A 2007 study by Görling et.al. did an investigation of the adoption of SPF records 1 year after the RFC was adopted. The study was very limited in scope, only investigating Swedish country code (.se) domains [21].

**Table 3.** SPF Adoption

| | **2007**[21] | **2015**[19] | **2017**[43] | **2016-2018** |
|---|---|---|---|---|
| Domains Considered | 385,862 | 1,000,000 | 100,000 | 336,963,348 (mean) |
| Domains with SPF | 6,286 | 401,356 | 53,365 | 41,432,865 (mean) |
| Domains with MX | 330,163 | 847,056 | - | 95,744,788 (mean) |
| % with SPF | 1.63% | 40.14% | 53.37% | 12.85% |
| % with MX and SPF | 1.9% | 47.38% | - | 42.78% |

They found that overall adoption was very low (See Table 3), and the majority of domains publishing an SPF record did so with a rule that failed to adequately enforce forged message origin, either using a neutral or soft fail qualifier on the `all` mechanism. Of the 1.4 billion TXT records examined in this study, 76% of them were SPF records.

In more recent work, Durumeric et al. [19] and Szalachowski and Perrig [43] both measured the appearance of SPF records in the most popular domains on the Internet (the Alexa top 1 million and top 100 thousand, respectively). Their results indicate that SPF usage over the most popular domains has significantly improved, with 40.14% of the top 1 million (and 53.37% of the top 100k) domains posting SPF records. Furthermore, Durumeric et al. show that in April 2015, 92% of Gmail's inbound email messages came from domains using SPF, which indicates that all of the major email providers (such as Yahoo, Outlook, and others) are employing SPF records to verify their email servers.

The measurements that we collected from a significantly larger set of domains demonstrate improvement in how SPF is used, but that this improvement is largely concentrated within the most popular domains online (see Figure 1). Adoption across the entire population of domains is still fairly low (12.85%). However, the adoption of domains with at least 1 MX record (meaning they are configured to accept email) increases the percentage to 42.78%, which is still lower than the 47.38% observed by Durumeric et al. in the top 1 million domains (See Table 3). This further reinforces the trend that the Durumeric study claims: while the most popular email domains are employing SPF records, adoption steadily decreases as more domains in the "long tail" of less popular services are considered. In our dataset, we also observed many domains with SPF records and no MX. A domain being used only for marketing may do this because it will send email but not receive it. Other domains may wish to indicate they will never send email and to treat any message from it as spam. Similar to SPF records are SenderID [29] records, which are limited to Microsoft Exchange Servers. These records are used much less, making up only 0.26% (3,696,073) of all TXT records observed.

Also of note is how the SPF qualifier for the `all` mechanism has changed over time. An SPF record may specify one of four policies for email: pass (for verified email messages), neutral (verification inconclusive), soft-fail (deliver but treat as "suspicious"), and hard-fail (do not deliver). The majority qualifier has

**Table 4.** SPF Operators

| | **2007**[21] | **2016-2018** |
|---|---|---|
| Domains with SPF | 6,286 | 41,432,865 (mean) |
| a | 6,192 (98.5%) | 14,100,368 (34.03%) |
| mx | 2,296 (36.5%) | 12,137,863 (29.30%) |
| ip4 | 1,573 (25.0%) | 10,839,731 (26.16%) |
| include | 710 (11.3%) | 18,371,446 (44.34%) |
| ptr | 350 (5.6%) | 2,758,898 (6.66%) |
| exist | 3 (0.05%) | 3 (0.00%) |
| ip6 | 0 (0.0%) | 7,049,348 (17.01%) |

**Table 5.** SPF Policies

| | **2007**[21] | **2015**[19] | **2016-2018** |
|---|---|---|---|
| Domains with SPF | 6,286 | 401,356 | 41,432,865 (mean) |
| Neutral | 3,430 (54.5%) | 80,394 (20.03%) | 7,515,050 (18.14%) |
| Soft fail | 775 (12.3%) | 226,117 (56.34%) | 21,264,822 (51.32%) |
| Hard fail | 1,233 (19.6%) | 84,801 (21.13%) | 15,965,363 (38.53%) |
| Pass | unknown | 10,045 (2.50%) | 106,606 (0.26%) |

shifted from neutral in 2007 (54.5%) to a soft-fail in 2015 [19] (56.34%), to a mixture of a hard (34.45%) or soft-fail (45.85%) in our study (see Table 5). This indicates an increased trust and reliance on SPF records to combat spam, but the contrast with the Durumeric study indicates that more popular domains favor the permissive nature of the soft-fail over rejecting delivery. The overall use of verification mechanisms has changed as well (see Table 4). In 2007, the most common mechanism type was `a` (98.5%), meaning single servers were responsible for sending email. In our study, the most common mechanisms are split between `include` (39.66%), `a` (30.41%), and `ip4` (23.36%). This indicates a shift away from single server solutions to the use of larger networks and hosted solutions to send email.

