What's in a Name? Exploring CA Certificate Control

Zane Ma[†] Joshua Mason[†] Manos Antonakakis[‡] Zakir Durumeric[§] Michael Bailey[†]

[‡]*Georgia Institute of Technology* [§]*Stanford University* [†]*University of Illinois at Urbana-Champaign*

Abstract

TLS clients rely on a supporting PKI in which certificate authorities (CAs)—trusted organizations—validate and cryptographically attest to the identities of web servers. A client's confidence that it is connecting to the right server depends entirely on the set of CAs that it trusts. However, as we demonstrate in this work, the identity specified in CA certificates is frequently inaccurate due to lax naming requirements, ownership changes, and long-lived certificates. This not only muddles client selection of trusted CAs, but also prevents PKI operators and researchers from correctly attributing CA certificate issues to CA organizations. To help Web PKI participants understand the organizations that control each CA certificate, we develop Fides, a system that models and clusters CA operational behavior in order to detect CA certificates under shared operational control. We label the clusters that Fides uncovers, and build a new database of CA ownership that corrects the CA operator for 241 CA certificates, and expands coverage to 651 new CA certificates, leading to a more complete picture of CA certificate control.

1 Introduction

Certificate Authorities (CAs) play a crucial role in modern web security by providing a scalable solution for identity verification. When HTTPS/TLS clients trust the certificates signed by a CA, they are relying on that CA's secure operations: issuance, revocation, key management, etc. Unfortunately, as demonstrated by years of CA mishaps and misconduct [6,54,58], not all CAs are trustworthy. When considering which certificates to trust, accurately identifying the CA operating a certificate is an imperative first step towards judicious and secure trust decisions.

As an example, in 2015, the CA WoSign came under public scrutiny for a series of operational issues [59]. One such problem was a misissuance bug that allowed owners of a subdomain (e.g., `evil.github.com`) to receive certificates for the base domain (i.e., `github.com`). This, along with WoSign's other transgressions, prompted root store operators to begin discussions about removing trust in WoSign certificates. In July 2016, a new discovery revealed that StartCom, a seemingly unaffiliated CA in Israel, was able to issue certificates signed by WoSign (a Chinese company). A deeper investigation of the incident eventually revealed that "the transaction which completed the chain to give WoSign 100% ownership of StartCom completed on November 1st 2015" [59]. Further evidence emerged that StartCom's CA certificates had likely been integrated with WoSign operations as early as December 2015 [56], when the removal of WoSign certificates from root stores appeared imminent. WoSign's stealthy acquisition of StartCom emphasizes the importance of transparency around operational CA control for a secure web. Distrust of WoSign certificates would have still allowed WoSign to stealthily issue trusted certificates through its StartCom CA certificates.

Unfortunately, today's CA certificates contain little reliable information about their operational control. The Subject field, which provides cryptographically attested identity in leaf certificates, does not provide identity or operational control guarantees in CA certificates. For example, consider the root certificate with subject "CN=Hotspot 2.0 Trust Root CA - 03; O=WFA Hotspot 2.0;C=US." The CA operator's identity (DigiCert) does not appear in this example subject name. Due to the laxness of CA certificate name requirements, Subject names are often based on branded product offerings or business partnerships, which provide limited utility for identifying CA control. To complicate matters further, CA certificates are long-lived, lasting up to 37 years[1], and may be exchanged or acquired through business transactions. Prior research on the CA ecosystem mischaracterizes CA certificate operators because they utilize certificate Subject Organization names for mapping CA certificates to CAs, despite their poor suitability for indicating operational control.

The best current solution for identifying CA certificate operators is the Common CA Database (CCADB) [55], a database run by Mozilla to store meta-information about CA

[1] https://crt.sh/?q=68409

certificates. CCADB provides a means for CA administrators to self-disclose their certificates, certificate policies, and audits. CCADB is a step towards improved CA transparency, and we first demonstrate its importance by applying CCADB to prior CA ecosystem results. Unfortunately, even CCADB is not meant to provide ground-truth CA operator information and can only serve as a proxy for CA certificate control.

This work augments CCADB's labeling of CA certificates and creates a new dataset that more accurately maps CA operators. To do so, we built Fides, which constructs a CCADB-independent understanding of shared CA control through three measurement perspectives: certificate issuance configurations, associated CA network infrastructure (revocation checking, chain building), and CA audit statements. One key contribution is the development of a novel fingerprinting technique that detects certificates generated by different configurations of CA issuance software. By grouping fingerprints into an *issuance profile*, we then correlate CA certificates with similar issuance practices. Fides applies the three CA operational vantage points across 2.9B certificates from Certificate Transparency (CT) and 1,266 CA audit statements. These measurements yield clusters of CA certificates with overlapping operations that likely fall under shared CA control.

No publicly available ground truth data exists, making evaluation of Fides's clusters difficult. As an alternative, we collected 28 bug reports that disclosed CA certificates under the scrutiny of Mozilla's root store maintainers as a pseudo ground truth dataset. Fides displays relatively high precision but low recall—it correctly clusters all 52 issuers that it can detect but only detects 31.8% of all CA issuers and 46.6% of operational CA issuers.

Finally, we generated a new dataset of CA certificate control by overlaying Fides's operational clusters with CCADB labels. We manually resolved conflicts arising from the administrative focus of CCADB, identifying 241 CA certificates where the CA operator disagrees with CCADB owner. Through cluster expansion, we also detail the CAs that control 654 previously unlabeled CA certificates. We open-sourced Fides's dataset of 6,849 CA certificates to enable future research and, ultimately, improve CA transparency [1].

2 Background and Motivation

TLS depends on a supporting Public Key Infrastructure (PKI), which provides a scalable mechanism for mapping network identifiers (e.g., domain names) to cryptographic keys. In this section, we outline the key parties in the Web PKI and their roles (Figure 1). We refer the reader to [22] for an in-depth introduction to the Web PKI.

2.1 Certificate Identity and Control

Certificates link an identity to a public key. For subscriber certificates, this identity is the domain or IP address that was validated during the issuance process. However, for CA certificates, the identity is the name of the organization. Specifically, the certificate Subject Common Name (CN) can be any unique string to help the CA identify the certificate, but the Subject Organization must contain the Subject CA's name or a doing-business-as (DBA) / fictitious business name [18].

CA certificates often live longer than CAs themselves, and a certificate's subject can be misleading in the case of a merger or acquisition, or if a CA decides to sell a root to another company. For example, as can be seen in Figure 2, Symantec/DigiCert and Comodo/Sectigo control two certificates that both appear to belong to UserTrust. UserTrust was an independent CA that transferred several of its root certificates to GeoTrust [76], which was acquired by VeriSign [53], then Symantec [3], and ultimately DigiCert [26]. UserTrust and its remaining root certificates were acquired by Comodo [24], which eventually rebranded as Sectigo [23]. While in some cases, it is possible to reassemble a CA certificate's history, many business transactions occur in private and there is often no paper trail that explicitly lays out the transfer of ownership/control of a CA certificate.

In most cases, we are most interested in who controls a CA certificate—the entity that has operational access to the cryptographic keys associated with a certificate and is responsible for the certificates issued by those keys. In this work, we consider a CA certificate operator to be the legal entity that controls the hardware security module (HSM) containing a CA certificate's private key. Intermediate certificates (if technically unconstrained) inherit the trust given to root certificates, but they don't necessarily inherit the same operator. Intermediate certificate control falls into three categories:

1. The intermediate is controlled by the root CA. This is common practice for all root CAs that wish to issue leaf certificates.

2. The intermediate is controlled by the root CA, but legally belongs to a subordinate CA. For example, Sectigo runs a white-labeled CA service for Web.com/Network Solutions [74], but Network Solutions owns the intermediate certificates issued by Sectigo [60].

3. The intermediate is controlled and owned by a subordinate CA. This often occurs when a new CA, such as Let's Encrypt, wants to bootstrap trust through an existing root CA [5].

This work focuses on understanding operational control, distinguishing scenarios 1–2 from scenario 3, rather than legal ownership, which requires more legal expertise.

2.2 User Agent Root Stores

Every User Agent (e.g., web browsers) that validates certificates ships a set of trusted "root" CA certificates that serve

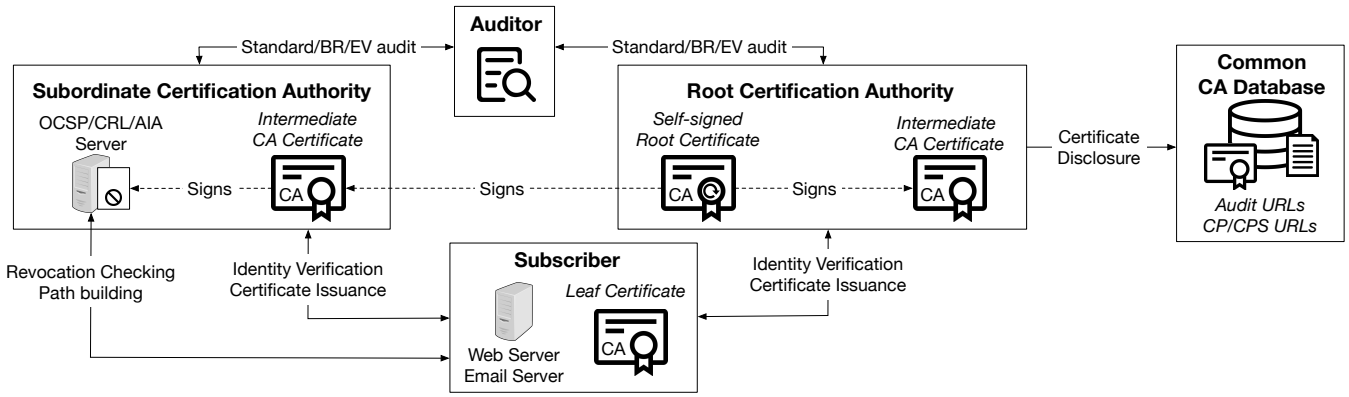

Figure 1: **CA PKI overview**—Root CAs can issue intermediate CA certificates for their own use or for independent subordinate CAs, which operate separate issuance and revocation/path building infrastructure. CAs are required by the NSS root store to disclose audit and policy document URLs for CA certificates through CCADB.

Symantec / DigiCert operated root c38dcb389593…

```
commonName          = UTN-USERFirst-NetworkApplications
orgUnitName         = http://www.usertrust.com
orgName             = The USERTRUST Network
localityName        = Salt Lake City
stateOrProvinceName = UT
countryName         = US
```

Comodo / Sectigo operated root 43f257412d44…

```
commonName          = UTN-USERFirst-Client Authentication and Email
orgUnitName         = http://www.usertrust.com
orgName             = The USERTRUST Network
localityName        = Salt Lake City
stateOrProvinceName = UT
countryName         = US
```

Figure 2: **Misleading Names**—The Subject fields of two roots previously operated by Symantec/DigiCert and Comodo/Sectigo illustrate that 1) the names in CA certificates do not reflect their operators, and 2) similar certificate names have no bearing on shared control.

