The Potential of Individualized Trusted Root Stores: Minimizing the Attack Surface in the Light of CA Failures

Johannes Braun
Technische Universität Darmstadt
jbraun@cdc.informatik.tu-darmstadt.de

Gregor Rynkowski
Technische Universität Darmstadt
grynkowski@cdc.informatik.tu-darmstadt.de

Abstract

The security of most Internet applications relies on underlying public key infrastructures (PKIs) and thus on an ecosystem of certification authorities (CAs). The pool of PKIs responsible for the issuance and the maintenance of SSL certificates, called the Web PKI, has grown extremely large and complex. Herein, each CA is a single point of failure for the security, leading to an attack surface, the size of which is hardly assessable.

This paper approaches the issue if and how the attack surface can be reduced in order to reduce the risk of relying on a malicious certificate. In particular we consider the individualization of the set of trusted CAs. We present a tool called Rootopia, which allows to assess the respective part of the Web PKI relevant for a user.

Our analysis of browser histories of 22 Internet users reveals, that the major part of the PKI is completely irrelevant to a single user. The attack surface can be reduced by more than 90%, which shows the potential of the individualization of the set of trusted CAs. Furthermore, all the relevant CAs reside within a small set of countries. Our findings confirm, that we unnecessarily trust in a huge number of CAs, exposing ourselves to unnecessary risks.

1 Introduction

Achieving the security targets confidentiality, data integrity and entity authentication is indispensable as more and more critical data is sent over the Internet in the context of many well known and extensively used applications like e-business, e-banking and e-government. Entity authentication and thus secure connection establishment builds on the underlying Web PKI. CAs certify the binding between the identity of an entity and its public key by issuing certificates. The certificates are then used during the SSL/TLS protocols to authenticate servers and to establish session keys.

Since the advent of the Internet, more and more such CAs have come into existence, forming a highly complex and interdependent ecosystem. Today, all of these CAs are equally trusted by users. This trust concept has been discussed and long since criticized in the scientific community [14–16] as it imposes severe security problems.

Concerning the Internet, this trust concept means that each of the CAs can issue a certificate for any domain, making each of the CAs a single point of failure. The compromise or misbehavior of a single one undermines the security of the entire system.

The security of the CAs does not only rely on cryptographic primitives and algorithms, which can be analyzed and checked for correctness. Security also relies on the quality and correctness of how the various organizations and entities conduct their work processes, which in turn are difficult to assess. The existing baseline requirements and audits are not capable to protect from failures and misbehavior. This is discussed by many researchers and in the past has often been shown by real world examples, which lead to the issuance of malicious certificates.

Recent security incidents [17, 18] have led to growing reservations in the scientific community about CAs. Furthermore, it drives research towards monitoring of the Web PKI and searching for alternatives (cf. Section 3 for details).

However, despite all the criticism against PKI, this is the only established solution we have right now [27]. Thus, an important goal is to improve the security of the current system. A natural approach to reduce risks is to reduce the attack surface.

The most obvious way to achieve this in case of the Web PKI is to reduce the number of trusted CAs. Yet, how can this be achieved? While globally each one of the CAs might be required for interoperability, this is clearly not the case when considering a individual user [7]. Thus, a user should only trust in the CAs he really needs for the applications he uses. Excluding all the non-required CAs from the personal set of trusted CAs

reduces the risks significantly. Compromises and misbehavior of those CAs then cannot threaten the respective user anymore.

The work at hand deals with the issue of the individual limitation of trusted CAs. We present a tool that allows to identify the set of CAs relevant to a specific user based on his browser history. We conduct a study and evaluate how the currently deployed Web PKI is observed from a user's point of view. We show, that the set of CAs relevant to a user is indeed highly dependent on its individual behavior. A thorough analysis and characterization of the individual view on the Web PKI can help to minimize the attack surface in the future. We show that there is an immensely high potential to improve the security by maintaining individualized root stores. Furthermore, we identify the challenges that have to be solved in order to fully realize individualized root stores.

We note, that a global limitation of the trusted CAs seems to be no viable solution. It would lead to interoperability problems and additional warnings whenever a certificate issued by an unknown CA is presented to the user. The problem with warnings is, that users get used to and tend to ignore them (see e.g. [21, 35]), even leading to a weakening effect. Furthermore, since browsing behavior of various users can vary quite strong, many different CAs are required and a global minimization of CAs cannot lead to an optimal solution.

The paper is organized as follows. First, we describe the Web PKI. We explain the problems with its current deployment and present related work. Afterward, we define our scope and describe our methodology. In Section 5 we present our tool, followed by the description of the setup for our study. In Section 7 we present the findings. Then we evaluate the implications for our goal, the minimization of the attack surface. We show limitations, present future work and end with a conclusion.

2 The Web PKI

In this section we explain the basics of PKIs relevant for this paper. We focus on the current deployment of the Web PKI relevant for SSL/TLS. The Web PKI is based on the X.509 standard [8], thus the acting CAs accordingly issue X.509 certificates. Among others, X.509 certificates have an issuer and a subject field and contain a public key. The issuer field contains the Distinguished Name (DN) of the certifying CA, while the subject field contains the DN of the entity, whose key is certified. To ensure this binding, certificates are digitally signed by the issuing CA.

The certificates are used during the SSL/TLS protocols to establish secure connections and in this context they are mostly used to authenticate web servers. Thus, at the beginning of a session, a web server presents its certificate to a client. The client must then examine the certificate to check if he communicates with the desired server. That means, on the one hand it checks whether the certificate was issued by a CA it trusts in and on the other hand, if the contained data identifies the communication partner. If so, the public key is extracted from the certificate and for example used to establish session keys to encrypt the communication. In order to check, i.e. validate the CA's signature on the certificate, the verifier must know the CA's public key.

The Web PKI is based on the concept of hierarchical PKIs. So-called *Root CAs* build the top of the infrastructure and sign certificates of *subordinate CAs* (Sub CAs). These Sub CAs in turn can issue certificates to CAs or *end entities* like domain owners, resulting in a so called certificate chain. An end entity's key can then be validated by any chain of valid certificates up to one of the Root CAs. The verifier only needs to know the public key of the Root CA. In a certificate chain, the subject of a certificate is either the issuer of the subordinate certificate or the end entity. The issuer of the first certificate must be a Root CA in order to be able to validate the certificate chain.

The Web PKI is far from being a simple hierarchical PKI. There exist many different Root CAs, which are included into so called trusted lists, or root stores. If a CA is included into the root store of a client, it is automatically trusted as are all of its Sub CAs. This leads to a highly complex system of CAs. So called cross signing – where CAs mutually issue certificates to each other – make the system even more complex.

A simplified example of the resulting Web PKI is depicted in Fig. 1. Here, an exemplary certificate chain exists from the Root CA R-CA_1 to the end entity EE_1, where the arrows represent certificates. The circular arrows represent self signed certificates, which are often issued by Root CAs to themselves in order to publish their keys. To validate EE_1's certificate, one only needs to know the key of R-CA_1. All other keys are shipped within the intermediary certificates. In this small example, it can also be seen, that it can be difficult to determine all the trusted CAs. For example if R-CA_3 were removed from the root store, its direct Sub CA S-CA_4 would still be trusted due to the additional chain from R-CA_2. However, as there exists no public repository of all the certificates, the existing chains are in general unknown to users until they are presented during connection establishment.

In the current deployment of the Web PKI, the decision which CAs are included into the root stores is left to browser and operating system vendors and relies on their specific policies [3, 4]. Interestingly, the number of CAs included into those root stores has been constantly growing over the last years. For example, the root store of

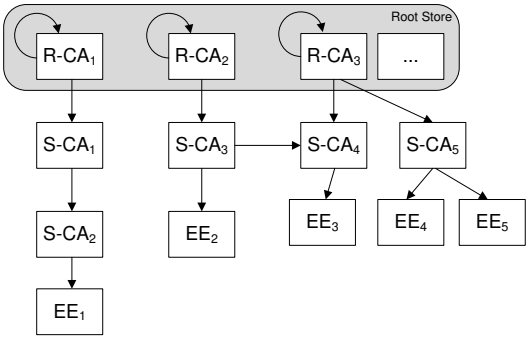

Figure 1: Example Web PKI

the Mozilla browser comes together with the NSS crypto library and contains about 160 CAs [23,30]. Another example is Microsoft's root store which contains about 264 CAs[1] [34], which are directly trusted. Additionally considering the Sub CAs, the number of trusted CAs rises to about 650 different organizations or 1,482 CAs according to data collected by the Electronic Frontier Foundation (EFF) [10,12].

Problems with the Web PKI The way how trust in the CAs is handled has severe drawbacks. Remember, each of the trusted CAs – either Root CA or Sub CA – can arbitrarily issue certificates for any domain. This makes each of the CAs a single point of failure for the whole system. Compromising or compelling a single one of the trusted CAs allows a potential attacker to impersonate as any web server, or to mount a man-in-the-middle attack on any SSL/TLS secured connection, thus opening doors for Internet fraud and surveillance. Moreover, every user of the Web PKI is affected by disruptions to the same extent. And to assume CAs as trusted third parties to be immune to compromises is simply wrong.

Over the last years, several incidents have been documented [6,17–20], with the Diginotar incident [20] being the most severe one, which finally led to the removal of the Diginotar CA from all trusted root stores. Actually, this was the first time that a CA was removed from the root stores. So far, compromises were compensated by revocation and black listing of maliciously issued certificates. Besides compromises, malpractices and evidence of governments that compelled CAs to cooperate have been reported [11, 15]. Even erroneously issued certificates can impose security holes as in the recent case of TURKTRUST [29], where CA certificates were issued instead of normal host certificates, giving the subject of those certificates the power to act as a fully trusted CA. With this in mind, it is in question if a user would trust in all of those CAs if he was aware of them.

[1] due to a silent update mechanism

We discuss related work dealing with the problems of the Web PKI in Section 3.