**DKIM Records** DKIM records are used to post public signing keys for a domain, which can then be fetched and used to verify incoming mail signed by the sending domain. As opposed to SPF or SenderID records, which must reside in the domain's apex, DKIM records are only fetched by an email MTA (Mail Transport Agent) once the message is received, and as directed by the DKIM-Signature header field. As such, these records are resolved less frequently than SPF, and are likely under represented when gathered using Active DNS. Because of this, DKIM records made up only 0.05% (657,458) of all records collected, representing a lower bound on potential deployment. Since these records are only relevant to domains that are configured to exchange email, we then considered the number of domains with MX records that also post an accompanying DKIM record (see Table 6). Even in this reduced set of domains, only 0.06% posted DKIM

**Table 6.** DKIM and DMARC Adoption

| | **2015**[19] | **2017**[43] | **2016-2018** |
|---|---|---|---|
| Domains Considered | 1,000,000 | 100,000 | 336,963,348 (mean) |
| Domains with MX | 847,056 | - | 95,744,788 (mean) |
| Domains with DKIM | - | 5,049 | 28,585 (mean) |
| Domains with DMARC | 8,890 | 7,361 | 33,224 (mean) |
| % with DKIM | - | 5.05% | 0.0085% |
| % with DMARC | 0.889% | 7.36% | 0.0098% |
| % with MX and DKIM | - | - | 0.0536% |
| % with MX and DMARC | 1.000% | - | 0.0439% |

**Table 7.** DMARC Policies and Notification Settings

| **Return Address** | **Count** | **Record Policy** | **Count** |
|---|---|---|---|
| rua only | 16250 (48.91%) | reject | 4801 (14.45%) |
| ruf only | 254 (0.76%) | quarantine | 2403 (7.23%) |
| both | 11891 (35.79%) | monitor | 276 (0.83%) |
| neither | 4828 (14.53%) | none | 24996 (75.23%) |
| Total | 33224 | Total | 33224 |

records along with an MX record (see Figure 1 for a monthly count). When compared to the Szalachowski and Perrig study [43], we again observed a pattern of deployment concentrated more in popular domains, with diminishing use in less popular services. While Durumeric et al. did not measure the appearance of DKIM records through the top 1 million domains, they did observe that 83% of the messages received by Gmail in April 2015 contained a DKIM signature, confirming that the most popular email service are cryptographically verifying email. The low occurrence rate of DKIM in less popular email domains could be related to inefficiency in running public-key cryptography at scale or a lack of email solutions that support it. Further passive collection of DNS records exchanged during SMTP sessions will be necessary to offer a definitive answer. Related to DKIM records are the fixed-length Base64 key signature records used by Microsoft's Active Directory in establishing a multi-domain trust. These records make up 0.64% (9,010,935) of all TXT records examined, and indicate that the use of DNS as a cryptographic key repository is highly application-dependent.

**DMARC Records** DMARC records are used to provide more granular control for email failing other verification methods (such as SPF and DKIM), as well as inform domain operators of how much email is failing checks. Across previous measurement studies, deployment was very low, with 7.36% of the top 100k domains and 0.889% of the top 1 million domains posting DMARC records. Our study further demonstrates the trend of diminishing deployment in unpopular domains, with only 0.05% (764,148) of all domains posting DMARC records.

Furthermore, in both our study and Durumeric et al., most of the DMARC record rules provide no rules for handling invalid email, but do request reports for domain owners about failure rates. DMARC allows for two possible reporting addressed to be listed: an `rua` address for reporting aggregate statistics about authentication checks, or an `ruf` address for sending forensic reports. As shown in Table 7, 84.7% of DMARC records were configured to at least receive aggregate reports, while only 22.51% of policies listed *any* policy for how to handle mail that does not pass authentication checks (reject message, quarantine message, or monitor message). This is likely due to two reasons. First, the DMARC website recommends initially rolling out DMARC policies with only reporting rules enabled [18]. Second, as observed by Durumeric et al., mailing lists frequently modify mail in transit, which will invalidate DKIM signatures and make publishing a DMARC reject policy problematic for popular mail services. This reserved approach to rolling out DMARC also mirrors early stages of SPF deployment, where neutral qualifiers for `all` mechanisms were the prevalent.

### 4.2 Domain Verification

These records are used when signing up for a SaaS cloud solution to handle some piece of an organization's infrastructure. Most of the non-standard record uses fall into this category. Since here are many applications using records for the same purposes, we examine records by category. In general, use of an implementing service requires proof of ownership of a domain by the customer when registering an account. The service generates a random alphanumeric string that the domain operator must insert in their DNS authoritative server as a TXT record. The service then checks for the existence of this record, and when it is found, the verification step is complete. These records all have some form of identifier, then a randomized value that follows a pattern established by the service. Our ability to verify that any particular record is associated with a service is somewhat limited, as public documentation providing this proof is sometimes not available. In addition, few of these services' provide documentation for when records can be removed. In total, 43 such applications we identified, comprising 15.6% of the TXT records we observed.

