

# Security Economics in the HTTPS Value Chain

Hadi Asghari\*, Michel J.G. van Eeten\*, Axel M. Arnbak<sup>+</sup> & Nico A.N.M. van Eijk<sup>+1</sup>

\* h.asghari@tudelft.nl, m.j.g.vaneeten@tudelft.nl

Delft University of Technology, Faculty of Technology Policy and Management

<sup>+</sup> a.m.arnbak@uva.nl, vaneijk@uva.nl

University van Amsterdam, Faculty of Law, Institute for Information Law

**Abstract.** Even though we increasingly rely on HTTPS to secure Internet communications, several landmark incidents in recent years have illustrated that its security is deeply flawed. We present an extensive multi-disciplinary analysis that examines how the systemic vulnerabilities of the HTTPS authentication model could be addressed. We conceptualize the security issues from the perspective of the HTTPS value chain. We then discuss the breaches at several Certificate Authorities (CAs). Next, we explore the security incentives of CAs via the empirical analysis of the market for SSL certificates, based on the SSL Observatory dataset. This uncovers a surprising pattern: there is no race to the bottom. Rather, we find a highly concentrated market with very large price differences among suppliers and limited price competition. We explain this pattern and explore what it tells us about the security incentives of CAs, including how market leaders seem to benefit from the status quo. In light of these findings, we look at regulatory and technical proposals to address the systemic vulnerabilities in the HTTPS value chain, in particular the EU eSignatures proposal that seeks to strictly regulate HTTPS communications.

**Keywords:** HTTPS, Cybersecurity, Internet Governance, Constitutional Values, E-Commerce, Value Chain Analysis, Security Economics, eSignatures Regulation, SSL, TLS, Digital Certificates, Certificate Authorities.

---

<sup>1</sup> This paper includes a few substantially revised and condensed sections of an earlier paper: A.M. Arnbak & N.A.N.M. van Eijk (2012). *Certificate authority collapse: regulating systemic vulnerabilities in the HTTPS value chain*. Presented at TPRC 2012: the research conference on communication, information and internet policy. Online at: [http://papers.ssrn.com/sol3/papers.cfm?abstract\\_id=2031409](http://papers.ssrn.com/sol3/papers.cfm?abstract_id=2031409).

For inspiration and comments, the authors would like to thank: Bernhard Amann, Ian Brown, Joris van Hoboken, Ralph Holz, Chris Hoofnagle, Kees Keuzenkamp, Samad Khatibi, Stephen Schultze, Ton Slewe, Christopher Soghoian, Sid Stamm, Peter Swire, Marcelo Thompsom, Jan Joris Vereijcken, Frederik Zuiderveen Borgesius and participants of TPRC 2012, a Berkman Center Lecture series Sept. 2012, 29c3, a UC Berkeley TRUST Seminar Jan. 2013 and a HKU Law & Tech Talk, Feb. 2013. The authors are solely responsible for this draft, comments are gratefully received at h.asghari@tudelft.nl and a.m.arnbak@uva.nl

## Table of Contents

<b>1. Introduction .....</b>	<b>3</b>
<b>2. The HTTPS Authentication Model .....</b>	<b>4</b>
2.1 HTTPS Communications .....	4
2.2 The Actor-Based Value Chain of the HTTPS Market.....	5
<b>3. Systemic Vulnerabilities.....</b>	<b>7</b>
3.1 Comparing Known CA Breaches .....	7
3.2 Systemic Vulnerabilities of the HTTPS Authentication Model .....	8
<b>4. Methodology.....</b>	<b>9</b>
4.1 The SSL Observatory Data .....	9
4.2 Market Data.....	10
4.3 Scope and Limitations.....	11
<b>5. The Market for SSL Certificates.....</b>	<b>12</b>
5.1 How Many Organisations Issue Certificates? .....	12
5.2 Market Shares of Certificate Types and Vendors.....	13
5.3 Market Prices of SSL Certificates .....	15
<b>6. Analysis of Market Incentives .....</b>	<b>18</b>
6.1 No Race to the Bottom.....	18
6.2 What is Being Sold in SSL Certificate Markets? .....	19
6.3 Incentives for Security .....	22
<b>7. Improving HTTPS Governance .....</b>	<b>23</b>
7.1 Regulatory Solutions.....	23
A) Security Requirements .....	25
B) Security Breach Notification.....	25
C) Liability.....	26
D) Chain of Trust Transparency.....	27
7.2 Technical Solutions .....	28
7.3 Evaluation .....	30
<b>8. Conclusion.....</b>	<b>31</b>
<b>References.....</b>	<b>32</b>

## 1. Introduction

Hypertext Transfer Protocol Secure (‘HTTPS’) has evolved into the de facto standard for secure web browsing. Through the certificate-based authentication protocol, web services and internet users protect valuable communications and transactions against interception and alteration by cybercriminals, governments and business. In only one decade, it has facilitated trust in a thriving global E-Commerce economy, while every internet user has come to depend on HTTPS for social, political and economic activities on the internet.

The HTTPS authentication model mediates the trust relationship between web site operators,<sup>2</sup> Certificate Authorities (‘CAs’) that issue SSL certificates, web browsers and end-users. For years security experts have sounded the alarm bells about several systemic vulnerabilities of HTTPS communications. A successful attack on HTTPS communication itself – so ignoring attack strategies like SSL stripping or taking over of one of the end points directly – requires a compromised certificate and the ability to modify IP traffic.<sup>3</sup> The 2011 security breach at Dutch CA DigiNotar exposed fundamental weaknesses in the design of the HTTPS authentication model to a global audience. Meanwhile, larger CAs such as Comodo, GlobalSign, Verisign and Trustwave have also suffered substantial breaches, while notably suffering less in their aftermath – an issue to which we return in the paper.

While serving as the de facto standard for secure web browsing, the security of HTTPS is broken in many ways. HTTPS authentication is by and large unregulated,<sup>4</sup> but it has become a top priority in telecommunications policy. European policymakers suggested a review of the EU Electronic Signatures Directive that pioneers a legal framework for HTTPS in June 2012, the European Parliament will vote about the proposal and several amendments in September 2013. Upon adoption, the proposed regulation acquires immediate binding force in the legal systems of 27 Member States. This will impact global HTTPS governance substantially. As we will later show, the CAs that operate within the EU jurisdiction appear to make up around 80% of the HTTPS market. Thus, the proposal is one to watch.

In light of regulatory and technical attempts to resolve the vulnerabilities of HTTPS, it is important to understand the incentives of the actors in the value chain, most notably the Certification Authorities, as their role has been put centre stage by in the recent attacks.

This paper first examines the HTTPS value chain (Section 2) and the systemic vulnerabilities of the technology as demonstrated by the breaches (Section 3). Next, we turn to the CAs and the market for SSL certificates. To better understand the security incentives under which they operate, we analysed the SSL Observatory dataset of certificates for HTTPS traffic. Section 4 outlines the methodology. In

---

<sup>2</sup> This group includes websites (HTTPS) and other services (such as POP/IMAP). For ease of reading, we mostly use ‘HTTPS’ and ‘web sites’ throughout the paper.

<sup>3</sup> Certificate compromise is extensively discussed in Section 3. Manipulating IP traffic may be achieved through a rogue hotspot, poisoning DNS/APR cache, malware or by accessing traffic at ISPs directly. Network providers, DNS servers and governments may have this type of access, while it is a relatively straightforward affair for cybercriminals.

<sup>4</sup> See Section 7.

Section 5 we explore what the data tells us about the number of CAs, the firms that own them, their market shares and the pricing strategies. This uncovers a surprising pattern: a highly concentrated market with very large price differences among suppliers and limited price competition. We then ask how we can explain the pattern that we have uncovered and what this tells us about the security incentives of CAs (Section 6). In light of these findings, we look at regulatory and technical proposals to address the systemic vulnerabilities in the HTTPS value chain, with special attention of the EU eSignatures proposal (Section 7). After this, we wrap up our analysis in a brief conclusion and discussion of the main findings (Section 8).

As far as the authors are aware of, this research project is the first in-depth multi-disciplinary analysis of HTTPS governance. Both descriptive and normative legal research is conducted, as well as value-chain and empirical analysis adopting security economics concepts to research the incentives in the HTTPS ecosystem. As such, this paper extends and deepens an earlier working paper on the same topic by two of its authors, that primarily analysed the legal and value-chain aspects of the HTTPS ecosystem (Arnbak & Van Eijk 2012). In addition to studying the SSL market, the methods of which we outline later, we employed desk research, actor-based value-chain analysis and numerous (non-structured) interviews with crucial stakeholders. We also benefitted greatly from expert comments received at conferences and workshops.

## **2. The HTTPS Authentication Model**

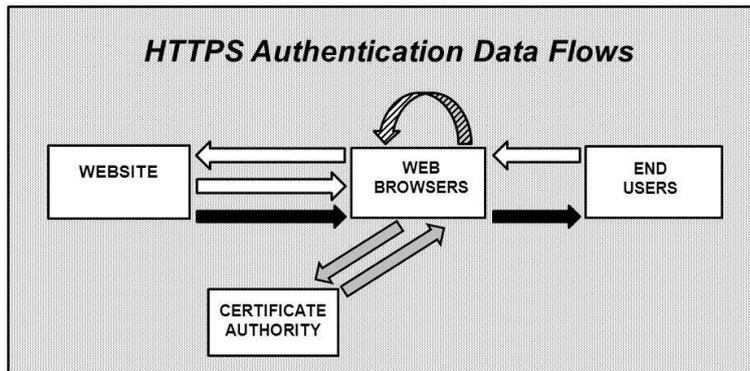
This section describes the HTTPS Authentication Trust Model, the HTTPS market and the actor-based HTTPS Authentication Value Chain, in order to gain insight in the interactions between its key stakeholders. For a more extensive description, we refer to earlier work by two of the authors (Arnbak & Van Eijk, 2012, section 2).

### *2.1 HTTPS Communications*

Essentially, HTTPS is a two-step process: first, a trust relationship (a ‘handshake’) is established between a website operator and an end-user. This is done with the help of an SSL certificate containing basic information for authentication purposes. If the web browser of the end-user trusts the certificate and the issuing CA, this authentication handshake succeeds. Secondly, successful authentication leads to a TLS/SSL encrypted channel between the website and browser, called a ‘tunnel’ (Anderson, 2008, p. 670), and the web browser will alert the user, for instance through depicting a padlock, a green address bar. If the SSL certificate or the issuing CA cannot be trusted, the web browser will show a security warning to the end user.<sup>5</sup> The handshake authentication thus serves as the stepping stone for the confidentiality and integrity that HTTPS seeks to deliver (Roosa & Schultze, 2010). The described data flows are visualised in Figure 1.

---

<sup>5</sup> See Arnbak and Van Eijk (2012) for a more detailed description.



- Data Flows: 4 Phases
1. *White* = HTTPS request and subsequent SSL Certificate offering
  2. *Pattern* = CA Root verification
  3. *Grey* = Certificate signature verification (OSCP)
  4. *Black* = 'Handshake' – authentication

**Figure 1: HTTPS Authentication Data Flows**

If a website operator seeks to provide HTTPS communications, it thus needs to obtain an SSL certificate from a CA. Basically, these SSL certificates are small computer files that might contain information on hostname (website), certificate owner (website), certificate issuer (CA), validity period and public key (Anderson, 2008, p. 672). The amount of information that SSL certificates provide depends on the type of certificate purchased by its owner. Domain Validated (DV) certificates can be acquired at low costs and may require a website operator to reply to an e-mail sent by the CA to a standard e-mail address in the WHOIS database for domain validation (CA/Browser-Forum, 2011). The various types of Organization Validated (OV) and Extended Validation (EV) certificates require more thorough validation by the CA, for example by phone, written letter or face-to-face, verifying both domain and the organization behind it – the end-point (CA/Browser-Forum, 2012). If validation succeeds, CAs sign the OV or EV certificate.