	Organization	#	Symantec Affiliation
Blacklisted Roots	Symantec	10	–
	VeriSign	14	Acquired by Symantec (2010) [3]
	TC TrustCenter	10	Acquired by Symantec (2010) [75]
	GeoTrust	8	Acquired by VeriSign (2006) [53]
	Equifax	4	Acquired by GeoTrust (2001) [2]
	UserTrust	1	GeoTrust partnership (2001) [76]
	Thawte	10	Acquired by VeriSign (1999) [34]
	RSA Data Sec.	1	Spun out VeriSign (1995) [33]
Whitelisted	Apple	6	Sub-CA intermediates
	Google	1	Sub-CA intermediates
	DigiCert	2	Cross-signed DigiCert roots
	DigiCert	2	Transition intermediates

Table 1: **Symantec Distrust**—Blacklisting of Symantec-controlled roots involved 58 root certificates [4] with 8 separate orgs. in their X.509 Subject field. These orgs. are linked through a scattered history of corporate spin-offs and acquisitions.

as the root of trust in the Web PKI. CAs rarely use root certificates to directly sign leaf "subscriber" certificates (e.g., the certificate for a website). Rather, root certificates sign intermediate CA certificates, which handle day-to-day signing of subscriber certificates. Root certificates can thus remain offline, protecting them from compromise. This is a necessary precaution due to the difficulty of updating root stores. While new roots are added/removed as CAs emerge, dissolve, or adopt new technology, it can take years for a new root to propagate to clients and become globally reliable. Intermediate CA certificates are also used to delegate trust to third parties.

Products typically have their own root stores or borrow the root store of another product; each product also has its own requirements for including a CA in its root store. For example, Mozilla requires roots to publicly disclose unconstrained intermediates in Common CA Database (CCADB) [55], a Mozilla-operated repository, while Microsoft does not. CAs demonstrate their compliance with root store requirements by publishing Certificate Policies (CP) and Certification Practice Statements (CPS) that describe how the CA operates. CAs then enlist a third-party accredited/licensed auditor to verify the CA's compliance with their own written policies and public standards like the CA/B Forum Baseline Requirements [18], and either the European Telecommunications Standards Institute (ETSI) criteria or the WebTrust criteria. The CP, CPS, and audit documents provide a detailed look at a CA's operations and management. Several browser operators, including Microsoft and Mozilla, require that root CAs register their CA certificates along with links to audit/CP/CPS documentation in CCADB.

2.3 Operational Consequences

CA ownership data is critical to root store operators as exemplified by the distrust of Symantec roots in 2017. Between 2009–2017, Symantec repeatedly misissued certificates [58], and as a result, Google Chrome [62], Mozilla [49], Apple [10],

and Microsoft [52], discontinued their trust in Symantec-issued certificates. Identifying Symantec-controlled CA certificates required significant manual investigation of CA audits, CA operational characteristics, and corporate ownership structure. In total, Chrome blacklisted 58 root certificates belonging to eight Subject Organizations, seven of which had no direct indication of Symantec ownership (Table 1). Intermediate CA certificates also require attribution because root CAs can delegate intermediate certificates to independent "subordinate" organizations. For example, even after all Symantec roots were distrusted, not all of their child intermediates were subject to distrust. Apple and Google both operated subordinate-CAs (sub-CAs) that chained to Symantec's roots, and these intermediates were explicitly whitelisted and exempt from distrust due to their independent operation.

2.4 Research Consequences

Since no methods have previously existed for mapping CA certificates to their operators, existing research has defaulted to the information available in X.509 chains when attempting to characterize the CA ecosystem. For example:

CA ecosystem. Studies of the CA and certificate ecosystem [21, 28, 44, 47, 77] have aggregated certificates based on their Subject Organization names. Figure 3 provides a brief comparison between the perspectives provided by Subject Org. names and CCADB, which maps more closely to CA operation as detailed in the following section. At the root certificate level, Subject Orgs. and CCADB labels are within the same order of magnitude. However, the number of Subject Orgs. for intermediate certificates significantly exaggerate the diversity of the CA ecosystem. Similarly, the number of CAs responsible for 50% of the CA ecosystem shrinks from 54, based on Subject Org., to 4, based on CCADB data. Reframing prior studies and performing future studies under the context of CA certificate control will lead to more accurate and actionable results.

BGP attacks on domain validation. Previous work in 2018 [11] identified that Symantec was vulnerable to BGP hijacking attacks during domain control verification. However, in 2017 DigiCert acquired Symantec's Website Security and PKI solutions, and Symantec's CA certificates had transitioned to DigiCert operational control by December 1, 2017 [69]. Future analysis could identify whether the reported issues were specific to Symantec and its web of CA acquisitions (Section 2.3), or if they were systemic throughout DigiCert, which is the largest CA by distribution of roots and intermediates.

Phishing certificates. A study from 2019 [50] identified the top ten issuers of phishing certificates, listing Let's Encrypt as the most common issuer at 34.4%, followed by cPanel at 22.2%, and RapidSSL at 9.1%. When we take CA operators into account, we find that cPanel is actually operated by

	Certs	Subject orgs	CCADB owners
All Trusted Roots	366	178	130
Microsoft Roots	354	176	130
Apple Roots	166	83	60
NSS Roots	147	72	52
All Trusted Intermediates	3,447	637	90
All Trusted Certs	3,813	685	132

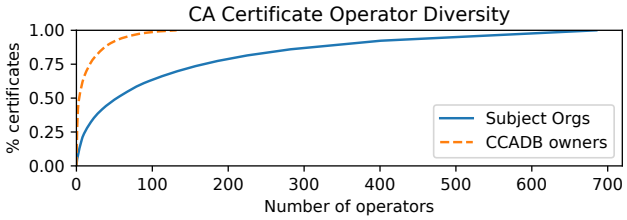

Figure 3: **CA Certificate Owner Perspectives**—Certificate Subject organization names exaggerate the size of the CA ecosystem. CCADB hints at more condensed CA certificate control, with a few major players.

Sectigo, and that Sectigo controls three (cPanel, COMODO RSA, COMODO ECC) of the top five most common issuers. Similarly, DigiCert operates four (RapidSSL TLS RSA, CloudFlare Inc ECC CA-2, DigiCert SHA2, RapidSSL CA) of the top ten phishing certificate issuers. Taken together, the top three CAs issuing the most phishing certificates would be Let's Encrypt (34.4%), Sectigo (32.8%), and DigiCert (18.9%). This CA control perspective reveals that phishing certificates are concentrated not just within Let's Encrypt, but Sectigo as well.

2.5 Potential Sources of Truth

CCADB—the CA certificate database that major browsers jointly maintain—presents an enticing alternative to the unreliable Subject fields found in CA certificates. Although CCADB was publicly accessible as early as 2016, academic PKI research has largely overlooked it. CCADB's low utilization likely stems from poor public awareness and its misalignment with CA operations.

CCADB does not currently have a field for the legal entity that manages each intermediate or root certificate. The closest field is "CCADB owner" field, which denotes the Salesforce account that is responsible for administrative reporting. For example, although DigiCert acquired QuoVadis and assumed control of all CA operations in January 2019 [27], both DigiCert and QuoVadis exist as separate CCADB owners as of July 2020. Historically, independent subordinate CAs were also disclosed under their root CA in CCADB (e.g. Apple intermediates were disclosed under DigiCert and Sectigo). To address this discrepancy, in March 2019 CCADB began reporting intermediate CA certificates with their own audit state-

ments that are not inherited from the parent certificate. Since independent audits often indicate independent operation, this reporting expanded the utility of CCADB for mapping CA certificate control. However, even CCADB's subordinate CA labels are not necessarily representative of actual CA certificate control. For example, as detailed in Section 4.1, Let's Encrypt's cross-signed certificates from IdenTrust are not disclosed as an independent subordinate CA, since they are listed under an IdenTrust audit (which states, ironically, that the cross-signs are not covered by the audit).

Second, not all CA certificates are disclosed through CCADB, since not all root store operators (e.g., Apple) require public disclosure. Furthermore, Mozilla only requires CCADB disclosure for technically unconstrained certificates, which allows for certificates to go unlabeled. Subsequent analysis in Table 2 reveals that 665 (20%) of trusted issuers (by Subject + SPKI) are missing from CCADB.

Audit, CP, and CPS documents supplied by CAs initially appear to provide the basis for identifying the organization that controls each CA certificate. However, this is often not possible in practice. For example, in the case of the Symantec distrust, Thawte, GeoTrust, and Symantec all submitted independent audits that did not indicate the relationship between the companies. Furthermore, audits are composed of plaintext English rather than structured data, which would require manual analysis for thousands of CA certificates to uncover ownership details.

3 Fides: A System for Uncovering Ownership

Building on the observation that certificates issued by the same organization are likely to be structured similarly, we introduce Fides, which we use to aggregate and label CA behavior through fingerprinting of certificate generation software, network infrastructure, and audit details. By clustering on these features, we can detect when CA certificates are controlled by the same party, and using CCADB to seed our labeling process, we build a new dataset that more accurately depicts CA certificate control.

3.1 Data Collection

We began our investigation by collecting 2.9B certificates available prior to July 1, 2020 from all CT logs trusted by Google Chrome or Apple [9,36]. We observed 121,482 unique CA certificates and then filtered out 117,106 ad hoc CA certificates issued by Google to check CT server uptime [37]. We also included 2,240 CA certificates that were disclosed in CCADB, but not present in CT. We labeled the CA certificate data with the trusted root certificates for Apple, Microsoft, and Mozilla NSS as of July 1, 2020 as well as revocation data from NSS OneCRL and Chrome CRLSets revocation lists.

Although CT provides a complete certificate chain for a given certificate (that leads to a trusted root for a given CT log), the presented certificate chain may represent just one of many valid certificate chains. For example, consider a chain of two certificates, *A* and *B*, where *A* is the issuer of *B*. Now consider a third certificate *C*, that has the same Subject+Subject Public Key Info (SPKI) as *A*. The chain of *A–B* does not preclude the possibility that *C* actually issued *B*, since *C–B* is a valid chain as well. This is a consequence of the flexible design of X.509 certificate chaining, which considers any certificates with the same Subject+SPKI (SSPKI) to be interchangeable parents. Given a child certificate, only the SSPKI of the parent certificate(s) can be determined, and we perform parent/child analysis at the granularity of unique SSPKI pairs.

As part of our data acquisition, we independently verified the certificate chains for each CA certificate and discovered 32 certificates containing signature algorithm inconsistencies[2]. We removed these certificates, yielding a final dataset of 2.9B trusted leaf certificates and 9,154 CA certificates accounting for 6,549 unique SSPKI pairs (Table 2).

3.2 Fingerprinting Leaf Certificates

We first analyze the ASN.1 structure of the leaf certificates that each CA certificate has signed and then identify clusters of consistent certificate structure. This is possible because CAs have considerable freedom in how the certificates they generate are structured, particularly for the fields in the Subject DN and data included in X.509 extensions. For example, one CA's certificate generation software might only create certificates with 2048-bit RSA public keys, while another may always include both HTTP- and LDAP-based revocation URLs. X.509 certificates are structured in an ordered tree, following the hierarchical ASN.1 data format. Each non-leaf ASN.1 node, including the root, represents a compound ASN.1 field that has one or more sub-fields. Each leaf node contains the value for a field, which can be a string, integer, OID, etc.