**SaaS Infrastructure and Applications** Many applications related to SaaS infrastructure require proof of domain ownership. Four of the top five domain verification applications (see Figure 2) fell into this category: GSuite (formerly Google Apps) with 5.4% (75,633,895), Microsoft Office365 with 1.8% (26,014,855), Zoho email hosting with 0.09% (1,221,180), and Microsoft Outlook email hosting with 0.06% (903,999). Over the past two years, Zoho has seen a consistent rise in the number of deployed domains, while Outlook verification records have steadily decreased. Both GSuite and Office365, while suffering a drop in the number of records in October 2016 (immediately after Google Apps were renamed to GSuite), have seen a steady increase in deployment since. Other less common applications included shared resources such as Cisco's WebEx, website

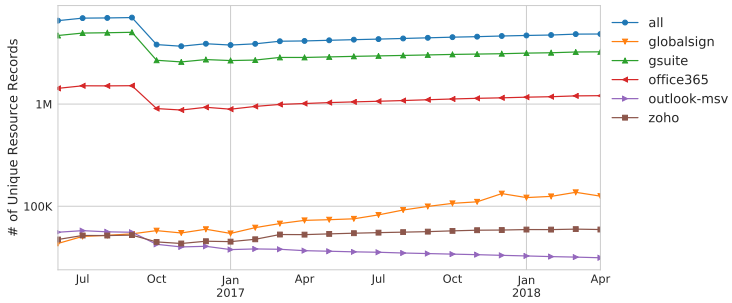


**Fig. 2.** The count over time of TXT records in our data set for the top five domain verification applications. Note the log-scale on the y-axis.

hosting, and vanity domains or custom URLs. Since these services are providing critical infrastructure to an organization, the service requires the domain operators to verify their control as part of the registration process. This is done to prevent fraudulent account registration or misuse of a trusted domain name. Additionally, we found one example of a service (StatusPage) requiring domain verification before emails from a custom domain could be sent. This is required in addition to the normal sender validation steps of adding an SPF and DKIM record. They require a domain verification step in order to set a custom `FROM` address in emails used to communicate a domain's service outages with customers [42]. The service also checks for this domain verification record every hour, so removal of it will cause the StatusPage alerting system to stop working.

**Security and Identity** We identified 12 SaaS products that provide security and identity services requiring domain verification. The third most frequent domain verification record we observed belonged to the cloud-based certificate authority (CA) Globalsign (0.14%, 1,910,055), which had a steadily increasing number of records observed over the collection period. We observed significantly smaller counts for services providing "Verified" status for accounts on services like Facebook or Docusign. We also identified a few services for providing or consuming SAML logins that required an additional domain verification step, such as the case with Adobe's IDP or Atlassian's Confluence Wiki. Since these services rely on validating identity, domain verification is used as a means to bootstrap trust from DNS control. Normally, the exchange of cryptographic keys and HTTP service endpoints (metadata) is sufficient to establish trust, but some services require additional verification. This use of DNS as a "trust anchor" has been studied in relation to expiring domains and stale records, and has been shown to cause problems [28, 9].

**Domain Scanning** There were 9 cloud services identified that required domain validation before they would scan a domain's infrastructure (0.01%, 129,378 records across all 9 applications). The validation is a necessary step to prevent

the service from being used to generate denial of service attacks or expose vulnerabilities to third parties. This includes cloud-based load testers such as Blitz.io and Loader.io, and vulnerability scanners like Detectify and Cloudpiercer. These services generally check the records on demand when the scan runs. Some scanners, such as Botify, perform Search Engine Optimization (SEO) by mimicking search engines' scanning algorithms in order to improve visibility. We also found scanners that look in repositories of stolen account credentials for data belonging to a particular organization, as is the case with "Have I Been Pwned". Since the information being retrieved is very sensitive, the service requires domain operators to prove ownership.

**Advertisement Monetization** We found 2 applications of domain verification as an important step when setting up a monetization agreement between a cloud-based service and a particular organization's domain. The Brave web browser blocks advertisements by default, but allows users to pay domain owners directly with cryptocurrency if the domain is registered with Brave. This registration requires domain verification. Additionally, Dailymotion is a video hosting platform that allows users to embed videos in other sites and collect revenue from the ads shown in those videos, but only once the domain operator has proven they own the domain in their Dailymotion account. These records only accounted for 0.001% (20,590) of the total dataset, indicating that this application is still not widely adopted.

Finally, we identified a pattern of fixed-length hexadecimal records with no identifier. As best we can determine, they are used for verification of domains being "parked" for advertising purposes. This is based off an analysis showing that these records frequently appear as the only TXT record for the domain (95%), and are frequently seen with an NS record pointing to a parking service. These records alone comprise 7.86% of the total number of TXT records collected, and sometimes are present multiple times for a single domain. This is consistent with other research being performed into domain parking, where domain operators wishing to park a domain for monetization are required to prove ownership during registration with the monetization service [5]. Over all the domain verification applications we observed, most of these are poorly documented and all reveal significant amounts of information about a domain's internal infrastructure.

### 4.3 Resource Location

A Resource Location record is one in which the information contained in the TXT record points to the location of some other application or service. Presumably this information is being passed in DNS because it is convenient, as the service owner wishes the clients to be able to locate the resource with only a DNS request. Only a few examples of using TXT records to reference another resource were identified. These were Red Hat's JBoss Fuse server (212), Symantec's Mobile Device Management product (4,909), Ivanti's Landesk Application

(3,536), and client configuration suggestions for Bittorent (1,516), for a total of 4 applications and less than 0.001% of the records gathered. We also observed several instances of human-readable location "comments" in use for domains with multiple data centers. While these records did not have any regular pattern, they are clearly being used as implicit resource location identifiers, and as such our count of resource location records should be treated as a lower bound.