3 Related Work

One research direction focuses on the general deployment and the complete landscape of the Web PKI. Several works approach that issue by scanning the Internet for SSL connections and collecting the certificates in order to analyze those. The EFF conducted a scan of the entire IPv4 space and downloaded the certificates. Results were presented in [9–12]. They focused on the certification infrastructure as for example the number and relationships between CAs. The authors of [23] scanned the most popular web pages and passively scanned the traffic of Internet users from different locations around the world to collect certificates, focusing on the quality of certification practices. The authors identify many problems and conclude that the Web PKI is in a sorry state. The ICSI SSL Notary [25] is another work which aims at the detection of malicious certificates by passively scanning the Internet traffic aggregating the observed certificates into a central database. From this database, they derived the Tree of Trust, which shows the observed CAs and and intermediate CAs as well as their certification paths. All these works have in common, that they show a global view on the Web PKI considered from different angles. However, it is not possible to analyze how and which parts of the Web PKI are seen and actually required by a single user.

A common issue considered in these works is that each CA is a single point of failure for the complete ecosystem. In [34] the authors consider the threat of Governments compelling CAs, as CAs have the power to make ultimate surveillance possible by enabling man-in-the-middle attacks. The paper [24] proposes a tool to detect a possible man-in-the-middle attack and track the location of the adversary.

Over the years, many tools have been proposed to enhance Internet security and counteract threats imposed by possible CA failures. The approaches can be distinguished into those enhancing the PKI system with additional measures and those completely bypassing the mechanisms of PKI. However, all these approaches define additional or alternative measures coming with their own disadvantages, instead of improving on the inherent weaknesses of the deployed PKI.

One technique, called certificate pinning, lets users store certificates to websites they browse frequently. The certificates are then reused whenever they access the website again. This can be realized by implementing the trust on first use (TOFU) approach, where a certificate is added to a local trust store when a website is accessed for the first time. Add-ons like Certificate Patrol [33] and

Trustbar [22] implement this approach and additionally support the user in deciding about the trustworthiness of a presented certificate, i.e. by showing certificate information. Unfortunately, such a trust decision requires the user to have PKI expertise to a certain extent. However, in any case, TOFU requires an adversary to be present during the first connection establishment. Another possibility to realize certificate pinning is to include keys or white-lists of CAs directly into software as done by the Chrome browser, where pinning is limited to Google services.

Notarial solutions like the above ICSI SSL Notary [25], Convergence [28] or Perspectives [5] maintain databases containing formerly observed certificates and can be queried to reconfirm the authenticity of a specific certificate, sometimes also involving consensus decisions of several independent notary servers. Notarial solutions can be used to enhance the security of the PKI system or completely bypass CAs, if a user fully relies on the notarial decisions. Yet, in some sense this approach only defers the trust requirements from the CAs to the notaries.

A recently quite often discussed alternative to the X.509 PKI is the binding of certificates – above all self-signed certificates – to Domain Name System (DNS) names using DNSSEC. However, the security then relies on the security of the DNS infrastructure. Furthermore, DNSSEC is still in its infancy.

Another project we shortly want to mention is called Monkeysphere [2], which tries to make CAs completely obsolete. It aims at the application of PGP certificates in the context of web server authentication. In PGP, users sign certificates of other users and in this case also of web servers. Basically, a certificate is considered trustworthy if it is signed by anyone the validating user trusts in. However, the approach comes with scalability problems inherited from PGP, and the requirement of users to become actively involved.

4 Scope and Methodology

In the work at hand, we consider the question if and how it is possible to individually minimize the set of trusted CAs and therewith the attack points that threaten the user. Is a PKI user really required to trust in the whole set of CAs? Or is a small subset of those CAs enough to keep all his applications working and his desired web pages reachable via https without additional browser warnings? The relevance of this research is shown by the problems with the Web PKI discussed in Section 2. The approach is a logical consequence which is shown by a short example: One can ask, why a user should trust in a CA, which is supposedly exclusively used for e-government applications in a foreign country. He will – with high probability – never use such services. Thus, why should he take the risk of trusting in that CA? The simple answer to this is that there is no reason for it.

With these considerations in mind, the set of actually required CAs seems to be highly dependent on the user's applications and browsing behaviors. To answer the question if and how the set of trusted CAs can be reduced, the parts of the Web PKI relevant to a specific user need to be identified. Furthermore, the evolution of this personal view on the Web PKI must be assessed, to see if the set of trusted CAs can be reduced to a minimal subset. Given the subset of required CAs is more or less constant, a reduction to that set of CAs once they are known is possible. On the other hand, frequent changes in the observed CAs would require further mechanisms to maintain those changes in order to enable the reduction without introducing unavailabilities and interoperability problems.

So far, a detailed study on the issue of individual user requirements concerning the Web PKI is not available. We are only aware of descriptions for advanced users [7] and experimental self studies [32], which are not generally applicable to all Internet users. However, these experiments show that the reduction of trusted CAs is a promising approach.

To approach the research questions, we analyze the individually required CAs based on the past browsing behavior by analyzing the user's browser histories. We aim at deriving the complete view on the observed part of the PKI and also include Sub CAs, as also a small number of Root CAs can still impose a huge overall number of CAs. We also trace how this personal view on the Web PKI develops over time.

It can easily be seen that this is a privacy sensitive topic. Therefore, we chose the approach to analyze the histories locally instead of collecting user histories for analysis. This enables the users on the one hand to control the data which they are about to hand out and on the other hand enables us to show parts of the individual results directly to interested users. For the data collection we implemented a tool which is run locally on the user's machine. It reads the user's browsing history and gathers information about the CAs relevant for the respective user based on the certificates that are obtained from servers when opening https connections. We collect this data from all participating users and aggregate it for further analysis. Thereby we derive user specific information as well as similarities and differences among user groups. To be able to group the participants in this study we collect metadata based on a questionnaire.

5 The Tool – Rootopia

Our tool is called Rootopia. It is implemented in Java 1.7 to be as platform independent as possible. It currently runs under Windows and Mac OS X, but will soon be available also for Linux. It can read history files from Mozilla Firefox or Google Chrome directly. Looking at Microsoft Internet Explorer (IE), we decided to use an existing external tool called IEHistoryView [1] due to compatibility reasons with different Windows and IE versions. IEHistoryView allows to read the IE history and store it in a text file, which can then be read by Rootopia. IEHistoryView can be downloaded out of Rootopia's menu to make the use as convenient as possible. With the possibility of supporting these three browsers we are currently able to serve over 90% of users [36].

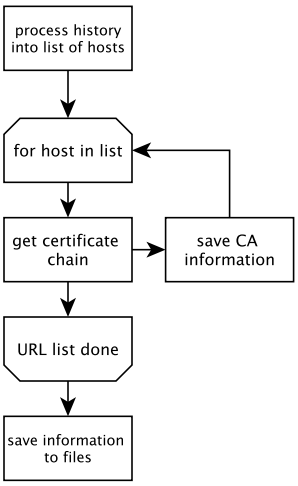

Figure 2: Rootopia Control Flow

5.1 Functionality

As we are interested in the CAs seen by the specific user, the first step is to extract the hosts to which https connections where established in the past. As Firefox and Chrome store their histories within SQLite Databases, the relevant entries can be queried directly. In the case of IE we filter the hosts contained in the exported text file for URLs starting with "https". The hosts are then filtered for multiple occurrences and sorted by the date when they where first accessed. For each host, the date of the first and the last visit is stored to draw conclusions on the dates when a related CA was observed. To gather information about the CAs involved in the certification of the respective host, the tool establishes an SSL connection to each of the hosts and retrieves the certificates provided by the host server. The certificates are then analyzed to identify the involved CAs and the corresponding certificate chains. Note that for each obtained certificate chain, the path validation provided by the Java Cryptography Architecture with the default TrustManager, which implements the standard X.509 path validation, is executed. Path validation is based on the root store provided by Java which contains 79 Root CAs. We only include certificate chains that could be validated in our analysis to ensure, that only valid CAs are counted. Note that this also leads to the exclusion of several CAs that are seen but e.g. only included within incomplete chains. Thus, our numbers might slightly underestimate the total number of observed CAs. On the other hand, if the CAs are only identified within chains that cannot be validated, a removal of those CAs from the root store does not change the user experience as an error is shown anyway. In case path validation fails, we store the certificates to evaluate the failures. We refer the reader to Section 7.5 for a discussion. Fig. 2 visualizes the steps. The collected data is stored into different files to be available for further investigation and comparison. We decided to store the data within CSV files, as these are on the one hand easily machine readable and on the other hand can be conveniently viewed and processed by major spreadsheet programs. This enables users to see and decide which data they provide for analysis counteracting privacy concerns. Additionally, our tool provides a visualization of the connections and dependencies between observed CAs. Furthermore, the tool shows a time course with the observed CAs so the user can see which CA has been seen for the first time at a specific date.

5.2 Simplifying Assumptions

The problem with rebuilding the view of a user on the part of the Web PKI he has seen so far is that the CA data is not available from prior interactions. Certificate chains obtained to establish a https connection are not stored. Some browsers, like Firefox, cache intermediate CAs from former visits. However, it is not possible to determine when the CA was first seen and how often. Thus, one cannot examine the development of the view and the importance of a CA for a user directly.

Hence, we assume that the certificates presented by a host are the same, or at least from the same CA as in the past. That means a certain host obtains its certificates (for example after an old certificate expires) from one and the same CA and uses the same certificate chain for each connection establishment. Thus, we retrieve the certificates from each host only once.

This is a slight simplification as huge server farms such as Google or Facebook (see also Section 8.1), are known to have certificates from different CAs. However,

as the assumption is true for most of the hosts, our analysis is close to reality. Furthermore, we tested the behavior by connecting to Amazon, Google, Facebook and Dropbox repeatedly over several days. The only host presenting different certificates during our tests was Dropbox. Furthermore, we executed Rootopia repeatedly on our own browser histories and always obtained the same set of CAs, which further justifies the assumption.

What we actually get from our approach is a current snapshot of the Web PKI seen by the user, i.e. we see all the CAs which are required at the moment of the analysis, to be able to establish all former https connections. Yet, assuming that a host sticks to one CA allows us to infer from this snapshot, how the view developed in the past.