We analyze a certificate's ASN.1 tree structure without leaf node values as our certificate fingerprint abstraction. This is because while certificate structure is relatively stable, the values within the structure are not. For example, a certificate's public key should be generated at random. Excluding leaf node values from an ASN.1 certificate focuses on differences caused by different certificate generation software/configuration, rather than differences arising from required high-entropy fields (e.g., serial number) or user input (e.g., Subject Name). The one general exception to this is enumerable values that are denoted by an Object Identifier (OID) ASN.1 node, which do not introduce high entropy. Extension types, for instance, are specified by an OID and are more indicative of software configuration rather than input diversity or required high-entropy fields.

[2]The signatureAlgorithm field in the TBS Certificate does not match the second signatureAlgorithm field after the TBS Certificate, or does not contain a known OID.

	Issuers	CCADB	Cert FPs	Cert URLs	Audits	Fides
All Trusted Roots	*359*	*354 (98.6%)*	*330 (91.9%)*	*296 (82.5%)*	*204 (56.8%)*	*343 (95.5%)*
Microsoft Roots	352	352 (100.0%)	325 (92.3%)	292 (83.0%)	203 (57.7%)	337 (95.7%)
Apple Roots	165	158 (95.8%)	164 (99.4%)	157 (95.2%)	123 (74.5%)	164 (99.4%)
NSS Roots	147	147 (100.0%)	144 (98.0%)	143 (97.3%)	128 (87.1%)	147 (100.0%)
All Trusted Intermediates	*3,058*	*2,375 (77.7%)*	*1,858 (60.8%)*	*1,783 (58.3%)*	*1,736 (56.8%)*	*2,583 (84.5%)*
Microsoft Intermediates	3,031	2,364 (78.0%)	1,844 (60.8%)	1,773 (58.5%)	1,725 (56.9%)	2,558 (84.4%)
Apple Intermediates	2,493	2,168 (87.0%)	1,502 (60.2%)	1,469 (58.9%)	1,522 (61.1%)	2,110 (84.6%)
NSS Intermediates	2,366	2,135 (90.2%)	1,381 (58.4%)	1,355 (57.3%)	1,564 (66.1%)	2,019 (85.3%)
All Trusted Certs	*3,338*	*2,673 (80.1%)*	*2,111 (63.2%)*	*2,007 (60.1%)*	*1,909 (57.2%)*	*2,847 (85.3%)*
All CA Certs	*6,549*	*4,845 (74.0%)*	*4,363 (66.6%)*	*3,898 (59.5%)*	*2,868 (43.8%)*	*5,613 (85.7%)*

Table 2: **Data Coverage**—Fides's combined datasets miss sixteen trusted root issuers (subject+SPKI), and include 84.5% of trusted intermediate CA issuers.

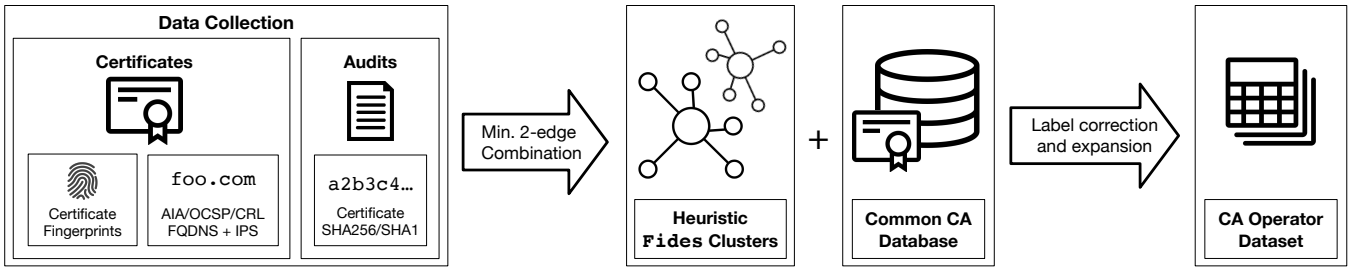

Figure 4: **Fides**—Integration of certificate- and audit-based data sources creates heuristic clusters that approximate shared CA certificate control. Combined with CCADB, Fides outputs a dataset of CA certificates and their operators.

Figure 5 provides a sample fingerprint, which demonstrates some of the certificate properties that it captures: the type of cryptographic keys as well as the type and order of extensions. Issuer Names are excluded from the fingerprint, since the absence or presence of the components are dictated by specific issuing certificates rather than certificate generation software. For extensions, X.509 certificates abstract extension data into an ASN.1 *octet string* field so that new and unknown extensions can be safely parsed. Extensions often have custom data formats that override the standard ASN.1 types, so we implement custom parsers for extensions to further increase the precision of our fingerprints. To encourage further research with this technique, we have open sourced our certificate ASN.1 fingerprinting tool [48].

Case study. To better understand the utility of certificate fingerprints, we performed a case study examining the certificates issued by three of the top CAs by issuance volume: Let's Encrypt, Sectigo, and DigiCert. Using CCADB as a rough approximation of CA certificate control, we find that Let's Encrypt has issued certificates with 66 distinct fingerprints while DigiCert (21,856) and Sectigo (23,576) have issued over three hundred times more fingerprints. This reflects the narrowness of Let's Encrypt's automated CA operations, which only issue domain-validated certificates from a single CA software, Boulder [46]. Sectigo, DigiCert, and Let's Encrypt fingerprints are disjoint, with the exception of 11 fingerprints that are shared between one Sectigo and two DigiCert CA issuers. These overlapping fingerprints occur in CA certificates labeled "TAIWAN-CA INC.", which suggests either the same issuance software/configuration between Sectigo and DigiCert that does not appear in any other DigiCert or Sectigo certificates, or the presence of an undisclosed sub-CA under the control of Taiwan CA. Fortunately, the DigiCert certificates expired in September 2016 and Sectigo certificate in May 2020, reducing the potential danger of improper disclosure in this particular case.

Looking at the top twenty most common fingerprints and their issuers amongst these three CAs (Figure 6), we observe that CAs are often responsible for more than one fingerprint. To best capture the operational nature of each issuer, we compare each issuer's *issuance profile*, which is the full set of fingerprints issued by an issuer. We use the following modified Jaccard similarity metric with a heuristic threshold of 0.5 to account for issuance profiles of different sizes:

$$J_{mod}(A,B) = \frac{|A \cap B|}{min(|A|,|B|)}$$

Figure 6 highlights three different operational snapshots. The DigiCert issuance profiles form two disjoint clusters: issuers with fingerprints 1–8 and 18, and issuers with fingerprints 9–17 and 19–20. Manual inspection of the certificates and audits in these two clusters reveal that the first cluster belongs to the "Citizen CA," which is the PKI used for Bel-

```
16 # x509 certificate root
   16 # TBS certificate
      0
         2
      2
      16 # Signature Algorithm
         6.1.2.840.113549.1.1.11 # sha256WithRSAEnc.
         5
      16 # Validity
         23
         23
      16 # Subject
         17
            16
               6.2.5.4.3 # Common Name
               19
      16 # Subject Public Key Info
         16
            6.1.2.840.113549.1.1.1 #rsaEncryption
            5
         3
      3 # Extensions
         16
            16
               6.2.5.29.15 # Key Usage
               1 # Critical
               4
                  3
                     100001 # digSig, keyCertSign
               16
                  6.1.3.6.1.5.5.7.1.1 # AIA
                  4
                     16
                        16
                           6.1.3.6.1.5.5.7.48.1 # OCSP
                           6
                        16
                           6.1.3.6.1.5.5.7.48.2 # Iss.
                           6
   16
      6.1.2.840.113549.1.1.11 # sha256WithRSAEnc.
      5
   3
```

Figure 5: **Sample certificate fingerprint**—Each node label is the ASN.1 universal tag type [61], with OID values added as a suffix for OID tag type 6. Extensions may override default ASN.1 tag types.

gium's electronic identity system. According to CA documents [19], Citizen CA is operated by Certipost, although a majority of intermediates are disclosed under DigiCert. The second DigiCert cluster contains an assortment of CA certificates operated by DigiCert itself. Sectigo's issuers display a patchwork of fingerprints with no clear clusters, which likely indicates diverse but shared issuance operations. Finally, Let's Encrypt demonstrates relatively restricted issuance across two issuers. The first issuer represents Let's Encrypt's X3 intermediate which was in operation during its introduction of support for elliptic curve public keys, pre-certificates, and OCSP Must-Staple extension. The second Let's Encrypt issuer is the retired X1 intermediate that issued a less diverse set of early certificates.

3.3 CA Network Infrastructure

CA network infrastructure (e.g., OCSP servers) can also hint at shared operation. We investigate the operational infrastructure used for online chain building (Authority Information Access (AIA) CA Issuer) and certificate revocation (CRL and OCSP). This infrastructure can be closely tied to the issuing CA certificate since child revocations are often signed by the issuing certificate, and AIA CA Issuer URLs provide copies of the issuing certificate. Other certificate fields that contain URLs, such as the Certificate Policies extension, do not relate directly to operational functions. In total, we extracted 2,334 FQDNs embedded within child certificates, which were composed of 991 OCSP names, 938 CRL names, and 800 AIA Issuers. We performed A-record DNS lookups for each FQDN, resulting in a total of 835 IPv4 addresses in 309 Autonomous Systems (ASes). Figure 7 presents the distribution of different network names and addresses across CA certificates. Approximately half of all FQDNs were only associated with a single CA certificate, which might suggest relatively isolated operations. However, the distribution of IPs have a much shorter tail, which indicates that many FQDNs share the same underlying IP addresses. To identify the shared IP addresses that are indicative of shared CA operation, we filtered out all ASes belonging to CDN networks, which can co-locate unrelated network services. After this filtering (11% IPs removed), we linked CA certificates with IPs within the same /24 subnet or with exact-match OCSP/CRL/AIA hostnames.

3.4 CA Audits

While certificate fingerprints and network infrastructure directly measure operational features to link CA certificates under shared control, CA audits provide a complementary data perspective. CA audits report on CA operations as examined by a third-party, professionally qualified auditor. In addition to disclosing a CA's conformance or deviation from its CA policies, CA audits often include a listing of certificates within the scope of the audit. CCADB retains a collection of all CA audits, collected from public data sources, which we downloaded, resulting in 1,266 PDF files. To convert audit PDF documents to text, we preprocessed all PDFs with Adobe Acrobat's OCR tool, since many documents contained full-page images, rather than actual text. Second, to extract text from the PDFs, we opened each file in Adobe Acrobat, selected all text, and copied it into a text file. This method was chosen after testing multiple PDF-to-text solutions (including `pdftotext`), which each had difficulty preserving the spatial relationships between text elements[3]. For the subset of PDFs that did not allow text extraction, we utilized Google Document's PDF to text conversion feature. We developed a set of simple regular expressions to extract the certificate SHA-256 fingerprints that were within scope of each audit document. As part of the extraction process we observed that some CA audits only contained alternate hashes, such as SHA1, which we developed regular expressions for as well.

