

Tracking Certificate Misissuance in the Wild

Deepak Kumar*, Zhengping Wang*, Matthew Hyder*, Joseph Dickinson*, Gabrielle Beck†, David Adrian†, Joshua Mason*, Zakir Durumeric*†‡, J. Alex Halderman†, Michael Bailey*

* University of Illinois Urbana-Champaign † University of Michigan ‡ Stanford University

Abstract—Certificate Authorities (CAs) regularly make mechanical errors when issuing certificates. To quantify these errors, we introduce ZLint, a certificate linter that codifies the policies set forth by the CA/Browser Forum Baseline Requirements and RFC 5280 that can be tested in isolation. We run ZLint on browser-trusted certificates in Censys and systematically analyze how well CAs construct certificates. We find that the number errors has drastically reduced since 2012. In 2017, only 0.02% of certificates have errors. However, this is largely due to a handful of large authorities that consistently issue correct certificates. There remains a long tail of small authorities that regularly issue non-conformant certificates. We further find that issuing certificates with errors is correlated with other types of mismanagement and for large authorities, browser action. Drawing on our analysis, we conclude with a discussion on how the community can best use lint data to identify authorities with worrisome organizational practices and ensure long-term health of the Web PKI.

I. INTRODUCTION

HTTPS depends on a supporting public key infrastructure (PKI) composed of hundreds of certificate authorities (CAs) that verify the identities of websites and issue digital certificates. To ensure compatibility between browsers and HTTPS-enabled websites, standards bodies like the IETF and CA/Browser Forum have developed policies that govern the digital certificates that CAs provide. Unfortunately, there is a long history of certificate authorities failing to adhere to accepted standards, due to both implementation errors and indifference. In this paper, we systematically analyze the errors that authorities make when constructing certificates and consider whether these errors can be used to predict more serious problems.

We begin by dissecting the policies set forth by RFC 5280 [14] and CA/Browser Forum Baseline Requirements [9]. We find that many aspects of certificate construction can be checked in isolation, and we codify these requirements in a set of 220 lints. We introduce and release ZLint, a Go-based linting framework that implements these checks and provides structured data on certificate construction and standards adherence.

To quantify misissuance (i.e., certificates with errors) in the Web PKI, we run ZLint on the 240 million browser-trusted certificates in Censys [17]. We find that misissuance is low

in aggregate. Only 0.02% of certificates violate one of the two standards in 2017; 3.3% do not adhere to community best practices. This is a significant improvement from 2012 when more than 12% of certificates contained errors and nearly one third violated community recommendations. However, while the global misissuance rate is low, this is predominantly due to a handful of large authorities that consistently issue certificates without error. The three largest CAs by organization—Let’s Encrypt, Comodo, and cPanel—signed 80% of the certificates in our dataset and have near-zero misissuance rates. Let’s Encrypt, the largest CA by number of certificates issued, has a particularly stellar incident rate. Of the 37 million certificates the CA has signed, only 13 contain errors. None have warnings.

The bulk of misissuance is due to two classes of authorities. The first class is mid-sized authorities that make a variety of errors in a small percentage of their certificates. The second class is a long tail of small authorities that make the same errors in every issued certificate. Nearly half of the organizations in our dataset misissue more than 10% of certificates, and seventeen have made errors in every certificate. More than half of the errors and warnings in ZLint are triggered at least once. Most often, authorities fail to fully populate the Subject Alternative Names extension, encode the wrong type of data in the extension, or include invalid DNS names. Beyond individual certificates, we find that many organizations struggle to properly maintain OCSP/CRL responders. During our three week test period, the OCSP responders for 73 organizations (10%) failed every health check.

Next, in order to determine whether Lint data can be used to predict more serious issues, we investigate the correlation between the organizations that issue certificates containing errors, OCSP/CRL endpoint uptime, and browser removal. We find that there is weak correlation between the organizations that issue certificates with errors and OCSP availability. For authorities that have issued more than 100K certificates, there is moderate to strong correlation between ZLint-identified misissuance and browser removal. Surprisingly, while there is discussion about large CAs with high error profiles, there is no correlation between the small authorities making errors and discussion in the community (e.g., in the Mozilla Developer Security Policy mailing list).

Our results indicate that large authorities are making progress in correctly issuing certificates. However, there remains a long tail of small authorities that fail to follow community standards and misissue most certificates. Most of these small authorities are not being actively discussed. We hope that by shedding light on these practices, we motivate the community to investigate struggling authorities and prompt discussion on whether lint data can be systemically used to help prevent future PKI incidents. Finally, by releasing ZLint, we hope to

help certificate authorities avoid making errors in the future.

II. BACKGROUND

HTTPS—and TLS more broadly—depend on a supporting public key infrastructure (PKI). The Web PKI that supports HTTPS on the public Internet consists of hundreds of certificate authorities (CAs)—organizations that user agents like browsers trust to verify the identities of websites and provide digital certificates. More than 400 organizations, ranging from commercial CAs to academic institutions, controlled browser-trusted signing certificates in 2013 [19]. The Web PKI was historically opaque as certificates and their issuers were often unknown until found in the wild. However, repeated compromise and anecdotes of negligence led to increased scrutiny and community initiatives to publicly log known certificates, analyze CA behavior, establish technical standards, and distrust abusive organizations. Our work builds on several of these initiatives:

Certificate Transparency Logs. Certificate Transparency (CT) is a Google-initiated effort to maintain public, cryptographically-verifiable ledgers of all browser-trusted certificates (logs) [27]. Originally started in 2013, CT logs initially contained certificates found primarily through Google web crawls and Internet-wide scanning. Since then, several large authorities (e.g., Let’s Encrypt and Symantec) have started logging certificates at the time of issuance [22], [45]. In 2017, Google Chrome announced plans to require CT logging for browser trust [41]; other browsers are expected to follow [33]. As a result, public certificate transparency servers have become a de facto data source for monitoring the PKI.

Internet-Wide Scanning. In 2010, research groups began to use Internet-wide scanning to identify trusted certificate authorities and to publish data sets of known certificates [19]–[21], [26]. These scans helped to identify the widespread delegation of signing credentials and uncover abuse. While much of the community now relies on CT servers for data, Vandersloot et al. recently found that the combination of CT and Internet-wide scanning provides the most comprehensive perspective of the PKI [47]. For our study, we use Censys [17], which aggregates certificates from publicly known CT servers and IPv4 scans of common protocols.

CA/Browser Forum. The CA/Browser Forum is a voluntary consortium of certificate authorities, browsers, and other PKI participants [9]. The forum maintains several binding technical guidelines. In June 2007, the CA/Browser Forum published their first standard, *Guidelines For The Issuance And Management Of Extended Validation (EV) Certificates* [10], which outlines the expectations associated with issuing EV certificates. In 2011, the Forum established a second standard, *Baseline Requirements for the Issuance and Management of Publicly-Trusted Certificates* [9], colloquially referred to as the *Baseline Requirements* or *BRs*. The CA/B Forum Baseline Requirements are binding for all member organizations and apply to certificates issued for web authentication. This enforcement is often performed by participating web browsers.

Common CA Database (CCADB). The Common CA Database (CCADB) is a Mozilla-led effort to establish a public database of trusted certificate authorities [32]. As part of its root store policy, Mozilla requires that roots submit unconstrained

intermediates to CCADB [31]. We use CCADB to help identify the owners of signing certificates.

Mozilla Developer Security Policy Mailing List. The Mozilla Developer Security Policy (MDSP) mailing list [1] is the primary public forum for discussing the Mozilla Root Program. It has become a de facto location for community members to disclose PKI issues and for browsers to solicit feedback from both CAs and the broader community. We use MDSP to identify which certificate authorities are being actively discussed.

Increased transparency has led browsers to distrust several authorities over the past three years. Chrome and Firefox distrusted WoSign and StartCom after they were found issuing backdated SHA-1 certificates [34], [49]. PROCERT was distrusted in 2017 after multiple incidents of misissuance [29]. Chrome announced plans in 2017 to gradually distrust Symantec roots after a long history of problems [30], [35].

A. Terminology

We use several PKI-specific terms throughout this work, which we define below:

- 1) *Subscriber Certificate:* Subscriber certificates are provided to end customers. In the web PKI, they are typically used to identify a website to a user agent. They are not CA certificates.
- 2) *Root Certificate:* A root certificate is the type of CA certificate shipped in a user agent’s root store and acts as the trust anchor in the PKI.
- 3) *Intermediate Certificate:* Intermediates are a type of CA certificate signed by a trust anchor. Authorities typically sign subscriber certificates using an intermediate instead of their root certificate. In some situations, trust anchors will provide an intermediate certificate to a third-party organization to sign their own certificates.
- 4) *Organization:* Many authorities use multiple intermediate CA certificates for signing subscriber certificates. To group these intermediate certificates together, we rely on the Organization field that CAs generally include in certificate subjects. When we refer to organizations in this work, we are referencing this field.
- 5) *Business Owner:* This is the business entity responsible for each intermediate in our dataset. A public record of business owners is operated by Mozilla as a part of the Common CA Database (CCADB) [32].
- 6) *NSS Root Store:* Mozilla Network Security Services (NSS) is the cryptographic library used by Firefox for performing TLS handshakes. The NSS root store contains the set of trust anchors trusted by Firefox.