5.3 Collected Data

Now we describe the data groups that are collected and explain the reasons and why this data is interesting. Table 1 shows the data sets we collected for each CA seen by a specific user. We identified CAs with their distinguished names (DN) extracted from the issuer and subject fields of the obtained certificates. "first seen" is the date, when – according to the user's history – for the first time a connection to any host was established where the CA was involved into the certificate chain. "last seen" analogously specifies the date when for the last time a connection was established to any host involving the CA. With that data it is possible to examine, how the view on the Web PKI changes over time according to the user's browsing behavior.

Data ID	Meaning
DN	identifies the CA
CA kind	specifies the role of the CA (Root, Root/Sub, Sub) in which the CA was observed
certificate	the certificate(s) certifying the CA
first seen	specifies the date when a CA was first seen
last seen	specifies the date when a CA was seen for the last time
Sub CAs	DNs of the CAs certified by the CA
Super CAs	DNs of the CAs certifying the CA
# hosts	number of hosts that have the CA in their certificate chain
# visits	total number TLS connections involving the CA
EE CA	boolean, specifies if a CA certifies end-entities

Table 1: Collected Data for each CA

The number of hosts and visits shows the relative importance of the respective CA for the user. The number of visits is the sum of the number of visits of the hosts related to the respective CA. Sub CAs and Super CAs represent the relationship between CAs. "CA kind" tells us how the respective CA appears in the different certificate chains. All Root CAs and Root/Sub CAs together define the absolutely minimal set of CAs that must be contained in a user's root store in order to be able to validate the certificate chains for all his previously accessed hosts. Root/Sub CAs are such CAs, that were seen both, sometimes as Root CA and sometimes as Sub CA. In contrast storing all CAs, where "EE CA" is true – which means that the respective CA has issued at least one end entity certificate – is the absolute minimal set of CAs that have to be trusted. All host certificates could be validated with a chain of length one, meaning that certification of Sub CAs could be completely ignored.

Thus, in the first case the number of stored CAs is minimal, in the second the number of trusted CAs is minimal. This is because in the first case, the CAs considered trustworthy may have an arbitrary number of Sub CAs which are transitively trusted, even if never needed by the user.

6 Analysis of the Web PKI - Setup

For the analysis we asked Internet users to participate. For the pilot study whose results are presented in Section 7, we mainly asked colleagues, students, friends and family members. The participants were asked to run the tool Rootopia on their browser histories. Afterward, we collected the generated data for further analysis. Before the participants handed over the collected data, the data was explained to the participants personally to enable them to understand the extend of data collection, and if required to refuse their data from being used. Besides that, we collected metadata using a questionnaire to be able to group the people into different categories. Within the questionnaire, we ask for different aspects which might have major influence on the browsing behavior and thereby on the viewed part of the PKI (see Section 6.1). The results are fully anonymized. Furthermore, for privacy reasons it is possible for a participant to deny the storage of the found host names.

6.1 Questionnaire

The questionnaire consists of the questions summarized in Table 2. The data is to be used to group the participants in order to analyze differences between user groups. The country of origin and country of residence are interesting, as people from different countries might be interested in different web pages due to language and social background. The interesting question here is if this has influence on the CAs, and in particular, on the countries where the CAs are located. Furthermore, PC and Internet usage give information on how intensively the PC and the Web is used. Intensive use may on the one hand lead to a bigger set of CAs and on the other hand, may lead to the complete set of required CAs in a much shorter time span. Furthermore, the use of business PCs is often restricted according to security policies of a company. Besides that, users that do e-Commerce, e-Banking and e-Government are more likely to often come in contact with secure connections. We wanted to figure out if this

has any implications on the seen CAs. IT security and general IT expertise may have implications on how people use the Internet in general. The last question refers to https enforcement tools like HTTPS Everywhere [13]. Those tools enforce https connections instead of http in many cases, which might have influence on the set of seen CAs. We analyzed the data according to these aspects. The results can be found in Section 7.

Criterion	Possible Answers
gender	male / female
country of origin	country name
country of residence	country name
PC usage	private / business / both
Internet usage	
e-commerce	yes / no
e-banking	yes / no
e-government	yes / no
hours per day	# of daily online hours
IT security expertise	expert / knowledgeable / some familarity / no familarity
general IT expertise	expert / knowledgeable / some familarity / no familarity
use of https tools	name of the tool, otherwise "-"

Table 2: Collected metadata per participant

6.2 Challenges

As mentioned before, Rootopia analyzes a user's browser history and therewith his browsing behavior. As this is a privacy sensitive task, we were confronted with several privacy and security concerns of the users. On the one hand users are not willing to hand out their browser histories. This is why we used the approach of a local analysis and provided the participants with all the data which is sent to us for further research. However, many users also feel uncomfortable with executing unknown programs on their PCs, and convincing people that no privacy sensitive data is extracted is not always an easy task. This has to be considered when conducting a broad scale study.

Another problem we were faced with is that many users regularly delete their browser histories, only leaving a short time span for analysis. In general this leads to an incomplete view on the required part of the CA ecosystem and allows only restricted conclusions.

7 Findings

We analyzed the browser histories of 22 persons. Four persons provided two histories, either from different browsers they use in parallel, or different PCs. We ended up with 26 history files. All persons currently live in Germany, but have different cultural backgrounds. The participants reach from IT experts to persons that only occasionally use a PC. We present some general data on the analyzed histories in Table 3. In the analysis we distinguished between true Root CAs and CAs that were seen both as Root and as Sub CAs (Root/Sub CAs). This occurred, as occasionally a (cross-) certificate from another Root CA in the root store certifying the respective CA was included into the chain and sometimes not. However, as both Root and Root/Sub CAs are required to be present in the root store to be able to validate each presented chain, we will in the following refer to the sum of them as the Root CAs if not explicitly distinguished between the two cases.

Interestingly, none of the users – even heavy users with a huge number of different https hosts – did see more than 22 different Root CAs, which is about 13.4% of the 164 CAs included in the Firefox root store. Furthermore, a maximum of 75 Sub CAs was reached. The absolute maximum of CAs in total seen by a single Internet user was 96, which is 6.5% of the 1,482 CAs found by the EFF. Even less CAs were found when only considering CAs that signed host certificates. Those CAs represent the minimum number of CAs that need to be trusted by a user to be able to verify all the certificates of the hosts he connected to. The maximum value of such host signing CAs was 68 or in other words 4.6% of the CAs observed by the EFF. Furthermore, the ratio of host signing CAs was in the span of 50%-75% of the total CAs found for the respective user and reached 63% on average. Note that our numbers might slightly underestimate the really required CAs as we did not consider certificate chains that could not be validated by the Java path validation with its internal root store and also some connections to hosts could not be established. Please see Section 7.5 for details on that issue.

Criterion	Average	Min	Max
Duration of analyzed period (months):	18	4	38
Total number of https hosts:	168	12	636
Total number of https connections:	18,475	162	159,882
Total number of Root CAs:	10	4	14
Total number of Root/Sub CAs:	4	0	8
Total number of Sub CAs:	36	11	75
Root + Root/Sub CAs:	14	4	22
# CAs that signed host certificates:	33	8	68

Table 3: Collected metadata per participant

Considering the total number of different Root and Sub CAs observed by the whole group of participants, namely the union of all sets of CAs, leads to 28 Root CAs and 145 Sub CAs (please find a list of all CAs in Appendix B). This shows that there is a high potential in limiting the number of trusted CAs. Furthermore, there is a high overlap in the CAs (i.e. CAs that were observed by several persons). The overlap is significantly higher for Root CAs than for Sub CAs, which can be seen as the set union of Root CAs is only 27% larger than the maximum number of Root CAs of a single user, while in the

case of Sub CAs the set union consists of twice the number of Sub CAs seen by a single user. However, the significant differences in the numbers for different users – reflected in the minimum and maximum values – shows, that true minima for a single user can only be reached by individualization. Still, grouping the users into dedicated user groups can lead to good results.

One influencing factor leading to a low number of different CAs is surely the fact, that there are few large CA companies with a high market share in the certification business. However, when considering the distribution we observed among those large players, it turns out that it is not according to the market shares from the Netcraft SSL Survey [31]. Most significant, VeriSign, Inc is involved into more than 20% of the certification paths relevant for our user group, while it has only around 6% of the market share in the Netcraft Survey. In contrast, Go Daddy with more than 20% of market share achieves only a rather low rate in our data, namely Go Daddy was a Root CA in only less then 4% of the certificate chains. This is another indication, that it highly depends on the individual browsing behavior of the users, which CAs are truly relevant for them. For a complete List of the CAs and their respective relevance, we refer the reader to Appendix B.

We also grouped the observed CAs by country, where it turned out that CAs from only 14 different countries were relevant for the considered user set (see Figure 7). The overwhelming majority of CAs is from the US (US) followed by Germany (DE), Great Britain (GB) and Belgium (BE). Considering the other countries, less than 5 CAs were observed from those and often only by very few users (cf. Section 7.2 for details). This shows, that it might even be viable to limit the number of trusted CAs based on the country they reside in.

In the following we present detailed results also considering different user groups, and how the views develop over time.

7.1 Temporal Evolution

In the following we discuss our findings concerning the development of the individual views on the Web PKI over time according to the dates when related URLs were accessed. It turns out that, in general, the number of observed CAs does not grow linear but shows restricted growth with high growth rates in the first few months. Considering Root CAs, the upper bound is reached after several months. However, growth rates depend on the intensity of Internet usage or rather on the number of https hosts a user connects to. Considering users with high numbers of https hosts the upper bound is reached faster than for users that only connect to https occasionally. For Sub CAs, the development is similar to the Root CAs, however, it is much less significant. Thus, the number of Sub CAs tends to keep growing over a long time. The temporal evolution of the number of Root and Sub CAs is shown in Figures 3 and 4. The figures include the data for the four users, with the highest number of different https hosts, averaged over the length of the analyzed time span. In contrast Figures 5 and 6 show the evolution of the view on the PKI for the ten histories containing the least number of https hosts. The number of CAs depicts how many CAs were seen during a https connection found in the user's history until the respective month.