## 2.2 The Actor-Based Value Chain of the HTTPS Market

Since the inception of the HTTPS authentication process with the advent of the Netscape browser in the 1990s, a vibrant market for HTTPS communications has emerged. This market involves roughly four direct stakeholders: i) website operators; ii) certificate authorities; iii) web browsers, and iv) end-users.

*Website operators* decide whether to deploy HTTPS or not. Deploying sends out a message that end-users can entrust the website with valuable information. If embedded content is a part of the revenue model of a website operator, which is the case with many websites, it has strong incentives not to deploy HTTPS (e.g., see: Arstechnica, 2011; Langley, 2012b). HTTPS implementation is hardly state of the art in terms of security. Vratonjic et al. found that ‘only 16% of the websites implementing HTTPS carry out certificate-based authentication properly’ (Vratonjic, Freudiger, Bindschaedler, & Hubaux, 2011). SSL Pulse, a project run by security firm Qualys, finds only 8% use EV certificates and less than 1% support the HTTP Strict Transport Security protocol (SSL-Pulse, 2013).

*Certificate Authorities* exist in three categories: Root CAs, intermediate/subordinate CAs and untrusted CAs. Root CAs are trusted by default by browsers, after they have solicited for such a status with the browsers and complied with the varying browser CA trust policies. Intermediate/subordinate CAs are either directly verified by one Root CA or part of a chain of trust of several intermediate CAs that ultimately ends with one Root CA. Many root CAs own multiple subordinate CAs that may partake in such a chain of trust and in that case enjoy default trust by browsers. Regardless of the level of HTTPS implementation in terms of security by a website owner, certificates of CAs not linked to a trusted Root CA and self-signed (by the owner of a website) certificates evoke the ‘untrusted connection’ security warning when they request an SSL connection to web browsers.

A crucial technical property of the HTTPS Authentication Model is that any CA can sign SSL certificates for any domain name. In other words, anyone can request a SSL certificate for [www.google.com](http://www.google.com) with any CA, even though this CA is not an organization Google itself has contracted to sign its SSL certificate. From the CA perspective, there are some institutional limits to issuing some types of certificates (e.g., validation procedures for EV certificates), but no technical ones. If one obtains this second certificate with a CA that has root status, browsers will react by trusting the second certificate by default. End-users will get the familiar HTTPS notification, without noticing whether their HTTPS communications are mediated by the Google-owned certificate or the second certificate. This ability to sign for any domain name has profound implications for the security of the HTTPS ecosystem, commonly referred to as the ‘weakest link’ problem – if one CA suffers a breach, the entire ecosystem is under attack (ENISA, 2011; Roosa & Schultze, 2010). On the other hand, it has spurred a flourishing CA industry over the last decade. We will return to this technical property throughout this paper. (Roosa & Schultze, 2010; Soghoian & Stamm, 2012)

*Web browser vendors* serve as the interface between website owners, CAs and the end-user. In determining whether CAs should be granted root status, browsers have developed different trust policies. This leads to a different number of root CAs per browser. We return to this in Section 5.

In the case of (or if there is reason to suspect) certificate or even CA compromise, swift trust revocation is essential to minimise the associated risk. For certificates, all major browsers employ Online Certificate Status Protocol (OCSP) responders. These are operated by CAs and let browsers check whether trust in a certain certificate has been revoked. For CA revocation, browsers need to alter aforementioned root CA lists and patch the browser software, which end-users subsequently need to update to take effect. An important drawback of OCSP effectiveness, is that its use by CAs is not mandatory and often overruled in order to maintain connectivity between a web service and users (Langley, 2011).

*End-users* have an interest in seeking HTTPS communications with websites, as it is their valuable information that is on the line. They depend to a large degree on security decisions made by the aforementioned stakeholders. Only a very small margin of technically savvy users might pursue an (indirect) relationship with CAs through browser preferences, for example by blocking all certificates provided by a certain CA (ENISA, 2011; Vratonjic et al., 2011).

There seems to be wide consensus that the average end-user cannot reasonably be expected to exert control over the HTTPS ecosystem (Bakos, Marotta-Wurgler, & Trossen, 2009; ENISA, 2011). The next section describes the systemic vulnerabilities of the HTTPS ecosystem in theory and in practise, based on above conceptualisation of the HTTPS value chain.

### 3. Systemic Vulnerabilities

Earlier work analyses the well-documented landmark breach at Dutch CA DigiNotar and other breaches at CAs extensively (Arnbak & Van Eijk, 2012 para 3.1). Here, we overview the known security breaches at CAs more generally in order to present several systemic vulnerabilities of the HTTPS ecosystem.

#### 3.1 Comparing Known CA Breaches

On Friday 2 September 2011, a nocturnal press conferences of the Dutch Minister of Internal Affairs marked the beginning of the DigiNotar affair. It was triggered by unauthorized access, reportedly by a hacker sympathizing with the government of Iran in mid July 2011, to the root CA capacity of DigiNotar. When the breach became public three months later, it emerged that in this long period of obscurity 531 false certificates had been created for widely used and highly sensitive domain names such as \*.google.com, \*.facebook.com, update.windows.com and \*.cia.gov (Fox-IT, 2011, p. 10). DigiNotar, a niche player in the global market with a strong presence in the niche for Dutch eGovernment services, had root status with all major browser vendors, so these corrupt SSL certificates would have been trusted by default. The forensic report illuminated that thirty critical updates had not been performed, logging was insufficient and no anti-virus protection was in place at the time of the intrusion (Fox-IT, 2012, pp. 62-63). Interestingly, the CA complied to existing regulations and had successfully passed several ESTI standardised periodic auditing procedures by renowned accounting firms for the issuance of EV certificates and Qualified signatures (ENISA, 2011). In Section 7 of the paper, we will return to these observations. The damage was probably enormous, but cannot be determined with certainty due to the unreliability of the log files. ENISA speaks of breached communications of ‘millions of citizens’, particularly connected to the \*.google.com certificate, and notes that some experts believe that the lives of Iranian activists have been put at risk (ENISA, 2011). Upon publication of the breach, the trust in the entire range of DigiNotar activities was revoked.

Of the less documented CA incidents, the range of breaches at market leading CA Comodo has probably received the most attention (InfoSecurity, 2011). The best documented breach at Comodo was the compromise of its ‘UTN-USERFirst-Hardware’ certificate. According to data analysis from its SSL observatory, EFF calculated that ‘85,440 public HTTPS certificates were signed directly by UTN-USERFirst-Hardware. Indirectly, the certificate had delegated authority to a further 50 intermediate CAs, collectively responsible for another 120,000 domains’ (EFF, 2011a).

We know that Verisign, another dominant CA, was hacked in 2010. The breach was only discovered by news agency Reuters (2012) in February 2012, after Security and Exchange Commission regulations mandate companies to notify investors of intrusions since October 2011. According to the Reuters reports, a former CTO claimed he had not learned of the intrusion until contacted by Reuters and said Verisign ‘probably can’t draw an accurate assessment’ of the damage, ‘given the time elapsed since the attack and the vague language in the SEC filing’ (Reuters, 2012).

The breach at CA GlobalSign, yet another market leading CA, is another example of poor security practises, as software running on a public-facing webserver was not updated.

Another instance involved CA Trustwave. It became public that it had used its root CA status to enable third parties to issue SSL server certificates for employee

monitoring purposes. Trustwave subsequently claimed that this is common practice among other root CAs (Computer-World, 2012). This illustrates the compelled-CA attack of Soghoian & Stamm in real life: CAs are in a unique position to enable surveillance of end-users (Soghoian & Stamm, 2010).

This section has not covered all publicly known CA breaches,<sup>6</sup> but several patterns emerge. Regardless of scale, CAs get breached. They are reluctant to inform both relevant authorities, customers and the general public (end-users) about these breaches. Security practises at DigiNotar and (to a lesser extent) GlobalSign proved to be below a ‘state of the art’ or ‘general industry practise’ level, while this cannot be established for the breach at Verisign and the multiple Comodo breaches. Given the current reluctance to report breaches, there is no way of knowing that larger CAs and the certificates they issue are more secure than small CAs. The question emerges why the trust in the entire CA practise of DigiNotar was revoked by the web browser vendors, while larger CAs dodged the bullet. Clearly, the root capacity of DigiNotar was severely breached, but pragmatic considerations may have played a larger role. ENISA argued in the aftermath of the DigiNotar breach that if a larger CA would suffer a similar security breach, trust revocation by browser vendors in its certificates would seriously impact web communications on a global scale: ‘it can even be argued that CAs of this size are too large to fail’ (ENISA, 2011). The decision to punish a small CA for bad security practises probably is considered to be less problematic than removing market leaders such as Comodo, GlobalSign or Verisign from the trusted root list.

### 3.2 Systemic Vulnerabilities of the HTTPS Authentication Model

‘Systemic vulnerabilities’ point towards those vulnerabilities that are inherent to the HTTPS ecosystem as opposed to incidental vulnerabilities that have occurred at a particular stakeholder during an isolated incident. For instance, the fact that DigiNotar employed one extremely weak password to secure all of its systems is not a systemic vulnerability, but the fact that the result of poor security practises at one marginal CA may undermine the security of the entire HTTPS ecosystem is.

The fact that any CA can vouch for any domain name, then, is probably the most important and widely recognised vulnerability. This characteristic makes all CAs in over fifty jurisdictions a ‘weakest link’ for potentially all HTTPS communications (see Section 5 for the numbers). As ENISA (2011) observes: ‘The security of HTTPS equates to the security of the weakest CA.’

The scenario’s for failure are manifold: any CA could facilitate or be a malicious actor engaging in cybercrime, or be a company monitoring its employees, or could be compelled by a state actor to enable mass surveillance of internet users (Soghoian & Stamm, 2012), or one of its administrators could simply have a bad day – forgetting updates, writing poor code or in his own right be coerced to cooperate in malicious activities.

The recurring information asymmetries are another striking systemic vulnerability. Organisations – including CAs and website operators – have strong

---

<sup>6</sup> Roosa and Schultze (2010, p. 5) report on other breaches. Furthermore, KPN/Getronics, StartSSL TurkTRUST and several other CAs have been breached in recent years.

incentives to conceal poor security practises and breaches. Reporting has a strong public interest dimension: a breach risks not only the untrustworthiness of the entire ecosystem, but also renders trust of end-users unjustified: end-users may disclose highly sensitive information based on erroneous assumptions of security.

From the viewpoint of web browser vendors, the interests of providing connectivity versus assuring trustworthiness may conflict. This is demonstrated in the overruling of OCSP responses and in browser management of root status. But browsers face the hard choice of rendering a large part of the HTTPS encrypted web inaccessible to its end-users upon a breach. If ENISA notes that major CAs are too big to fail (ENISA, 2011), the weakest link phenomenon is even more worrying. In other words, CA scale is a risk vector when it comes to security: a breach may compromise more communications, but revocation is more complicated.

The current regulatory regime and auditing obligations have proved quite ineffective. The qualified certificate practises of DigiNotar were strictly regulated and passed the periodic audits based upon regulation and internationally recognised industry standards. The perceived security that the current auditing schemes should deliver is another systemic vulnerability of HTTPS (Roosa & Schultze, 2010).

In all this, damages associated with security breaches are pushed downstream by the stakeholders towards end-users, even though end-users cannot reasonably be held accountable to evaluate security practises in the current HTTPS authentication model. A common practise for CAs is to disclaim liability for losses suffered as a cause of reliance in certificates (Roosa & Schultze, 2010; Vratonjic et al., 2011).