III. ZLINT: A CERTIFICATE LINTER

In order to programmatically detect flaws in certificates, we built ZLint—a linter that checks a certificate for conformance with RFC 5280 [14] and CA/B Forum Baseline Requirements [9]. We manually dissected the two standards and built a set of standalone *lints* that check a certificate for problems in isolation. We note that the Baseline Requirements standardize many aspects of issuance that cannot be verified in an isolated environment. For example, it is not possible

```

func (l *InvalidCertificateVersion) Execute(
    cert *x509.Certificate) *LintResult {
    if cert.Version != 3 {
        return &LintResult{Status: Error}
    }
    return &LintResult{Status: Pass}
}

func init() {
    RegisterLint(&Lint{
        Name: "e_invalid_certificate_version",
        Description: "Certificates MUST be of
            type X.509 v3",
        Source: CABFBaselineRequirements,
        Citation: "CABF BR 7.1.1",
        EffectiveDate: util.CABV130Date,
        Lint: &InvalidCertificateVersion{},
    })
}

```

Code Block 1: **Example Lint**—Lints are self-contained Go functions that check for adherence with technical standards. This lint checks that a certificate uses the correct X.509 version.

to determine whether the destination domain was correctly validated. ZLint consists of 220 lints and has 95% coverage of certificate-related BR clauses and 90% of RFC clauses.¹ We note that the CA/B Forum Baseline Requirements [9] only applies to ServerAuth certificates, that is, certificates used for TLS server authentication.

Lints can be one of several severity levels, including NOTICE, WARNING, and ERROR, which correspond to varying requirement levels in standards documents. For example, the failure to adhere to an RFC MUST clause maps to an ERROR and disregarding a SHOULD clause maps to a WARNING. ZLint does not contain any warning or error lints that do not correspond to a specific RFC or BR clause.² We label any certificate containing an error as misissued. For an example, consider the following clause from BR §7.1.1: Certificates MUST be of type X.509 v3. In this situation, we would produce an error (and consider the certificate misissued) if it is any other X.509 version. We show the corresponding lint in Code Block 1.

We note that while RFC 5280 is a static document, the Baseline Requirements continually evolve. New requirements typically are not retroactive, and as a result, not all BR clauses apply to every certificate. To ensure that we fairly grade historical certificates, we additionally encode an “effective date” in each lint and do not use it to score certificates issued prior to that date.

ZLint is not the first certificate linter. We were inspired by certlint, which was released by Peter Bowen in early 2016 [7] and has been used to uncover numerous PKI issues (e.g., [37], [38]). Several other linters have since started in parallel, including X.509 Lint [36] and GlobalSign certlint [25]. There are several architectural differences between ZLint and certlint. First, because we use ZLint to grade authorities, we

¹For a complete list, see <https://zmap.io/zlint/coverage>.

²We make one exception in mapping standards clauses to lint severity. The common name field is no longer recommended by the BRs, but is included in nearly every subscriber certificate as many user agents do not correctly parse certificates without a common name. We exclude the lint from our set of warnings such that every certificate does not flag a warning.

Issuer	Certificates	w/ Errors [†]	w/ Warnings [†]
Let’s Encrypt	37 M (61%)	13 (0.0%)	0 (0.0%)
Comodo	6.7 M (11%)	3,219 (0.0%)	7,902 (0.1%)
cPanel	4.7 M (7.8%)	131 (0.0%)	1 (0.0%)
Symantec	2.8 M (4.6%)	23,053 (0.8%)	2.7 M (99.9%)
GeoTrust Inc.	1.9 M (3.2%)	5,694 (0.3%)	1.9 M (99.9%)
GoDaddy	1.6 M (2.7%)	38,215 (2.0%)	5,186 (0.3%)
GlobalSign	1.2 M (1.9%)	837 (0.0%)	237 (0.0%)
Other	4.5 M (7.3%)	67 K (1.5%)	1.3 M (28.7%)
Total	61 M (100%)	140 K (0.2%)	6 M (9.9%)

TABLE I: **Largest Authorities**—We show the misissuance rates for authorities (by organization field) that have issued at least 1M certificates. [†] Percentages are based on the total certificates issued by the authority, not the certificates in the associated severity class.

only include lints that are directly based on a published standard (i.e., RFC 5280 and CA/B Forum BRs). We do not include other community best practices as warnings or errors. Second, we restrict lints to their effective periods. Third, instead of producing an text-based errors of problems to investigate, ZLint produces structured data that can be used for analysis. Fourth, we implemented ZLint as a standalone Go library. ZLint can validate 327K certificates per core hour, a 2235% speedup over certlint, which can process 14K certificates per hour. We hope that by implementing ZLint in a performant language, we can encourage CAs to integrate linters into their issuance processes.

We have released ZLint under the Apache 2.0 license,³ and it has been integrated into the two leading certificate search engines, crt.sh [44] and Censys [17]. The Censys Team has agreed to provide long term maintenance for the project, and we have already begun to see contributions from several popular certificate authorities.

IV. MISISSUANCE TODAY

We characterize misissuance by running ZLint on NSS-trusted certificates in Censys [17]. Unless specified otherwise, the numbers we present in this work are for the certificates that were valid on July 23, 2017. In total, our dataset contains 61 M current certificates and 170 M historically trusted certificates. These certificates were signed by 1320 CA certificates, 618 organizations, and 64 CCADB business owners. ZLint has been integrated into Censys and the data we use in this paper can be found in the normal Censys certificate dataset.

A. A Long Tail of Misissuance

There is a small amount of aggregate misissuance in 2017. Only 0.02% of certificates contain errors (i.e., are misissued), while 3.3% have warnings. This represents a significant improvement from 2012, when more than 12% of certificates contained errors (Figure 1, Figure 2). Though the aggregate misissuance rate is low, this is largely due to a small number of large authorities consistently producing well-constructed certificates. Three authorities—Let’s Encrypt, Comodo, and cPanel—signed for 80% of the certificates in our dataset, and have near-zero misissuance rates (Table I).

³ZLint is available at <https://github.com/zmap/zlint>.




Fig. 1: **Historical Misissuance**—Except for a few exceptions, error and warning rates have steadily decreased over time. The 2010 spike was caused by Thawte issuing a large number of certificates without the Authority Key Identifier (AKID) field. A similar event occurred in late 2014 when Symantec issued certificates without a Subject Key Identifier (SKID).

Year	Certificates	w/ Warnings	w/ Errors
2012	2.3 M	609,646 (26.3%)	287,454 (12.4%)
2013	2.9 M	860,481 (29.9%)	240,943 (8.4%)
2014	3.6 M	933,358 (26.2%)	101,631 (2.9%)
2015	7.0 M	1.8 M (25.4%)	35,419 (0.5%)
2016	50 M	3.4 M (6.7%)	24,008 (0.04%)
2017	102 M	3.4 M (3.3%)	23,207 (0.02%)

Fig. 2: **Misissuance Rates**—The fraction of misissued certificates has reduced significantly over the past five years.




Fig. 3: **Percent Misissuance by Organization Size**—We show a CDF of the percent misissuance in organizations for three sizes, those that issue <1000 certificates, those that issue between 1000 and 100 K certificates, and those that issue >100 K certificates. Larger organizations misissue a smaller fraction of their total certificates compared to smaller organizations.

Let’s Encrypt, the largest CA by number of certificates issued, has a particularly stellar misissuance rate. Of the 37 million certificates the CA has signed, only 13 contain errors.⁴ No certificate issued by Let’s Encrypt contains warnings.

On the flip side, nearly half of authorities (48% by organization) misissue at least 10% of certificates (Figure 4). These authorities tend to be small, issuing fewer than 1,000 certificates a piece (Figure 3, Figure 5). Seventeen organizations (2.8%) have misissued all certificates. In most of these cases, the same errors are present on each certificate within an organization. For

⁴The 13 certificates that Let’s Encrypt misissued contain invalid UTF-8 encoding for internationalized domain names and have since been revoked [2].

example, all Nestle certificates contain the same error: failure to include a certificate policy statement (CPS). While this might not seem surprising, it does contrast with the authorities responsible for the most misissuance by raw number. We show the organizations with the highest misissuance rates in Table II.

The authorities responsible for the largest number of misissued certificates tend to be mid-sized and misissue a small number of certificates relative to their size. Their misissuances inconsistently violate various standards, but the bulk of their errors are associated with a single type of problem. In the most extreme example, GoDaddy is responsible for nearly 40 K misissued certificates (28.6% of all misissuances), but only 2.4% of GoDaddy certificates contain errors.

We note that our investigation likely underestimates misissuance for companies that have acquired smaller authorities because we group CA certificates using the embedded Organization field. For example, at the time our study in July 2017, Symantec Corporation owned four of the authorities that misissued the largest number of certificates: VeriSign, GeoTrust, Thawte, and Symantec [11], [15], [46]. In another example, WoSign owned StartCom [49]. To understand the impact this has on our results, we consider misissuance rates based on CA Owner in CCADB [32]. This field, however, does not always indicate the party responsible for a specific intermediate certificate. For example, IdenTrust appeared as the owner for Let’s Encrypt until they established their own, independent trust anchor. However, it may identify cases not surfaced by grouping intermediates by their embedded organization field.

For the most part, we identify the same set of authorities with the greatest misissuance, but with slightly differing misissuance rates. For example, we find that DigiCert is responsible for 7,203 (1%) misissued certificates. The approach identified one additional, significant CA Owner, Deutsche Telekom, which misissued nearly 20% of all of its certificates—just over 10K certificates in total. There were 59 intermediates in our dataset that were not registered in CCADB at the time of our investigation.