Both Ivanti Landesk and Symantec's MDM assist with the management and identification of mobile devices on a corporate network. The endpoint of the master server must be made publicly available in a TXT record. In the case of Symantec's MDM, this may have been in error, as none of the endpoints tested responded to external `curl` HTTP probes, and several were using private IP space IPv4 addresses. In the case of Landesk, this seems to be deliberate. All endpoints tested here responded with either empty HTTP replies, or in a few cases, service health information, without providing any login credentials.

Red Hat's JBoss Fuse product is used in the development of API driven microservices. One of the connectors offered by the product is a JDBC endpoint, allowing applications to interact with information stored in a database using XML and SOAP. Although we were unable to locate any documentation recommending it as a best practice, several domains were identified publishing the location of their JBoss Fuse Server in TXT records. A few of the JBoss Fuse endpoints responded to external HTTP `curl` probes, but did not return any useful information. In all of these resource location applications, the exact location and product used is made publicly available and therefore trivializes the information gathering step of an attack.

One interesting use case of Resource Location TXT records in Bittorrent client automatic configuration. This is done in order to prevent misconfigured torrent trackers that publish a URL as the tracker location from flooding websites with Bittorrent requests. This information was formalized in BEP 34 [33]. Domains running Bittorrent trackers are able to publish TXT records so that clients can discover the actual location of the tracker service located at that domain. This is also useful if the tracker has to change ports or locations after torrent files have been distributed, solving an automatic discovery problem for an otherwise peer-to-peer network. This record format can also be used to inform clients no tracker is present in the domain.

### 4.4 Unknown Records

The remaining 8% of TXT records that belonged to unidentified applications presented a challenge for manual analysis, as many of the records appear without a discernible pattern and appear in very few domains. However, we wanted to ascertain whether the text in these records tended to be structured or random, as well as their potential for use in amplification attacks. To identify unusual trends in these records, we computed the Shannon entropy over characters in each record and length of each record in characters. We then compared these results to known use cases to determine how they differ from typical TXT usage. To perform the comparison, we categorized all records into bucketed ranges of

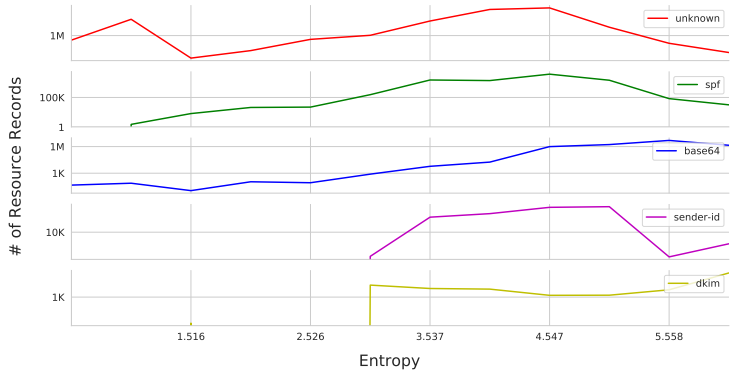


**Fig. 3.** The five representative applications with the lowest Shannon entropy. All other applications had entropy of approximately 3.5-6. Note the log-scale on the y-axis.

entropy and record length, with each bucket containing records with entropy within approximately 0.5 and record lengths within approximately 45 characters.

**Entropy** We observed that most known applications yielded records with entropy in the range of 3-6 (see Figure 3), with some domain verification records containing random tokens falling in a higher range (such as Google site verification records). However, when compared with our collection of unknown records, we noted a large spike in the number of low entropy records when compared with known applications. Upon investigating the contents of these low-entropy records, our analysis showed either records containing `‘‘~’’` or double quotes surrounding a string of a single repeated character (e.g., `‘‘nnn’’`). While we were unable to identify any use for the repeated character strings, we speculate that the `‘‘~’’` records could be a misconfiguration of SPF, which uses the tilde to indicate a “soft fail”. Beyond these examples, we found that the unknown records tended to follow the entropy trends of the rest of our known applications.

**Record Length** The vast majority of records we collected had a length of 500 characters or less, with all known applications (except SPF) having a maximum record length around 1,000 characters (see Figure 4). For SPF records, we found a consistent decrease in the number of records as the length increased, but still found a large number of samples containing long lists of allowed IP addresses, up to a maximum of approximately 3,750 characters.

In the unknown record category, we observed a surprising number of very large TXT records up to approximately 4,500 characters long. All of the records we inspected in this set belong to domains in the `.tel` TLD [44]. Originally designed as a DNS repository for telephone contact information, the TLD now allows domain owners to store arbitrary files in DNS as TXT records. The service offers storage of PDF documents, images, and even the ability to host a

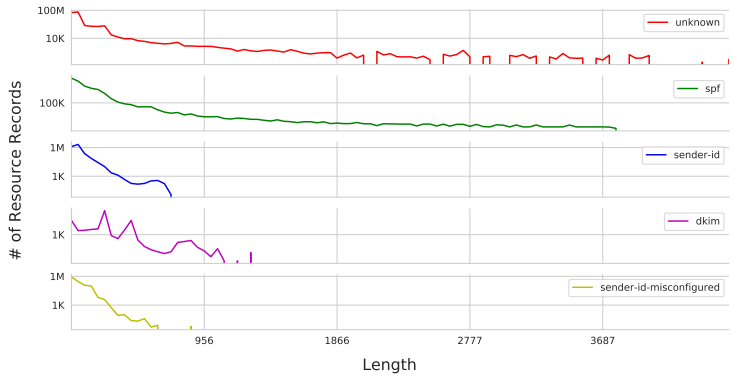


**Fig. 4.** The five representative applications with the longest records. All other applications had records below 950 characters. Note the log-scale on the y-axis.