These conceptual considerations provide guidance into our empirical research of the HTTPS certificate market. We first outline the methodology. Then we describe the dominant properties of this market by looking at the number of CAs, the firms that own them, their market shares and the pricing strategies. We then ask how we can explain the market properties and what this tells us about the security incentives of CAs. In light of these findings, we reflect on the governance of the HTTPS value chain.

## 4. Methodology

The empirical part of this study builds primarily upon two datasets: the Electronic Frontier Foundation's SSL Observatory data and a custom set of market prices for the different offerings of certificate authorities.

### 4.1 *The SSL Observatory Data*

The SSL Observatory is a project that investigates the certificates used to secure all of the sites encrypted with HTTPS on the Web.<sup>7</sup> The Observatory scanned the full IPv4 address space for publicly visible webservers running HTTPS, over a course of several weeks. All certificates returned by these servers were saved along with some metadata. This amounts to 4-6 million certificates, out of which only a portion is considered as valid by browsers (having a valid certificate chain, not being expired,

---

<sup>7</sup> See: <https://www.eff.org/observatory>

etc.). After filtering out the invalid, the Observatory dataset provides approximately 1.5 million SSL certificates.

The dataset is the most comprehensive of its kind, but has one major drawback: its age. The version we accessed was the final public release of December 2010. For this reason, we explored other available datasets, to no avail. They were either from the same period, the collection methodology made them less comprehensive or the data was not readily accessible to other researchers.<sup>8</sup>

From the SSL Observatory data, we generated a list of certificate authorities using several standard queries (e.g., looking at *basicConstraints* or the *issuer* field). This results in approximately 1,100 CAs. The self-signed CAs on this list were matched using fingerprints to the Microsoft Root Certificate Program list (Microsoft, 2012) and to the Mozilla source file that has the roots in it<sup>9</sup>. The matching allows us map the CAs to the owning organisation information kept by the root stores.<sup>10</sup>

Next, we identified certificate types. EV certificates can be determined via the existence of certain policy object identifiers (OIDs) in the *Certificate Policies* field. These object identifiers (OIDs) are extracted from the Chromium browser source file. Distinguishing between DV & OV certificates is tricky and can turn into art - we adapted the heuristic algorithm suggested by (Hurst, 2012). The gist of the algorithm is to see whether the certificate *subject* field contains data that can identify an organisation, using city and state fields as extra hints. The determined types were crosschecked by looking at the percentage of DV/OV/EV certificates each CA had issued, as a majority of owners issue only one type of certificate per CA.

## 4.2 Market Data

A dataset was built that maps each CA to its market name, product offerings and prices. The starting point was generating a list of all CAs that had issued more than 500 certificates. The majority of these are subordinate CAs, for which we used web search to determine the owners. In most cases, this was straightforward. In some instances, we had to make an educated guess based on the results of web searches on CA and owner names.

Current product and price information is taken from the owner's website.<sup>11</sup> A number of these vendors do not provide prices, and some only on request. To illustrate: the website of *Secure Business Services*, which has 3000 certificates in the

---

<sup>8</sup> This included the SSL Landscape project at TU München (<https://pki.net.in.tum.de/node/8>), the Berkeley ICSI Certificate Notary (<http://notary.icsi.berkeley.edu/>) and a few others. The former unfortunately also dates back to March 2011; the latter is unfortunately not readily accessible due to privacy considerations, except in a highly aggregated graph which we have actually used for triangulation purposes. Other projects have similar limitations.

<sup>9</sup> See: <http://www.mozilla.org/projects/security/certs/included/>

<sup>10</sup> The matching was not perfect, with several fingerprints not being found – typically for new and retired roots; in some of these cases, we made inferences using the subject field.

<sup>11</sup> A number of other smaller known brands that have issued less than 500 certificates were looked into as well to generate some insight in the long tail of CAs.

dataset, gives neither prices nor an option to request a quote.<sup>12</sup> The contact telephone provided on the site does not work either. We have skipped such vendors, accounting together for 2% of the market.

To make prices comparable, we standardized them to the extent possible as follows: (i) US dollar prices are used if available; if not, we convert using current rates; (ii) VAT is added when explicitly excluded; (iii) we only include prices of certificates with a one year validity period; (iv) all discounts including multi-year, bulk, as well as various bundled offerings, are ignored; (v) reseller pricing is ignored. Most SSL vendors have partner programs and their resellers often set lower prices – in one case, down to a fourth (€49 versus €12). It is not possible to tell from the certificates which ones have been bought via a reseller, so we cannot factor this in.

We considered wildcard and UCC certificates separately, and given the three DV/OV/EV types, this yields in total eight price categories.<sup>13</sup> Different brands of a vendor are also considered separately when the certificates can be technically distinguished, e.g. Symantec/Verisign, Symantec/Thawte, Symantec/GeoTrust and Symantec/RapidSSL. In the end, 98% of the SSL certificates in the dataset are mapped to a brand and prices are available for 96% of the certificates.

### 4.3 *Scope and Limitations*

Several limitations need to be taken into account. The first is a matter of scope: the SSL Observatory collects only certificates of publicly visible web servers. This fits with our analysis, which focuses on HTTPS. We should point out, however, that the data does not capture the digital certificates used by back-end systems, personal email certificates, and other use cases. One point of comparison comes from Verisign's annual report. In December 2009, they had a 1.2 million installed base of SSL certificates ('business authentication services'), and an unspecified number of 'user authentication services' certificates (Verisign, 2010, p. 49). In the Observatory data, Verisign has approximately 663 thousand certificates, which is a significant proportion of all server certificates.

A second limitation is the time mismatch between market shares and prices. Market shares, having been calculated from the Observatory data, are from December 2010, while we have gathered the price data in February 2013. This can be overcome when updated Observatory data is released.<sup>14</sup>

A third limitation stems from price accuracy. We already mentioned resellers offering different and lower prices; prices can also be lower due to discounts; they can also be several times higher when considering that certificate typically needs to be installed on multiple servers, and several brands put limits on this. We have aimed to standardize the prices as much as possible.

We attempted to triangulate the SSL Observatory data with aggregated statistics available from two other sources: the Certificate Notary and the NetCraft

---

<sup>12</sup> <http://www.securebusinessservices.com>

<sup>13</sup> These are: single-domain DV, multi-domain DV, wildcard DV, single-domain OV, multi-domain OV, wildcard OV, single-domain EV and finally multi-domain EV certificates. EV certificates do not support wildcards.

<sup>14</sup> Retrieving consistent pricing data from the past was generally not feasible, but we have a brief look at historical prices by the end of Section 5.

SSL Survey. The patterns were similar – neither conclusive fits, nor large discrepancies could be noted.<sup>15</sup>

## 5. The Market for SSL Certificates

### 5.1 How Many Organisations Issue Certificates?

The question of how many organisations can issue certificates seems straightforward, but it has been the source of speculation and controversy. The X.509 standard specifies a structure composed of root certificate authorities, intermediate certificate authorities and end entities (IETF, 2005). Root CAs are trusted directly by the end applications; they typically certify intermediate CAs (also known as subordinate CAs), who in turn certify other intermediates or issue certificates for end entities. Browser and OS vendors have their own policies for determining which CAs to include in their root stores; such is the case with as those of Microsoft, Mozilla and Apple (see: Apple, 2013; Microsoft, 2009; Mozilla, 2013). Software can also use roots provided by the underlying operating system, like Google Chrome does.

Looking at the certificates marked as valid in the SSL Observatory dataset, we see approximately 1,100 issuing CAs. A company or organisation can own and operate multiple root and intermediate CAs – for reasons such as operational procedures, redundancy, security, branding, or as a consequence of acquisitions.<sup>16</sup> We map root CAs to their owning entities by looking at the details provided on vendor root stores. Root stores have regular audit requirements, keeping their lists up to date. Table 1 shows the results of matching the Microsoft and Mozilla root stores with the EFF dataset. Microsoft supports more root CAs than Mozilla, especially among governmental owners. Root CAs are located in 43 countries. Intermediate CAs add another 11 countries to this list.<sup>17</sup> Determining the ownership of the intermediate CAs is more complicated, as a base list to compare against does not exist. A portion of them are owned by the same organisations owning the roots; others are separate entities. We have mapped this manually for all intermediate CAs that have at least 500 certificates in the dataset (93 CAs are above the threshold) to their respective owners and, in the case of firms with multiple brands, to each brand.

---

<sup>15</sup> The Observatory data has more certificates, which is to be expected as it based on a full scan of the IPv4 space. The Notary data is based on certificates in active use in US networks monitored by the Notary project. The list of CAs and roots match to a large degree, but not perfectly, as they are from different points in time. The NetCraft SSL Survey, a recognized industry report, is available for a fee. We have compared their summary graph with our own data (see: NetCraft, 2012). The results are consistent.

<sup>16</sup> One interesting case is the DFN hierarchy, used by the academic network in Germany. In the DFN-PKI scheme, each institution has its own signing CA, resulting in more than 250 subordinate CAs. The private key for all of them is kept centrally at the DFN, and not given to the institutions. In practical terms, all these CAs fall under one organisation.

<sup>17</sup> Based on country data provided in the CA's certificate subject; looking at the where the owners are actually located results in slightly different counts.

Searching the web, we connected these CA names to their owners. CAs with similar names were then also identified as belonging to the same owners, bringing the total mapped to 134. Finally, we separately tagged the 261 CA names from the DFN-Cert hierarchy.

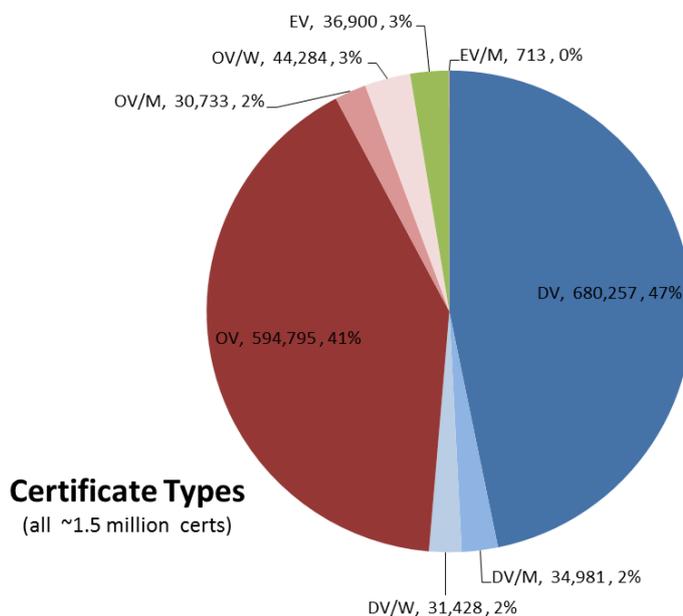
**Table 1**

Root store	Root owners (organizations)	Percentage governmental	Root CAs	CAs under hierarchy	Hierarchy level
Microsoft	116 (89 in dataset)	36%	333 (173 in dataset)	1096 in dataset	Median 1, Max 4
Mozilla	61 (56 in dataset)	20%	158 (130 in dataset)	907 in dataset	Median 1, Max 4

In summary, although it is very difficult to come up with an exact number for the total number of organisations issuing certificates, we can provide a reasonable estimate. There are already over one hundred owners for the root CAs alone (Table 1). Mapping intermediate CAs with the aforementioned criteria adds only 24 additional owners to the count, bringing the total to 140. Our impression is that mapping the whole population of CAs would bring the total number of owners to somewhere between 200 and 300.

## 5.2 Market Shares of Certificate Types and Vendors

Figure 2 shows the distribution of the different certificate types. DV and OV each hold around half of the total. The figure also shows the number of domains each certificate was issued for: a single domain, multiple domains or a wildcard.