B. What Do Authorities Get Wrong?

Authorities make a variety of errors when issuing certificates (Table III). Of the 179 error lints and 41 warning lints in ZLint, 97 (54.2%) and 25 (61%) were triggered at least once, respectively. The most common error (by misissued certificates) is that authorities fail to correctly populate the Subject Alternate Names (SANs) extension. This error accounts for 70 K misissued certificates (31%); 94 authorities have made the error at least once. 45 authorities similarly fail to include the SAN extension. Most of the failures to include the extension are the responsibility of GoDaddy, who issued 94% of the 40K certificates with this error in 2012–2013. This error results in GoDaddy being responsible for the largest number of misissuances in our dataset. Symantec was responsible for the second most misissued certificates (23K), of which 95% failed to properly encode internationalized domain names in the common name field. WoSign, StartCom, and VeriSign failed to include the authority key identifier extension in 11.7 K, 9.8 K, and 9.8 K certificates respectively.

For all of the organizations that issued at least 10K certificates in our dataset, one error accounts for at least 90%

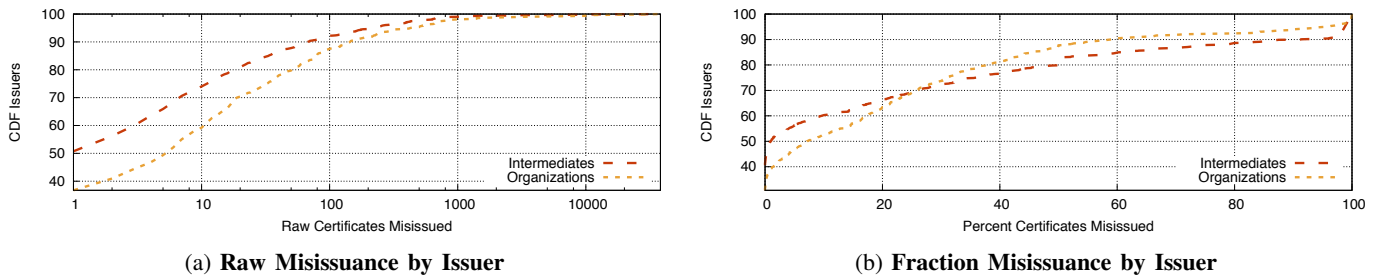


Fig. 4: **Misissuance Rates**—Most Intermediates, Organizations, Business Owners, and Root CAs misissue some fraction of certificates on the public Internet.

Organization	Misissued	Organization	Misissued	Organization	Misissued
Nestle (1)	968 100%	Consortici Catalunya (2)	1,117 58.8%	GoDaddy.com (3)	38,215 2.4%
PSCProcert (1)	39 100%	RHRK (2)	1,171 35.6%	Symantec Corp. [†] (22)	23,053 0.8%
Giesecke and Devrient (1)	18 100%	KPN Corporate BV (2)	1,933 34.5%	StartCom Ltd. [‡] (17)	11,617 2.1%
Unizeto Sp. z o.o. (1)	18 100%	DFN-Verein (5)	1,689 29.8%	WoSign CALtd. [‡] (39)	9,849 5.0%
CertiPath LLC (1)	9 100%	Universitaet Stuttgart (1)	1,830 29.2%	VeriSign [†] (10)	9,835 23.1%
Helsana Gruppe (1)	8 100%	AC Camerfirma S.A. (1)	2,725 25.9%	GeoTrust Inc. [†] (22)	5,694 0.3%
Chunghwa Telecom Co. (1)	7 100%	VeriSign (10)	42,622 23.1%	Comodo Ltd. (30)	3,219 0.1%
TSCP Inc. (1)	5 100%	Trend Micro Inc (1)	6,374 19.8%	DigiCert (43)	2,597 0.1%
Dell Inc. (1)	4 100%	AlphaSSL (1)	3,848 17.2%	Thawte [†] (4)	1,751 0.4%
DigitPA (1)	2 100%	Uni Erlangen Nuernberg (1)	1,115 14.4%	TERENA (9)	1,405 1.7%

(a) Highest Misissuance Rate

(b) Highest Misissuance Rate (>1K certificates)

(c) Most Misissued Certificates

TABLE II: **Organizations with Highest Misissuance**⁵—We show the organizations that misissued the most certificates by fraction or raw count. The number in parentheses is the number of intermediates each organization is responsible for. [†] Owned by Symantec Corp. [‡] Owned by Qihoo 360 Technology Co. Ltd.

Error	Source	Certificates [†]	Warning	Source	Certificates [†]
Subject CN not from SAN	BR §7.1.4.2	70 K (31.0%)	SKID missing	RFC 5280 §4.2	5.6 M (90%)
SAN extension missing	BR §7.1.4.2	39 K (17.3%)	ExtKeyUsage not critical	RFC 5280 §4.2	260 K (4.3%)
Invalid character in DNSName	BR §7.1.4.2	30 K (13.6%)	Explicit Text not UTF-8	RFC 5280 §4.2	184 K (3.0%)
AKID missing	RFC 5280 §4.2	30 K (13.4%)	Policy contains NoticeRef	RFC 5280 §4.2	67 K (1.1%)
SAN email field present	BR §7.1.4.2	12 K (5.2%)	AIA missing CA URL	BR §7.1.2.3	41 K (0.7%)
Invalid TLD in DNSName	BR §7.1.4.2	6.5 K (2.9%)	ExtKeyUsage Extra Values	BR §7.1.2.3	15 K (0.3%)

TABLE III: **Most Common ZLint Errors and Warnings**

of all errors for that authority. However, we also note that 63% of the authorities that issued at least 100K certificates made at least ten errors. This contrasts with organizations that misissue nearly all their certificates, as they generally misissue all of their certificates in the exact same way. The lints that the largest number of authorities triggered are that the email field is present in the SAN extension (44% of organizations), the SAN is not properly populated (15.2%), and DNSName entries are not well-formed (13.4%).

C. Warnings

In addition to making errors, authorities often fail to adhere to recommendations in community standards (i.e., ZLint warnings). Unlike errors, there are a handful of large players who do not follow community recommendations (Figure 6). For example, Symantec, GeoTrust, and Thawte—three large authorities owned by Symantec at the time of our study—trigger warnings for 99% of their certificates. All of these fail to include the subject key identifier (SKID) extension

in end-entity certificates. Despite accounting for 90% of the warnings triggered by ZLint, this warning is only triggered by 23 (3.7%) authorities. The second most common warning is that the extended key usage extension is not marked critical, which appears in 260 K certificates and is triggered at least once in 408 (66%) authorities. HostPoint AG triggers this warning the most—it appears in 62 K of their certificates. They are followed closely by WoSign, which issues 61 K certificates with this warning.

⁵An early version of this paper erroneously listed D-Trust GmbH, Freistaat Bayern, and FNMT-RCM as having high misissuance rates in Table II. This inconsistency was due to a bug in the Go ASN.1 parser, which prevented ZLint from identifying secondary CRL endpoints in some certificates.

In addition, KPN Corporate Market B.V, Microsec Ltd., Trustis Ltd., and TUBITAK were listed with high misissuance rates in Table VIII due an improper check of the length field in the `ExplicitText` field.

Both errors have been corrected in this version of the paper and in ZLint. We thank D-TRUST and Microsec Ltd. for bringing these issues to our attention.

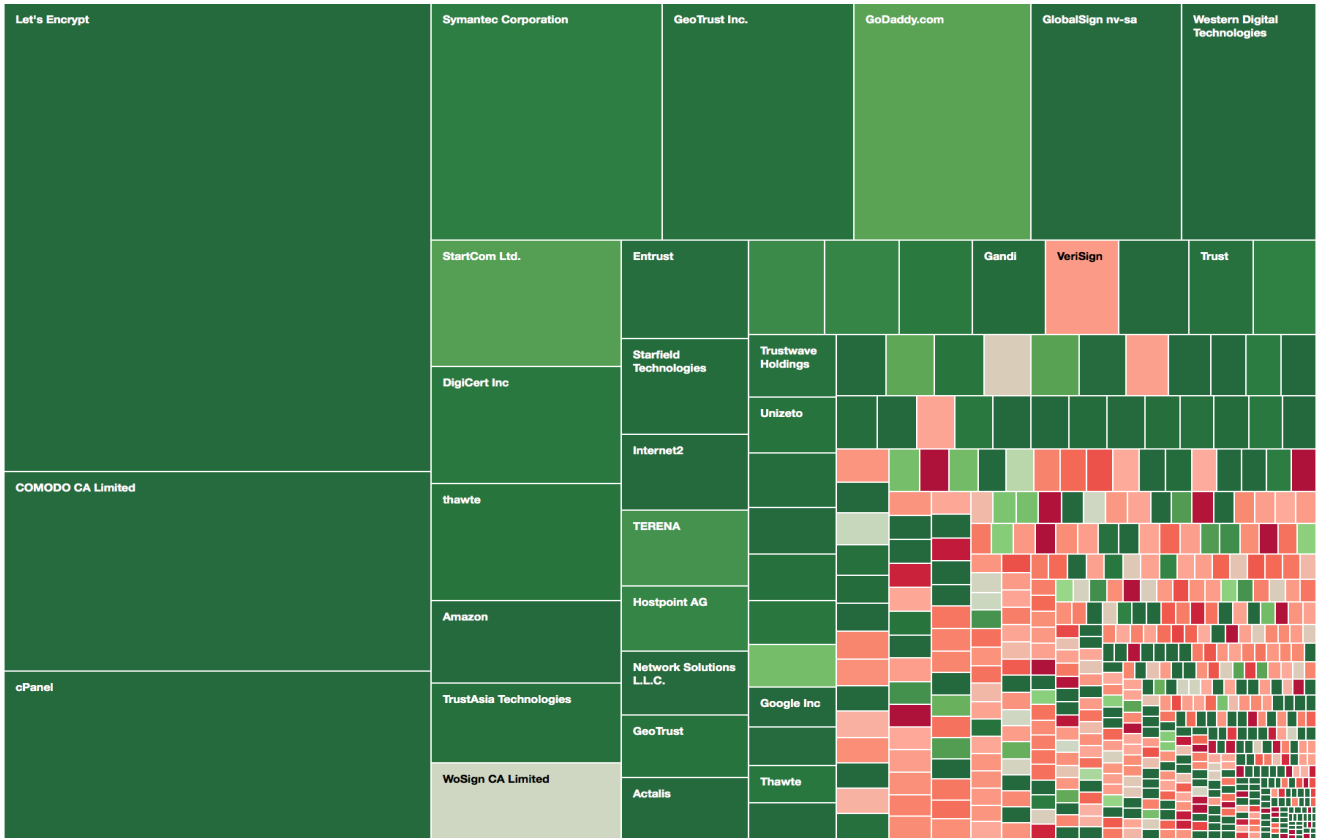


Fig. 5: Percent ZLint Errors by Total Certificates Issued—Large certificate authorities generally issue certificates with fewer ZLint errors than smaller authorities.