[3]PDFs prioritize universally consistent rendering and specify the location of text elements, rather than their logical grouping, and many solutions extract text from left-to-right, top-to-bottom. This caused issues for tables with wrapped text columns, as an example.

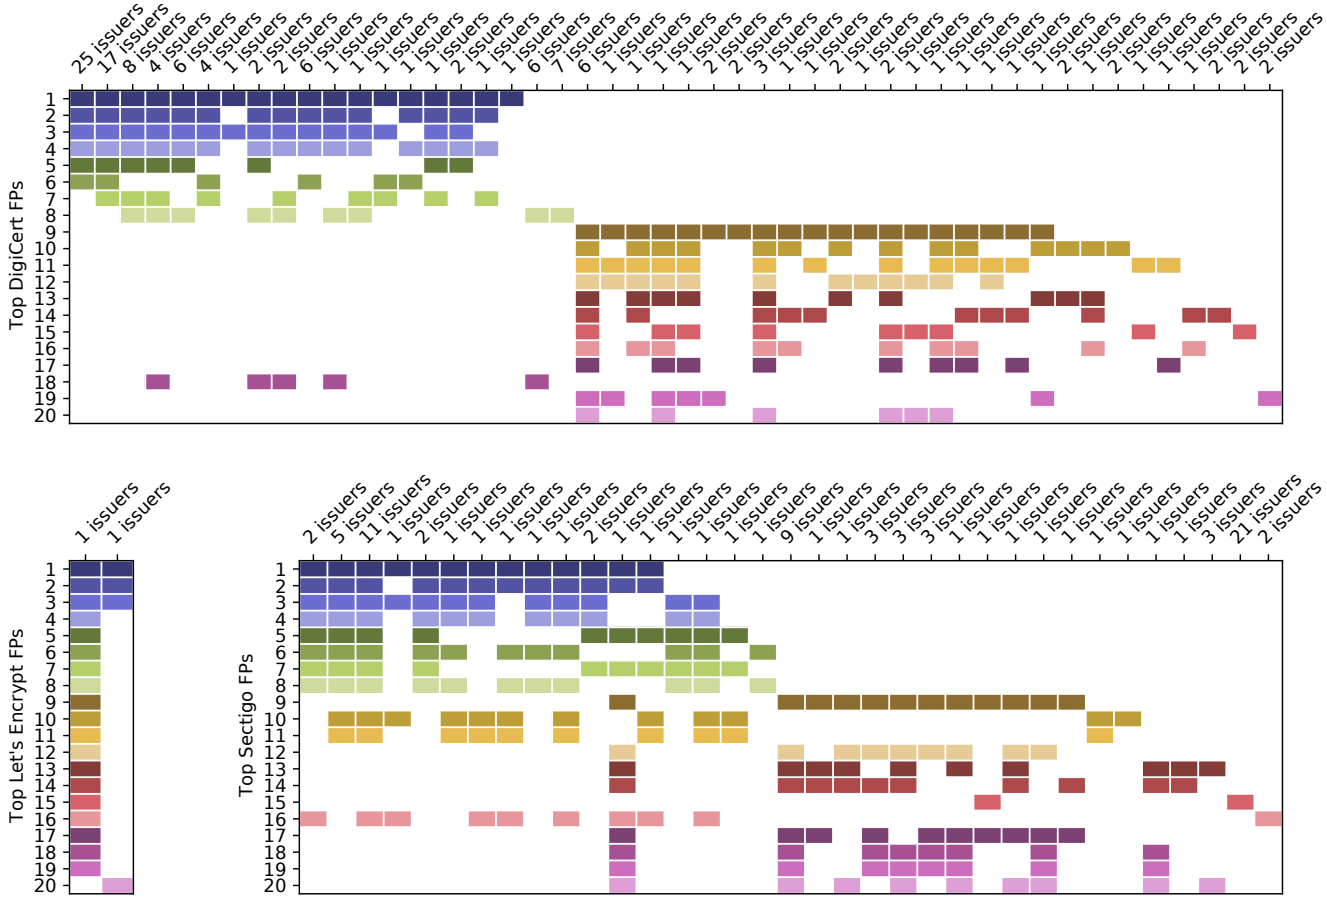

Figure 6: **Top Fingerprint Issuance**—The top twenty issued fingerprints for each of the top three CAs by volume reflect differences in certificate issuance. Let's Encrypt employs only two issuers that contain overlapping issuance profiles, while DigiCert and Sectigo display a multitude of issuers with varying overlap. DigiCert's issuers generate two disjoint sets of fingerprints that reflect independent CA operations (Certipost and DigiCert).

3.5 Combining Perspectives

We took a conservative approach to combining the audit, network, and certificate-based techniques discussed previously. Only certificates that were associated through at least two perspectives are considered to have common CA operation. In general, we expect well connected CA certificates—those that issue similarly fingerprinted child certificates, fall within scope of the same audits, and/or utilize the same network infrastructure—to belong to the same operational control. After applying this criteria, Fides generated 320 clusters containing 2,599 CA issuers, with clusters ranging in size from 2–696 CA certificates. We subsequently utilized these clusters to identify discrepancies between CCADB owner labels and Fides's heuristic clusters of shared operational control.

To better understand the contribution of each perspective, we measured the co-occurrence of the three techniques (Table 3). Audit disclosure is the largest source of linkages between CA certificates (388k), followed by overlapping network infrastructure (188k); however, these represent noisier data that did not match other sources. The differing rates of co-occurrence for each perspective indicates their precision. Certificate fingerprints (99% overlap) have relatively high precision, whereas audit and network features are less precise ($\leq$ 39% overlap). The combination of these diverse perspectives provides a more complete picture of CA operations and can guide manual investigation of CA certificate control.

3.6 Limitations

The novelty of this work yields its primary limitation: little ground truth data exists to evaluate the accuracy of Fides. We recognize this limitation and work to reduce its impact. Fides's results do not constitute ground truth data; instead we use its multi-layered aggregation of perspectives to identify higher-level certificate control inconsistencies within CCADB and point develop a new dataset that better aligns with CA operator transparency.

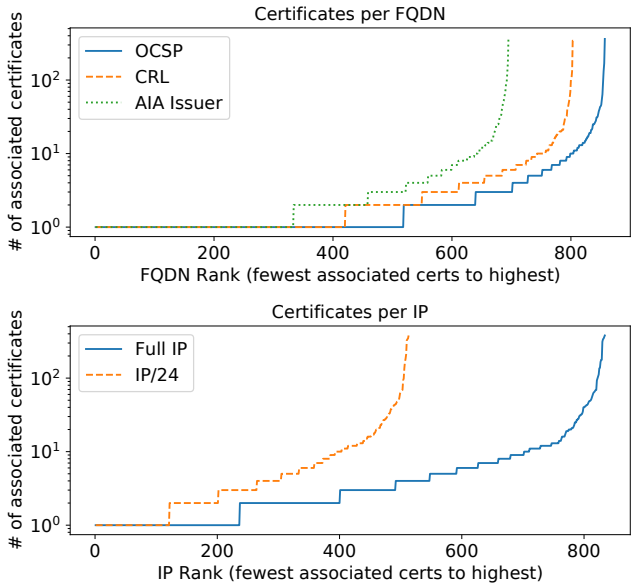

Figure 7: **Certificate URL and IPs**—Distribution of network infrastructure used for OCSP, CRL, and AIA Issuer URLs in CA certificates.

	Cert FPs	Network	Audit
Cert FPs	–	28.8k (15%)	4.4k (1%)
Network	28.8k (98%)	–	47.9k (12%)
Audit	4.4k (15%)	47.9.6k (26%)	–
Shared	29.1k (99%)	72.6k (39%)	48.2k (12%)
Total	29.6k (100%)	188.2k (100%)	387.9k (100%)

Table 3: **Edge Co-occurrence by Perspective**—The co-occurrence rates of each technique highlight their individual precision, led by certificate fingerprints. Network infrastructure and audits account for the most CA associations, as indicated by shared edge counts.

The certificate- and document-based properties that we observe are not guaranteed indicators of CA certificate association. It is possible, for example, that two independent CAs coincidentally issue certificates with the same ASN.1 fingerprints and share AIA Issuer/OCSP/CRL infrastructure. To help mitigate these false positive associations—independent CA operations that are wrongly associated—we take a conservative approach for each individual perspective. For example, we designed certificate fingerprints to err on the side of high precision (87,009 unique fingerprint-profiles across 4,376 issuer Subject+SPKI), and our work combines unique perspectives to strengthen confidence in common CA control.

On the flip side, there are likely false negative associations as well—common CA control that is not identified by our methods. Such instances can arise as a result of complex CA operations, CAs with a single CA certificate, omissions in audit documentation, undisclosed certificates, etc. Discovery of missing associations proves to be a challenge, but will improve as the CA community moves towards increased public disclosure and documentation.

	Bug Reports	Correct	Issuers	Correct
All Issuers	28	3 (10.7%)	150	48 (32.0%)
Active Issuers	22	7 (31.8%)	103	48 (46.6%)

Table 4: **Fides evaluation**—Fides's ability to identify previously undisclosed CA certificates ranges from 32% of all issuers, to 47% when excluding inactive issuers that do not issue certificates published in CT.

3.7 Evaluation

Without ground truth data on who controls CA certificates, it is difficult to directly evaluate whether Fides correctly identifies their owners. Instead, we evaluate whether Fides is able to correctly detect ownership in the cases where there is a Mozilla Bugzilla ticket that establishes clear ownership details. Mozilla's root store policy requires timely CCADB disclosure of all technically capable CA certificates. When community members notice the lack of proper disclosure, issues are created to investigate and disclose ownership in CCADB. We identified 28 instances between May 2014 to July 2019 of delayed or invalid disclosure bugs (Appendix A). These bug reports contain certificates that were manually examined by root store maintainers prior to CCADB disclosure and we use them as approximate ground truth data to evaluate Fides.

We extracted the CA certificates and issuers (Subject+SPKI) associated with each bug report and manually determined CA ownership based on bug details. We then mapped each CA certificate to its corresponding Fides cluster. Since Fides clusters are initially unlabeled—they group operationally similar certificates without identifying an CA operator—we manually labeled the clusters that contained CA certificates found in bug reports. This manual identification entails looking at the individual components of Fides: certificate fingerprints, network infrastructure, and audit statements to assign a likely CA operator for each target. We then compared the manually identified CA operator(s) with the CA disclosed in each bug report. For instance, bug report #1503638 contains a CA certificate found in a cluster of 11 certificates that are covered by WISEKey audits and utilize shared network infrastructure (e.g., `ocsp.wisekey.com`).