**Figure 2**

Table 2 shows the percentage of top sites, based on the Alexa ranking, that are using SSL certificates in general, and EV in particular. Although the higher ranking sites have higher HTTPS adoption, they do not differ significantly in terms of using EV over OV or DV, despite the fact that browsers provide more explicit trust signals with EV.

**Table 2**

<b>Top domains (Alexa ranking)</b>	<b>Percentage that have an SSL certificate (Dec 2010)</b>	<b>Percentage of which is EV (Dec 2010)</b>
<b>Top 1000</b>	35.3%	6.8%
<b>Top 10k</b>	25.2%	11.3%
<b>Top 100k</b>	15.2%	10.7%
<b>Top 500k</b>	5.0%	8.5%

In the next step we set out to map the market shares of the certificate authority owners and brands. Figure 3 shows the total market share of each CA, for the combined OV/DV/EV submarkets. Around 98% of all the certificates in the SSL Observatory dataset are accounted for. The results indicate a highly concentrated market: three vendors – Symantec, GoDaddy and Comodo – hold more than three quarters of the market share.

To test whether higher ranking websites chose similar vendors or not, Figure 4 provides the distribution of certificates used by the top-thousand and top-hundred-thousand domains. Although individual market shares differ, the concentration and overall pattern is the same as that for the total set of domains.<sup>18</sup> The largest difference is the FIRM-OWN-CA subgroup in the top-thousand domains, as companies such as Google, Facebook and Microsoft issue certificates from their own intermediate CAs.

To assess the degree of market concentration, we calculated the Herfindahl-Hirschman Index (HHI). The results are shown in Table 3. The scores are above 2,500 which indicates a highly concentrated market (DOJ & FTC, 2010).

**Table 3**

	<b># Firms in Set</b>	<b>HHI-4</b>
<b>All certificate types</b>	23+	2729
<b>DV market</b>	11+	3739
<b>OV market</b>	21+	2862
<b>EV market</b>	12+	5343

---

<sup>18</sup> The Spearman rank coefficient shows a high similarity between the sets: rho=0.75/sig=0.00 between total set of domains & top1k; and rho=0.94/sig=0.00 between total set of domains & top100k.



As can be seen in Figures 5, 6 and 7, high-priced brands enjoy large market shares, in the case of EV even the bulk of the market share.

**Table 4**

Certificate type	Min price	Max price	Average (std. dev.)
DV	\$0	\$249	\$81 (74)
OV	\$38	\$1172	\$258 (244)
EV	\$100	\$1520	\$622 (395)

The situation is comparable when we look at the long tail of market shares. Here, we often encounter smaller CAs with specific geographic markets. We took a closer look at five brands (Table 5). Instead of competing with the market leaders on price, they seem to be focused on reaping similar profits from their local customers, i.e., adopting a niche market strategy.

**Table 5**

Company (Country)	Approx. certificate price (OV)	Global Market Share
Etisalat (U.A.E.)	\$326	0.0%
Netlock (Hungary)	\$102	0.1%
RBC (Russia)	\$107	0.0%
TürkTrust (Turkey)	\$196	0.0%
CERTUM (Poland)	\$155	0.1%
Comodo	\$131	12%

In sum: the market shares shows limited signs of price competition. How about price pressure over time? In a market with high fixed cost and low marginal cost, a consistent price decrease would be the predicted result. We could not systematically establish whether prices have gone down, because we weren't able to accurately reconstruct price levels in earlier years. We checked the pages for twelve of the bigger brands in 2009, using the Internet Archive<sup>19</sup>. The comparison shows no definitive trend of increase or decrease: four of the brands maintained the same price; three brands increased prices for DV/OV certificates, while decreasing the price of EV certificates; of the remaining, four dropped prices and one increased them. This mixed result does not signal strong price competition.

---

<sup>19</sup> <http://archive.org>

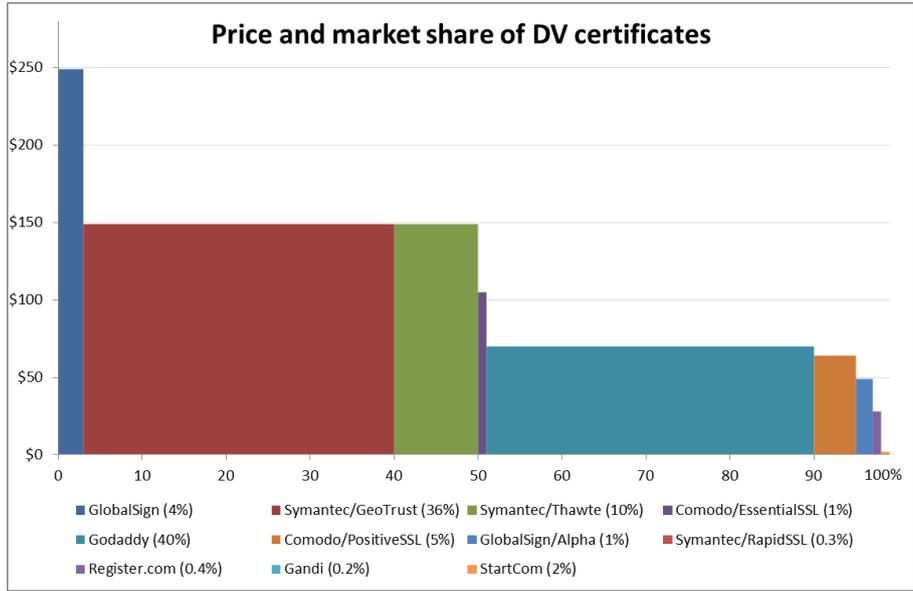


Figure 5

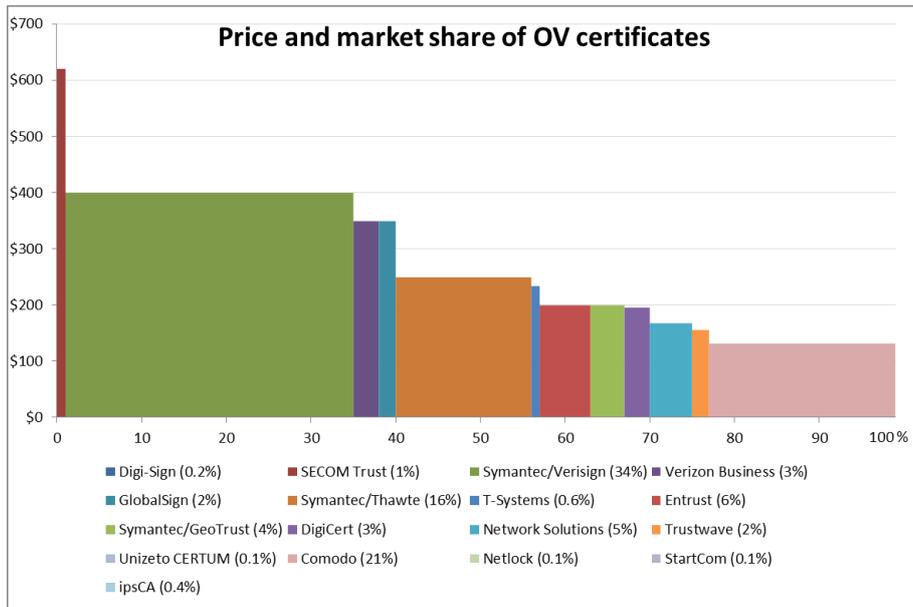


Figure 6

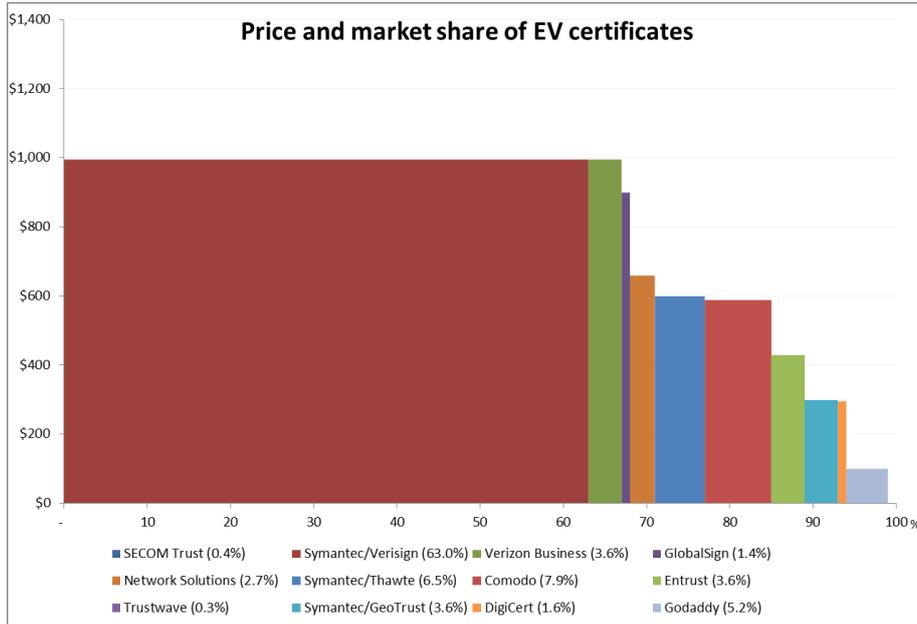


Figure 7

## 6. Analysis of Market Incentives

The empirical data has revealed a pattern that requires an explanation: notwithstanding the fact that certificates of one type are technically perfect substitutes, each submarkets is highly concentrated, with very large price differences among suppliers and limited price competition. How can this be explained? In one sentence: because this market is not driven by the sale of the certificates themselves, but by the services and reputations signals bundled with the certificates.

### 6.1 No Race to the Bottom

Before we analyse the empirical pattern in more detail, we first want to highlight the fact that it falsifies a much-repeated claim about CAs, namely that they compete in a race to the bottom. Various researchers and industry observers have claimed that such a race exists in this market and some associate this with the poor security practices at DigiNotar and other compromised CAs (Kelkman, 2013; Mills, 2011; Roosa & Schultze, 2010, p. 6; Vratonjic et al., 2011, pp. 30-32).

At first glance, such a race is indeed what one would expect. The certificates of one type are perfect substitutes. This would suggest that the market is completely commoditized. Also, buyers can't meaningfully distinguish secure from less secure offerings. There are strong information asymmetries between the CAs and the buyers. More importantly, any CA can issue a certificate for any domain, which means that the security of SSL to prevent man-in-the-middle is determined by the weakest link in the market – i.e., the most insecure CA. In other words, buying from a supposedly

more secure CA cannot protect the site owner against the threat of an attacker fraudulently signing his domain with a certificate from a compromised CA.

The combination of these two conditions – a completely commoditized market in which buyers have no way of telling which offering is more secure – should have produced a ‘race to the bottom’: a market dominated by fierce competition pushing prices towards marginal cost, with perverse incentives for security (Anderson, 2008, p. 223; Shapiro & Varian, 1998, pp. 19-52).

The data, however, clearly suggests otherwise. We see market concentration, but not because dominant players leverage their increasing returns to scale to compete on price. There seems to be very little price pressure at work, in fact, especially in the market for EV certificates. The most expensive suppliers have large market shares, leaving only marginal shares for the cheapest ones. Even RapidSSL, the comparatively cheap ‘fighter brand’ of market leader Symantec, captures less than a 0.5% share of the DV certificate market.

One explanation for the lack of price competition could be the existence of entry barriers. It is unclear, however, what these barriers would entail exactly. It takes a substantial investment to get a root into the root stores of the leading browser and OS vendors. But there is a large group of CAs that are already present in those stores and that are cheaper than the dominant players. It does not seem to be a successful strategy. At the tail end of the market, where certificates are sometimes 5 to 10 times cheaper than those from the market leaders, we see that low prices have, by and large, only attracted minor market shares. There is one notable exception: GoDaddy, the hosting provider. Its cheaper DV certificates have captured 40% of the market, perhaps aided by the fact that they can bundle them with its huge hosting business. This stands in stark contrast to the market for EV certificates. Here GoDaddy’s price is among the lowest prices in the market – and 10 times cheaper than the market leaders – and it has managed to capture only around 5% of the market. So the presence of entry barriers cannot really explain this pattern.