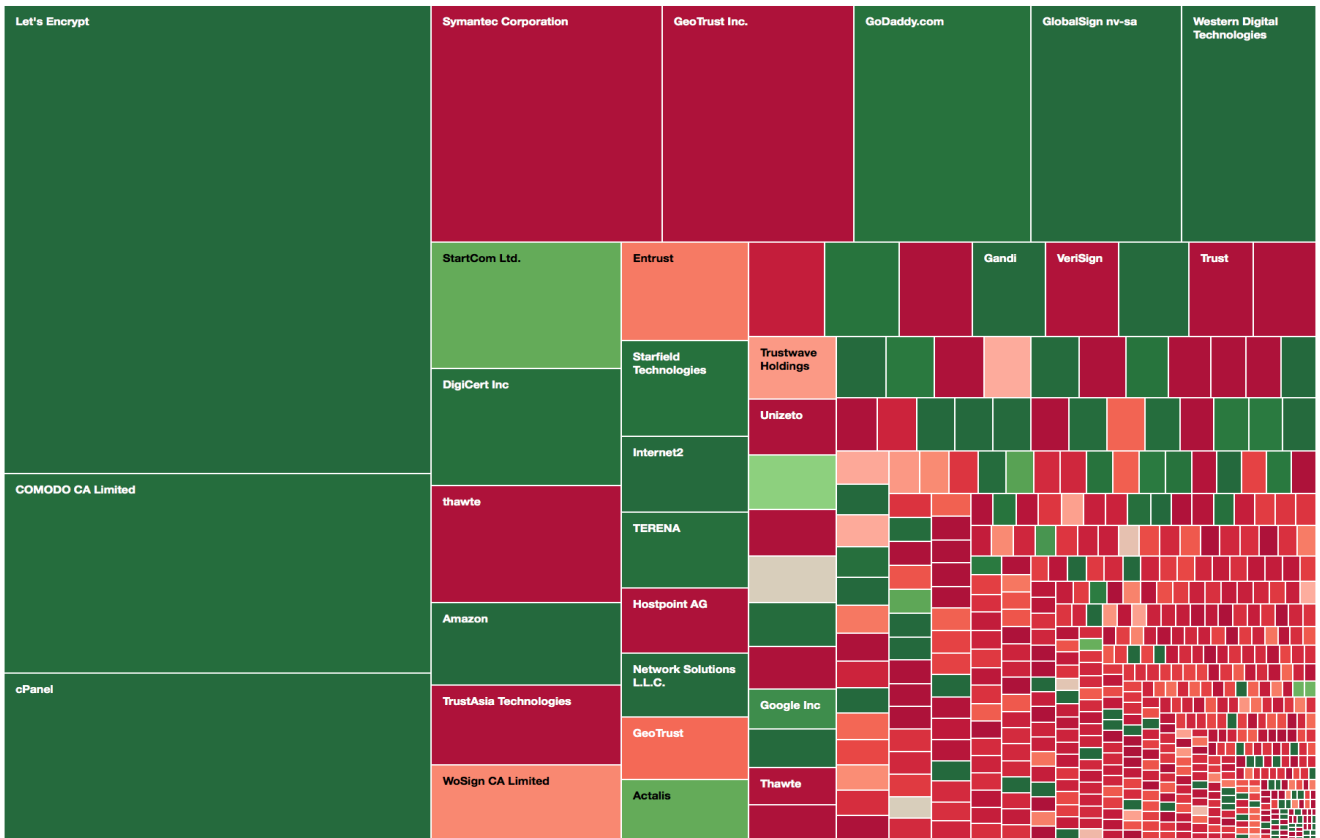


Fig. 6: Percent ZLint Warnings by Total Certificates Issued—ZLint warnings are more dispersed throughout the ecosystem, affecting both large and small players. A handful of large players, Symantec, GeoTrust, thawte, and TrustAsia, all issue more than 95% of their certificates with ZLint warnings.

Organization	Intermediates	Largest Offending Intermediate	% of Issued	% of Errors
GoDaddy Inc.	3	Go Daddy Secure Certification Authority	9%	99.8%
Symantec Corp.	22	Symantec Basic DV SSL CA - G2	60%	94.2%
StartCom Ltd.	17	StartCom Class 1 DV Server CA	83%	76.5%
WoSign CA Ltd.	39	WoSign CA Free SSL Certificate G2	58%	59.0%
VeriSign	10	VeriSign Class 3 International Server CA - G2	32%	96.1%
GeoTrust Inc.	22	RapidSSL SHA256 CA	52%	49.1%
Comodo CA Ltd.	29	COMODO RSA Domain Validation Secure Server	22.7%	85.5%
DigiCert Inc.	43	DigiCert SHA2 Secure Server CA	52.4%	55.5%
Thawte	12	thawte DV SSL CA - G2	30.4%	26.0%
TERENA	9	TERENA SSL CA 3	53.1%	53.2%

TABLE IV: **Intermediate contribution to Organizational Misissuance**—We show the intermediate that contributes the most to misissuance per organization. In 80% of the organizations that issue more than 10K certificates, the majority of misissued certificates are generated by just one intermediate.

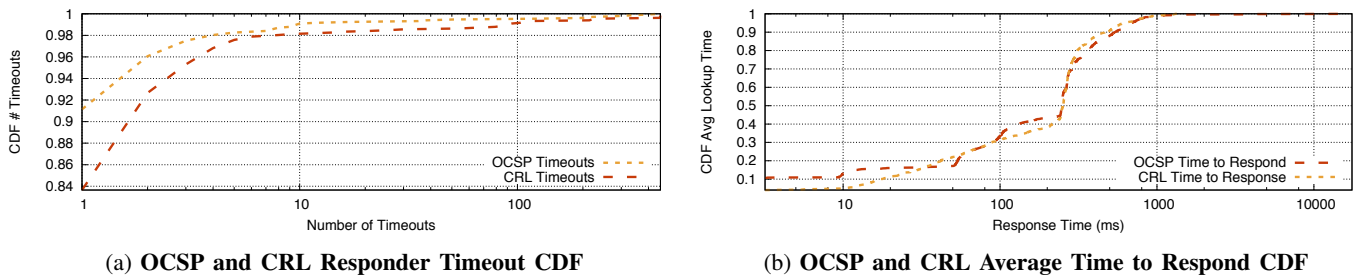


Fig. 7: **OCSP, CRL Timeouts and Average time to Respond**—We show the average response time as well as the number of timeouts for the OCSP responders and CRL distribution points found in certificates in our dataset. While the majority of organizations adhere to community standards, there are still a handful of poorly managed and broken services encoded in trusted certificates.

D. Intermediate Diversity

While it is difficult to externally observe what causes mid-to-large sized authorities to misissue a large number of certificates, we do observe one pattern. Table IV shows the “largest offending intermediate” in raw ZLint error output for organizations that issue more than 10K certificates. We find that in 80% of these organizations, the majority of ZLint error certificates can be traced back to exactly one intermediate. For example, the majority of *COMODO*’s misissued certificates are generated by the intermediate *COMODO RSA Domain Validation Secure Server*, which accounts for 85.5% of the error certificates. We posit that this is because intermediates within an organization are using entirely disparate codebases or infrastructure to issue certificates. This has implications for PKI organizations moving forward—as rules continue to grow and change, maintenance of such processes may become unwieldy, and the surface area for a programming error to manifest may increase. Ultimately, oversight of an organization is harder when there are disparate issuance processes.

V. BEYOND INDIVIDUAL CERTIFICATES

While ZLint can systematically identify problems in the construction of individual certificates, there are many aspects of CA operation that cannot be verified in isolation. These range from correctly validating domain ownership to generating consistent OCSP responses. We measure two supporting

services—OCSP responders and CRL distribution points—to see if they correlate with certificate misissuance.

A. OCSP Responders

OCSP (Online Certificate Status Protocol) is an Internet protocol that allows authorities to revoke issued certificates [39]. Per BR §4.9, authorities are required to continually operate OCSP responders that allow user agents to check for revocation. We measured for adherence with following requirements:

- 1) **Endpoint Availability** Certificate authorities must maintain an online OCSP responder (BR §4.9.10). We check for baseline responsiveness every hour during this period using a 60 second timeout.
- 2) **Regular Updates** Per BR §4.9.10, OCSP responders must update revocation responses at least once every four days. We check that responses are not older than this.
- 3) **No Response for Unknown Certificates** Authorities must not return a `GOOD` response for certificates they did not issue. We construct and send an OCSP request for the serial number `deadbeefdeadbeef`, which we found no CA had issued in our dataset.