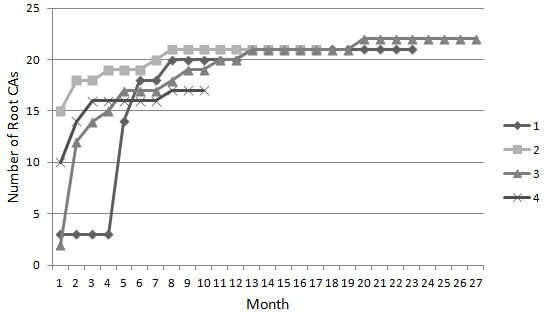

Figure 3: Temporal Evolution: Root CAs - high number of https hosts

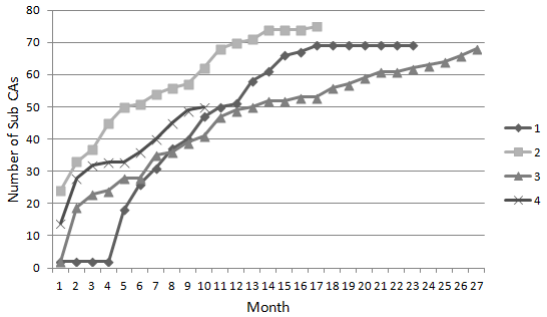

Figure 4: Temporal Evolution: Sub CAs - high number of https hosts

For the users that use https less intensively, it takes a much longer time span until the number of CAs tends towards an upper bound, but the same tendency is observable. Besides that, the upper bounds seem to lie strictly below the ones observed for heavy users.

On the other hand, there are users, that only connect to a very limited number of hosts but where the upper bounds on CAs are reached after very few months. This can be seen best in one data set, where the maximum of 4 Root CAs is reached after 3 months and is constant afterward (16 months). The picture for Sub CAs is nearly the same in that data set. Having a closer look shows, that the data belongs to a person doing e-banking and e-commerce, but besides that only occasionally surfs the Internet (a fact, which was identified during a personal

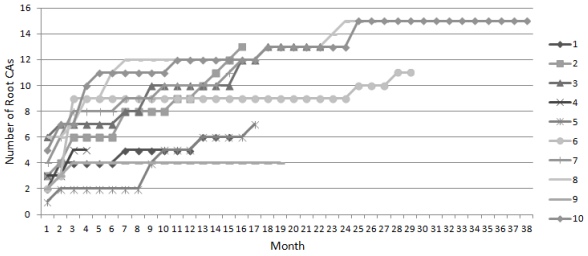

Figure 5: Temporal Evolution: Root CAs - low number of https hosts

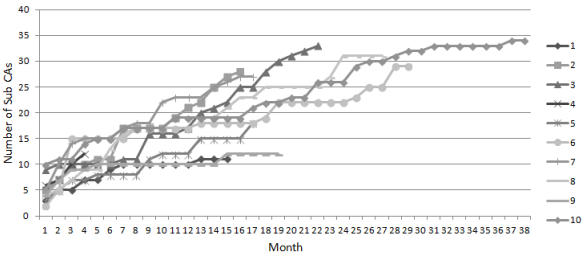

Figure 6: Temporal Evolution: Sub CAs - low number of https hosts

discussion).

To summarize our findings on the development over time, we state that it is not possible to give a concrete number of months after which all relevant CAs have been seen and the number of CAs stagnates. This is highly depended on the individual browsing behavior. However, in many cases due to the regular deletion of the histories, these are not long enough to derive the upper bound and the set of relevant CAs for the respective user. Yet, in general, our observations show that the number of CAs tends towards an upper bound significantly below the total number of existing CAs. This in turn shows the potential for the security gain by limiting the number of CAs. For completeness, please find the temporal evolution of the number of CAs for all analyzed data sets in Appendix A.

7.2 CA countries

As stated above, most of the observed CAs are from the US. Figure 7 shows the observed countries and the number of CAs including all data sets. The second most observed country in our set of participants is Germany. However, this is also a user group dependent outcome and results from the set of analyzed histories. A large number of participants are either from the scientific community or students at a university. Building two groups, the first containing people with academic background and the second one without, shows that the occurrence

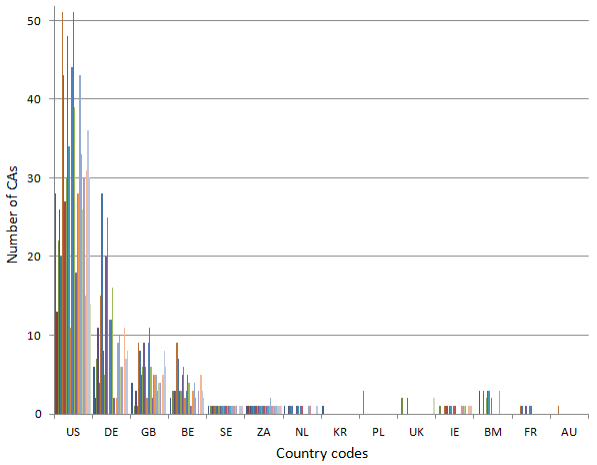

Figure 7: Distribution of CA countries, different colors represent different users

of German CAs is much less for the second group. The percentage of German CAs is on average 18.3% of all observed CAs per user in the first, and only 7.1% in the second group. It results from the fact, that most universities have their own CAs, certified by the DFN Root CA. Those CAs are completely irrelevant for the non-academic users. The distribution of CAs over the other countries did not change significantly. Besides that we grouped the data into users that originate from Germany and those who do not. Yet, interestingly this did not have significant effects on the distribution over the countries. However, when considering single users, the relevant CA countries can depend on the country of origin as we observed it for a user from Poland (PL). A grouping into different countries of residence would be interesting, yet could not be done with our data set and thus is left for future work.

Considering all data sets, there are country codes that were observed for most of the participants. These are SE, ZA, NL, and IE. However, the observed CA was always one and the same. We collected these CAs in Table 4.

DN	Country
EMAILADDRESS=premium-server@thawte.com, CN=Thawte Premium Server CA, OU=Certification Services Division, O=Thawte Consulting cc, L=Cape Town, ST=Western Cape, C=ZA	ZA
CN=TERENA SSL CA, O=TERENA, C=NL	NL
CN=AddTrust External CA Root, OU=AddTrust External TTP Network, O=AddTrust AB, C=SE	SE
CN=Baltimore CyberTrust Root, OU=CyberTrust, O=Baltimore, C=IE	IE

Table 4: Important CAs from rarely observed countries

For the remaining countries (KR, PL, UK, BM, FR, AU) no fix pattern is observable with FR and BM observed most often.

7.3 Relevance of Respective CAs

To measure the relevance of a CA for a user, we counted the number of hosts related to the respective CA. Interestingly, the percentage of Sub CAs that are related to only one host lies between 20% and 60% and is about 43% on average, which shows that Internet users observe many CAs where the relevance of each one is really low. Thus it is highly questionable if the benefits for the user by trusting into those CAs counterbalances the imposed risks, not speaking about the CAs a user never observes. As it might occur that a single host is accessed extremely often by one user and thus the related CA becomes more relevant to him, we also measured the number of visits, namely taking into account how often a host was accessed. As expected, the number of Sub CAs only observed during a single connection is lower. But still, rates of up to 38% of the total number of CAs for single users are reached and are 17.5% on average. That shows that many of the CAs are only observed by chance.

Furthermore, our data shows that a user observes the CAs most relevant for him during the first months, while CAs which are found later are less relevant, both either measured by the relative number of hosts or visits. As mentioned above we averaged the CA's relevance over all users. It can be seen, that there is a strong correlation between the number of occurrences in the data sets of a CA, i.e. how many users observed a CA, and the averaged relevance of the respective CA. The numbers can be found in Appendix B. From this findings we conclude, that building user groups and taking the CAs which most users of that group have in common can be a good starting point to set up an individualized root store for e.g. a user where no history data is available.

7.4 Number of CAs and Overlaps

We computed the union set of CAs for different user groups. To identify the similarity of the views on the Web PKI within a group, we computed overlaps in the CAs, namely how many users have how many CAs in common. If not differently specified, in the following with overlap we mean the ratio of CAs that all group members have in common.

The group of the four users with most https hosts leads to 25 Root CAs and 108 Sub CAs. With 64% the overlap in the Root CAs is twice the overlap for Sub CAs (31%). That shows, that the set of Root CAs relevant to a user is much less dependent on the individual browsing behavior. This also holds for the other groupings we analyzed and is as expected, as the total number of existing Root CAs is nearly ten times smaller than the number of Sub CAs. Comparing the CA sets of the group of heavy users with the union set of all data sets, it shows that the heavy users make up for 89% of all Root CAs and 74% of the Sub CAs. Thus, most of the CAs required by the other users are also seen by the heavy users.

When comparing the groups of academic and non-academic users, the first observe significantly more CAs (27 vs. 19 Root CAs and 140 vs. 63 Sub CAs). This seems to result from the fact, that all the heavy users are also part of the academic group. The overlaps in the academic group are higher than in the non-academic group.

However, to really do a fine grained grouping, more data sets are required. High overlaps were achieved only for the group of heavy users, thus an interesting remaining question is if this results from the fact that the heavy users reach a set of CAs that satisfy the requirements of most of the users or if the grouping resulted into a good match of browsing behavior. There are indications for both. The observation, that the remaining users do not need too many additional CAs speaks in favor of the first. On the other hand, the heavy users are all from the same scientific working group, and have comparable backgrounds.

7.5 Path Validation Errors

As mentioned above, we only considered valid certificate chains in our evaluation of the user's views on the PKI guaranteeing that no invalid CA certificates (due to revocation, expiry or other reasons) are considered. Yet, we also did a quick check on the certificate chains obtained from hosts where the connection failed due to path validation errors. On average, about 10% of the connections failed due to path validation errors, which is strictly below the the numbers of failed chains observed by other studies [23]. This mainly comes from the fact, that Firefox as well as Chrome do not store URLs in the history in case path validation fails and if no exception is added by the user. Thus, many of the invalid paths are filtered by the browsers and therefore do not occur in our analysis. Most fails result from Java's smaller root store and less robust path validation or from CAs manually added by the users.Nevertheless, we found some interesting reasons when checking the failed paths.

We only present a short overview on the reasons. A rather large set of failed path validations resulted as the Root CA was not included within the Java root store, like StartCom. In other cases, which mostly occurred with servers from universities, was that they only presented the host certificate without any certificate chain. This results from the fact that universities often use their own CAs to sign their host certificates, but do not present a chain up to a Root CA, which requires their users to manually add these CAs to their root stores. Interestingly, we often found those university's CAs within another chain which could be validated. Thus, several servers just do

not provide the complete chain even if there exists one.