In total, we examined 28 bug reports that spanned 150 issuers. Fides correctly identified the operators of 32% of 150 issuers and was able to correctly label all issuers in only 3 of the 28 bug reports. This is in part because 47 unidentified issuers were *not operational* and had not issued any certificates, which prevented us from fingerprinting them. Excluding these, Fides was able to correctly identify 47% of operational CA

issuers, and correctly label all issuers in 7 of 22 bug reports. For example, for #1499585, Fides is able to provide more detailed information than the bug report itself. DigiCert disclosed 14 issuers, most of which are identified by Fides as DigiCert, but one of which is correctly identified by Fides as Cybertrust Japan, an independent sub-CA. As a whole, Fides has relatively high precision, since all 48 issuers within Fides appeared in the correct cluster (i.e., no certificates for CA A occurred in a cluster dominated by CA B), but low recall, only accounting for about half of unlabeled issuers.

4 Towards CA Transparency

In this section, we utilize Fides to uncover CA certificates that have incorrect or misleading ownership data in CCADB. We first label CA certificates with their CCADB owner, cluster CA certificates using Fides, and detect when clusters have conflicting CCADB owners and when unlabeled CA certificates belong to a labeled cluster. We manually investigate these incongruities by examining audit documentation and operational features, and construct a new dataset that aligns more closely with CA certificate operational control.

4.1 Labeling Fides Nodes

CCADB tracks individual CA certificates (i.e,. by SHA-256 fingerprint) whereas Fides tracks issuers by SSPKI. When we group CCADB certificates by SSPKI (Appendix 9), we find 39 issuers (110 certificates) that map to more than one CCADB owner. 31 SSPKIs contain certificates that are revoked, expired, or properly disclosed as sub-CAs (i.e., different controlling owner). For the remaining 8 SSPKIs (20 certificates), CCADB presents control ambiguity with a single key appearing to have multiple CCADB owners. This includes 4 Let's Encrypt intermediates that are cross-signed by IdenTrust and disclosed under the IdenTrust CCADB owner, rather than the Internet Security Research Group (ISRG). The cross-signed certificates are also disclosed in the IdenTrust audit, which explicitly declares "the cross-signed certificates are not controlled by IdenTrust" [71]. This misrepresentation of CA certificate control is not a violation of disclosure policies—it merely highlights limitations of CCADB's record model. We manually identify a single CCADB owner for each of the ambiguous issuers by examining audits and certificate subjects. As part of this process, we discovered an improperly disclosed Subordinate CA by Camerfirma [17], which added to the growing list of compliance issues that recently prompted discussion about Camerfirma's fitness for Mozilla's root store [78].

After resolving the 39 SSPKI ownership conflicts (correcting 64 certificates), we find 557 CA certificates not present in CCADB that share an SSPKI with a CCADB-labeled certificate. Because a shared SSPKI represents the same cryptographic keying material, we expand CCADB labels to these CA certificates. We characterize these newly labeled certificates in Section 4.4. Having resolved SSPKI ownership inconsistencies, we next identify more subtle discrepancies and absences in two ways.

4.2 Multi-operator Clusters

The presence of multiple CCADB owners within a single Fides cluster points to likely misalignment between CCADB labels (including sub-CA reporting) and CA certificate control as automatically inferred by Fides. We identified eleven such instances (Figure 5) and then manually inspected each cluster to determine the root causes of mismatch, described below. Two are false positives, and the remaining nine clusters—comprised of 728 issuers across 581 certificates—point to a range of discrepancies between CCADB labels and CA control. By resolving the issues described below, we correct the labels for 125 issuers and 136 CA certificates.

White-label sub-CA. Cluster 2, the second largest cluster, contains CA certificates belonging to both Sectigo and Web.com. Web.com is properly disclosed as a sub-CA of Sectigo, but unlike other disclosed sub-CAs (e.g., Apple, which is also a sub-CA of Sectigo), Web.com exhibits the same operational features as its parent CA, Sectigo. Several Web.com and Sectigo issuers share the same AIA infrastructure (`crt.usertrust.com`) and their OCSP/CRL infrastructure utilize identical IP addresses. Furthermore, 391 out of 540 Web.com issuance fingerprints overlap with Sectigo generated fingerprints, suggesting a shared certificate issuance pipeline. Audit reports corroborate these findings, indicating that Sectigo controls cross-signed certificates for Web.com and that both CAs operate in the same locations globally [30, 31]. Most sub-CAs operate independent of their parent CA, but Web.com appears to utilize a white-label CA service provided by Sectigo [74]. Fides automatically spotlights the shared operations of Sectigo and Web.com, which should be treated as closely intertwined participants in the CA ecosystem, despite their differentiation in CCADB. We further update the Fides dataset to indicate that all Web.com CA certificates are effectively operated by Sectigo.

Undisclosed control. Six clusters with multiple CCADB labels constitute undisclosed CA certificate control. Cluster 4, which contains CA certificates used for Belgium's electronic ID cards, contains the largest number of misrepresented CA certificates. Although three root certificates are disclosed as a sub-CA of DigiCert called Certipost NV/SA (which runs the Belgian Citizen CA), all of the intermediates under those roots contain only a DigiCert CCADB label. All operational features, including third-party audits [63], point to Certipost control of these CA certificates. Cluster 60 also displays incomplete sub-CA disclosure: 2 PKIoverheid intermediates are disclosed as a Digidentity sub-CA, but their child intermediates are labeled as PKIoverheid, contradicting audit records [16].

Cluster	CA1: # issuers (certs)	CA2: # issuers (certs)	Shared Features					Outcome
			CRL	OCSP	AIA	Cert FP	Audit	
2	Sectigo: 313 (382)	Web.com: 6 (14)	✓	✓	✓	✓	✓	White-label sub-CA.
4	DigiCert: 109 (110)	Certipost: 19 (21)	✓	✓	✓	✓	✓	Undisclosed control.
6	GlobalSign: 75 (118)	Google: 23 (33)	✓	✓	✓	✓	✓	False positive.
21	GoDaddy: 9 (19)	Amazon: 2 (7)	✓	✓	✓	-	✓	False positive.
60	Digidentity B.V.: 3 (4)	PKIoverheid: 2 (2)	-	✓	-	-	✓	Undisclosed control.
64	DigiCert: 2 (4)	Sectigo: 1 (1)	✓	-	-	✓	-	Undisclosed third-party.
67	TC TrustCenter: 2 (3)	DSV GmbH: 1 (1)	-	-	✓	✓	-	Undisclosed control.
94	Deutsche Telekom: 2 (2)	DigiCert: 1 (1)	-	✓	-	✓	-	Undisclosed control.
183	StartCom: 1 (1)	Certinomis: 1 (1)	-	✓	-	✓	-	Undisclosed control.
212	E-Tugra: 1 (1)	e-tugra: 1 (1)	-	✓	-	✓	-	Clerical error.
252	E-Tugra: 1 (1)	e-tugra: 1 (1)	-	✓	-	✓	-	Clerical error.

Table 5: **Multi-operator clusters**—11 clusters have clashing CCADB labels. Green/red cells represent correct/incorrect match between CCADB owner labels and CA operational control. CCADB labels misrepresent the control of 125 issuers across 136 certificates.

Cluster 67 contains two root certificates that CCADB labels as TC TrustCenter, and one root certificate that CCADB labels as DSV GmbH. All three utilize the same AIA issuer (`www.trustcenter.de`), contain the name "TrustCenter", and generate a shared set of globally-unique issuance fingerprints. The evidence suggests that TC TrustCenter controls all three roots. Cluster 94 and 183 represent similar cases between Deutsche Telekom / DigiCert and StartCom / Certinomis. As previously discussed in Section 3.2, Taiwan CA (TWCA) appears to operate cluster 64 based on certificate fingerprinting. The issuers in cluster 64 also share CRL infrastructure (`sslserver.twca.com.tw`), further suggesting TWCA operated as an undisclosed sub-CA of both Sectigo and DigiCert. In this instance, a third-party not indicated by CCADB labels actually operated the CA certificates within the cluster.

Clerical error. Clusters 212 and 252 provide an example of CCADB clerical error. CCADB contains two distinct variations of the Turkish SSL provider, E-Tugra (10 CA certs) and e-tugra (6 CA certs), which suggests two distinct CCADB administrator accounts for a single CA. The explanation for this behavior is unknown. While this is the only administrative quirk that emerges from Fides cluster analysis, a manual investigation of CCADB's Subordinate CA owners reveals further clerical quirks. 7 sub-CAs contain inconsistent naming such as alternate spellings (e.g., "Quo Vadis" versus "QuoVadis") or syntactic differences (e.g., "DigitalSign – Certificadora Digital, SA" versus "DigitalSign –Certificadora Digital, S.A."). These may seem like minor details that manual inspection can clarify, but we note that CAs may have very similar names, as is the case with SSLCOM and SSL.com, and sloppiness can lead to misidentification.

False positives. Fides falsely grouped two clusters of CA certificates. The first, cluster 6, contained CA issuers labeled by CCADB as GlobalSign (76 issuers) and Google (23 issuers). Fides detected shared OCSP infrastructure, audits, and certificate fingerprints between two GlobalSign issuers[4] and two issuers[5] labeled as Google Trust Services (GTS) by CCADB. In actuality, these CAs are currently operated independently, but Fides mistakenly clusters GTS and GlobalSign issuers because they were historically operated by GlobalSign. GTS acquired two GlobalSign roots in 2016 [40], but Fides's chronology unawareness leads to the false positive grouping of CA operation. A very similar root acquisition occurred between Amazon Trust Services (ATS) and GoDaddy [43], leading to the second false positive clustering. To better address these scenarios, future work can incorporate chronologically differentiated operational profiles to detect transitions in certificate control.

4.3 Minority unlabeled clusters

We identified 17 Fides clusters (Table 6) where a minority of nodes are unlabeled, and a supermajority (more than 70%) of nodes share the same CCADB owner label. In total, Fides labeled 94 certificates spanning 84 issuers. Due to insufficient audit data and CCADB metadata for these newly-labeled CA certificates, we cannot properly assess the accuracy of these new labels, and false positives such as those identified in Section 4.2 could exist. To reduce these possibilities, we chose a conservative 30% threshold of unlabeled nodes. Fides's CA operator labels represent a best-effort guess for CA certificates that would otherwise have no CA control information available. We further examine these previously unlabeled certificates in Section 4.4, alongside the SSPKI expanded labels from Section 4.1.

[4] *GlobalSign PersonalSign 2 CA - SHA256 - G3* and *GlobalSign EC Administration CA2*

[5] *GlobalSign ECC Root CA - R4* and *GlobalSign EC Administration CA1*

Cluster	Primary Operator	Unlabeled Iss. (Certs)	Unlabeled %
2	Sectigo	7 (8)	2.1%
3	DigiCert	7 (8)	3.8%
4	Certipost s.a./n.v.	41 (41)	24.3%
5	DigiCert	7 (10)	6.2%
7	Asseco	3 (4)	4.5%
8	HARICA	2 (2)	3.6%
13	Entrust	2 (2)	8.3%
15	SwissSign AG	2 (2)	10.5%
16	SecureTrust	1 (2)	5.6%
28	Gov. of Hong Kong	1 (1)	11.1%
36	DigiCert	2 (2)	28.6%
38	DigiCert	2 (2)	28.6%
41	IdenTrust	1 (1)	14.3%
42	Cybertrust Japan	2 (3)	28.6%
57	GlobalSign	1 (4)	20.0%
67	TC TrustCenter	1 (1)	25.0%
69	KIR S.A.	1 (1)	25.0%
17 clusters	*14 operators*	*83 (94)*	–

Table 6: **Minority unlabeled clusters**—94 CCADB-undisclosed certificates appear in 17 clusters with a super-majority (>70%) of known issuers. Undisclosed, Fides-clustered CA certificates occur across a range of CA operators.