We performed these checks for the 1,419 responders in our dataset every hour from the University of Illinois Urbana-Champaign campus between September 1–20, 2017.

Given the large number of organizations that misissue certificates, we were surprised to find that 75% of respon-

Organization	Certificates	Rank	% Error	% Warn	All Authorities		CAs w/ >1,000 Certificates		
					Error Percentile	Warn Percentile	Error Percentile	Warn Percentile	
Wosign									
StartCom	536,175	9	2.2%	2.7%	42%	18.2%	73%	43%	
Wosign	195,513	14	5.04%	33.1%	48%	23%	81%	54%	
Symantec									
Symantec	2,767,815	4	0.83%	99.9%	39%	85%	64%	80%	
GeoTrust	1,913,352	5	0.3%	99.5%	36%	84%	55%	75%	
Thawte	459,435	11	0.38%	99.7%	37%	83%	57%	76%	
Verisign	42,622	27	23.07%	98%	67%	83%	90%	72.5%	
PROCERT	39	364	100%	0%	97%	12%	n/a	n/a	

TABLE V: **Distrusted Certificate Authorities**—We show the three major organizations that have lost browser trust compared to the rest of the ecosystem. In 2016 and 2017, all three organizations were misissuing certificates at a rate 2-8x worse than the remainder of trusted organizations.

ders (associated with 89.8% of intermediates and 94.2% of organizations) were available for all 454 checks during our two week measurement period. Just over 9% of responders (134) responders were offline for all checks. These belonged to 73 organizations. Beyond the responders that were never available, there is a large variance in responder uptime (Figure 7a). Baseline Requirements §4.10.2 stipulates that OCSP responders must respond in under 10 seconds under normal operating conditions. As can be seen in Figure 7b, the median response time is 247 milliseconds, but 2 (0.2%) responders have median response times above 10 seconds. Both belonged to WISeKey. In addition, 53 responders had worst-case response times greater than 10 seconds.

There were 25 responders (11 organizations) that incorrectly provided a GOOD response for an unknown certificate. These belonged to smaller authorities, with an median issuance of 24 certificates. The largest of these authorities was Swisscom Smaragd CA 2, which issued 571 certificates in our dataset. Only 10 responders (7 organizations) did not refresh validity periods every four days. These were similarly small authorities, issuing a median 22 certificates. In general, we note a positive trend in OCSP responder health from 2012 when only 8.27% of OCSP queries took less than 100ms [43]. In our experiment, just over 33% of our OCSP requests required 100ms, and in general, there are only a handful of responders in the wild that behave incorrectly.

B. CRL Maintenance

While OCSP is the preferred mechanism for revoking certificates, some CAs continue to support CRLs. When CAs decide to provide CRL-based revocation, they must guarantee the uptime of CRL distribution points per BR §4.10. There are 3.4K unique CRLs, which are associated with 90% of intermediates and 95% of organizations in our dataset.

We observe a smaller number of CRL distribution points that time out for every request compared to OCSP (6 servers, 0.2%). The median average response time is 252 milliseconds, which is well under the BR requirement of 10 seconds (Figure 7b). However, 2 CRL distribution points had worst-case response times of 15 seconds. The rate of the curve increases rapidly after the median point. In other words, the latter 50% of CRLs are far significantly worse than the former 50%. We find that the latter 50% of CRLs generally belong to much

smaller organizations that issue a median 37 certificates. This is similar with our earlier finding that smaller organization misissue a larger fractions of their certificates.

VI. MISISSUANCE AS A PREDICTOR

While adhering to technical standards is innately important for guaranteeing correct behavior with user agents, it is also natural to ask if the misissuance that ZLint detects can act as a predictor for more severe problems like incorrectly validating domain ownership.

In the past year, Google Chrome and Mozilla Firefox took removal steps for three certificate authorities: WoSign, Symantec, and PROCERT [29], [35], [49]. The browsers did not revoke these authorities solely based on failure to adhere to the Baseline Requirements [49]. However, these issuers had some of the worst ZLint misissuance rates in their corresponding classes. Of the 16 authorities with more than 100 K certificates in our dataset, Symantec and Wosign controlled five of the six organizations with the highest misissuance rates (Table V). These were 5–24 times higher than the median misissuance rate for this class of authorities. The sixth organization in the equivalence class is GoDaddy, who misissued 38 K (2.4%) certificates. We manually investigated the certificates GoDaddy issued and find that 95% were issued in late 2012 to early 2013. All suffered from the same problem: missing the required SAN extension. Since the issue was resolved in 2013, GoDaddy’s misissuance rate has been 0.1%, in line with other authorities like Let’s Encrypt and Comodo.

For large authorities with over 100K certificates, there is strong correlation between organization misissuance rate (i.e. percent of certificates containing errors) and browser action. The Point-Biserial Correlation between organization misissuance and browser action is 0.71 ($p = 0.002$). However, we note that while several intermediate organizations were distrusted by browsers, these were subsidiaries of only two companies. We manually categorize the issuer organizations for these top CAs by parent company and find moderate correlation of (0.58, $p = 0.04$). This correlation breaks down when run against all authorities, because there are a large number of very small authorities who have misissued all certificates. 37 organizations misissue more than 90% of their certificates and 18 organizations have a 100% misissuance rate. As such,

	Errors	Warnings	MDSP Mentions	OCSP Uptime	CRL Uptime
Errors	–	0.26 (< 0.01)	0.26 (0.03)	0.10 (< 0.01)	0.07 (0.01)
Warnings	0.26 (< 0.01)	–	-0.19 (0.06)	0.19 (< 0.01)	0.17 (< 0.01)
MDSP Mentions	0.26 (0.03)	-0.19 (0.06)	–	-0.53 (< 0.01)	-0.17 (0.09)
OCSP Uptime	0.10 (< 0.01)	0.19 (< 0.01)	-0.53 (< 0.01)	–	0.59 (< 0.01)
CRL Uptime	0.07 (0.01)	0.17 (< 0.01)	-0.17 (0.09)	0.59 (< 0.01)	–

TABLE VI: **Correlation coefficients and p-values between various measurements.**—We show the correlation between misissuance, mismanagement, and the public profile of a CA. Values with correlation coefficients from 0.1 to 0.3 are weakly correlated, 0.3 to 0.5 are moderately correlated, and greater than 0.5 are strongly correlated.

there is no correlation for situations like PROCERT where 39 (100%) of certificates are misissued.

Beyond the correlation with browser action, there is weak correlation between organizations that sign certificates with warnings and organizations that sign certificates with errors.⁶ There is strong correlation (0.59) between OCSP and CRL mismanagement, but only weak correlation between misissuance (as defined by ZLint) and OCSP mismanagement. This result lends some credence to the notion that misissuance can serve as an indicator for larger organizational mismanagement issues, as PKI members have previously suggested [42].

A. Community Discussion of Misissuance

Given the number of the authorities with exceedingly high misissuance rates, it begs the question: is the community aware of and addressing these authorities’ behavior? To investigate this, we investigated the correlation between mismanagement and misissuance with discussion on the Mozilla Developer Security Policy (MDSP) mailing list. The MDSP mailing list is the primary public forum for discussing web PKI issues and for reporting misissuance [1]. As such, MDSP acts as a “public profile” of a CA. We used the number of times a CA was mentioned doing something unexpected as a proxy for “CA Notoriety” in the community. We crawled the MDSP forum from the beginning of 2015 and searched for links to certificates that were misissued (e.g., a link to crt.sh). We augmented this list by manually inspecting posts for other kinds of mentions that did not include certificate links. In total, we classified 6,069 posts and comments, identifying a total 630 CA mentions. These misissued certificates are associated with 130 intermediate certificates and 83 organizations.

Surprisingly, there is only a weak correlation between ZLint errors (i.e., misissuance) and MDSP discussion of misissuance (Table VI). However, we find that many organizations are mentioned infrequently on MDSP, and as a result, rank ordering in the latter 50% of the dataset is fairly unstable. If we only consider the larger half of the organizations, we observe strong positive correlation (0.88, $p \ll 0.01$). In other words, large authorities that misissue certificates are discussed on MDSP, but there is a lack of conversation around small organizations. This may be because community members are only discussing misissued certificates when they are first logged in certificate

⁶We use the Spearman’s rank correlation test to determine correlation between types of mismanagement. According to Cohen’s guidelines [13], values with absolute correlation coefficients from 0.1–0.3 are weakly correlated, values from 0.3–0.5 are moderately correlated, and values greater than 0.5 are considered strongly correlated.

transparency servers and many of the small authorities have only issued a handful of publicly-known certificates.

Last, we note that discussion is skewed towards misissuance rather than management of other resources like CRLs and OCSP, despite these being a critical part of the Baseline Requirements. This may be because misissued certificates can be found during ex post facto investigation rather than requiring continuous monitoring.

VII. DISCUSSION

It is not immediately clear what the PKI community should do with lint results, or the correlation between misissuance and other types of mismanagement. In this section, we discuss potential routes forward as well as other observations from our investigation.

A. Community Response

Given the positive correlation between ZLint-identified misissuance, other types of operational mismanagement, and more serious concerns, we argue that the PKI community should use lint results along with other data sources (e.g., OCSP monitoring) to identify authorities with worrisome operational practices. While many of the correlations we discussed are not strong under all circumstances, we note that we are not arguing for browsers to take direct action based on lint results. Rather, we argue that lint data can be used to indicate where the community should focus more attention.