simple website in DNS as a combination of TXT and NAPTR records, which is then rendered through their free "Telhosting" service. This use of DNS as an arbitrary key-value data store is a significant departure from the intended use of the service, and presents a serious security risk for amplification attacks. Given that the TLD is signed using DNSSEC, an attacker with the ability to automate domain registration and upload of files to `.tel` could potentially flood a target with these large TXT records and the accompanying DNSSEC records.

## 5 Security Implications

### 5.1 Information Leakage

One of the biggest take-aways from this study is that the presence of particularly formatted TXT records conveys information about how a domain does business. A TXT record's key often identifies the requesting service by name, and in cases where it does not, it is easy to track down this information. This leads to a situation where a third party looking to infiltrate an organization can use DNS as an information gathering tool. The aggregate of this information can provide insight into services a domain is using that an outsider would not otherwise be privy to. All categories of TXT records display this leakage to some degree.

Protocol Enhancement records are intended to be publicly accessible in order to serve their function for enhancing another use case. As such, the disclosure here may not be as damaging, but remains useful to attackers. By looking at a domain's SPF record, an interested party can determine if mail is sent using a cloud service, on premises servers, or some combination. Knowing the origin of sent mail from cloud providers could be used in the creation of phishing attacks, allowing an attacker to masquerade as a cloud provider in an attempt to steal unwary users' sensitive information. This could also include information about

```
"adobe-idp-site-verification=b3947e8b-e9ab-4a78-93df-e1d11a06e155"
"google-site-verification=rrX-JT3vh3S6nvjDwRIlLPyr3m9_6aMp00lU5jHSQGA"
"pardot_82202_*=69be9d019c63cbfa7e4816f0774d2e4d829c2c10bdf8c5450320cdddb2f1a493"
"dsUs6X1AHF2QnsOeFRlq97nF4u+DGclGBD3vPQTlgK1lVwZ3vgbi0BNuwogA0KRdzg2RSAJcq92sG+YIwB8AKQ=="
"v=spf1 ip4:198.102.61.0/24 ip4:198.102.62.0/24 ip4:198.102.63.0/24 ip4:198.102.32.0/24
     ip4:162.209.25.132 include:_spf.salesforce.com include:spf-00151a02.pphosted.com
     include:aspmx.pardot.com include:amazonses.com ip4:204.93.64.116 ip4:204.93.64.117
     ip4" ":192.250.208.112 ip4:192.250.208.113 include:sendgrid.net
     include:spf.workfront.com ~all"
"docusign=5a8c5ea6-9539-457f-8930-a2021aca99ac"
"docusign=905a8e17-2e3a-4fb0-9bf3-99ec8b1ec979"
"amazonses:dY6kgeXGidLNGtmdsUxmjPnMcEXJ+U5FIKNDH+JxAE8="
"adobe-sign-verification=b6d1b80570e07516da53ff616b5b41d"
```

**Fig. 5.** TXT records for an example domain from our dataset.

software being used to handle email, such as a SenderID record or Base64 Active Directory Federation record indicating use of Microsoft Exchange.

Domain Verification records convey the most information about how a domain operates because many of them contain the product name requesting verification in the TXT record. Documentation about when these records can be removed, if ever, is far from comprehensive, an example of neglect for known security best practices [16]. As a result, these records are left in DNS indefinitely. This information allows third parties to profile a domain's cloud service usage. Of the 43 distinct use cases of Domain Verification we found, only 10 indicated how long the record was needed, with 4 stating the record *could be removed after the initial verification*. These records can sometimes identify services in use the public would otherwise not know about, such as the "citrix-verification-code" record indicating GoToMeeting use. This information can be used to craft targeted phishing attacks for users of that domain, as well as giving an attacker additional vectors for service disruption or caches of sensitive information.

Resource Location records not only convey information about product usage, but also provide a specific endpoint for an attacker to target. In our analysis, the targets of these location identifiers have varying accessibility from the public Internet based on the service they are for. The records for Symantec's MDM suite all had targets that did not respond to `curl` probes. However, some of the JBoss Fuse records and all of the Ivanti Landesk records pointed to servers that did respond. In the case of Ivanti Landesk, depending on the mobile device operating system they were servicing, the response was either blank or contained service health information. This behavior was observed across all domains using this product, including Ivanti itself, so this appears to be the desired behavior. This service is used in management of other devices and is a very attractive target to be used as command-control for more nefarious purposes.