4.4 CA operator dataset

Building off of CCADB-labeled clusters, and merging in our investigation of owner ambiguities (Section 4.1) and discrepancies between CCADB labels and Fides's operational clusters, we develop a new dataset that more accurately describes the organizations that control each CA certificate. The dataset corrects the administrative CCADB labels of 241 CA certificates by resolving multiple CCADB owner conflicts within a single SSPKI or Fides cluster. Through SSPKI and cluster expansion, the dataset also extends coverage to 651 CA certificates beyond CCADB disclosure, which is limited by CA self-reporting and the fact that not all root stores require CA certificate disclosure. In total, Fides improves or extends coverage for 208 trusted CA certificates, or 6.2% of all 3,338 CA certificates trusted by Microsoft, Apple, or NSS (Table 7). We hope that this dataset, which we provide open-source [1], enables improved CA research and CA trust decision making. Below, we investigate the potential explanations for Fides's findings and CCADB's shortcomings.

Fides Relabeled Guided by manual analysis, Fides identifies 241 CA certificates where CCADB labels disagree with operational features. The conflicting DigiCert/Certipost cluster accounts for nearly half (114) of these instances. Excluding these certificates, we find twenty CAs that act as the CCADB administrator for a CA certificate they do not operate, indicating that for many CAs, CCADB owner labels signal administrative responsibility rather than operational control. In many cases, these CA certificates are disclosed as sub-CAs, but disclosure is often incomplete, as detailed in Section 5. About a quarter (64 out of 241) of Fides relabeled CA certificates resulted from conflicting CCADB owners for a shared SSPKI. Although conflict resolution requires manual investigation, CCADB could add an automated notification or require a sub-CA label when a single SSPKI maps to certificates with multiple CCADB owners.

Fides Newly Labeled Fides automatically assigned labels to 651 CA certificates not present in CCADB. In the absence of ground truth data for newly labeled certificates, we tracked ten CA certificates that were added to CCADB between July 2020 and February 2021. All ten CA certificate had CCADB labels that matched the independently generated Fides labels (including four Web.com /Sectigo certificates). As an additional confirmation of these new Fides labels, we examined the 62 CA certificates that appeared in audits. We manually identified the CA operator in each audit and found that 60 out of 62 (96.7%) Fides-labeled operators match audit records. The two certificates with erroneous labels occur because two DigiCert cross-signs of MULTICERT certificates contain a DigiCert CCADB label. In this instance, Fides propagates a CCADB label that does not match CA certificate control. Future work classifying the CAs described in CA audits could provide additional consistency checks to further improve Fides's accuracy.

Why were these newly labeled certificates not included in CCADB? Only 75 certificates are trusted by NSS, and only 20 had not expired before February 2017 when NSS mandated CCADB disclosure [57]. NSS does not require disclosure of technically constrained CA certificates (6) or those without TLS server authentication capabilities, which applies to the remaining 13 certificates. Fides does not discover improperly undisclosed NSS-trusted CA certificates, suggesting general compliance with the Mozilla Root Store disclosure policies. However, as described in Section 4.1 we do discover improperly disclosed CA certificates.

For CA certificates trusted by Apple or Microsoft, Fides expands coverage of CA operators by 209 certificates. 141 of these certificates are expired, limiting their utility, but can provide data for historical CA behavior studies. The 68 remaining certificates improve public understanding of active CA operation, especially for the six CA certificates (4 Sectigo, 1 DigiCert, 1 TrustFactory) with unconstrained server authentication capabilities. Because each unconstrained CA certificate is a single point of widespread failure (i.e., a compromised CA certificate can impersonate most domains[6]), comprehensive transparency of the CA certificates wielded by each CA can help attribute suspicious behavior or mitigate more serious issues when they occur.

[6]Exceptions for HSTS, preload, and CAA.

	Total Iss. (Certs)	Trusted Iss. (Certs)	Valid Iss. (Certs)
CCADB	4,845 (6,195)	2,673 (2,961)	3,457 (4,077)
Relabeled	189 (241)	85 (90)	103 (121)
New label	404 (651)	90 (115)	130 (164)
Fides	4,928 (6,846)	2,707 (3,076)	3,490 (4,241)

Table 7: **CCADB/Fides Comparison**—Fides yields a CA operator dataset that corrects CA operator labels for 90 trusted CA certificates from 85 issuers, and extends coverage by 115 trusted CA certificates. Fides improves/increases coverage for 6.1% of all 3,338 trusted CA certificates.

5 Discussion

While Fides can detect inconsistencies between CCADB ownership labels and operational practices, it is not a long-term solution. Its heuristics are not perfect, and while its precision is high, its recall is low. We hope that Fides sheds light on the poor state of affairs, quantifying how certificate subjects poorly reflect CA ownership and showing how CCADB does not currently address controlling ownership. True transparency requires changes to existing CA procedures and root store requirements. Below, we explore potential solutions:

CCADB Structured Data. At the moment, CCADB plays a critical role in the PKI ecosystem: it provides a mechanism for CA certificate data to update independent of the actual certificate itself. CCADB provides mutability to CA certificates. Because the frequency of CA certificate control changes outpaces the frequency of CA certificate replacement, current CA certificates must divorce their names (stored in the certificate) from their identity (stored outside of the certificate). CCADB is a natural location to track who controls each CA root and intermediate certificate. While in some cases we can infer certificate control from CCADB record owners and uploaded audits, the data is not easily accessible. Adding explicit fields for ownership details would allow both root store operators and researchers to better track CA behavior, and would additionally provide data compare against regular Fides runs. This proposal is the simplest to implement, but would require careful auditing and consistent monitoring to protect against error-prone or even nefarious self-reporting. User agents can also enforce more stringent CCADB inclusion policies to help remove trust dependencies on CAs that have refused to submit details to CCADB.

Increased Intermediate Restrictions. While trust anchors are long-lived and shipped with user agents, intermediate CA certificates do not need to be. User agents can require that intermediate CA certificates contain up-to-date ownership details, similar to the requirements for Extended Validation (EV) certificates, and to restrict their change of ownership. Because leaf certificates are signed by intermediates rather than a trust anchor, this would allow users to always identify the entity that signed the certificate used when accessing a website. User agents could further limit the validity period of intermediate certificates to disincentivize transfer of ownership and reduce the impact of changes in certificate control.

Reconsider Root CA Labels. Today, user agents already ignore some details about trust anchors, including their validation periods. We should consider whether we should also ignore included trust anchor subject names and to instead ship these details with the root store. As it stands, labels on roots are misleading for a significant fraction of CAs, and browser-supplied labels could provide more up-to-date ownership details (e.g. as extracted from CCADB).

These proposals are orthogonal to the development and deployment of Fides, which can help identify user errors and suspicious CA practices. Future development of Fides can verify the consistency between CA documentation/audits claims and externally measurable behavior. For example, by extending Fides to include the IP addresses from which CAs deliver certificates to Subscribers, we could automatically check the accuracy of the operational locations disclosed in CA documentation/audits. Further development of Fides's certificate fingerprinting techniques can also identify the CA software that different CAs use, leading to better discovery and remediation of certificate issuance problems, such as the widespread usage of 63-bit serial numbers due to an EJBCA bug [12].

6 Related Work

The CA ecosystem has received extensive examination from security researchers. Prior work can be grouped into two categories: the security and properties of issued certificates and the correctness of certificate validation. Our work focuses on the former, since research investigating certificate validation issues [15, 20, 35, 72] is not germane to this study.

Initial work by Holz et al. in 2011 and Durumeric et al. in 2013 focused on the acquisition of certificate data, revealing a fractured ecosystem fraught with problematic certificates, untrusted chains, re-used certificates, and the aftermath of known issuer compromise [6, 28, 39]. Chung et al. performed a similar study in 2016, but instead focused on the sources and uses of invalid certificates. A meta-study of the certificate ecosystem in 2016 [77] found that scanning IPv4 address space for certificates only captured a fraction of the overall certificate ecosystem, and that Certificate Transparency (CT) [45] contained a predominant, and proliferating, share of certificates. In 2018, Chrome [73] and Apple [8] began requiring CT inclusion for all future trusted certificates, paving the way for strict CT enforcement by other browsers. This study uses CT as the authoritative source of certificates in the PKI ecosystem.

Our work is an application that extends the transparency originally intended by CT, which was explicitly designed for domain owners to detect misissued certificates and for public

auditors to expose certificates that are not compliant with the Baseline Requirements [45]. Several works have used CT as a source for discovering phishing DNS names [42, 51, 66, 70], while others have focused on the privacy implications of domain exposure through CT [67, 70].

To combat the issue of insecure and problematic certificates, the CA/Browser (CA/B) Forum established a set of binding *Baseline Requirements* (BRs) in 2011 [18]. The BRs mandate secure as well as hygienic certificate issuance practices (e.g., the subject distinguished name must not have a leading whitespace). Several "linting" tools have been developed to check certificate compliance with the BRs [13, 68]. Although the certificate hygiene BRs do not have direct security consequences, Kumar et al. demonstrated that poor certificate hygiene is strongly correlated with instances of certificate insecurity [44]: issuers that don't run a pristine certificate issuance operation are more likely to make security mistakes. Hiller et al. detailed the cross-signing complexity of the web PKI [38], and noted the difference between "internal" and "external" cross-signs, but did not discuss the issue of CA issuer control. To our knowledge, this work is the first to address the challenge of CA identification and improve the state-of-the-art.

One contribution of this work is the certificate fingerprinting technique used to link related issuance operations. Fingerprinting techniques have been previously used in the context of SSL/TLS. Ristić first described fingerprinting the Client Hello messages exchanged in SSL/TLS handshakes [65], and this fingerprinting approach interception [29]. Client Hello fingerprinting has also been used to identify TLS clients for a range of purposes [7, 14, 41, 64], leading to recent work that masks client fingerprints to avoid detection [32]. The most closely related work, by Delignat-Lavaud et al. in 2014, created certificate templates that used manually selected fields to create certificate profiles [25]. While some overlapping features are captured by both techniques, there are two key differences. First, the use of the Issuer field as a high-importance clustering feature assumes a one-to-one mapping between issuer and certificate generation process. Our work invalidates this assumption. Second, the authors performed their study on 1.4M certificates, while this study investigates 2.9B certificates in an ecosystem that has grown prolifically in recent years.