Rather than a market around a commoditized product competing on price and locked into a race to the bottom, the empirical pattern suggests that this is in fact a market with highly differentiated products that can be sold at dramatically different prices. In one sense, it is good news that the market is not driven by a race to the bottom, given the perverse security incentives associated with such a race. It does beg the question of how sellers have managed to differentiate their products and what this tells us about the security incentives that operate in the market.

## 6.2 *What is Being Sold in SSL Certificate Markets?*

If the certificates themselves are perfect substitutes, then how are suppliers differentiating their products to allow for the large price differences? In short: by bundling them with additional services. This becomes visible when we look at the marketing tactics used in the retail channels for SSL certificates.

CAs go out of their way to suggest that their offerings are different from those of its competitors. This has resulted in a rather baroque set of selling points on which they try to differentiate their products. We won’t attempt to discuss them all, but rather focus on the main ones and then conceptually summarize the main differentiation strategies.

Some selling points are straightforward, such as the percentage of all internet users whose browsers will accept the certificate. There is no real differentiation here, however. All brands included in our overview (Figures 6-8) are included in the dominant trust stores and therefore have a near-complete browser coverage measured

in terms of internet users. Another selling point is the speed with which the certificate will be issued. Faster is seen as better. Most CAs promise to hand over DV certificates in minutes and EV certificates in a matter of days or even hours. We didn't find meaningful differences among the brands.

CAs and resellers also stress the security 'features' of their certificates, such as its key length and the encryption level it supports, even though these features are virtually the same across all CAs and the security problems with SSL have had nothing to do with breaking the encryption. As with browser coverage and speed, these features do not really differentiate the products on offer.

Another security-related tactic is leveraging the reputation of a CA brand. The market leaders all offer the buyers a seal to put on their site, indicating to site visitors that the site is secured by that specific brand. There are significant differences among brand reputations, if only in terms of name recognition, so this feature can account for a part of the price differentiation.<sup>20</sup>

Some CAs also bundle security services with the certificate, such as monitoring whether the buyer's domain is hosting malware or phishing sites. Another bundled tool supposedly scans whether the buyer's site handles credit card data in compliance with PCI standards.

Arguably the most incomprehensible differentiating tactic is the 'warranty' on which some CAs compete. The warranty is not for the buyer, but for the end users who suffer fraud when using a site that was secured by an SSL certificate from the CA that should not have been issued in the first place. It is a rather mindboggling exercise trying to understand in real world terms how this warranty would work and how it would benefit the buyer of the certificate. To illustrate: in the case of DigiNotar such a warranty seems to only come into play if DigiNotar had been the official supplier of certificates for Google and the Iranian victims would have suffered some sort of fraud. Unsurprisingly, as far as we can tell there are no cases where a CA actually paid damages to end users under this warranty. Still, the idea seems to be that it would function as a trust signal to third parties – i.e., the warranty provides the visitors of the buyer's site with extra assurance that it is safe to conduct business with the buyer, because they can hold the CA liable if it turns out that it is not really the site of the buyer. Of course, in reality end users have never heard of these warranties, the information about the warranty amounts is not available to them, let alone that they know which CA offers higher warranty amounts. In fact, end users rarely know, or care about, what CA actually issued the certificate in the first place. This has not stopped the CAs from competing the warranty amounts, where higher amounts supposedly demonstrate more secure or trustworthy certificates.

In addition to these selling points marketed in retail channels, there are also strategies that specifically target enterprise customers. These are much less visible to outside observers and we currently only have some anecdotal evidence on this from conversations with enterprise buyers. We encountered three additional differentiating features, each of which help explain the price differentiation and the dominance of the current market leaders.

---

<sup>20</sup> That said, we also found rather forced attempts to differentiate. CAs stress the difference between static and dynamic seals – which says nothing more than whether the seal is a static picture or an animated one with a bit of dynamic information, such as the current date.

First, and perhaps foremost, is the provision of enterprise-level certificate management services. One IT security manager of a multi-national firm explained how valuable support services are for the management, billing and reporting related to certificates. They employ certificates in thousands of domains for tens of different legal entities across many countries. Each entity faces different requirements in terms of billing languages, methods and periods, tax rules, reporting processes and more. What set the market leaders apart is the extensive and integrated back-end support for meeting these requirements. Smaller suppliers offered no such services.

Second, they bought from the market leaders because the reputations of the main brands functioned as a sort of a liability shield towards their corporate leadership, shareholders and regulators, in case something would go wrong. It is a variant on the old adage: 'Nobody ever got fired for buying IBM'.

The third benefit from buying from market leaders is less explicit and a bit counter-intuitive. Enterprise buyers understand that security in this market is a weakest-link problem. They also understand that three of the four market leaders got hacked in recent years and are therefore not immune to the threat that brought down DigiNotar. This all suggests that there might not be any real security benefits from buying from them. The attacks have also demonstrated something else, however: these CAs are less likely to be thrown out of the root stores.

To put it differently: the market leaders are, in a sense, too big to fail. Browser and OS vendors will be extremely reluctant to remove them from the root store. This can actually be a benefit to the CA's customers, because it provides them with better business continuity. The collapse of DigiNotar has underlined the value of this advantage. For the government and business customers of DigiNotar, the breach was in essence a crisis of availability (continuity), not of confidentiality or integrity. Tens of thousands of certificates had to be found and replaced in about a week. During that time, government representatives publicly acknowledged that they faced the threat of a large-scale 'blackout' of governmental services (NRC, 2011). That scenario is unlikely for the customers of the too-big-too-fail market leaders. Of course, buyers can still switch away from those suppliers if they choose to, but they can do so under less time pressure.

So, to sum up, what are buyers actually buying in this market and how can this explain the pattern of high concentration and of high price differentiation? The certificates themselves are perfectly substitutable, but CAs differentiate via:

- bundled security services, such as scans of the buyer's site for malware or PCI compliance;
- enterprise certificate management services, such as support for the management, billing, and reporting around large numbers of certificates;
- brand reputation as a liability shield against the buyer's organizational superiors, shareholders, regulators or others who may hold the buyer accountable in the face of security issues;
- trust or security signals aimed at third parties, most notably end users, such as brand reputation, site seals, warranty amounts and, in a sense, the high price of a certificate itself signals security;
- higher continuity in case of security failures at the CA, because of the unlikelihood of its root status being revoked by browser and OS vendors.

The technical artefact of a certificate is a perfectly substitutable information good, but in light of these features, one could argue that what CAs sell in practice is a subscription-based service. Subscription services are less substitutable and can thus be more effectively differentiated in the market.

The fact that some of the ‘security’ features of these services do not really provide actual security, does not change this. Knowledgeable buyers probably understand that buying from the market leaders does not actually increase the security of their HTTPs service. After all, the security of HTTPs is a weakest-link problem and thus determined by the weakest CA. Moreover, the reputation of the market leaders does not necessarily mean they are actually more secure, as the large CAs have also proven vulnerable to attack and have not always been transparent about this. Even when a buyer understands this, it still makes sense to buy from the market leaders rational. Enterprise support, a liability shield, security signals to third parties and better continuity insurance are all valuable.

The price differences among certificates are large in absolute terms, but they are modest when compared to other cost components. Saving several hundreds of dollars is a marginal gain in light of the cost of installation, perceived trustworthiness and better support. Furthermore, the price of a certificate will typically be amortized over millions or even billions of clicks. Even when compared to a company’s own intermediate CA, which can issue free certificates, the price difference is that significant. In the words of the respondent, self-issued certificates are ‘not as cheap as you would hope’. There are still substantial costs related to the need for dedicated and trained staff for certificate management and the time spent by other business units involved in billing and reporting.

All these considerations reinforce the choice to buy from the market leaders, i.e., they strengthen concentration in the market and differentiate them enough from competitors to charge substantially higher prices.

### 6.3 *Incentives for Security*

Now that we better understand what the market is actually selling, what does this tell us about the security incentives at work? Given that the market leaders successfully differentiate their products via, among other things, security-related features, there appears to be a significant willingness-to-pay for security among buyers. But does this willingness-to-pay translate into actual security incentives? In other words, can CAs attract more customers or charge higher prices by investing more in security? This is not at all clear. Two classic problems affect the proper alignment of incentives: information asymmetry and externalities.

The information asymmetry prevents buyers from knowing what CAs are really doing. Buyers are paying for the perception of security, for a liability shield and for trust signals to third parties. None of these correlate verifiably with actual security. Given that CA security is largely unobservable to buyers, their demand for security does not necessarily translate into strong security incentives for CAs.

The incentive problem is exacerbated by the negative externalities that are the result of the weakest-link security of the system. The failure of a single CA impacts the whole ecosystem, not just that CA’s customers. All other things being equal, these interdependencies would undermine the incentives of CAs to invest, as the security of their customers also depends on the efforts of all other CAs.

The most powerful incentive for security seems to be reputation effects. Given that the market leaders leverage their reputation to charge higher prices and capture a larger market share, does this make them more sensitive to the reputation damage caused by breaches? Again: not necessarily. Yes, they have more of a reputation to lose compared to smaller, lesser-known brands. But they also are less threatened by the ultimate reputation effect: being removed from the root stores of browser and OS vendors and, as almost unavoidable consequence, going into

bankruptcy. The fact that the market leaders are more or less too-big-to-fail provides a perverse incentive to browser and OS vendors to keep them in the root store even at high cost. To phrase it differently: those vendors have to trade off availability of a large portion of the web against the confidentiality and integrity of the communications of the specific domains that are attacked.

Ironically, the security problems that have plagued the HTTPS ecosystem over the past few years may in fact benefit the market leaders, even though they themselves were partially to blame for these problems. The breaches have increased the demand for security and this demand seems to latch onto whatever security signals are available, regardless of their relationship to actual security. It seems reasonable to assume that post DigiNotar, buyers felt the pressure to shift from smaller CAs towards the larger, more ‘trusted’ brands.<sup>21</sup> The security problems also appear to have led enterprise customers to strategies of redundancy – i.e., encrypting connections using two certificates from two suppliers instead of one – which, again, would benefit the market leaders.

All of this may impact the attempts to fix the systemic vulnerabilities of the system. The current incentive structures seems quite favourable for the dominant players, which might make them reluctant, or at least less eager, to push for adoption of one of the proposed technical solutions. This is not to suggest that they will act against them, but rather that the status quo works quite well for them – perhaps even more so because of recent breaches. We should keep this in mind during the last part of this paper, where we discuss possible improvements in HTTPS governance.

## **7. Improving HTTPS Governance**

In Sections 7.1 and 7.2, we analyse the current regulatory and technical solutions to the systemic vulnerabilities of HTTPS. We evaluate the different regulatory and technical proposals, for improving HTTPS governance, along with possible alternatives, in section 7.3.

### *7.1 Regulatory Solutions*

In terms of regulation, and government policy in general, the HTTPS authentication model is by and large untouched in both the US and the EU (Roosa & Schultze, 2010; Van Eijk, 2012; Voulon, 2012).<sup>22</sup> This is bound to change in the near future. Interestingly, the United States and the European Union seem to have opted for a different approach. The European Commission (2012a) has proposed a new ‘eSignatures Regulation’, which contains several provisions that, if enacted in its current form, will impact the HTTPS ecosystem globally. The US National Institute for Standards and Technology (NIST), on the other hand, is opting for a multi-

---

<sup>21</sup> We hope to test this hypothesis in the future by analysing the 2013 or 2013 dataset of the SSL Observatory and looking at the changes in market shares since the DigiNotar collapse once it becomes available.