Consider one example domain in our dataset with a diverse set of TXT records shown in Figure 5. This domain has 9 records covering 8 separate applications. By looking at the types of TXT records, we can gain a lot of insight into how this domain does business. In the SPF record, there are a mix of `ip4` and `include` directives, so it is reasonable to conclude that the domain hosts

some of their mail infrastructure locally and leverages cloud services for others. The include statements have targets for Salesforce Pardot and Amazon SES, so we can say that those services are probably part of this solution. This is backed up by the presence of Salesforce Pardot and Amazon SES domain verification records. A Base64 record indicating the use of Microsoft Exchange (`dsUs6...`) is present, so that is most likely what they are using as their on-premises email solution. There is an Adobe IDP Site Verification record, so we know this is the domain's SAML Single Sign-On solution for cloud services. There is a record for GSuite domain verification indicating that this domain uses some features of GSuite. Records for Adobe Sign and Docusign indicate those services are being used to digitally sign documents. In previous record retrievals for this zone, there were also records for 2 different VOIP solutions; Citrix Verification Code and LogMeIn OpenVoice. They could be using these in conjunction for internal communication, or phasing one out in favor of the other.

Given this information, an attacker has many options for crafting convincing spear phishing messages or searching for application-specific vulnerabilities. Starting with the protocol enhancement records, the presence of Microsoft Exchange records indicates a program to search for any known remote vulnerabilities. Exposed servers in the ranges contained in the SPF record could lead to exploitation of these vulnerabilities for Exchange servers there. Unfortunately, because the information contained in an SPF record needs to be public to assist in spam prevention, these IP ranges can not be hidden. Other preventative techniques, like restrictive firewall rules limiting port access to the mail servers could be leveraged to limit the attack surface.

The domain verification records, however, could be hidden, removed, or obfuscated. Their presence and clear formatting in the domain inform third parties of details about the domain's SaaS usage. An email attempting to phish credentials appearing to come from any of the services using these verification records will appear more genuine. This is because we know that, at least at some point, these services were in active use by the domain. Using social engineering to obtain access to some services, like the Adobe IDP, Adobe Sign or Docusign, would be particularly appealing. These services could contain sensitive credentials or important documents. If domain verification records were removed shortly after the verification step was performed, or did not contain service specific identifiers, none of this information would be available in profiling operations. Additionally, if the verification records were hidden in unguessable subdomains, their presence would remain unknown to most third parties and would still be usable by the service requesting them.

### 5.2 Service Hijacking

Reliance on information stored in DNS to control the security or functionality of an application opens that application up to service hijacking if cache poisoning is used to alter the content of that record. The DNS server being poisoned would need to be the recursive resolver of the client accessing DNS, making this a very targeted proposition. However, blind trust in the integrity of DNS records

to validate other processes is a dangerous assumption to make, particularly in cases where domains are not DNSSEC signed or resolvers are not validating responses. In our dataset, domains publishing TXT records also publish `RRSIG` records only 6% of the time. This includes a dramatic increase in the number of TXT records appearing with an accompanying RRSIG record in January 2017 (see Figure 1). After closer inspection, we found that 97% of the new DNSSEC verified records that appeared that month belonged to the Western Samoa (.ws) top-level domain. This TLD is managed by a single registrar, so the sharp increase in RRSIG records would indicate a change in signing policy by this particular registrar. In total, our observed occurrence of DNSSEC-verified domains is higher than a recent estimate of global DNSSEC deployment which places the rate of signed zones closer to 1% [12], but it it is still very low.

In the case of Protocol Enhancement records, this could be used to bypass the protections offered entirely. In DKIM records, content validation could be bypassed by injecting a cache poisoned record with a fraudulent public key to make it seem genuine. RFC 4871 directly identifies this possibility as a flaw in the design [4]. In section 8.4, the authors state "DKIM is only intended as a 'sufficient' method of proving authenticity. It is not intended to provide strong cryptographic proof about authorship or contents." The RFC goes on to recommend that any domain using DKIM as part of its email validation plan should also use DNSSEC to sign the domain.

Domain Verification records have varying degrees of vulnerability, depending on use case. The most damaging use of cache poisoning would be in records used for signing up for a cloud-based CA or for domain scanning services. In the case of the CA, if a forged verification record appeared in the CA's recursive resolver, an account could be created allowing trusted certificates to be issued for a domain the attacker does not own [10]. For scanning services, forged records could turn these platforms into on-demand denial of service generators.

For Resource Location records, a forged record could change the target of clients to redirect them to a malicious server instead. In remote management solutions like Symantec MDM or Ivanti Landesk, this could be a management server that instead functions as command-control for malware installation or forcing clients to participate in botnets.

## 6 Conclusion

DNS TXT records are used for a wide variety of purposes, some formal and many more application specific and informal. We have identified 52 structured usage patterns of TXT records covering 92% of TXT records from a representative set of domains, but others likely exist. These patterns generally fall into 3 application categories: Protocol Enhancement, Domain Validation, and Resource Location. We have observed that the use of Protocol Enhancement records has increased dramatically in the last 10 years, as evidenced by the widespread use of SPF records. The majority of the informal use cases fall into the Domain Validation or Resource Location categories, and generally suffer from poor documentation

about when they can be removed. The use of these records exposes business information about the organization publishing them, allowing for easier profiling by attackers or identifying specific targets to exploit. Furthermore, we identified an unconventional use of DNS as a key-value data store, which represents a severe threat for potential amplification attacks. Both domain operators and service owners have a responsibility to make sure the openness of DNS TXT records is not abused, and our work shows that carelessness in their use can lead to significant public information leakage and vulnerability to attacks.