Other interesting cases of broken certificate chains resulted from a wrong ordering of the certificates (although the chain was complete), additional certificates obtained together with the chain but which were completely unrelated to the rest of the chain, or multiple occurrences of one and the same certificate within the chain. Besides that we found several self-signed certificates.

In summary, when also considering CAs contained in the failed chains, this would increase the numbers of CAs by about 10% on average per user. While changing the absolute numbers, the inclusion of such chains has only small effects on our general results. Thus, expanding our analysis to those chains is left for future work, and requires a more thorough analysis of the failures.

8 Evaluation and Future Work

In this section we describe limitations of our approach and discuss how the data can be used to minimize the set of trusted CAs in practice. We also present interesting research questions for future work.

8.1 Limitations

Here we summarize some limitations of our approach. While we can draw a good picture of a user's past view on the Web PKI it is not possible to predict the future. Our data shows, that the number of CAs approaches a certain bound after some time. However, new CAs do occur after long time periods. Thus, derived views might always lack some CAs.

Furthermore, the question remains how to deal with hosts that have their certificates issued by several different CAs. If we consider Google as an example, the issuers of certificates for Google currently include Verisign, Google Internet Authority, Equifax, GeoTrust, and DigiCert [26], which shows that the number of different CAs can in practice be quite large. One possibility could be to adjust this manually to complete the views on the Web PKI. However, further research is needed to identify those hosts and their actual behavior, as we have discussed above, for many of those hosts the same certificate chains where presented during repeated connections.

Another limitation in the approach evolves from the fact, that many users delete their history – partly or completely – quite often. In such cases, it is not possible to derive the CAs relevant to the user. Furthermore, the browser histories do not contain the CAs relevant to applications installed on the user's system. Thus, we miss those CAs that might be relevant for e.g. software updates.

8.2 Minimizing the Attack Surface

The results from our analysis of the Web PKI show that, in case of individualization of the root stores, the number of CAs can be reduced drastically for the respective user. However, due to the discussed limitations, the identified set of CAs cannot be used for the individualization without further considerations. The applicability also depends on the usage scope of the root store that is about to be limited. As far as a root store exclusively used by a browser – like in the case of Firefox – is concerned, the data generated by Rootopia builds a good starting point. But still, further research is needed on the question how to deal with CAs that are observed later, since the tool can only provide a snapshot of the current situation. Our results also show that this can occur at any time, even though the probability shrinks the longer the analyzed browsing history is. A possibility would be to only accept new CAs for new hosts, however this requires to keep browsing histories, or at least to maintain a list of formerly observed https hosts. Another promising approach could be to combine the reduction of the set of trusted CAs with a sophisticated update mechanism that does further security checks whenever a new CA is observed.

Furthermore, the results show that a limitation of the total number of CAs (by explicitly limiting the Sub CAs) is much more challenging than a reduction of the Root CAs. This is because the set of relevant Root CAs is much more limited, while Sub CAs are often observed with only one single host, and the set of relevant Sub CAs keeps growing in many cases without approaching a maximum within the analyzed time periods.

Considering root stores that reside within the operating system, the limitation is still problematic. A multitude of different applications rely on those root stores and thus, those dependencies must be considered to prevent these applications from stopping to work properly. From our point of view, several separated root stores dedicated to different purposes have advantages over a central all-purpose root store in respect to the minimization of the attack surface.

8.3 Future Work

While the number of participants in our first study was enough to show the relevance of our approach, a larger study is required to identify potential differences between several user groups. Here, the views from different countries is of special interest. Furthermore, a large number of participants is relevant to the possibility to group users into different categories. An interesting question is, if group profiles can be derived and how they can be applied to define the set of required CAs for users where

browser histories are not available. Even if such group profiles might overestimate the set of actually required CAs, this will still lead to a significant reduction of the risks.

Another interesting research question is to consider the view on the Web PKI within certain contexts, e.g. e-banking, e-commerce or general web surfing. This would allow to make the decision if a CA is considered trustworthy context dependent. Assigning CAs to a certain context can further reduce the set of trusted CAs, limiting the impact of malfunctions of the system.

Besides that, Rootopia is to be extended by a browser Plugin that allows the automatized reduction of the root store of the Firefox browser. However, many browsers, e.g. Chrome and IE, use the root store of the operating system. We will examine possibilities to detect CAs required by the installed applications to get a broader view on the required PKI parts and to identify possibilities for a minimization of these root stores. Thereby, interdependencies with root store updates as for example applied by Microsoft need to be considered.

9 Conclusion

In this work we showed that the risk to be affected by CA failures is unnecessarily high. We presented a tool that allows to derive and assess the personal requirements of Internet users based on their browser histories. It turned out, that the individual views on the Web PKI tend towards a fixed individual set of CAs. The temporal evolution described in Section 7.1 actually shows different courses, thus confirming that the set depends on a user's individual browsing behavior. Our analysis revealed, that a reduction of trusted CAs by more than 90% is possible without restricting the respective user in his daily Internet use. Furthermore, there are big differences in the relevance of the CAs, which leaves further room for improvement. Also, a limitation based on the countries the CAs reside in is promising. The CAs we observed for different users originate from a rather small set of countries. On the other hand, it turned out that it is a challenging task to completely define the set of relevant CAs for an individual user. One problem is the unavailability of sufficient data about the user's browsing history. In such cases, grouping users and deriving group profiles can help to provide a starting point for the limitation. Further research is needed to define such profiles. On the other hand, mechanisms are needed to deal with CAs that are newly observed and interdependencies between the root store under consideration and applications apart from browsers need to be considered. We conclude that the individualization of root stores bears huge security improvements. However, the realization of those improvements remains challenging.

References

[1] IEHistory View. http://download.cnet.com/IEHistoryView/3000-2381_4-10448770.html.

[2] The monkeysphere project. http://web.monkeysphere.info/.

[3] Microsoft root certificate program, 2009. http://technet.microsoft.com/en-us/library/cc751157.aspx.

[4] Mozilla ca certificate policy, 2013. http://www.mozilla.org/projects/security/certs/policy/.

[5] CARNEGIE MELLON UNIVERSITY. Perspectives Project. http://perspectives-project.org//, visited July 2012.

[6] COMODO. The Recent RA Compromise. http://blogs.comodo.com/it-security/data-security/the-recent-ra-compromise/, visited Nov. 2011.

[7] CONETRIX.COM. How-to Limit the Number of Certificate Authorities Your Browser Trusts, 2011.

[8] COOPER, D., SANTESSON, S., FARRELL, S., BOEYEN, S., HOUSLEY, R., AND POLK, W. RFC 5280 – Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile. RFC 5280 (Proposed Standard), 2008.

[9] ECKERSLEY, P., AND BURNS, J. An Observatory for the SSLiverse. Defcon 18, July 2010. https://www.eff.org/files/DefconSSLiverse.pdf.

[10] ECKERSLEY, P., AND BURNS, J. Is the SSLiverse a Safe Place? 27C3, 2010. https://www.eff.org/files/ccc2010.pdf.

[11] ECKERSLEY, P., AND BURNS, J. The (Decentralized) SSL Observatory. Invited talk at 20th USENIX Security Symposium, August 2011.

[12] The EFF SSL Observatory. https://www.eff.org/observatory.

[13] ELECTRONIC FRONTIER FOUNDATION. HTTPS Everywhere. https://www.eff.org/https-everywhere. visited July 2012.

[14] ELLISON, C., AND SCHNEIER, B. Ten Risks of PKI: What You're Not Being Told About Public Key Infrastructure. *Computer Security Journal 16*, 1 (2000), 1–7.

[15] GUTMANN, P. Pki: it's not dead, just resting. *Computer 35*, 8 (aug 2002), 41 – 49.

[16] GUTMANN, P. *Engineering Security*. 2013. Book draft available online at http://www.cs.auckland.ac.nz/~pgut001/pubs/book.pdf.

[17] H ONLINE. Fake google certificate is the result of a hack, 2011. http://h-online.com/-1333728.

[18] H ONLINE. Flame – oversights and expertise made for windows update worst case scenario, 2012. http://h-online.com/-16142341.

[19] H ONLINE. Attack on Israeli Certificate Authority. http://h-online.com/-1264008, visited Nov. 2011.

[20] H ONLINE. Fake Google certificate is the result of a hack. http://h-online.com/-1333728, visited Nov. 2011.

[21] HERLEY, C. So long, and no thanks for the externalities: the rational rejection of security advice by users. In *Proceedings of the 2009 workshop on New security paradigms workshop* (New York, NY, USA, 2009), NSPW '09, ACM, pp. 133–144.

[22] HERZBERG, A., SINAI, A., DVORKIN, A. A., SHARFI, I., JBARA, A., VAITSMAN, A., AND TOV, R. TrustBar: Re-establishing Trust in the Web, 2006.

[23] HOLZ, R., BRAUN, L., KAMMENHUBER, N., AND CARLE, G. The ssl landscape: a thorough analysis of the x.509 pki using active and passive measurements. In *Proceedings of the 2011 ACM SIGCOMM conference on Internet measurement conference* (New York, NY, USA, 2011), IMC '11, ACM, pp. 427–444.

[24] HOLZ, R., RIEDMAIER, T., KAMMENHUBER, N., AND CARLE, G. X.509 forensics: Detecting and localising the ssl/tls men-in-the-middle. In *ESORICS* (2012), pp. 217–234.

[25] ICSI. The ICSI Certificate Notary, 2013. `http://notary.icsi.berkeley.edu/`.

[26] IMPERIALVIOLET. Public key pinning, 2011. `http://www.imperialviolet.org/2011/05/04/pinning.html`.

[27] LAURIE, B. Seven and a Half Non-risks of PKI: What You Shouldn't Be Told about Public Key Infrastructure. `https://www.apache-ssl.org/7.5things.txt`.

[28] MARLINSPIKE, M. Convergence. `http://convergence.io/`, visited July 2012.

[29] MICROSOFT SECURITY RESPONSE CENTER. Security Advisory 2798897 , 2012. `http://technet.microsoft.com/en-us/security/advisory/2798897`.