<sup>22</sup> Industry standards are formulated, amongst others, by the American Bar Association, the American Institute of Certified Public Accountants, the CA/Browser Forum and ETSI.

stakeholder solution and organizing a series of workshops aimed at non-regulatory policy and technical resolutions to overcome the systemic vulnerabilities.<sup>23</sup>

The proposed EU Regulation on ‘electronic identification and trust service for electronic transactions in the internal market’ will replace the 1999 Electronic Signatures Directive. The ordinary legislative procedure will be followed, meaning that the definitive contents of the Regulation are to be negotiated between the Council and European Parliament. The Parliament is set to vote on the proposal and a series of amendments to it in September 2013. Once enacted, the Regulation acquires the status of binding legislation in all 27 EU Member States,<sup>24</sup> whereas the more common Directives need subsequent implementation by the Member States.

With regard to its scope, the EU proposal targets ‘trust service providers’, a concept that includes CAs issuing SSL certificates for HTTPS communications.<sup>25</sup> Other critical HTTPS stakeholders – browser vendors and website operators – remain unregulated. We learn from section 5.2 that roughly 80% of the CA market is controlled by a limited number of companies and that these organizations appear to fall within EU jurisdiction.<sup>26</sup> This makes regulation in itself a solution to consider and deconstructs conventional wisdom that internet regulation should not be considered because laws are inherently local, whereas ‘the internet’ is a global communications system.

A fundamental legal requirement for HTTPS governance is that any regulatory proposal needs to apprise constitutional values such as legality, privacy and communications freedom. In the aftermath of the DigiNotar breach, economic interests prevailed over legality (that government action must be based in law) and confidentiality interests. The EU proposal falls short with regard to providing solid normative guidance on how to balance the underlying values of information security – availability, confidentiality, integrity – even though fundamental rights frameworks positively require legislative bodies to do so (Arnbak & Van Eijk, 2012).

In the following, we subsequently discuss governance proposals on security requirements, security breach notification, liability – typical remedies suggested by security economics to reduce information asymmetries and externalities – and on instituting ‘chain of trust transparency’. Earlier work by two of the authors discussed these aspects in more legal detail, as well as other important dimensions of the Regulation, such as underlying normative values, the scope of governance, supervision and auditing (Arnbak & Van Eijk, 2012).

---

<sup>23</sup> See: [http://www.nist.gov/itl/csd/ct/ca\\_workshop.cfm](http://www.nist.gov/itl/csd/ct/ca_workshop.cfm)

<sup>24</sup> Cf. art. 288 Treaty on the Functioning of the European Union.

<sup>25</sup> Art. 3[7] sub 12 explicitly refers to ‘website authentication’, while the Impact Assessment details that this refers to the issuance of SSL certificates (European-Commission, 2012b, pp. 86-88).

<sup>26</sup> A quick scan of the legal documentation of these CAs tells us that all major CAs have offices within the European Union. See: GeoTrust (2013); Comodo (2013); GlobalSign (2013); GoDaddy (2013); Symantec (2013); Thawte (2013); Entrust (2013); with the notable exception of CA DigiCert, that has a market share of about 5%. Roughly 15% at the tail-end of the market hasn’t been researched.

## A) Security Requirements

The EU proposal introduces a new obligation for CAs on security requirements. CAs need to implement ‘appropriate technical and organisational measures to manage the risks posed to the security of the trust services they provide [...] having regard to the state of the art’, according to art. 15[1]. ‘In particular, measures shall be taken to prevent and minimise the impact of security incidents and inform stakeholders of adverse effects of any incidents.’ Failure to comply will cause the CA to be liable for any direct damages on the basis of art. 9[1] (discussed in section C). Bearing in mind previous breaches and the critical role of CAs in HTTPS communications, it clearly makes sense to mandate CAs to have state of the art security practises in place. On the other hand, as long as the weakest link problem hasn’t been solved (technically), it only takes on CA to be breached to undermine the security of the entire ecosystem. Another real challenge that is not addressed by this EU proposal lies with website HTTPS implementation (see section 2.1). A value chain approach towards regulation would have exposed this important aspect.

The specific security requirements are not summed up in the Regulation, while both the European Commission and national supervisory bodies are granted executive power to adopt delegated acts and issue binding instructions on the basis of art. 15[4] to 15[6]. This provides flexibility for regulators and enforcers to adapt security requirements in line with best practises. But balancing of different private and public interests is equally important. Notably, recital 26 mentions that the security requirements should serve ‘to boost user trust in the single market’, rather than to protect the integrity and confidentiality of trust services. The recital seems to imply that security requirements are there to keep up appearances with users (boost trust), rather than effectively contributing to securing HTTPS communications and the systems it relies on. As observed before, we see a prevailing economic rationale, rather than one concerned with the broader underlying interests of information security and constitutional values.

## B) Security Breach Notification

A recurring characteristic of the CA breaches discussed in Section 3 is the tendency of CAs to conceal these for both browsers, websites, authorities and the public. Indeed, strong incentives exist to do so. The associated lack of transparency complicates threat and vulnerability modelling and consequent informed HTTPS governance responses, thus should be prioritized in any regulatory initiative.

Security breach notifications should, at least in theory, help minimise the damage after a breach has occurred and provide incentives for organisations to invest in information security upfront. The EU proposal introduces such an SBN in art. 15[2]. CAs are to notify relevant authorities of a breach of security or a loss of integrity ‘where feasible within 24 hours’, if the breach ‘has a significant impact on the trust service provided and on the personal data maintained therein’. If disclosure of the breach is in the public interest, relevant authorities may inform the public or require the CA to do so. Cross-border breaches should be notified by the authorities to the relevant supervisory bodies in other Member States and to ENISA. Supervisory bodies are to report on the notifications to the European Commission and ENISA (art. 15[3]).

There appears to be broad consensus that breach notifications are an appropriate measure to relieve the HTTPS ecosystem of perceived trust in an succeeded authentication, where the validity of the authentication is unwarranted

(ENISA, 2011). It is telling that the security breach at Verisign only became public two years after the incident and through an indirect way, when Security and Exchange Commission (2011) regulations mandated companies to notify investors of intrusions since October 2011.

SBN legislation is not in itself a silver bullet in augmenting security levels. Much of the impact will, again, depend on the details. The authority to flesh out the details is, as with the security requirements provision, delegated to the European Commission (art. 15[6]) and supervisory bodies (art. 15[4]). Breach notifications are of paramount importance for the well-functioning of HTTPS communications and addressing information asymmetry. As notifications enable trust revocation by browser vendors, security measures to be taken by other CAs, transition to secure certificates by website operators and security measures by end-users, a strict regime for notifications – which types of breaches should be made public by default, for instance – is defensible.

Experiences with SBN legislation in the United States also suggest that notifications need to be complemented with pro-active and punitive enforcement to be effective. Enforcement should be pro-active in order to avoid non-compliance (for example, as a part of yearly audits). If this fails to materialize, strong incentives exist not to notify breaches at all, at the expense of the well-intentioned companies that take security and the interests of customers seriously (Winn, 2009, p. 33). Enforcement also should have a punitive element, in addition to reputational costs for CAs. This could be achieved by effectuating liability when security breaches remain unreported, a logical framework not included in the current Commission proposal. In Section 6.3, we have argued that reputation losses might not affect major CAs. If notifications are not supported with a scheme of meaningful sanctions, they may only impact smaller market players who risk being thrown out of root stores for non-reporting. Another lesson from the US experience is to consider avoiding ‘safe harbors’, instances in which companies are exempted from notification, for encrypted data. Winn notes that this creates ‘perverse incentives to invest in mitigating harms after they occur instead of prevention’ (Thaw, 2011; Winn, 2009, p. 3).

### C) Liability

Currently, liability for security breaches is disclaimed across the HTTPS value chain and transferred to relying parties such as end-users. This practise is explicitly approved of in industry self-regulation policies (CA/Browser-Forum, 2011). In art. 9[1], the EU proposal introduces a new liability regime for ‘trust service providers’ that puts an end to this practise. It provides for ‘entitlement to compensation of damage caused by any negligent trust service provider for failure to comply with security good practices which result in a security breach which has a significant impact on the service.’ Article 15[1] on security practises constitutes the threshold for effectuating liability, but failing to notify a security breach does not (in the current proposals).

The breaches at CAs are indeed a concern to HTTPS communications and point to substantial negative externalities associated with a breach at one isolated CA, as the entire HTTPS ecosystem is at risk of being compromised. A liability regime may incentivise CAs under EU jurisdiction to take security more seriously. In addition, the burden of proof lies with the CA and that may lead to investment in proper logging functions (unlike those DigiNotar had in place).

The proposal has serious drawbacks, however. The proposed liability regime doesn’t appreciate the dynamics of the HTTPS authentication value chain. Art. 9 has

two dimensions, one of which is not thought through in the proposal: ‘negligence’ and ‘any direct damage’. Regardless of the security practises and intentions of one individual CA (‘negligence’), no single company is able to stand in for the consequences for the entire HTTPS ecosystem or a specific target of an attack (‘any direct damage’) once its systems are breached. Again: consider DigiNotar, with its an annual budget of a few million US Dollars, whereas rogue certificates were issued for activities of Google, Facebook, Skype, cia.gov, etc. (see Section 3.2). In such a scenario, liability for any given CA not seems unreasonable, and outright harmful: DigiNotar went bankrupt in the aftermath of its breach. One would assume that this leads to liability circumvention through creating subsidiary special purposes companies that bear full liability and can be easily be filed for bankruptcy – in the same way that DigiNotar quickly went under, while its parent company Vasco so far escaped unscathed.

Introducing liability regimes for CAs operating in the EU may have other undesirable effects. A liability regime might raise entry barriers and favour incumbent CAs who are in a relatively strong position to shield themselves from liability. Conversely, small CAs will think twice before doing business with large corporations processing vast amounts of sensitive data, or might not even enter that market at all.

HTTPS value chain analysis suggests alternative approaches, in which liability is spread across the value chain according to the risk associated with certain activities. CAs have their share in this risk, but are mostly unaware what value a sold certificate should protect, whereas website owners know what kinds of sensitive information they are dealing with (online banking, E-Commerce, private communications, etc.). Another aspect that would deserve attention in the context of liability, is the option for CAs and other stakeholders to pass on liability to information technology producers such as software developers, who in many cases ‘are in a better position than database owners to fix problems with information security’ (Winn, 2009).

#### D) Chain of Trust Transparency

The last policy proposal we discuss is instituting so-called ‘chain of trust transparency’. In section 2.2, we described the value-chain dynamics of chains of trust, for example when a Root CA issues signs a subordinate CA which can then sell certificates on the market, or a subordinate CA that buy itself into a chain of trust with a Root CA in order have its certificates trusted by default by web browser vendors.

For reasons of trust, security and privacy, discussed in section 3.2, it is crucial to know what organizations are behind the signing of a certificate. While most website operators end end-users may benefit from HTTPS deployment in itself, chain of trust transparency is both key to accountability of CAs and is particularly important for certain groups, such as banks, health institutions and political organizations that structurally engage in sensitive communications. Without knowledge of the organizations behind the signing of one certificate, these groups could be subject to systematic monitoring of their sensitive communications unnoticed (Soghoian & Stamm, 2010).

The subordinate CAs that operate within these chains of trust are hardly known today. Transparency is non-existing at the institutional level and only starting to emerge through various (research) projects. In section 5.2, we have also made a first assessment of these chains of trust, and put some initial numbers to the number of CAs that actually operate on the HTTPS market. Several add-ons to web browsers,

such as CertPatrol for Firefox,<sup>27</sup> provide insight into the chain of trust of an individual certificate for tech savvy end-users.