## References


1. Agar, R.J.M.: The domain name system (DNS): Security challenges and improvements. Tech. rep., Royal Holloway, University of London (2010)
2. Agten, P., Joosen, W., Piessens, F., Nikiforakis, N.: Seven months' worth of mistakes: A longitudinal study of typosquatting abuse. In: Proceedings of the Network and Distributed System Security Symposium (NDSS) (2015)
3. Akamai: Security bulletin: Crafted dns text attack. `https://www.akamai.com/us/en/multimedia/documents/state-of-the-internet/dns-txt-amplification-attacks-cybersecurity-threat-advisory.pdf` (2014)
4. Allman, E., Callas, J., Delany, M., Libbey, M., Fenton, J., Thomas, M.: Domainkeys identified mail (dkim) signatures. RFC 4871, RFC Editor (2007), `http://www.rfc-editor.org/rfc/rfc4871.txt`
5. Alrwais, S.A., Yuan, K., Alowaisheq, E., Li, Z., Wang, X.: Understanding the dark side of domain parking. In: USENIX Security Symposium (2014)
6. Amann, J., Gasser, O., Brent, L., Carle, G., Holz, R.: Mission Accomplished? HTTPS Security after DigiNotar. In: Proceedings of the ACM Internet Measurement Conference (IMC) (2017)
7. Barnes, R., Hoffman-Andrews, J., McCarney, D., Kasten, J.: Draft: Automatic certificate management environment (ACME). `https://www.ietf.org/id/draft-ietf-acme-acme-18.txt` (2019)
8. Bellis, R.: DNS transport over TCP - implementation requirements. RFC 5966, RFC Editor (2010), `http://www.rfc-editor.org/rfc/rfc5966.txt`
9. Borgolte, K., Fiebig, T., Hao, S., Kruegel, C., Vigna, G.: Cloud Strife: Mitigating the Security Risks of Domain-Validated Certificates. In: Proceedings of the Network and Distributed System Security Symposium (NDSS) (2018)
10. Brandt, M., Dai, T., Klein, A., Shulman, H., Waidner, M.: Domain Validation++ For MitM-Resilient PKI. In: Proceedings of the ACM Conference on Computer and Communications Security (CCS) (2018)
11. Bushart, J., Rossow, C.: DNS Unchained: Amplified Application-Layer DoS Attacks Against DNS Authoritatives. In: Proceedings of the International Symposium on Research in Attacks, Intrusions, and Defenses (RAID) (2012)
12. Chung, T., van Rijswijk-Deij, R., Chandrasekaran, B., Choffnes, D., Levin, D., Maggs, B.M., Mislove, A., Wilson, C.: A longitudinal, end-to-end view of the DNSSEC ecosystem. In: USENIX Security Symposium (2017)
13. Chung, T., van Rijswijk-Deij, R., Choffnes, D., Levin, D., Maggs, B.M., Mislove, A., Wilson, C.: Understanding the role of registrars in DNSSEC deployment. In: Proceedings of the ACM Internet Measurement Conference (IMC) (2017)
14. Cisco: Cisco umbrella populatiry list, september 26 2017. `http://s3-us-west-1.amazonaws.com/umbrella-static/top-1m-TLD-2017-09-26.csv.zip`

15. Dagon, D., Provos, N., Lee, C.P., Lee, W.: Corrupted DNS resolution paths: The rise of a malicious resolution authority. In: Proceedings of the Network and Distributed System Security Symposium (NDSS) (2008)
16. Dietrich, C., Krombholz, K., Borgolte, K., Fiebig, T.: Investigating System Operators' Perspective on Security Misconfigurations. In: Proceedings of the ACM Conference on Computer and Communications Security (CCS) (2018)
17. Dinaburg, A.: Bitsquatting: DNS hijacking without exploitation. In: Proceedings of BlackHat Security (2011)
18. DMARC.org: Dmarc overview. `https://dmarc.org/overview/`
19. Durumeric, Z., Adrian, D., Mirian, A., Kasten, J.: Neither Snow Nor Rain Nor MITM... An Empirical Analysis of Mail Delivery Security. In: Proceedings of the ACM Internet Measurement Conference (IMC) (2015)
20. Foster, I.D., Larson, J., Masich, M., Snoeren, A.C., Savage, S., Levchenko, K.: Security by Any Other Name: On the Effectiveness of Provider Based Email Security. In: Proceedings of the ACM Conference on Computer and Communications Security (CCS) (2015)
21. Görling, S.: An overview of the sender policy framework (spf) as an anti-phishing mechanism. Internet Research **17**(2), 169–179 (2007)
22. Herzberg, A., Shulman, H.: DNSSEC: Security and availability challenges. In: IEEE Conference on Communications and Network Security (CNS). pp. 365–366. IEEE (2013)
23. Hu, H., Wang, G.: End-to-End Measurements of Email Spoofing Attacks. In: USENIX Security Symposium (2018)
24. Kaminsky, D.: Black ops 2008: It's the end of the cache as we know it. Black Hat USA (2008)
25. Kintis, P., Miramirkhani, N., Lever, C., Chen, Y., Romero-Gómez, R., Pitropakis, N., Nikiforakis, N., Antonakakis, M.: Hiding in plain sight: A longitudinal study of combosquatting abuse. In: Proceedings of the ACM Conference on Computer and Communications Security (CCS) (2017)
26. Kountouras, A., Kintis, P., Lever, C., Chen, Y., Nadji, Y., Dagon, D., Antonakakis, M., Joffe, R.: Enabling network security through active DNS datasets. In: Proceedings of the International Symposium on Research in Attacks, Intrusions, and Defenses (RAID) (2016)
27. Le, T., Van Rijswijk-Deij, R., Allodi, L., Zannone, N.: Economic incentives on DNSSEC deployment: Time to move from quantity to quality. In: IEEE/IFIP Network Operations and Management Symposium (NOMS) (2018)
28. Lever, C., Walls, R., Nadji, Y., Dagon, D., McDaniel, P., Antonakakis, M.: Domainz: 28 registrations later measuring the exploitation of residual trust in domains. In: IEEE Symposium on Security and Privacy (SP) (2016)
29. Lyon, J., Wong, M.: Sender id: Authenticating e-mail. internet engineering task force (ietf). RFC 4406, RFC Editor (2006), `http://www.rfc-editor.org/rfc/rfc4406.txt`
30. M. Kucherawy, E., E. Zwicky, E.: Domain-based message authentication, reporting, and conformance (dmarc). RFC 7489, RFC Editor (2015), `http://www.rfc-editor.org/rfc/rfc7489.txt`
31. MacFarland, D., Shue, C., Kalafut, A.: Characterizing optimal DNS amplification attacks and effective mitigation. In: International Conference on Passive and Active Network Measurement (PAM) (2015)
32. Mockapetris, P.: Domain names - implementation and specification. RFC 1035, RFC Editor (1987), `http://www.rfc-editor.org/rfc/rfc1035.txt`