This is slightly different from how lint results are being used today. Currently, conversation looks to focus on a specific incident or a single class of error. In most cases, these are errors that larger authorities are making. While addressing these errors will likely lead to better certificate quality, we encourage the community to look beyond individual misissuances and to focus attention on the authorities that are failing to address sustained misissuance.

There is slight risk that constant attention on lints will train authorities to simply pass linters rather than focus on the larger issues at hand. This could remove one of a small number of externally visible indicators of CA behavior before a more serious incident occurs. However, given that the CA/B Forum Baseline Requirements are an evolving document and that authorities are *continuing* to misissue certificates despite the emergence of tools like certlint, this seems unlikely. We encourage authorities to integrate tools like ZLint and certlint into their issuance processes. We expect that this will provide a low-pass filter on organizations looking to reduce their error

profile. Lint errors should further only make up one component of a certificate authority’s public profile. Other aspects such as CRL and OCSP availability as well as CAA adherence can and should be used as indicators. Root stores might consider setting strict limits before authorities are disqualified.

A critical step in focusing attention on the most worrisome authorities is independent, long-term data collection and reporting. Certificate Transparency and services like crt.sh are a huge step forward in addressing problems in the web PKI. However, the ecosystem still needs additional monitors that track other aspects of behavior such as OCSP/CRL uptime and better reporting tools to track CA behavior over time.

B. Small Authorities

As has been noted by many prior studies, there are a worrisome number of small certificate authorities that can sign certificates for nearly any website [5], [19], [47], [48]. In most cases, this concern has been theoretical, based on the nature of the controlling organizations and the concept of “least privilege.” However, as can be seen in Figure 5, most of these small authorities are struggling to follow community standards and correctly issue valid certificates.

PROCERT is one of seventeen authorities that have misissued *every* certificate valid at the time of our study. 37 authorities have misissued 90% of certificates and 76 authorities have errors in at least half of their certificates. Like PROCERT, the authorities with the worst issuance profiles tend to be small. Of the organizations who have issued more than half of their certificates incorrectly, the majority (77%) have issued fewer than 100 certificates and only two have issued more than 1,000 certificates. Unfortunately, as discussed in the previous section, while the misbehavior of medium and large authorities has been discussed, many of these very small authorities have gone unmentioned on public forums like MDSP, but deserve additional attention.

C. Lack of Authority Transparency

CCADB is a significant step forward for tracking who controls CA certificates. However, we find that it is still missing 59 (4.4%) of the intermediates in our investigation. All misissued at least one certificate. We were further surprised to find that while CCADB oftentimes contains the information needed to identify the CA owner (e.g., audit statements), there is no discrete field that indicates an intermediate’s operator. For example, Let’s Encrypt is listed as being owned by both IdenTrust and ISRG, because IdenTrust served as their trust anchor before Let’s Encrypt established their own root certificate. We expect that making this data explicit will ensure more accurate measurement and enable new types of monitoring.

VIII. RELATED WORK

Our understanding of the certificate ecosystem has largely been informed by active probing of HTTPS sites and Internet-wide scanning of the IPv4 space. Holz et al. published a study in 2011 that focused on a lack of standardization across leaf certificates found through active and passive measurements [26]. The Electronic Frontier Foundation (EFF) operates the SSL Observatory, which aggregates certificates based on frequent

IPv4 scans [21]. In 2013, Durumeric et al. analyzed the state of the HTTPS PKI, focusing on certificate collected through aggregate scans of the IPv4 space [19]. Unfortunately, the ecosystem is becoming increasingly opaque to scanning. VanderSloot et al. showed that scanning can no longer be used as the sole source for certificate measurement, and that we must instead leverage a variety of perspectives to further our understanding. As a result, this study utilizes both Censys [17] and Certificate Transparency Logs [27], which cover 99.4% of the certificates observed by VanderSloot et al [47].

In addition to these measurements, there is a long history of work from both academia and industry that have measured and improved the HTTPS and certificate ecosystem. For example, Akhawe et al. published two studies on how to improve the effectiveness of HTTPS browser warnings [4], [5]. In addition, researchers have focused on certificate and PKI errors from both the perspectives of the CA and the browser, and in contexts outside of HTTPS [3], [6], [18], [48]. Recent work from Porter Felt et al. studied the increased adoption of HTTPS, and offers insight into improving adoption further [23]. Finally, there have been a number of papers that investigate the state of revocation in the Web PKI [28], [50].

Our work is not the first to measure certificate compliance with community standards. In particular, Delignat-Lavaud et al. investigated certificate conformance to the BRs on an aggregate level in 2014 [16]. Our work draws on their initial analysis, but provides an update on the ecosystem in 2017. Further, we released ZLint as an open source project, and deployed our linter in production environments to provide actionable feedback to the PKI community. Finally, we focus on using compliance data as a predictor for other security related problems with CAs, which was not present in the original work.

The idea of building a certificate linter was drawn from prior community efforts to build similar tools. In early 2016, AWS Labs published a tool called certlint [7], which checks a myriad of important fields on certificates for conformance. Shortly after, a tool named x509lint was released [36], which broadly checks the correctness of fields in certificates. Unfortunately, we found using these tools to be a challenge, as they lacked consistency and did not entirely cover the clauses in the standards.

In addition to large scale measurements, there has also been focused work on testing the quality of certificate validation in popular libraries [8], [12], [24]. These techniques often employ a combination of fuzzing, differential testing, and symbolic execution to determine if there are code paths in validation libraries that are at odds with standards, or worse, introduce bugs that can be exploited by attackers. Sivakorn et al. used black-box testing on SSL/TLS libraries in order to check the correctness of the hostname validation process in a variety of client-side libraries [40]. Our work sits on the opposite end of these studies—we investigated the quality of the certificates and the organizations that issue them rather than the quality of the libraries and applications that verify them.

IX. CONCLUSION

In this work, we introduced ZLint, a linter that checks certificates for conformance with published technical standards. We ran ZLint against browser-trusted certificates in Censys

and characterized misissuance in the Web PKI. Misissuance has drastically reduced over the past five years, and only 0.2% of certificates contain errors in 2017. However, there remains a long tail of authorities that continue to misissue certificates. 295 authorities misissue at least 10% of certificates and 18 authorities made errors in all of their certificates. Most of the organizations making widespread errors are small CAs and have issued fewer than 1 K certificates. We find that there is correlation misissuance, other types of mismanagement, and in some cases, with browser action. This lends credence to the idea that lint errors can help identify authorities with worrisome operational practices. However, we find that many of the authorities with the worst misissuance rates have not been discussed publicly. We conclude with a discussion about how the community might better use lint data to strengthen the Web PKI.

ACKNOWLEDGEMENTS

The authors thank Jonathan Rudenberg and Rob Stradling for their help and feedback. This work was supported in part by the National Science Foundation under awards CNS 1530915, CNS 1518741, CNS 1409505, and CNS 1518888.