[30] MOZILLA. Mozilla CA Certificate Store - BuiltInCAs , 2012. `http://www.mozilla.org/projects/security/certs/`.

[31] NETCRAFT. Netcraft SSL Survey, 2011. `http://news.netcraft.com/ssl-survey/`.

[32] NETSEKURE.ORG. Results after 30 days of (almost) no trusted CAs, 2010. `http://netsekure.org/2010/05/results-after-30-days-of-almost-no-trusted-cas/`.

[33] PSYC. Certificate Patrol. `http://patrol.psyced.org/`.

[34] SOGHOIAN, C., AND STAMM, S. Certified lies: Detecting and defeating government interception attacks against ssl. Tech. rep., Indiana University Bloomington - Center for Applied Cybersecurity Research, 2010.

[35] SUNSHINE, J., EGELMAN, S., ALMUHIMEDI, H., ATRI, N., AND CRANOR, L. F. Crying wolf: An empirical study of ssl warning effectiveness. available online at `http://static.usenix.org/event/sec09/tech/full_papers/sunshine.pdf`.

[36] W3C. Browser statistics and trends, 2013. `http://www.w3schools.com/browsers/browsers_stats.aspl`.

A Temporal Evolution

The Figures 8 and 9 show the temporal evolution of the user's views on the Web PKI for all analyzed histories.

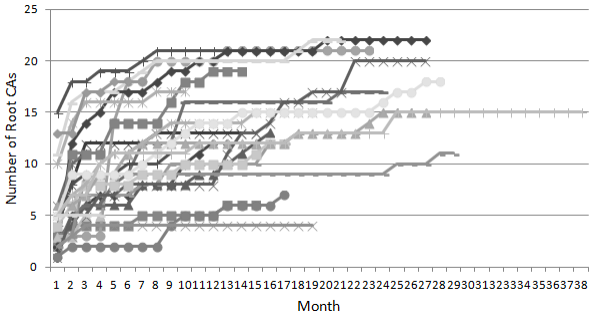

Figure 8: Temporal Evolution: Root CAs

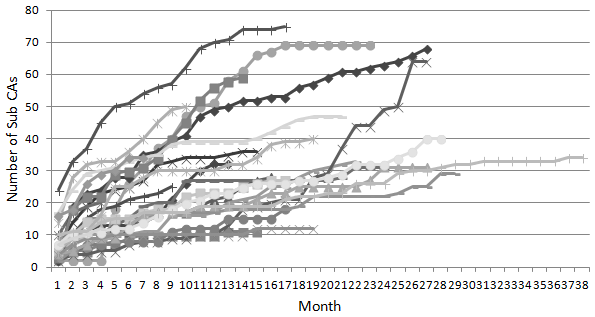

Figure 9: Temporal Evolution: Sub CAs

B Observed CAs

Table 5 shows all CAs that where observed as Root or Root/Sub CAs in one of the analyzed user data sets. "Relevance by data sets" and "Relevance by percentage of hosts" are indicators for the relevance to the user group whose data was analyzed. The former shows the number of user data sets that contained the respective CA, the latter shows with how many hosts the respective CA was observed in percent of the total hosts averaged over all users for which that CA was observed. Tables 6-9 show the same information for all observed Sub CAs.

DN	Relevance by data sets	Relevance by percentage of hosts
OU=Equifax Secure Certificate Authority, O=Equifax, C=US	26	16,65%
EMAILADDRESS=premium-server@thawte.com, CN=Thawte Premium Server CA, OU=Certification Services Division, O=Thawte Consulting cc, L=Cape Town, ST=Western Cape, C=ZA	26	11,52%
OU=Class 3 Public Primary Certification Authority, O="VeriSign, Inc.", C=US	26	20,54%
CN=AddTrust External CA Root, OU=AddTrust External TTP Network, O=AddTrust AB, C=SE	23	9,62%
CN=GeoTrust Global CA, O=GeoTrust Inc., C=US	23	7,42%
CN=GlobalSign Root CA, OU=Root CA, O=GlobalSign nv-sa, C=BE	22	4,00%
CN=VeriSign Class 3 Public Primary Certification Authority - G5, OU="(c) 2006 VeriSign, Inc. - For authorized use only", OU=VeriSign Trust Network, O="VeriSign, Inc.", C=US	22	20,09%
CN=GTE CyberTrust Global Root, OU="GTE CyberTrust Solutions, Inc.", O=GTE Corporation, C=US	21	4,01%
CN=Entrust.net Secure Server Certification Authority, OU=(c) 1999 Entrust.net Limited, OU=www.entrust.net/CPS incorp. by ref. (limits liab.), O=Entrust.net, C=US	20	3,51%
CN=Deutsche Telekom Root CA 2, OU=T-TeleSec Trust Center, O=Deutsche Telekom AG, C=DE	18	10,02%
OU=Go Daddy Class 2 Certification Authority, O="The Go Daddy Group, Inc.", C=US	18	3,82%
CN=TC TrustCenter Class 2 CA II, OU=TC TrustCenter Class 2 CA, O=TC TrustCenter GmbH, C=DE	16	2,10%
CN=DigiCert High Assurance EV Root CA, OU=www.digicert.com, O=DigiCert Inc, C=US	16	4,35%
CN=Entrust.net Certification Authority (2048), OU=(c) 1999 Entrust.net Limited, OU=www.entrust.net/CPS_2048 incorp. by ref. (limits liab.), O=Entrust.net	14	1,89%
CN=Baltimore CyberTrust Root, OU=CyberTrust, O=Baltimore, C=IE	14	2,93%
CN=thawte Primary Root CA, OU="(c) 2006 thawte, Inc. - For authorized use only", OU=Certification Services Division, O="thawte, Inc.", C=US	12	11,46%
CN=UTN-USERFirst-Hardware, OU=http://www.usertrust.com, O=The USERTRUST Network, L=Salt Lake City, ST=UT, C=US	11	3,18%
OU=VeriSign Trust Network, OU="(c) 1998 VeriSign, Inc. - For authorized use only", OU=Class 3 Public Primary Certification Authority - G2, O="VeriSign, Inc.", C=US	8	0,65%
CN=QuoVadis Root Certification Authority, OU=Root Certification Authority, O=QuoVadis Limited, C=BM	5	0,56%
EMAILADDRESS=info@valicert.com, CN=http://www.valicert.com/, OU=ValiCert Class 2 Policy Validation Authority, O="ValiCert, Inc.", L=ValiCert Validation Network	5	5,88%
OU=Starfield Class 2 Certification Authority, O="Starfield Technologies, Inc.", C=US	5	0,71%
CN=QuoVadis Root CA 2, O=QuoVadis Limited, C=BM	4	1,01%
CN=TC TrustCenter Universal CA I, OU=TC TrustCenter Universal CA, O=TC TrustCenter GmbH, C=DE	2	2,58%
CN=UTN - DATACorp SGC, OU=http://www.usertrust.com, O=The USERTRUST Network, L=Salt Lake City, ST=UT, C=US	2	0,78%
CN=Certum CA, O=Unizeto Sp. z o.o., C=PL	1	1,69%
CN=VeriSign Class 3 Public Primary Certification Authority - G3, OU="(c) 1999 VeriSign, Inc. - For authorized use only", OU=VeriSign Trust Network, O="VeriSign, Inc.", C=US	1	0,31%
CN=GeoTrust Primary Certification Authority, O=GeoTrust Inc., C=US	1	0,42%
CN=America Online Root Certification Authority 1, O=America Online Inc., C=US	1	0,88%

Table 5: Root CAs found for 26 Browser Histories

DN	Relevance by data sets	Relevance by percentage of hosts
CN=Thawte SSL CA, O="Thawte, Inc.", C=US	26	9,40%
CN=VeriSign Class 3 Extended Validation SSL CA, OU=Terms of use at https://www.verisign.com/rpa (c)06, OU=VeriSign Trust Network, O="VeriSign, Inc.", C=US	25	6,09%
CN=VeriSign Class 3 Secure Server CA - G3, OU=Terms of use at https://www.verisign.com/rpa (c)10, OU=VeriSign Trust Network, O="VeriSign, Inc.", C=US	25	5,34%
CN=Google Internet Authority, O=Google Inc, C=US	24	13,43%
CN=VeriSign Class 3 International Server CA - G3, OU=Terms of use at https://www.verisign.com/rpa (c)10, OU=VeriSign Trust Network, O="VeriSign, Inc.", C=US	24	4,30%
CN=VeriSign Class 3 Extended Validation SSL SGC CA, OU=Terms of use at https://www.verisign.com/rpa (c)06, OU=VeriSign Trust Network, O="VeriSign, Inc.", C=US	23	6,67%
CN=RapidSSL CA, O="GeoTrust, Inc.", C=US	23	3,60%
OU=www.verisign.com/CPS Incorp.by Ref. LIABILITY LTD.(c)97 VeriSign, OU=VeriSign International Server CA - Class 3, OU="VeriSign, Inc.", O=VeriSign Trust Network	22	2,04%
CN=GeoTrust SSL CA, O="GeoTrust, Inc.", C=US	21	4,44%
CN=DigiCert High Assurance CA-3, OU=www.digicert.com, O=DigiCert Inc, C=US	20	3,32%
CN=COMODO High-Assurance Secure Server CA, O=COMODO CA Limited, L=Salford, ST=Greater Manchester, C=GB	20	3,02%
CN=Akamai Subordinate CA 3, O=Akamai Technologies Inc, C=US	20	2,22%
CN=COMODO Certification Authority, O=COMODO CA Limited, L=Salford, ST=Greater Manchester, C=GB	19	2,12%
CN=Thawte DV SSL CA, OU=Domain Validated SSL, O="Thawte, Inc.", C=US	19	1,93%
CN=DFN-Verein PCA Global - G01, OU=DFN-PKI, O=DFN-Verein, C=DE	18	9,34%
SERIALNUMBER=07969287, CN=Go Daddy Secure Certification Authority, OU=http://certificates.godaddy.com/repository, O="GoDaddy.com, Inc.", L=Scottsdale, ST=Arizona, C=US	18	3,18%
CN=GlobalSign Organization Validation CA - G2, O=GlobalSign nv-sa, C=BE	18	1,59%
CN=thawte Extended Validation SSL CA, OU=Terms of use at https://www.thawte.com/cps (c)06, O="thawte, Inc.", C=US	18	1,55%
EMAILADDRESS=tud-ca@hrz.tu-darmstadt.de, CN=TUD CA G01, O=Technische Universitaet Darmstadt, L=Darmstadt, ST=Hessen, C=DE	16	7,80%
CN=TC TrustCenter Class 2 L1 CA XI, OU=TC TrustCenter Class 2 L1 CA, O=TC TrustCenter GmbH, C=DE	16	1,78%
CN=DigiCert High Assurance EV CA-1, OU=www.digicert.com, O=DigiCert Inc, C=US	16	1,49%