33. Neij, F., Norberg, A., Brown, C.: Bep 34: DNS tracker preferences. http://www.bittorrent.org/beps/bep_0034.html
34. Nikiforakis, N., Balduzzi, M., Desmet, L., Piessens, F., Joosen, W.: Soundsquatting: Uncovering the use of homophones in domain squatting. In: International Conference on Information Security (ISC) (2014)
35. Nikiforakis, N., Van Acker, S., Meert, W., Desmet, L., Piessens, F., Joosen, W.: Bitsquatting: Exploiting bit-flips for fun, or profit? In: Proceedings of the international conference on World Wide Web (WWW) (2013)
36. Osterweil, E., Ryan, M., Massey, D., Zhang, L.: Quantifying the operational status of the DNSSEC deployment. In: Proceedings of the ACM Internet Measurement Conference (IMC) (2008)
37. Pearce, P., Jones, B., Li, F., Ensafi, R., Feamster, N., Weaver, N., Paxson, V.: Global measurement of DNS manipulation. In: USENIX Security Symposium (2017)
38. van Rijswijk-Deij, R., Sperotto, A., Pras, A.: DNSSEC and its potential for DDoS attacks. In: Proceedings of the ACM Internet Measurement Conference (IMC) (2014)
39. Scheitle, Q., Hohlfeld, O., Gamba, J., Jelten, J., Zimmermann, T., Strowes, S.D., Vallina-Rodriguez, N.: A Long Way to the Top: Significance, Structure, and Stability of Internet Top Lists. In: Proceedings of the ACM Internet Measurement Conference (IMC) (2018)
40. Scheitle, Q., Mislove, A., Carle, G., Chung, T., Hiller, J., Gasser, O., Naab, J., van Rijswijk-Deij, R., Hohlfeld, O., Holz, R., Choffnes, D.: A First Look at Certification Authority Authorization (CAA). ACM SIGCOMM Computer Communication Review **48**(2), 10–23 (2018)
41. Schlitt, W., Wong, M.W.: Sender policy framework (spf) for authorizing use of domains in e-mail, version 1. RFC 4408, RFC Editor (2006), http://www.rfc-editor.org/rfc/rfc4408.txt
42. Statuspage: DNS configuration requirements. https://help.statuspage.io/knowledge_base/topics/domain-ownership
43. Szalachowski, P., Perrig, A.: Short paper: On deployment of DNS-based security enhancements. In: Financial Cryptography and Data Security (2017)
44. Telnames Limited: .tel. https://www.do.tel/ (2019)
45. Wander, M.: Measurement survey of server-side DNSSEC adoption. In: Proceedings of the Network Traffic Measurement and Analysis Conference (TMA) (2017)
46. Wang, Y.M., Beck, D., Wang, J., Verbowski, C., Daniels, B.: Strider typo-patrol: Discovery and analysis of systematic typo-squatting. SRUTI **6**, 31–36 (2006)
47. Weaver, N., Kreibich, C., Paxson, V.: Redirecting DNS for ads and profit. In: USENIX Workshop on Free and Open Communications on the Internet (FOCI) (2011)
48. Zdrnja, B., Brownlee, N., Wessels, D.: Passive monitoring of DNS anomalies. In: Conference on Detection of Intrusions and Malware & Vulnerability Assessment (DIMVA). pp. 129–139. Springer (2007)
49. Zmijewski, E.: Accidentally importing censorship. https://dyn.com/blog/fouling-the-global-nest/ (March 2010)