7 Conclusion

In this work, we analyzed the ownership and control of CA certificates. We showed that embedded ownership data is often inaccurate due to mergers/acquisitions, business transactions, and record keeping failures. To scalably identify discrepancies between certificates, audit records, and operational practices, we introduced Fides, which empirically tracks CA behavior and clusters CA certificates with shared operational fingerprints. Our dataset draws attention to the administrative, rather than operational, focus of CCADB, which is the best existing delineation of CA operations. We found 241 CA certificates where CCADB labels diverge from CA control. Fides also automatically labeled an additional 651 CA certificates that were not disclosed in CCADB. In addition to promoting CA operational transparency, the Fides dataset has also identified or corroborated several CA disclosure issues. To help future studies accurately characterize the Web PKI, we are releasing our dataset of 6,846 CA certificates, their operational fingerprints, and CA operator labels [1].

Acknowledgments

The authors thank Ryan Sleevi and the anonymous reviewers for providing insightful feedback on various parts of this work. This work was supported in part by the Yunni & Maxine Pao Memorial Fellowship and a gift from Google, Inc.

References

[1] Fides source code/data. https://github.com/zzma/ca-transparency.

[2] GeoTrust acquires Equifax's digital certificate business. https://www.bizjournals.com/atlanta/stories/2001/09/24/daily17.html.

[3] Notice regarding Symantec acquisition of Verisign's authentication businesses. https://www.verisign.com/en_US/verisign-repository/symantec/index.xhtml.

[4] Symantec roots. https://chromium.googlesource.com/chromium/src/+/master/net/data/ssl/symantec/README.md.

[5] J. Aas. Let's Encrypt is trusted. https://letsencrypt.org/2015/10/19/lets-encrypt-is-trusted.html.

[6] H. Adkins. An update on attempted man-in-the-middle attacks. https://security.googleblog.com/2011/08/update-on-attempted-man-in-middle.html.

[7] B. Anderson, S. Paul, and D. McGrew. Deciphering malware's use of TLS (without decryption). *arXiv preprint arXiv:1607.01639*, 2016.

[8] Apple. Apple's Certificate Transparency policy. https://support.apple.com/en-us/HT205280.

[9] Apple. Current Apple CT logs. https://valid.apple.com/ct/log_list/current_log_list.json.

[10] Apple. Information for website operators about distrusting Symantec certificate authorities. https://support.apple.com/en-us/HT208860.

[11] H. Birge-Lee, Y. Sun, A. Edmundson, J. Rexford, and P. Mittal. Bamboozling certificate authorities with BGP. In *27th USENIX Security Symposium (USENIX Security)*, 2018.

[12] J. Bohm. General issues that came up in the DarkMatter discussion(s). https://groups.google.com/g/mozilla.dev.security.policy/c/7WuWS_20758/m/erK0-f0GCwAJ.

[13] P. Bowen. Certlint. https://github.com/awslabs/certlint.

[14] L. Brotherston. Stealthier attacks and smarter defending with TLS fingerprinting. *DerbyCon*, 2015.

[15] 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.

[16] BSI. Digidentity B.V. audit. https://www.digidentity.eu/assets/files/terms/20200123-ETS043-411-1.pdf, 2020.

[17] Bugzilla. Camerfirma: Failure to abide by section 8 of Mozilla policy: Unauthorized, improperly disclosed subordinate CA. https://bugzilla.mozilla.org/show_bug.cgi?id=1672029.

[18] CA/Browser Forum. Baseline requirements. https://cabforum.org/baseline-requirements-documents/.

[19] Certipost. Belgian certificate policy & practice statement for eID PKI infrastructure Citizen CA. http://repository.eid.belgium.be/downloads/citizen/archive/en/CITIZEN_CA_2018.pdf.

[20] S. Y. Chau, O. Chowdhury, E. Hoque, H. Ge, A. Kate, C. Nita-Rotaru, and N. Li. Symcerts: Practical symbolic execution for exposing noncompliance in X.509 certificate validation implementations. In *38th IEEE Symposium on Security and Privacy*, 2017.

[21] T. Chung, Y. Liu, D. Choffnes, D. Levin, B. M. Maggs, A. Mislove, and C. Wilson. Measuring and applying invalid SSL certificates: the silent majority. In *16th ACM Internet Measurement Conference*, 2016.

[22] J. Clark and P. C. Van Oorschot. Sok: SSL and HTTPS: Revisiting past challenges and evaluating certificate trust model enhancements. In *34th IEEE Symposium on Security and Privacy*, 2013.

[23] Comodo. Comodo CA is now Sectigo. https://comodosslstore.com/sectigo.

[24] Comodo. Secure faxing. Secure e-mail. Secure backup - anywhere, anytime. https://www.comodo.com/news/press_releases/12_01_04.html.

[25] 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 (NDSS)*, 2014.

[26] DigiCert. DigiCert completes acquisition of Symantec's Website Security and related PKI solutions. https://www.digicert.com/news/digicert-completes-acquisition-of-symantec-ssl/.

[27] DigiCert. Digicert completes purchase of QuoVadis, expands European presence and TLS, PKI offerings. https://www.digicert.com/news/pr/digicert-completes-purchase-of-quovadis-ssl/.

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

[29] Z. Durumeric, Z. Ma, D. Springall, R. Barnes, N. Sullivan, E. Bursztein, M. Bailey, J. A. Halderman, and V. Paxson. The security impact of HTTPS interception. In *24th Network & Distributed Systems Symposium*, 2017.

[30] Ernst & Young LLP. Sectigo: Report of independent accountants. https://bug1472993.bmoattachments.org/attachment.cgi?id=9078178.

[31] Ernst & Young LLP. Web.com: Report of independent accountants. https://www.cpacanada.ca/generichandlers/CPACHandler.ashx?attachmentid=230861.

[32] S. Frolov and E. Wustrow. The use of TLS in censorship circumvention. In *26th Network & Distributed Systems Symposium*, 2019.

[33] Funding Universe. RSA Security Inc. history. http://www.fundinguniverse.com/company-histories/rsa-security-inc-history/.

[34] Funding Universe. VeriSign, Inc. history. http://www.fundinguniverse.com/company-histories/verisign-inc-history/.

[35] M. Georgiev, S. Iyengar, S. Jana, R. Anubhai, D. Boneh, and V. Shmatik. The most dangerous code in the world: Validating SSL certificates in non-browser software. In *19th ACM Conference on Computer and Communications Security*, 2012.

[36] Google. Chrome compliant CT logs. https://www.certificate-transparency.org/known-logs.

[37] P. Hadfield. Comment on retrieving, storing and querying 250m+ certificates like a boss. https://medium.com/@hadfieldp/hey-ryan-c0fee84b5c39.

[38] J. Hiller, J. Amann, and O. Hohlfeld. The boon and bane of cross-signing: Shedding light on a common practice in public key infrastructures. In *27th ACM Conference on Computer and Communications Security*, 2020.

[39] 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.

[40] R. Hurst. Google Trust Services roots. https://groups.google.com/g/mozilla.dev.security.policy/c/1PDQv0GUW_s/m/ErxcjAcFDwAJ.

[41] M. Husák, M. Cermák, T. Jirsík, and P. Celeda. Network-based HTTPS client identification using SSL/TLS fingerprinting, 2015.

[42] P. Kintis, N. Miramirkhani, C. Lever, Y. Chen, R. Romero-Gómez, N. Pitropakis, N. Nikiforakis, and M. Antonakakis. Hiding in plain sight: A longitudinal study of combosquatting abuse. In *24th ACM Conference on Computer and Communications Security*, 2017.

[43] J. Kozolchyk. How to prepare for AWS's move to its own certificate authority. https://aws.amazon.com/blogs/security/how-to-prepare-for-aws-move-to-its-own-certificate-authority/.

[44] D. Kumar, Z. Wang, M. Hyder, J. Dickinson, G. Beck, D. Adrian, J. Mason, Z. Durumeric, J. A. Halderman, and M. Bailey. Tracking certificate misissuance in the wild. In *39th IEEE Symposium on Security and Privacy*, 2018.

[45] B. Laurie, A. Langley, and E. Kasper. Certificate Transparency. RFC 6962, 2013.

[46] Let's Encrypt. Boulder: An ACME-based certificate authority, written in Go. https://github.com/letsencrypt/boulder.

[47] 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*, 2015.

[48] Z. Ma. asn1-fingerprint. https://github.com/zzma/asn1-fingerprint.

[49] G. Markham. Mozilla's plan for Symantec roots. https://groups.google.com/forum/#!msg/mozilla.dev.security.policy/FLHRT79e3XE/riCrpXsfAgAJ.

[50] U. Meyer and V. Drury. Certified phishing: Taking a look at public key certificates of phishing websites. In *15th Symposium on Usable Privacy and Security*, 2019.

[51] U. Meyer and V. Drury. Certified phishing: Taking a look at public key certificates of phishing websites. In *15th Symposium on Usable Privacy and Security*, 2019.

[52] Microsoft Secure Blog Staff. Microsoft partners with DigiCert to begin deprecating Symantec TLS certificates. https://www.microsoft.com/security/blog/2018/10/04/microsoft-partners-with-digicert-to-begin-deprecating-symantec-tls-certificates/.

[53] R. Miller. VeriSign to buy GeoTrust, combining top SSL providers. https://news.netcraft.com/archives/2006/05/17/verisign_to_buy_geotrust_combining_top_ssl_providers.html.

[54] Mozilla. CA/closed incidents. https://wiki.mozilla.org/CA/Closed_Incidents.

[55] Mozilla. Common CA Database. https://www.ccadb.org/.

[56] Mozilla. WoSign and StartCom. https://docs.google.com/document/d/1C6BlmbeQfn4a9zydVi2UvjBGv6szuSB4sMYUcVrR8vQ/edit.

[57] Mozilla Wiki. CA/Root store policy archive. https://wiki.mozilla.org/CA/Root_Store_Policy_Archive.

[58] Mozilla Wiki. CA:Symantec issues. https://wiki.mozilla.org/CA:Symantec_Issues.

[59] Mozilla Wiki. CA:WoSign issues. https://wiki.mozilla.org/CA:WoSign_Issues.

[60] Network Solutions, LLC. Network Solutions certification practice statement. https://assets.web.com/legal/English/CertificationPracticeStatement.pdf.

[61] Objective Systems. ASN.1: Listing of universal tags. https://obj-sys.com/asn1tutorial/node124.html.

[62] 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.

[63] PWC. Certipost independent assurance report. https://bug1461443.bmoattachments.org/attachment.cgi?id=9149485, 2020.

[64] A. Razaghpanah, A. A. Niaki, N. Vallina-Rodriguez, S. Sundaresan, J. Amann, and P. Gill. Studying TLS usage in Android apps. In *13th International Conference on emerging Networking EXperiments and Technologies*, 2017.

[65] I. Ristic. HTTP client fingerprinting using SSL handshake analysis. https://blog.ivanristic.com/2009/06/http-client-fingerprinting-using-ssl-handshake-analysis.html.

[66] R. Roberts, Y. Goldschlag, R. Walter, T. Chung, A. Mislove, and D. Levin. You are who you appear to be: A longitudinal study of domain impersonation in TLS certificates. In *26th ACM Conference on Computer and Communications Security*, 2019.