In a recent amendment to its CA policy, web browser vendor Mozilla seeks to enhance chain of trust transparency. Amongst others, Mozilla requires that subordinate CA certificates ‘either be technically constrained or be publicly disclosed and audited’.<sup>28</sup> Subordinate CAs, in other words, must either be constrained to only issue certificates for a (small set of) domain name(s) – on internal networks, for example – or their chain of trust must be publicly disclosed and audited, basically holding subordinate CAs to the same standard as root CAs and making a root CA accountable for all the certificates it signs. Existing subordinate CA certificates have to comply by 15 May 2014. So far, the requirement is not part of the EU proposal.

## 7.2 Technical Solutions<sup>29</sup>

A host of technical solutions to the systemic vulnerabilities of the current system are currently being developed. Among the most prominent proposals are Convergence (Convergence, 2011), Perspectives (Perspectives, 2011), DANE (IETF, 2012b), Sovereign Keys (EFF, 2011b), Certificate Transparency (IETF, 2013a; Transparency, 2012), Public Key Pinning (IETF, 2012a)<sup>30</sup> and TACK (IETF, 2013b; TACK, 2012).

Within the confines of the paper we cannot discuss the technical merits of the different approaches and their complicated and evolving implementations, not to mention the combinations in which they could potentially coexist or reinforce each other. From the perspective of governance, however, we can make several general observations:

- all proposals can function on top of the current CA system, leaving it in place as is or even depending on it, while a subset of proposals can in principle replace it (e.g., Perspectives, DANE, Sovereign Keys);
- all proposals can follow more or less incremental adoption paths alongside the current process, albeit that some paths are a lot more difficult than others and that all need support from browsers;
- all proposals attempt to solve the weakest-link problem by introducing another source of authority to either check the CA – i.e., to check whether the certificate that is validated through the normal SSL/TLS process, and that resolves back to a CA with root status, is indeed the correct certificate – or to replace the CA;

---

<sup>27</sup> See: <https://addons.mozilla.org/en-us/firefox/addon/certificate-patrol/>

<sup>28</sup> See: Mozilla (2013), art. 8: ‘all certificates that are capable of being used to issue new certificates, and which directly or transitively chain to a certificate included in Mozilla’s CA Certificate Program, MUST be operated in accordance with Mozilla’s CA Certificate Policy and MUST either be technically constrained or be publicly disclosed and audited.’

<sup>29</sup> We are grateful to Bernhard Amann (ICSI, Berkeley) for his comments on this Section.

<sup>30</sup> Public Key Pinning has a client-based configuration where the browser comes pre-shipped with some keys, and is actively being used by Chrome. It is however not scalable, an issue that the sever-based configuration, designed as an extension to the HTTP aims to tackle. Here we discuss the latter.

- all proposals reduce the information asymmetry of buyers and users versus the CAs by more rapidly and systematically uncovering suspect certificates and those who issued them.

The new sources of authoritative information about certificates vary. Convergence and Perspectives introduce notaries; while DANE, Sovereign Keys, Public Key Pinning and TACK use different approaches to make the domain owner himself the authority; Certificate Transparency locates authority in public and auditable logs of all issued certificates.

We have to keep in mind that none of these solutions are anywhere close to moving from experimental designs to large-scale adoption. That said, they do seem promising in terms of addressing the current weaknesses, especially the weakest-link problem, for which regulatory solutions appear ineffective. In the longer run, therefore, these options are preferable. This makes it relevant to assess how they relate to the incentives of the actors in the value chain.

In the previous section, we argued that the insecure status quo is not per se bad for the business of the market leaders, on the contrary. In light of this, one might assume that CAs are not particularly keen on actively helping any of these proposals along, especially the ones that theoretically could make them obsolete at some point in the future. In practice, however, there are examples of CAs involved in the development, such as DigiCert and Comodo who are experimenting with Certificate Transparency (Langley, 2012a). Other proposals, such as DANE, point to the fact that they require non-trivial activities on the side of the domain owner and that these activities may be done by their CA. What these hints suggest is the fact that for CAs, these proposals do not so much crowd out their services in the current SSL market, at least not in the immediate future. They may in fact create new markets for complementary services to support HTTPS long before these solutions may start cannibalizing their current business model.

The incentives of the browser vendors also come into play. Each proposal is intensely debated in terms of its impact on browser performance. Given the huge scale of their deployment, browsers have to operate under wildly different and sometimes very unfavourable conditions. Many of the proposals struggle with how to perform the proposed checks on certificates under some of those conditions, such as the filtering in place in many access networks. These are the same reasons why OCSP never worked as intended. As a first step, some level of browser support could be delivered via extensions that third parties can develop and users can install on their own – as Convergence has done for Firefox. Any form of large-scale adoption, however, requires default support by the browser vendor. Vendors have been active in this area, especially Google and Mozilla. Chrome, for example, already supports public key pinning for certain opt-in domains. While none of these solutions are easy to scale, it does suggest there are benefits for early adopters. This would mitigate the problem of positive externalities that many security solutions struggle with: adoption imposes immediate cost on actors, while the security benefits are in the future and will only occur if and when a critical mass of actors have adopted it.

As far as users and domain owners are concerned, the incentives regarding adoption of new solutions are not that straightforward. Yes, both parties have a stake in securing their HTTPS communication, but there are also costs associated with the complications that the new solutions generate in actual use. Whether these costs are worth it depends on the kind of threats they want defend themselves against. It seems that your average cybercriminal is not interested in breaching a CA and manipulating network traffic. Financially attractive information can also be acquired through more cost-effective attacks (Florencio & Herley, 2011; Langley, 2013). From previous

attacks and breaches, it appears that state-sponsored attackers and large corporations, rather than profit-driven criminals, are more likely to engage in the more complex MITM attacks in the realm of the SSL authentication model.<sup>31</sup> For some user groups and some domains, this threat might make early adoption of a new solution attractive.

### 7.3 Evaluation

How do the different technical and regulatory solutions deal with the market failures of information asymmetry and negative externalities? Our analysis suggests the regulatory options they do not seem promising in this respect.

CA liability could theoretically internalize the externalities, but the potentially wide-ranging consequences of a breach mean that they are quickly beyond the scope of an individual firm. In all likelihood, such provisions will lead to liability-avoiding legal arrangements, such as putting liabilities in separate legal entities. Also, the fact that the leading CAs are all too big to fail reduces their exposure to liability, just like it has for the banks in the financial sector. Ineffective liability provisions also undermine the security requirements that the proposal introduces for CAs, as they depend in large part on the liability as the means of enforcement.

Security breach notification requirements might help reduce the information asymmetry, but not in the weak form that the EU proposal entails. Breach notification is also often discussed as a possible security incentive via reputation effects. In light of the recent breaches, this seems unlikely to happen in this market. The market leaders have appeared relatively impervious to reputation damage.

In other words, the EU proposal does not seem too helpful in aligning the incentives in the CA market towards securing the HTTPS value chain. When it comes to the overcoming the weakest-link security of the current model, the technical proposals offer hope. They also promise to reduce the information asymmetry, in the sense of quickly and systematically exposing suspect certificates and their issuers.

Of course, none of the technical proposals are available at scale for the immediate future. This raises the issue of the transition. It is interesting to note that most of the proposed solutions appear not to depend on the cooperation of the CAs. They are in the hands of the domain owners, end users, and browser vendors – and among the latter we have already seen active support and development. Particularly, the institutional policy suggested by Mozilla to either technically constrain or publicly disclose chains of trust should be extended throughout the certificate issuance value-chain (Roosa & Schultze, 2013, p. 9), and included in the EU Regulation .

That said, most solutions would scale up sooner if the CAs have the incentives to help put a new system in place. The market is highly concentrated and four dominant firms can reach most of the buyers of SSL certificates. The incentives of the dominant CAs seem mixed. Given the willingness-to-pay for security in the market, we expect the market leaders to thrive because of, rather than in spite of, the recent breaches. Does that mean they oppose or ignore the proposed technical solutions? That doesn't appear to be the case. Most new solutions will function on top of the existing market for the foreseeable future. Their implementation also often

---

<sup>31</sup> And indeed, from the very recently disclosed breach at CA Turktrust. See Schultze (2013).

requires non-trivial actions by the domain owners. Taken together, this suggests that the CAs might see new revenue models open up with services for domain owners that complement, rather than cannibalize their existing models.

## 8. Conclusion

HTTPS has become the de facto standard for securing web communications, and its systemic vulnerabilities are a reasonable concern to policy-makers on both sides of the Atlantic. A better understanding of the interactions and incentives between website operators, certificate authorities, browsers vendors and end-users is vital to inform policy responses. This paper contributes in this regard, both on the conceptual and empirical level.

The breaches at CAs – including those at DigiNotar, Comodo, Verisign, GlobalSign and Trustwave – have exposed different systemic vulnerabilities of HTTPS: the security of the entire ecosystem depends on the security of one of the several dozens of CAs (‘weakest-link’), browsers are not really able to revoke trust in major CAs (‘too big to fail’), most website operators choose not to offer HTTPS or implement it poorly, and while end-users cannot reasonably be expected to check the security of certificates, they currently bear the liability of security incidents. Notably, all breached CAs complied to security auditing schemes in place in regulatory instruments such as the 1999 EU eSignatures Directive (as far as applicable) and the various CA policies of browser vendors. And all breached CAs known to date initially managed to conceal these security incidents, which begs the question how many CA breaches have gone unnoticed to the general public.

Our empirical analysis has uncovered intriguing results. We found that the market is highly concentrated, with very large price differences among suppliers and limited price competition. Market leaders differentiate their offerings partially via security features: their reputation enables them to offer security signals – though some of these signals are absurd and none of them correlate verifiably with actual security – and a limited liability shield. In other words, the current vulnerabilities may actually benefit rather than hurt the dominant CAs.

In terms of solutions, the EU has opted for a regulatory response, while the US seeks resolve in multi-stakeholder and technical approaches. In general, the technical solutions are more promising than the regulatory ones when it comes to the most urgent problem that needs to be solved, namely the weakest-link security of the current ecosystem. However, none of these proposals are anywhere close to large-scale implementation, let alone adoption. Regulation can play a role and territorial law evasion seems difficult with over 80% of the HTTPS market owned by a small number of CAs, that appear to fall within European jurisdiction.

In general, regulation should provide normative guidance on the balancing of confidentiality, integrity and availability of HTTPS. Also, specific regulatory measures need to allocate responsibilities throughout the value chain and may include internal security requirements paired with proportionate liability provisions and meaningful security breach notifications coupled with proactive enforcement and punitive elements, as previous breaches has shown that reputation damage associated with breaches does hardly effect major consolidated CAs. Chain of trust transparency, requiring technical name constraints or full disclosure of certificate hierarchy and audits associated with certificate issuance, should be included in any policy responses, given the associated interests for high-end security seeking website operators and end-users.

While the EU eSignatures proposal pioneers regulatory provisions in most of these areas, its details fail to achieve these policy goals. As such, the current proposal will reinforce market dominance of large CAs as well as current systemic vulnerabilities, such as the weakest-link problem. Instead of improving the security incentives in the market, it creates new long-term institutional dependencies on the actors whose roles should be limited from a security perspective.