Table 6: Sub CAs found for 26 Browser Histories

DN	Relevance by data sets	Relevance by percentage of hosts
CN=COMODO Extended Validation Secure Server CA 2, O=COMODO CA Limited, L=Salford, ST=Greater Manchester, C=GB	16	1,35%
CN=thawte Primary Root CA, OU="(c) 2006 thawte, Inc. - For authorized use only", OU=Certification Services Division, O="thawte, Inc.", C=US	14	12,22%
CN=GeoTrust DV SSL CA, OU=Domain Validated SSL, O=GeoTrust Inc., C=US	14	1,40%
CN=GlobalSign Domain Validation CA - G2, O=GlobalSign nv-sa, C=BE	14	1,30%
CN=GeoTrust Extended Validation SSL CA, OU=See www.geotrust.com/resources/cps (c)06, O=GeoTrust Inc, C=US	13	1,89%
CN=PositiveSSL CA 2, O=COMODO CA Limited, L=Salford, ST=Greater Manchester, C=GB	13	1,65%
CN=Entrust Certification Authority - L1C, OU="(c) 2009 Entrust, Inc.", OU=www.entrust.net/rpa is incorporated by reference, O="Entrust, Inc.", C=US	13	1,46%
CN=GeoTrust Primary Certification Authority, O=GeoTrust Inc., C=US	12	2,01%
CN=USERTrust Legacy Secure Server CA, O=The USERTRUST Network, L=Salt Lake City, ST=UT, C=US	12	1,53%
CN=InCommon Server CA, OU=InCommon, O=Internet2, C=US	12	1,51%
CN=COMODO Extended Validation Secure Server CA, O=COMODO CA Limited, L=Salford, ST=Greater Manchester, C=GB	12	0,93%
CN=Microsoft Internet Authority	11	2,41%
CN=TeleSec ServerPass CA 1, OU=Trust Center Services, O=T-Systems International GmbH, C=DE	11	2,08%
CN=TERENA SSL CA, O=TERENA, C=NL	11	2,05%
CN=Microsoft Secure Server Authority, DC=redmond, DC=corp, DC=microsoft, DC=com	11	1,84%
CN=DPWN Root CA R2 PS, OU=IT Services, O=Deutsche Post World Net, DC=com	10	1,52%
CN=DPWN SSL CA I2 PS, OU=I2 PS, O=Deutsche Post World Net	10	1,52%
CN=PositiveSSL CA, O=Comodo CA Limited, L=Salford, ST=Greater Manchester, C=GB	10	0,96%
CN=UTN-USERFirst-Hardware, OU=http://www.usertrust.com, O=The USERTRUST Network, L=Salt Lake City, ST=UT, C=US	9	2,51%
CN=Network Solutions Certificate Authority, O=Network Solutions L.L.C., C=US	9	0,71%
CN=Thawte SGC CA - G2, O="Thawte, Inc.", C=US	8	1,22%
CN=AlphaSSL CA - G2, O=AlphaSSL	8	0,81%
CN=EssentialSSL CA, O=COMODO CA Limited, L=Salford, ST=Greater Manchester, C=GB	8	0,79%
EMAILADDRESS=ca@zivit.de, CN=ZIVIT CA - G01, OU=Betrieb, O=Zentrum fuer Informationsverarbeitung und Informationstechnik, C=DE	8	0,67%
CN=VeriSign Class 3 Secure Server CA - G2, OU=Terms of use at https://www.verisign.com/rpa (c)09, OU=VeriSign Trust Network, O="VeriSign, Inc.", C=US	8	0,65%
CN=Cybertrust Public SureServer SV CA, O=Cybertrust Inc	8	0,64%
CN=COMODO SSL CA, O=COMODO CA Limited, L=Salford, ST=Greater Manchester, C=GB	7	1,00%
EMAILADDRESS=pki@h-da.de, CN=Hochschule Darmstadt, O=Hochschule Darmstadt, L=Darmstadt, C=DE	7	0,70%
CN=GlobalSign Domain Validation CA, O=GlobalSign nv-sa, OU=Domain Validation CA, C=BE	7	0,61%
CN=WebSpace-Forum Server CA, O="WebSpace-Forum, Thomas Wendt", C=DE	7	0,58%
CN=DigiCert High Assurance EV Root CA, OU=www.digicert.com, O=DigiCert Inc, C=US	6	3,42%
CN=COMODO SSL CA 2, O=COMODO CA Limited, L=Salford, ST=Greater Manchester, C=GB	6	1,06%
CN=QuoVadis Global SSL ICA, OU=www.quovadisglobal.com, O=QuoVadis Limited, C=BM	6	0,83%
CN=Cybertrust Global Root, O="Cybertrust, Inc"	6	0,65%
CN=TC TrustCenter Class 4 Extended Validation CA II, OU=TC TrustCenter Class 4 L1 CA, O=TC TrustCenter GmbH, C=DE	6	0,62%
CN=MSIT Machine Auth CA 2, DC=redmond, DC=corp, DC=microsoft, DC=com	5	1,27%
CN=UTN - DATACorp SGC, OU=http://www.usertrust.com, O=The USERTRUST Network, L=Salt Lake City, ST=UT, C=US	5	1,19%
SERIALNUMBER=10688435, CN=Starfield Secure Certification Authority, OU=http://certificates.starfieldtech.com/repository, O="Starfield Technologies, Inc.", L=Scottsdale, ST=Arizona, C=US	5	0,59%
CN=GlobalSign, O=GlobalSign, OU=GlobalSign Root CA - R2	5	0,48%
CN=DFN-Verein-GS-CA - G02, OU=Geschaeftsstelle, O=DFN-Verein, C=DE	5	0,47%
CN=VeriSign Class 3 Public Primary Certification Authority - G5, OU="(c) 2006 VeriSign, Inc. - For authorized use only", OU=VeriSign Trust Network, O="VeriSign, Inc.", C=US	4	26,86%
CN=Entrust Root Certification Authority, OU="(c) 2006 Entrust, Inc.", OU=www.entrust.net/CPS is incorporated by reference, O="Entrust, Inc.", C=US	4	0,61%
CN=Entrust Certification Authority - L1E, OU="(c) 2009 Entrust, Inc.", OU=www.entrust.net/rpa is incorporated by reference, O="Entrust, Inc.", C=US	4	0,61%
CN=Gandi Standard SSL CA, O=GANDI SAS, C=FR	4	0,30%
CN=Vodafone (Corporate Domain 2009), O=Vodafone Group, C=UK	3	1,68%
CN=Vodafone (Corporate Services 2009), O=Vodafone Group, C=UK	3	1,68%
CN=GlobalSign Organization Validation CA, O=GlobalSign, OU=Organization Validation CA	3	0,56%
CN=GlobalSign Extended Validation CA - G2, O=GlobalSign nv-sa, C=BE	3	0,50%
EMAILADDRESS=ca-btu@tu-cottbus.de, CN=BTU-CA (G01 2008), OU=Rechenzentrum, O=Brandenburgische Technische Universitaet Cottbus, L=Cottbus, ST=Brandenburg, C=DE	3	0,48%
EMAILADDRESS=ca@pki.tu-dortmund.de, CN=TU Dortmund CA - G01, OU=ITMC, O=Technische Universitaet Dortmund, C=DE	3	0,41%
CN=COMODO High Assurance Secure Server CA, O=COMODO CA Limited, L=Salford, ST=Greater Manchester, C=GB	3	0,39%
EMAILADDRESS=pki@hu-berlin.de, CN=HU-CA, O=Humboldt-Universitaet zu Berlin, C=DE	3	0,36%
CN=Zertifizierungsstelle der TUM, O=Technische Universitaet Muenchen, C=DE	3	0,29%
CN=Trusted Secure Certificate Authority, O=Trusted Secure Certificate Authority, C=US	2	1,70%
CN=TC TrustCenter Class 3 L1 CA IX, OU=TC TrustCenter Class 3 L1 CA, O=TC TrustCenter GmbH, C=DE	2	1,29%
EMAILADDRESS=pki-admin@uni-potsdam.de, CN=Universitaet Potsdam CA - G01, O=Universitaet Potsdam, L=Potsdam, C=DE	2	1,07%
CN=EuropeanSSL Server CA, O=EUNETIC GmbH, C=DE	2	1,02%
CN=Register.com CA SSL Services (OV), O=Register.com, C=US	2	0,94%
EMAILADDRESS=rubca@ruhr-uni-bochum.de, CN=Ruhr-Universitaet Bochum CA, O=Ruhr-Universitaet Bochum, L=Bochum, ST=Nordrhein-Westfalen, C=DE	2	0,80%
EMAILADDRESS=mpg-ca@mpg.de, CN=MPG CA, O=Max-Planck-Gesellschaft, C=DE	2	0,63%
CN=DFN-Verein CA Services, OU=DFN-PKI, O=DFN-Verein, C=DE	2	0,62%

Table 7: Sub CAs found for 26 Browser Histories (cont.)