[67] R. Roberts and D. Levin. When Certificate Transparency is too transparent: Analyzing information leakage in HTTPS domain names. In *18th ACM Workshop on Privacy in the Electronic Society (WPES)*, 2019.

[68] K. Roeckx. X509lint. https://github.com/kroeckx/x509lint.

[69] J. Rowley. https://groups.google.com/d/msg/mozilla.dev.security.policy/_EnH2IeuZtw/AdZvpzGJAwAJ.

[70] Q. Scheitle, O. Gasser, T. Nolte, J. Amann, L. Brent, G. Carle, R. Holz, T. C. Schmidt, and M. Wählisch. The rise of Certificate Transparency and its implications on the internet ecosystem. In *18th ACM Internet Measurement Conference*, 2018.

[71] Schellman & Company, LLC. IdenTrust - 2019 WebTrust for CAs. https://www.cpacanada.ca/generichandlers/CPACHandler.ashx?attachmentid=236834.

[72] S. Sivakorn, G. Argyros, K. Pei, A. D. Keromytis, and S. Jana. Hvlearn: Automated black-box analysis of hostname verification in SSL/TLS implementations. In *38th IEEE Symposium on Security and Privacy*, 2017.

[73] 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.

[74] R. Sleevi. Disclosure and CP/CPS for cross-signed roots. https://groups.google.com/d/msg/mozilla.dev.security.policy/89iF_4Ovpwg/zboFW5c6DwAJ.

[75] Symantec. Symantec acquires PGP and GuardianEdge. http://eval.symantec.com/mktginfo/enterprise/other_resources/b-pgp_guardianedge_acq_faq.en-us.pdf.

[76] USERTrust Inc. Web archive: GeoTrust and USERTrust offering free SSL service to businesses worldwide. https://web.archive.org/web/20020802152743/http://www.usertrust.com/index.asp?key=news/free-ssl-service.

[77] 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.

[78] B. Wilson. Summary of Camerfirma's compliance issues. https://groups.google.com/g/mozilla.dev.security.policy/c/dSeD3dgnpzk.

A Historic Disclosure Issues

Issue #	Date	# Certs	# Issuers	Owner	Description
1012744	2014-05-19	8	8	Firmaprofesional	Publicly disclosed subordinate CA certificates from Firmaprofesional
1013081	2014-05-20	6	6	ACCV	ACCV publicly disclosed subordinate CA certificates
1016347	2014-05-27	3	3	ACEDICOM	Publicly disclosed subordinate CA certificates from ACEDICOM
1017583	2014-05-29	40	27	GlobalSign	Public disclosure of GlobalSign Subordinate CAs
1018158	2014-05-30	9	9	GRCA	GRCA publicly disclosed subordinate CA certificates
1309707	2016-10-12	7	6	WoSign	Distrust new certs chaining up to current WoSign/StartCom roots
1367842	2017-05-25	3	3	TurkTrust	TurkTrust: Non-audited, non-technically-constrained intermediate certs
1368171	2017-05-26	2	2	Firmaprofesional	Firmaprofesional: Non-audited, non-technically-constrained intermediate certs
1368176	2017-05-26	7	5	DigiCert	DigiCert: Non-audited, non-technically-constrained intermediate certs
1368178	2017-05-26	1	1	Symantec	Symantec: Non-audited, non-technically-constrained intermediate cert
1373452	2017-06-15	3	2	TrustID	Identrust TrustID Subordinate CA - Revocation Notification
1386891	2017-08-02	2	2	StartCom	Certinomis: Cross-signing of StartCom intermediate certs, and delay in reporting it in CCADB
1432608	2018-01-23	15	4	Gov. of Portugal (SCEE)	Add EC Raiz Estado Cross Certificates to OneCRL
1451950	2018-04-05	2	2	Gov. of Portugal (SCEE)	DigiCert: Intermediate Cert(s) not disclosed in CCADB
1451953	2018-04-05	4	4	TeliaSonera	TeliaSonera: Intermediate Cert(s) Not Disclosed in CCADB
1455119	2018-04-18	2	2	Firmaprofesional	Firmaprofesional: Undisclosed Intermediate certificate
1455128	2018-04-18	4	2	Certicamara	Certicamara: Undisclosed Intermediate certificates
1455132	2018-04-18	13	13	SwissSign	SwissSign: Undisclosed Intermediate Certificates
1455137	2018-04-18	1	1	T-Systems	T-Systems: Undisclosed Intermediate certificate
1464359	2018-05-25	1	1	Firmaprofesional	Firmaprofesional: Undisclosed Intermediate certificate SDS
1497700	2018-10-09	1	1	DocuSign/Keynectis	DocuSign/Keynectis: Undisclosed Intermediate certificate
1497703	2018-10-09	2	2	SECOM	SECOM: Undisclosed intermediate certificates
1499585	2018-10-16	26	21	DigiCert	Digicert: Undisclosed CAs -Federated Trust CA-1
1503638	2018-10-31	1	1	WISeKey	WISeKey: Failure to disclose intermediate in CCADB
1542082	2019-04-04	1	1	IdenTrust	Identrust: Failure to disclose Unconstrained intermediate Within 7 Days
1563573	2019-07-04	22	22	DigiCert	DigiCert: Failure to disclose Unconstrained Intermediate within 7 Days
1563574	2019-07-04	2	2	SECOM	SECOM: Failure to disclose Unconstrained Intermediate within 7 Days
1563575	2019-07-04	1	1	TeliaSonera	Telia: Failure to disclose Unconstrained Intermediate within 7 Days
Total: 28	–	186	150	21	–

Table 8: **CCADB disclosure issues**—28 resolved disclosure issues provide an approximate ground truth dataset of 150 issuers (186 certificates) for Fides evaluation.

B Issuers with multiple CCADB Owners

Issuer (Subject+SPKI)	CCADB owners	# Certs	Details
ec38da6:MULTICERT SSL CA 005	AC Camerfirma, S.A. | MULTICERT	2	Undisclosed / unaudited MULTICERT sub-CA
49d8519:Starfield Services Root CA - G2	Amazon Trust Services | GoDaddy	3	Undisclosed Amazon sub-CA
98ac41c:StartCom Class 3 OV Server CA	StartCom | WoSign	2	Undisclosed StartCom sub-CA
5e87566:Belgium Root CA4	Certipost s.a./n.v. | DigiCert	3	Undisclosed Certipost sub-CA
d42c25d:Let's Encrypt Authority X1	IdenTrust | ISRG	3	Undisclosed ISRG sub-CA
dafa2be:Let's Encrypt Authority X2	IdenTrust | ISRG	3	Undisclosed ISRG sub-CA
78d2913:Let's Encrypt Authority X3	IdenTrust | ISRG	2	Undisclosed ISRG sub-CA
fdeacfa:Let's Encrypt Authority X4	IdenTrust | ISRG	2	Undisclosed ISRG sub-CA
6ee23dd:SSL.com EV Root CA RSA R2	Asseco | SSL.com	3	Disclosed SSL.com sub-CA
39904e6:SSL.com Root CA RSA	Asseco | SSL.com	2	Disclosed SSL.com sub-CA
51b64a7:UCA Global G2 Root	Asseco | Shanghai Elec. CA	2	Disclosed SHECA sub-CA
fa2de6c:GTS Root R1	GlobalSign | Google Trust Services	2	Disclosed GTS sub-CA
2da3659:DigiCert High Assurance EV Root CA	DigiCert | Entrust	6	Expired Entrust cross-sign
4098e01:Network Solutions CA	Sectigo | Web.com	8	Expired Sectigo cross-sign
df6609e:Government CA	Certipost s.a./n.v. | DigiCert	2	Expired / undisclosed DigiCert cross-sign
67d2813:Government CA	Certipost s.a./n.v. | DigiCert	2	Expired / undisclosed DigiCert cross-sign
219718a:Federal Bridge CA 2013	DigiCert | IdenTrust | US Federal PKI	3	All revoked
4c76dcf:Actalis Authentication CA G2	Actalis | DigiCert	3	All revoked
1fb3270:Certipost E-Trust Primary Normalised CA	Certipost s.a./n.v. | DigiCert	2	All revoked
ed0fa26:ECRaizEstado	DigiCert | Gov. of Portugal (SCEE)	6	Revoked DigiCert cross-sign
c23714e:Belgium Root CA2	DigiCert | GlobalSign	3	Revoked GlobalSign cross-sign
5c78ccd:WellsSecure Public Root CA	DigiCert | Wells Fargo Bank N.A.	2	Revoked DigiCert cross-sign
1fc94be:WellsSecure Public Root CA 01 G2	DigiCert | Wells Fargo Bank N.A.	3	Revoked DigiCert cross-sign
5451b03:AffirmTrust Commercial	Entrust | SwissSign AG	2	Revoked SwissSign cross-sign
8b7b0ab:AffirmTrust Networking	Entrust | SwissSign AG	3	Revoked SwissSign cross-sign
69286df:GlobalSign Root CA	GlobalSign | Google Trust Services	2	Revoked GlobalSign transfer
1ac0e91:GlobalSign	GlobalSign | Google Trust Services	3	Revoked GlobalSign cross-sign
e125939:CA of WoSign	Sectigo | StartCom | WoSign	6	Revoked StartCom/Sectigo cross-sign
5b5804f:CA of WoSign G2	Asseco | WoSign	3	Revoked Asseco cross-sign
676cf22:CA Wotong Root Certificate	StartCom | WoSign	3	Revoked StartCom cross-sign
b53b021:Microsoft Azure TLS Issuing CA 06	DigiCert | Microsoft Corporation	2	Revoked DigiCert cross-sign
4002521:Microsoft Azure ECC TLS Issuing CA 01	DigiCert | Microsoft Corporation	2	Revoked DigiCert cross-sign
c0f4b26:Microsoft Azure TLS Issuing CA 05	DigiCert | Microsoft Corporation	2	Revoked DigiCert cross-sign
0c6cbcf:Microsoft Azure TLS Issuing CA 02	DigiCert | Microsoft Corporation	2	Revoked DigiCert cross-sign
04026ad:Microsoft Azure TLS Issuing CA 01	DigiCert | Microsoft Corporation	2	Revoked DigiCert cross-sign
6664c4c:Microsoft Azure ECC TLS Issuing CA 02	DigiCert | Microsoft Corporation	2	Revoked DigiCert cross-sign
52929fe:Microsoft Azure ECC TLS Issuing CA 06	DigiCert | Microsoft Corporation	2	Revoked DigiCert cross-sign
e6b1b8a:Microsoft Azure ECC TLS Issuing CA 05	DigiCert | Microsoft Corporation	2	Revoked DigiCert cross-sign
7712fbc:MULTICERT SSL CA 001	AC Camerfirma, S.A. | MULTICERT	3	Revoked AC Camerfirma cross-sign

Table 9: **Issuers with multiple CCADB owners**—39 issuers (Subject+SPKI) across 110 certificates have ambiguous CCADB owners, which reflect CCADB's unsuitability for mapping CA certificate control.