## References

- Anderson, R. J. (2008). *Security Engineering: A Guide to Building Dependable Distributed Systems*: Wiley.
- Apple. (2013). Updated Apple Root Certificate Program. Retrieved 1 Feb 2013, from [http://www.apple.com/certificateauthority/ca\\_program.html](http://www.apple.com/certificateauthority/ca_program.html)
- Arnbak, A., & Van Eijk, N. (2012). Certificate Authority Collapse: Regulating Systemic Vulnerabilities in the Https Value Chain.
- Arstechnica. (2011). Https Is Great: Here's Why Everyone Needs to Use It (So We Can Too). Online at: <http://arstechnica.com/business/2011/03/https-is-great-here-is-why-everyone-needs-to-use-it-so-ars-can-too/> doi:
- Bakos, Y., Marotta-Wurgler, F., & Trossen, D. (2009). *Does Anyone Read the Fine Print? Testing a Law and Economics Approach to Standard Form Contracts*. Paper presented at the Testing a Law and Economics Approach to Standard Form Contracts (October 6, 2009). CELS 2009 4th Annual Conference on Empirical Legal Studies Paper.
- CA/Browser-Forum. (2011). Baseline Requirements for the Issuance and Management of Publicly-Trusted Certificates, Version 1.0, Effective 1 July 2012. Online at: [http://www.cabforum.org/Baseline\\_Requirements\\_V1.pdf](http://www.cabforum.org/Baseline_Requirements_V1.pdf).
- CA/Browser-Forum. (2012). Guidelines for the Issuance and Management of Extended Validation Certificates (version 1.4, effective 29 May 2012 ed.). Online at: [http://www.cabforum.org/Guidelines\\_v1\\_4.pdf](http://www.cabforum.org/Guidelines_v1_4.pdf)
- Comodo. (2013). Updated Contact Comodo. Retrieved 2013-06-02, from <http://www.comodo.com/contact-comodo/contact-us.php>
- Computer-World. (2012). Trustwave Admits Issuing Man-in-the-Middle Digital Certificate; Mozilla Debates Punishment. Online at: [http://www.computerworld.com/s/article/9224082/Trustwave\\_admits\\_issuing\\_man\\_in\\_the\\_middle\\_digital\\_certificate\\_Mozilla\\_debates\\_punishment](http://www.computerworld.com/s/article/9224082/Trustwave_admits_issuing_man_in_the_middle_digital_certificate_Mozilla_debates_punishment) doi:
- Convergence. (2011). Updated Convergence - Details. Retrieved 2013-03-01, from <http://convergence.io/details.html>
- DOJ, & FTC. (2010). Horizontal Merger Guidelines: U.S. Department of Justice & the Federal Trade Commission. Online at: <http://www.justice.gov/atr/public/guidelines/hmg-2010.html>.
- EFF. (2011a). Iranian Hackers Obtain Fraudulent Https Certificates: How Close to a Web Security Meltdown Did We Get? Online at: <https://www.eff.org/deeplinks/2011/03/iranian-hackers-obtain-fraudulent-https>
- EFF. (2011b). Updated The Sovereign Keys Project. Retrieved 2013-03-01, from <https://www.eff.org/sovereign-keys>
- ENISA. (2011). Operation Black Tulip: Certificate Authorities Lose Authority, Version 2, Dec. 2011. Online at: <http://www.enisa.europa.eu/media/news-items/operation-black-tulip>.
- Entrust. (2013). Updated Contact Us - Office Locations. Retrieved 2013-06-02, from <http://www.entrust.com/contact/offices.htm>

- European-Comission. (2012a). Com(2012) 238/2. Regulation of the European Parliament and of the Council on Electronic Identification and Trust Services for Electronic Transactions in the Internal Market.
- European-Comission. (2012b). Swd(2012) 135. Comission Staff Working Paper Impact Assessment Accompanying the Proposal for a Regulation of the Ruoepan Parliament and of the Council on Electronic Identification and Trust Services for Electronic Transactions in the Internal Market.
- Florencio, D., & Herley, C. (2011). *Where Do All the Attacks Go?* Paper presented at the Workshop on Economics of Inofrmation Security (WEIS) 2011. Online at: <http://research.microsoft.com/pubs/149885/WhereDoAllTheAttacksGo>.
- Fox-IT. (2011). Diginotar Certificate Authority Breach, 5 Sep. 2011. Online at: <http://www.rijksoverheid.nl/documenten-en-publicaties/rapporten/2011/09/05/diginotar-public-report-version-1.html>
- Fox-IT. (2012). Black Tulip – Report of the Investigation into the Diginotar Certificate Authority Breach, 13 Aug. 2012. Online at: <http://www.rijksoverheid.nl/documenten-en-publicaties/rapporten/2012/08/13/black-tulip-update.html>
- GeoTrust. (2013). Updated Geotrust - About Us. Retrieved 2013-06-02, from <http://www.geotrust.eu/en/about+us/>;
- GlobalSign. (2013). Updated Globalsign International. Retrieved 2013-06-02, from <https://nl.globalsign.com/en/about+globalsign/international/>
- GoDaddy. (2013). Updated About Go Daddy. Retrieved 2013-06-02, from <https://www.godaddy.com/newscenter/about-godaddy.aspx?ci=9079>
- Hurst, R. (2012). How to Tell Dv and Ov Certificates Apart. September 11, 2012. Online at: <http://unmitigatedrisk.com/?p=203>
- IETF. (2005). Internet X.509 Public Key Infrastructure: Certification Path Building, Rfc 4158, September 2005. Online at: <http://tools.ietf.org/html/rfc4158>.
- IETF. (2012a). Public Key Pinning Extension for Http, Draft-Ietf-Websec-Key-Pinning-04. Online at: <http://tools.ietf.org/html/draft-ietf-websec-key-pinning-04>.
- IETF. (2012b). Rfc 6698 - the Dns-Based Authentication of Named Entities (Dane), Transport Layer Security (Tls) Protocol: Tlsa. Online at: <http://tools.ietf.org/html/rfc6698>.
- IETF. (2013a). Certificate Transparency, Draft-Laurie-Pki-Sunlight-07. Online at: <http://tools.ietf.org/html/draft-laurie-pki-sunlight-07>.
- IETF. (2013b). Trust Assertions for Certificate Keys, Draft-Perrin-Tls-Tack-02.Txt. Online at: <http://tools.ietf.org/html/draft-perrin-tls-tack-02>.
- InfoSecurity. (2011). 31 March 2011 Comodo Admits Two More Registration Authorities Hacked. Online at: <http://www.infosecurity-magazine.com/view/16986/comodo-admits-two-more-registration-authorities-hacked>.
- Kelkman, O. M. (2013). Dnssec Musings - Diginotar, Dane and Deployment: NLnet Labs. Online at: [http://conference.apnic.net/\\_\\_data/assets/pdf\\_file/0005/58901/dnssec-diginotar-dane\\_1361864377.pdf](http://conference.apnic.net/__data/assets/pdf_file/0005/58901/dnssec-diginotar-dane_1361864377.pdf).
- Langley, A. (2011). Browsers Have Never Been Able to Make Ojsp Lookups Blocking, and Therefore Ojsp Is Basically Useless for Security. Online at: <http://www.imperialviolet.org/2011/11/29/certtransparency.html>
- Langley, A. (2012a). Certificate Transparency. Online at: <http://www.imperialviolet.org/2012/11/06/certtrans.html>
- Langley, A. (2012b). Living with Htpps Online at: <http://www.imperialviolet.org/2012/07/19/hope9talk.html>

- Langley, A. (2013). World Crypto 2013. Online at: <http://www.imperialviolet.org/2013/01/13/rwc03.html>
- Microsoft. (2009). Updated Microsoft Root Certificate Program, Updated January 15 2009. Retrieved 1 Feb 2013, from <http://technet.microsoft.com/en-us/library/cc751157.aspx>
- Microsoft. (2012). Updated Windows and Windows Phone 8 Ssl Root Certificate Program (Member Cas); Updated December 2012. Retrieved 01 Feb 2013, from <http://social.technet.microsoft.com/wiki/contents/articles/14215.windows-and-windows-phone-8-ssl-root-certificate-program-member-cas.aspx>
- Mills, E. (2011). August 30, 2011 Google Users in Iran Targeted in Ssl Spoof. *CNET*. Online at: [http://news.cnet.com/8301-27080\\_3-20099421-245/google-users-in-iran-targeted-in-ssl-spoof/](http://news.cnet.com/8301-27080_3-20099421-245/google-users-in-iran-targeted-in-ssl-spoof/).
- Mozilla. (2013). Updated February 14, 2013 Mozilla Ca Certificate Policy, Version 2.1, . from <http://www.mozilla.org/projects/security/certs/policy/>
- NetCraft. (2012). Updated Retrieved 01 Feb 2013, from <http://www.netcraft.com/internet-data-mining/ssl-survey/>
- NRC. (2011). 10 September 2011 Mooie Dag Voor Een Black-out' ('Nice Day for a Black-out'), *NRC Handelsblad*, pp. 14-15.
- Perspectives. (2011). Updated Perspectives Project - What Is Perspectives? Retrieved 2013-03-01, from <http://perspectives-project.org/>
- Reuters. (2012). Key Internet Operator Verisign Hit by Hackers. Online at: <http://www.reuters.com/article/2012/02/02/us-hacking-verisign-idUSTRE8110Z820120202> doi:
- Roosa, S. B., & Schultze, S. (2010). The "Certificate Authority" Trust Model for Ssl: A Defective Foundation for Encrypted Web Traffic and a Legal Quagmire. *Intellectual Property & Technology Law Journal*, 22(11), 3.
- Roosa, S. B., & Schultze, S. (2013). Trust Darknet: Control and Compromise in the Internet's Certificate Authority Model (April 11, 2013). Online at: <http://ssrn.com/abstract=2249042>
- Schultze, S. (2013). Turktrust Certificate Authority Errors Demonstrate the Risk of "Subordinate" Certificates. Online at: <https://freedom-to-tinker.com/blog/sjs/turktrust-certificate-authority-errors-demonstrate-the-risk-of-subordinate-certificates/>
- SEC. (2011). Commission Cf Disclosure Guidance: Topic No. 2. Cybersecurity. Online at: <http://www.sec.gov/divisions/corpfin/guidance/cfguidance-topic2.htm>
- Shapiro, C., & Varian, H. (1998). *Information Rules*: Harvard business school press.
- Soghoian, C., & Stamm, S. (2012). Certified Lies: Detecting and Defeating Government Interception Attacks against Ssl (Short Paper) *Financial Cryptography and Data Security* (pp. 250-259): Springer.
- SSL-Pulse. (2013). Updated Survey of the Ssl Implementation of the Most Popular Web Sites. Retrieved 2013-03-01, from <https://www.trustworthyinternet.org/ssl-pulse/>
- Symantec. (2013). Updated About Symantec - Locations. Retrieved 2013-06-02, from <https://www.symantec.com/en/uk/about/profile/locations/>
- TACK. (2012). Updated Tack, for Pinning. Retrieved 2013-03-01, from <http://tack.io/>
- Thaw, D. (2011). *Characterizing, Classifying, and Understanding Information Security Laws and Regulations, Forthcoming Ph.D. Dissertation*. University of California, Berkeley.

- Thawte. (2013). Updated Over Ons - Samenwerking En Persoonlijke Ondersteuning. Retrieved 2013-06-02, from <http://www.thawte.nl/nl/over+ons/>
- Transparency. (2012). Updated Certificate Transparency - Home. Retrieved 2013-03-01, from <http://www.certificate-transparency.org/>
- Van Eijk, N. (2012). Diginotar: Lessons to Be Learnt. *Ars Aequi*, 61(2), 80-82.
- Verisign. (2010). Verisign Annual Report 2009 Online at: <http://files.shareholder.com/downloads/VRSN/2358873860x0x365048/ea1e2339-4582-4149-bf73-5391991cc3c1/>.
- Voulon, M. B. (2012). Toezicht Op Certification Service Providers (Csps). *Computerrecht*, 2012/1.
- Vratonjic, N., Freudiger, J., Bindschaedler, V., & Hubaux, J.-P. (2011). *The Inconvenient Truth About Web Certificates*. Paper presented at the Workshop on Economics of Information Security (WEIS) 2011.
- Winn, J. (2009). Are 'better' security Breach Notification Laws Possible? *Berkeley Technology Law Journal*, 24.