DN	Relevance by data sets	Relevance by percentage of hosts
EMAILADDRESS=gwdg-ca@gwdg.de, CN=Universitaet-Goettingen CA, O=Georg-August-Universitaet Goettingen, L=Goettingen, ST=Niedersachsen, C=DE	2	0,59%
EMAILADDRESS=ca@d-nb.de, CN=DNB-CA, O=Deutsche Nationalbibliothek, L=Frankfurt am Main, C=DE	2	0,56%
CN=GlobalSign Extended Validation CA, O=GlobalSign, OU=Extended Validation CA	2	0,47%
CN=QuoVadis Root CA 2, O=QuoVadis Limited, C=BM	2	0,46%
EMAILADDRESS=ca@kit.edu, CN=KIT-CA, OU=Steinbuch Centre for Computing, O=Karlsruhe Institute of Technology, L=Karlsruhe, ST=Baden-Wuerttemberg, C=DE	2	0,44%
EMAILADDRESS=pki@uni-regensburg.de, CN=Uni Regensburg CA - G01, O=Universitaet Regensburg, L=Regensburg, ST=Bayern, C=DE	2	0,33%
EMAILADDRESS=ca@rrze.uni-erlangen.de, CN=FAU-CA, OU=RRZE, O=Universitaet Erlangen-Nuernberg, L=Erlangen, ST=Bayern, C=DE	2	0,33%
EMAILADDRESS=rum-ca@rz.uni-mannheim.de, CN=RUM-CA-G Zertifizierungsinstanz, OU=Rechenzentrum, O=Universitaet Mannheim, L=Mannheim, ST=Baden-Wuerttemberg, C=DE	2	0,32%
CN=GlobalSign Primary Secure Server CA, OU=Primary Secure Server CA, O=GlobalSign nv-sa, C=BE	2	0,28%
CN=GlobalSign ServerSign CA, OU=ServerSign CA, O=GlobalSign nv-sa, C=BE	2	0,28%
CN=Network Solutions DV Server CA, O=Network Solutions L.L.C., C=US	2	0,21%
EMAILADDRESS=pki@tu-dresden.de, CN=TU Dresden CA - G02, OU=ZIH, O=Technische Universitaet Dresden, C=DE	2	0,21%
CN=Experian Root CA, CN=AIA, CN=Public Key Services, CN=Services, CN=Configuration, DC=experian, DC=local	2	0,19%
CN=Experian Issuing CA 1, CN=AIA, CN=Public Key Services, CN=Services, CN=Configuration, DC=experian, DC=local	2	0,19%
CN=Fraunhofer Root CA 2007, OU=Fraunhofer Corporate PKI, O=Fraunhofer, C=DE	2	0,18%
EMAILADDRESS=pki-ca@bundestag.de, CN=Deutscher Bundestag CA - G01, OU=Deutscher Bundestag, O=Deutscher Bundestag, C=DE	2	0,18%
CN=GeoTrust Global CA, O=GeoTrust Inc., C=US	1	25,00%
CN=VeriSign Class 3 Secure Server CA, OU=Terms of use at https://www.verisign.com/rpa (c)05, OU=VeriSign Trust Network, O="VeriSign, Inc.", C=US	1	3,13%
EMAILADDRESS=ca@uni-wuerzburg.de, CN=UNIWUE-CA - G01, O=Universitaet Wuerzburg, C=DE	1	2,94%
CN=Certum Trusted Network CA, OU=Certum Certification Authority, O=Unizeto Technologies S.A., C=PL	1	1,69%
CN=Certum Extended Validation CA, OU=Certum Certification Authority, O=Unizeto Technologies S.A., C=PL	1	1,69%
CN=SGTRUST CERTIFICATION AUTHORITY, O=SGssl, C=KR	1	1,45%
EMAILADDRESS=ca@rz.uni-saarland.de, CN=CA Universitaet des Saarlandes, O=Universitaet des Saarlandes, L=Saarbruecken, ST=Saarland, C=DE	1	1,26%
EMAILADDRESS=caadmin@uni-bonn.de, CN=Universitaet Bonn CA, OU=Hochschulrechenzentrum, O=Universitaet Bonn, L=Bonn, ST=Nordrhein-Westfalen, C=DE	1	1,11%
CN=adidas Global Intermediate CA 01, O=adidas AG, C=DE	1	1,03%
CN=adidas EMEA Issuing CA 01, O=adidas AG, C=DE	1	1,03%
CN=Universitaet Bremen CA, O=Universitaet Bremen, L=Bremen, ST=Bremen, C=DE	1	1,03%
EMAILADDRESS=jgu-ca@uni-mainz.de, CN=JGU CA - G01, O=Johannes Gutenberg-Universitaet Mainz, L=Mainz, ST=Rheinland-Pfalz, C=DE	1	0,92%
CN=AOL Member CA, O=America Online Inc., L=Dulles, ST=Virginia, C=US	1	0,88%
EMAILADDRESS=pki@hs-mannheim.de, CN=HS Mannheim CA, O=Hochschule Mannheim, C=DE	1	0,78%
EMAILADDRESS=pki@smi.sachsen.de, CN=Sachsen Global CA, OU=Saechsisches Staatsministerium des Innern, O=Freistaat Sachsen, L=Dresden, ST=Sachsen, C=DE	1	0,46%
EMAILADDRESS=caadmin@fernuni-hagen.de, CN=FernUniversitaet in Hagen Global CA, OU=Zentrum fuer Medien und IT, O=FernUniversitaet in Hagen, L=Hagen, ST=Nordrhein-Westfalen, C=DE	1	0,46%
EMAILADDRESS=zertifizierungsstelle@nw.neclab.eu, CN=NECLAB-CA, OU=NEC Laboratories Europe, O=NEC Europe Ltd., C=DE	1	0,46%
EMAILADDRESS=ca@uni-ulm.de, CN=Global-Uni-Ulm-CA, O=Universitaet Ulm, C=DE	1	0,46%
EMAILADDRESS=pki@uni-marburg.de, CN=Uni Marburg CA - G02, OU=Hochschulrechenzentrum, O=Universitaet Marburg, C=DE	1	0,46%
EMAILADDRESS=ca@rwth-aachen.de, CN=RWTH Aachen CA, O=RWTH Aachen, C=DE	1	0,46%
EMAILADDRESS=pki@uni-kiel.de, CN=Uni Kiel CA - G02, OU=Rechenzentrum, O=Universitaet Kiel, L=Kiel, ST=Schleswig-Holstein, C=DE	1	0,46%
CN=Intel External Basic Policy CA, O=Intel Corporation, C=US	1	0,43%
EMAILADDRESS=hrz-ra@uni-bielefeld.de, CN=CA der Universitaet Bielefeld - G02, O=Universitaet Bielefeld, C=DE	1	0,42%
EMAILADDRESS=camaster@uni-koeln.de, CN=UniKoeln CA, O=Universitaet zu Koeln, L=Koeln, C=DE	1	0,42%
EMAILADDRESS=ca@fh-muenster.de, CN=FH Muenster CA - G01, OU=Datenverarbeitungszentrale, O=Fachhochschule Muenster, L=Muenster, ST=Nordrhein-Westfalen, C=DE	1	0,36%
CN=Thawte SGC CA, O=Thawte Consulting (Pty) Ltd., C=ZA	1	0,36%
CN=Oracle SSL CA, OU=Class 3 MPKI Secure Server CA, OU=VeriSign Trust Network, O=Oracle Corporation, C=US	1	0,31%
CN=Network Solutions EV Server CA, O=Network Solutions L.L.C., C=US	1	0,22%
CN=Cybertrust SureServer Standard Validation CA, O=Cybertrust Inc	1	0,22%
CN=Intel External Basic Issuing CA 3A, O=Intel Corporation, C=US	1	0,22%
CN=Intel External Basic Issuing CA 3B, O=Intel Corporation, C=US	1	0,22%
C=BE, O=GlobalSign nv-sa, OU=RootSign Partners CA, CN=GlobalSign RootSign Partners CA	1	0,21%
CN=Deutsche Telekom CA 5, OU=Trust Center Deutsche Telekom, O=T-Systems Enterprise Services GmbH, C=DE	1	0,21%
CN=Fraunhofer Service CA 2007, OU=Fraunhofer Corporate PKI, O=Fraunhofer, C=DE	1	0,21%
EMAILADDRESS=pki@dagstuhl.de, CN=Schloss Dagstuhl - LZI GmbH CA - G01, OU=IT-Abteilung, O=Schloss Dagstuhl - LZI GmbH, L=Wadern, ST=Saarland, C=DE	1	0,21%
EMAILADDRESS=pki@unibw.de, CN=UniBwM CA-G01, O=Universitaet der Bundeswehr Muenchen, L=Muenchen, ST=Bayern, C=DE	1	0,21%
CN=Cybertrust SureServer EV CA, O=Cybertrust Inc	1	0,21%
EMAILADDRESS=fhw-ca@itc.fh-wiesbaden.de, CN=FHW-CA, OU=IT-Center, O=Fachhochschule Wiesbaden, L=Wiesbaden, ST=Hessen, C=DE	1	0,20%

Table 8: Sub CAs found for 26 Browser Histories (cont.)

DN	Relevance by data sets	Relevance by percentage of hosts
EMAILADDRESS=pki@fraunhofer.de, CN=Fraunhofer Service CA - G01, OU=Fraunhofer Corporate PKI, O=Fraunhofer, L=Muenchen, ST=Bayern, C=DE	1	0,20%
EMAILADDRESS=pki@bsb-muenchen.de, CN=BSB-CA, OU=Bayerische Staatsbibliothek, O=Bayerische Staatsbibliothek, L=Muenchen, ST=Bayern, C=DE	1	0,20%
EMAILADDRESS=ca@uni-frankfurt.de, CN=UNI-FFM CA, O=Johann Wolfgang Goethe-Universitaet, L=Frankfurt am Main, ST=Hessen, C=DE	1	0,20%
EMAILADDRESS=pki@tu-bs.de, CN=Technische Universitaet Braunschweig CA, O=Technische Universitaet Braunschweig, L=Braunschweig, ST=Niedersachsen, C=DE	1	0,16%
CN=Fraunhofer User CA 2007, OU=Fraunhofer Corporate PKI, O=Fraunhofer, C=DE	1	0,16%
CN=Dell Inc. Enterprise CA, O=Dell Inc.	1	0,16%
CN=Dell Inc. Enterprise Issuing CA1, O=Dell Inc.	1	0,16%
CN=SecureTrust CA, O=SecureTrust Corporation, C=US	1	0,16%
CN=AusCERT Server CA, OU=Certificate Services, O=AusCERT, C=AU	1	0,16%

Table 9: Sub CAs found for 26 Browser Histories (cont.)