REFERENCES

- [1] Mozilla dev security policy. <https://groups.google.com/forum/#!forum/mozilla.dev.security.policy>.
- [2] J. Aas. 2017.08.10 let's encrypt unicode normalization compliance incident. https://groups.google.com/forum/#!searchin/mozilla.dev.security.policy/nfkc%7Csort:relevance/mozilla.dev.security.policy/nMxaxhYb_iY/AmjCI3_ZBwAJ.
- [3] M. E. Acer, E. Stark, A. P. Felt, S. Fahl, R. Bhargava, B. Dev, M. Braithwaite, R. Sleevi, and P. Tabriz. Where the wild warnings are: Root causes of Chrome HTTPS certificate errors. In *24th ACM Conference on Computer and Communications Security*, 2017.
- [4] D. Akhawe, B. Amann, M. Vallentin, and R. Sommer. Here's my cert, so trust me, maybe? understanding TLS errors on the web. In *22nd International World Wide Web Conference*, 2013.
- [5] D. Akhawe and A. P. Felt. Alice in warningland: A large-scale field study of browser security warning effectiveness. In *22nd USENIX Security Symposium*, 2013.
- [6] J. Amann, O. Gasser, Q. Scheitle, L. Brent, G. Carle, and R. Holz. Mission accomplished? HTTPS security after DigiNotar. In *17th ACM Internet Measurement Conference*, 2017.
- [7] P. Bowen. Certlint. <https://github.com/awslabs/certlint>.
- [8] C. Brubaker, S. Jana, B. Ray, S. Khurshid, and V. Shmatikov. Using frankencerts for automated adversarial testing of certificate validation in SSL/TLS implementations. In *35th IEEE Symposium on Security and Privacy*, 2014.
- [9] CA/Browser Forum. Baseline requirements documents. <https://cabforum.org/baseline-requirements-documents/>.
- [10] CA/Browser Forum. EV SSL certificate guidelines. <https://cabforum.org/extended-validation/>.
- [11] T. Callan. Verisign completes acquisition of geotrust. <https://www.symantec.com/connect/blogs/verisign-completes-acquisition-geotrust>.
- [12] S. Y. Chau, O. Chowdhury, E. Hoque, H. Ge, A. Kate, C. Nita-Rotaru, and N. Li. Syncerts: Practical symbolic execution for exposing noncompliance in X.509 certificate validation implementations. In *38th IEEE Symposium on Security and Privacy*, 2017.
- [13] J. Cohen. Statistical power analysis for the behavioral sciences, 1998.
- [14] D. Cooper, S. Santesson, S. Farrell, S. Boeyen, R. Housely, and P. W. Internet X.509 public key infrastructure certificate and certificate revocation list (CRL) profile. Technical report, 2008.
- [15] Dealbook. Symantec acquires verisign for 1.28 billion. <https://dealbook.nytimes.com/2010/08/10/symantec-acquires-verisign-for-1-28-billion>.
- [16] A. Delignat-Lavaud, M. Abadi, A. Birrell, I. Mironov, T. Wobber, and Y. Xie. Web PKI: Closing the gap between guidelines and practices. In *21st Network & Distributed Systems Symposium*.
- [17] Z. Durumeric, D. Adrian, A. Mirian, M. Bailey, and J. A. Halderman. A search engine backed by Internet-wide scanning. In *22nd ACM Conference on Computer and Communications Security*, 2015.
- [18] Z. Durumeric, D. Adrian, A. Mirian, J. Kasten, E. Bursztein, N. Lidzborski, K. Thomas, V. Eranti, M. Bailey, and J. A. Halderman. Neither snow nor rain nor MITM...: An empirical analysis of email delivery security. In *15th ACM Internet Measurement Conference*.
- [19] Z. Durumeric, J. Kasten, M. Bailey, and J. A. Halderman. Analysis of the HTTPS certificate ecosystem. In *13th ACM Internet Measurement Conference*, 2013.
- [20] Z. Durumeric, E. Wustrow, and J. A. Halderman. ZMap: Fast Internet-wide scanning and its security applications. In *22nd USENIX Security Symposium*, 2013.
- [21] EFF. SSL observatory, 2010.
- [22] L. Encrypt. Chain of trust. <https://letsencrypt.org/certificates/>.
- [23] A. P. Felt, R. Barnes, A. King, C. Palmer, C. Bentzel, and P. Tabriz. Measuring HTTPS adoption on the web. In *26th USENIX Security Symposium*, 2017.
- [24] M. Georgiev, S. Iyengar, S. Jana, R. Anubhai, D. Boneh, and V. Shmatikov. The most dangerous code in the world: Validating SSL certificates in non-browser software. In *19th ACM Conference on Computer and Communications Security*, 2012.
- [25] Globalseign. Globalseign certlint. <https://github.com/globalseign/certlint>.
- [26] R. Holz, L. Braun, N. Kammenhuber, and G. Carle. The SSL landscape: A thorough analysis of the X.509 PKI using active and passive measurements. In *11th ACM Internet Measurement Conference*, 2011.
- [27] B. Laurie, A. Langley, and E. Kasper. Certificate transparency. Technical report, 2013.
- [28] Y. Liu, W. Tome, L. Zhang, D. Choffnes, D. Levin, B. Maggs, A. Mislove, A. Schulman, and C. Wilson. An end-to-end measurement of certificate revocation in the web's PKI. In *15th ACM Internet Measurement Conference*.
- [29] Mozilla. CA:PROCERT issues. https://wiki.mozilla.org/CA:PROCERT_Issues.
- [30] Mozilla. CA:symantec issues. https://wiki.mozilla.org/CA:Symantec_Issues.
- [31] Mozilla. CCADB policy. <http://ccadb.org/policy>.
- [32] Mozilla. Common CA database. <http://ccadb.org>.
- [33] Mozilla. PKI:CT. <https://wiki.mozilla.org/PKI:CT>.
- [34] D. O'Brien. Final removal of trust in WoSign and StartCom certificates. <https://groups.google.com/a/chromium.org/forum/#!topic/net-dev/FKXe-76GO8Y>.
- [35] D. O'Brien, R. Sleevi, and A. Whalley. Chrome plan to distrust Symantec certificates. <https://security.googleblog.com/2017/09/chromes-plan-to-distrust-symantec.html>.
- [36] K. Roeckx. X509lint. <https://github.com/kroeckx/x509lint>.
- [37] J. Rudenberg. Certificate with metadata-only subject fields. <https://groups.google.com/forum/#!topic/mozilla.dev.security.policy/Sae5lpT02Ng>.
- [38] J. Rudenberg. Certificates with invalidly long serial numbers. https://groups.google.com/forum/#!topic/mozilla.dev.security.policy/b33_4CyJbWI.
- [39] S. Santesson, M. Myers, R. Ankney, A. Malpani, S. Galperin, and C. Adams. X.509 internet public key infrastructure online certificate status protocol - OCSP. <https://tools.ietf.org/html/rfc6960>.
- [40] S. Sivakorn, G. Argyros, K. Pei, A. D. Keromytis, and S. Jana. Hylearn: Automated black-box analysis of hostname verification in SSL/TLS implementations. In *38th IEEE Symposium on Security and Privacy*, 2017.
- [41] R. Sleevi. Certificate transparency in chrome - change to enforcement date. https://groups.google.com/a/chromium.org/forum/#!msg/ct-policy/sz_3W_xKBNY/6jq2ghJXBAAJ.
- [42] R. Sleevi. Certificates with invalidly long serial numbers. https://groups.google.com/forum/#!topic/mozilla.dev.security.policy/b33_4CyJbWI.
- [43] E. Stark, L.-S. Huang, D. Israni, C. Jackson, and D. Boneh. The case for prefetching and prevalidating TLS server certificates. In *19th Network & Distributed Systems Security Symposium*, 2012.

- [44] R. Stradling. crt.sh. <https://crt.sh>.
- [45] Symantec. Certificate transparency for Symantec SSL certificates. <https://knowledge.symantec.com/support/ssl-certificates-support/index?page=content&id=AR2177>.
- [46] Thawte. About Thawte. <https://www.thawte.com/about/>.
- [47] B. VanderSloot, J. Amann, M. Bernhard, Z. Durumeric, M. Bailey, and J. A. Halderman. Towards a complete view of the certificate ecosystem. In *16th ACM Internet Measurement Conference*, 2016.
- [48] N. Vratonjic, J. Freudiger, V. Bindschaedler, and J.-P. Hubaux. The inconvenient truth about web certificates. In *Economics of Security and Privacy III*. 2013.
- [49] K. Wilson. Distrusting new WoSign and StartCom certificates. <https://blog.mozilla.org/security/2016/10/24/distrusting-new-wosign-and-startcom-certificates/>.
- [50] L. Zhang, D. Choffnes, D. Levin, T. Dumitras, A. Mislove, A. Schulman, and C. Wilson. Analysis of SSL certificate reissues and revocations in the wake of heartbleed. In *14th ACM Internet Measurement Conference*, 2014.

	Organization	Errors	Warnings	Certificates	% Error	% Warn		Organization	Errors	Warnings	Certificates	% Error	% Warn
1	Let's Encrypt	13	0	36832906	0.00%	0.00%	51	KEYNECTIS	33	4437	4440	0.74%	99.93%
2	COMODO CA Ltd.	3219	7902	6753799	0.05%	0.12%	52	Globe Hosting	0	0	4281	0.00%	0.00%
3	cPanel	131	1	4698298	0.00%	0.00%	53	Dreamcommerce S.A.	5	4262	4262	0.12%	100.00%
4	Symantec Corporation	23053	2764783	2767815	0.83%	99.89%	54	Intermediate Certificate	0	3246	3877	0.00%	83.72%
5	GeoTrust Inc.	5694	1903908	1913352	0.30%	99.51%	55	AlphaSSL	662	0	3848	17.20%	0.00%
6	GoDaddy.com	38215	5186	1623395	2.35%	0.32%	56	FBS Inc	21	0	3807	0.55%	0.00%
7	GlobalSign nv-sa	837	237	1155449	0.07%	0.02%	57	Oracle Corporation	0	1	3784	0.00%	0.03%
8	Western Digital Tech.	0	0	944987	0.00%	0.00%	58	Japan Registry Services	0	3721	3721	0.00%	100.00%
9	StartCom Ltd.	11617	14612	536175	2.17%	2.73%	59	DOMENY.PL sp. z o.o	3	3	3702	0.08%	0.08%
10	DigiCert Inc	2597	1174	464448	0.00%	0.25%	60	HydrantID	1	1848	3662	0.03%	50.46%
11	thawte	1751	457997	459435	0.38%	99.69%	61	Buypass AS-983. . .	7	4	3081	0.23%	0.13%
12	Amazon	5	0	227056	0.00%	0.00%	62	FIDUCIA GAD IT AG	9	3054	3054	0.29%	100.00%
13	TrustAsia Technologies	533	210149	210359	0.25%	99.90%	63	Wells Fargo	3	19	3034	0.10%	0.63%
14	WoSign CA Limited	9849	64714	195513	5.04%	33.10%	64	AffirmTrust	17	20	3008	0.57%	0.66%
15	Entrust	229	57487	143282	0.16%	40.12%	65	Foerderung Deutschen Forschungsnetzes	0	1	2911	0.00%	0.03%
16	Starfield Technologies	111	411	136946	0.08%	0.30%	66	AC Camerfirma S.A.	705	421	2725	25.87%	15.45%
17	Internet2	2	32	85444	0.00%	0.04%	67	Rede Nacional de Ensino e Pesquisa - RNP	0	0	2576	0.00%	0.00%
18	TERENA	1405	183	83789	1.68%	0.22%	68	NetLock Kft.	127	377	2571	4.94%	14.66%
19	Hostpoint AG	716	62975	62975	1.14%	100.00%	69	TBS INTERNET	6	5	2288	0.26%	0.22%
20	Network Solutions L.L.C.	38	7	59989	0.06%	0.01%	70	EUNETIC GmbH	2	0	2136	0.09%	0.00%
21	GeoTrust	215	28354	59113	0.36%	47.97%	71	KPN B.V.	0	808	1957	0.00%	41.29%
22	Actalis S.p.A.[...]	17	1503	55548	0.03%	2.71%	73	Universitaet Stuttgart	534	1534	1830	29.18%	83.83%
23	Cybertrust Japan Co.	643	41608	47301	1.36%	87.96%	74	Plex	0	0	1759	0.00%	0.00%
24	QuoVadis Limited	511	177	44576	1.15%	0.40%	75	QuoVadis Trustlink BV	218	830	1704	12.79%	48.71%
25	Microsoft Corporation	296	44554	44554	0.66%	100.00%	76	DFN-Verein	504	1085	1689	29.84%	64.24%
26	Gandi	34	16	43689	0.08%	0.04%	77	Fraunhofer	0	518	1687	0.00%	30.71%
27	VeriSign	9835	41769	42622	23.07%	98.00%	78	Max-Planck-Gesellschaft	207	1298	1673	12.37%	77.59%
28	CloudFlare	2	14	41006	0.00%	0.03%	79	Swiss Government PKI	4	1623	1623	0.25%	100.00%
29	Trust Provider B.V.	94	33522	33522	0.28%	100.00%	80	TeliaSonera	51	355	1478	3.45%	24.02%
30	STRATO AG	311	32398	32406	0.96%	99.98%	81	Karlsruhe Inst. of Tech.	45	982	1338	3.36%	73.39%
31	Trustwave Holdings	75	6724	28078	0.27%	23.95%	82	Digi-Sign Limited	0	1	1284	0.00%	0.08%
32	Unizeto Technologies S.A.	76	22147	22195	0.34%	99.78%	83	Telecom Italia Trust Tech- nologies S.r.l.	59	28	1226	4.81%	2.28%
33	TAIWAN-CA	30	907	21603	0.14%	4.20%	84	Regionales Hochschulrechen- zentrum Kaiserslautern	417	935	1171	35.61%	79.85%
34	home.pl S.A.	4	15578	15578	0.03%	100.00%	85	Consorci Administracio Oberta de Catalunya	657	4	1117	58.82%	0.36%
35	Verizon Enterprise Sol.	38	1012	15276	0.25%	6.62%	86	Univ. Erlangen-Nuernberg	160	589	1115	14.35%	52.83%
36	USERTRUST Network	77	15	14240	0.54%	0.11%	87	Site Blindado S.A.	0	7	1068	0.00%	0.66%
37	SECOM Trust Systems	451	13087	13111	3.44%	99.82%	88	GANDI SAS	1	4	1067	0.09%	0.37%
38	Google Inc	0	164	11288	0.00%	1.45%	89	IZENPE S.A.	146	1016	1053	13.87%	96.49%
39	Corporation Service Co.	5	2	10746	0.05%	0.02%	90	DHIMYOTIS	0	0	1024	0.00%	0.00%
40	Thawte	51	9587	9709	0.53%	98.74%	91	Volusion	0	676	1013	0.00%	66.73%
41	nazwa.pl sp. z o.o.	3	9318	9318	0.03%	100.00%							
42	SSL.com	2	2	8631	0.02%	0.02%							
43	SwissSign AG	216	53	8588	2.52%	0.62%							
44	CertCenter AG	41	8447	8447	0.49%	100.00%							
45	T-Systems Intl. GmbH	566	1137	8160	6.94%	13.93%							
46	SecureCore	183	4	7946	2.30%	0.05%							
47	Natl. Inst. of Informatics	4	7911	7913	0.05%	99.97%							
48	Trend Micro Inc	1261	25	6374	19.78%	0.39%							
49	KDDI Web Comm. Inc.	0	6357	6357	0.00%	100.00%							
50	Nijimo	0	4615	4615	0.00%	100.00%							

TABLE VII: **Top Issuing Authorities**—We show the authorities that have issued more than 1000 trusted certificates in our dataset, along with their misissuance rates.

	Organization	Errors	Warnings	Certificates	% Error	% Warn		Organization	Errors	Warnings	Certificates	% Error	% Warn
1	Nestle	968	968	968	100.00%	100.00%	45	OpenTrust	12	14	17	70.59%	82.35%
1	PSCProcert	39	0	39	100.00%	0.00%	46	Helmholtz-Zentrum Berlin	29	38	42	69.05%	90.48%
1	Giesecke and Devrient	18	18	18	100.00%	100.00%	47	KBC Group	4	3	6	66.67%	50.00%
1	Unizeto Sp. z o.o.	18	2	18	100.00%	11.11%	47	Hochschulservicezentrum Baden-Wuerttemberg	2	3	3	66.67%	100.00%
1	CertiPath LLC	9	8	9	100.00%	88.89%	49	GeoForschungsZentrum Pots- dam	39	52	59	66.10%	88.14%
1	Helsana Gruppe	8	8	8	100.00%	100.00%	50	Georg-Eckert-Institut	11	14	17	64.71%	82.35%
1	Chunghwa Telecom Co.	7	0	7	100.00%	0.00%	51	Fachhochschule Neu-Ulm	54	81	85	63.53%	95.29%
1	TSCP Inc.	5	5	5	100.00%	100.00%	52	Hochschule fuer Film und Fernsehen Konrad Wolf	18	21	29	62.07%	72.41%
1	Dell Inc.	4	4	4	100.00%	100.00%	53	Baltimore	75	86	123	60.98%	69.92%
1	DigitPA	2	2	2	100.00%	100.00%	54	QuoVadis Trustlink Switzer- land Ltd.	3	3	5	60.00%	60.00%
1	Firmaprofesional SA	2	2	2	100.00%	100.00%	54	The Go Daddy Group	3	2	5	60.00%	40.00%
1	STRAC	2	2	2	100.00%	100.00%	56	Digidentity B.V.	144	238	242	59.50%	98.35%
1	LAWtrust	1	1	1	100.00%	100.00%	57	Consorti Catalunya	657	4	1117	58.82%	0.36%
1	LSU Health System	1	1	1	100.00%	100.00%	58	Firmaprofesional S.A.	25	41	43	58.14%	95.35%
1	SAFE-Biopharma	1	1	1	100.00%	100.00%	59	Hochschule Esslingen	18	30	31	58.06%	96.77%
1	Akademie LPD	1	1	1	100.00%	100.00%	60	InfoCert SpA	29	51	51	56.86%	100.00%
1	Ubizen nv	1	0	1	100.00%	0.00%	61	Universitaet Giessen	95	156	170	55.88%	91.76%
19	ACCV	840	840	844	99.53%	99.53%	62	Hochschulbibliothekszentrum NRW	54	83	98	55.10%	84.69%
20	Cybertrust Inc	320	322	322	99.38%	100.00%	63	GAD EG	24	41	44	54.55%	93.18%
22	Migros	88	89	89	98.88%	100.00%	64	TURKTRUST	6	0	11	54.55%	0.00%
23	KPN Corporate Market B.V.	55	56	56	98.21%	100.00%	65	Universitaet Greifswald	122	188	225	54.22%	83.56%
24	U.S. Government	126	46	129	97.67%	35.66%	66	Deutsches Biomasse- ForschungsZentrum	10	16	19	52.63%	84.21%
25	Xingzheng Yuan	32	1	33	96.97%	3.03%	67	SecureTrust Corporation	11	10	21	52.38%	47.62%
26	VISA	80	7	83	96.39%	8.43%	68	Bundesamt fuer Verbraucher- schutz	12	17	23	52.17%	73.91%
27	Secure Business Services	50	0	52	96.15%	0.00%	69	Fachhochschule Augsburg	29	46	57	50.88%	80.70%
28	CNNIC	21	21	22	95.45%	95.45%	70	Universitaet Marburg	89	155	175	50.86%	88.57%
29	Sonera	17	16	18	94.44%	88.89%	71	Hochschule fuer angewandte Wissenschaften Fachhochschule Coburg	12	23	24	50.00%	95.83%
30	agentschap	76	76	81	93.83%	93.83%	71	Digital Signature Trust Co.	9	8	18	50.00%	44.44%
31	Bechtel Corporation	14	15	15	93.33%	100.00%	71	Suva	3	3	6	50.00%	50.00%
32	Cybertrust	14	10	15	93.33%	66.67%	71	Institut fuer Weltwirtschaft an der Universitaet Kiel	1	2	2	50.00%	100.00%
33	DOUGLAS Holding AG	18	20	20	90.00%	100.00%	71	CRYPTONEO	1	2	2	50.00%	100.00%
34	Intel Corporation	664	738	738	89.97%	100.00%	71	Japan Certification Services	1	0	2	50.00%	0.00%
35	SCEE	17	17	19	89.47%	89.47%							
36	Fachhochschule Gelsenkirchen	14	16	16	87.50%	100.00%							
37	RUAG Holding Ltd	7	0	8	87.50%	0.00%							
38	Leibniz-Institut fuer Neuro- biologie Magdeburg	6	6	7	85.71%	85.71%							
39	Vodafone Group	790	543	937	84.31%	57.95%							
40	AC Camerfirma SA	77	89	92	83.70%	96.74%							
41	Banca d'Italia	111	134	134	82.84%	100.00%							
42	KIR S.A.	8	8	10	80.00%	80.00%							
43	ORC PKI	3	4	4	75.00%	100.00%							
44	Wells Fargo WellsSecure	5	0	7	71.43%	0.00%							

TABLE VIII: **Highest Misissuance Rate Organizations**—We show the authorities with the highest misissuance rate in our dataset, sorted by their total issuance.