Certified Lies: Detecting and Defeating Government Interception Attacks Against SSL

Christopher Soghoian
csoghoian@gmail.com

Sid Stamm
sidstamm@gmail.com

"Cryptography is typically bypassed, not penetrated."
— Adi Shamir [47]

"Just because encryption is involved, that doesn't give you a talisman against a prosecutor. They can compel a service provider to cooperate."
— Phil Zimmerman [48]

ABSTRACT

This paper introduces the *compelled certificate creation attack*, in which government agencies may compel a certificate authority to issue false SSL certificates that can be used by intelligence agencies to covertly intercept and hijack individuals' secure Web-based communications. Although we do not have direct evidence that this form of active surveillance is taking place in the wild, we show how products already on the market are geared and marketed towards this kind of use—suggesting such attacks may occur in the future, if they are not already occurring. Finally, we introduce a lightweight browser add-on that detects and thwarts such attacks.

1. INTRODUCTION

Consider a hypothetical situation where an American executive is in France for a series of trade negotiations. After a day of meetings, she logs in to her corporate webmail account using her company-provided laptop and the hotel wireless network. Relying on the training she received from her company's IT department, she makes certain to look for the SSL encryption lock icon in her web browser, and only after determining that the connection is secure does she enter her login credentials and then begin to upload materials to be shared with her colleagues. However, unknown to the executive, the French government has engaged in a sophisticated man-in-the-middle attack, and is able to covertly intercept the executive's SSL encrypted connections. Agents from the state security apparatus leak details of her communications to the French company with whom she is negotiating, who use the information to gain an upperhand in the negotiations. While this scenario is fictitious, the vulnerability is not.

The security and confidentiality of millions of Internet transactions per day depend upon the Secure Socket Layer (SSL)/Transport Layer Security (TLS) protocol. At the core of this system are a number of *Certificate Authorities* (CAs), each of which is responsible for verifying the identity of the entities to whom they grant SSL certificates. It is because of the confidentiality and authenticity provided by the CA based *public key infrastructure* that users around the world can bank online, engage in electronic commerce and communicate with their friends and loved ones about the most sensitive of subjects without having to worry about malicious third parties intercepting and deciphering their communications.

While not completely obvious, the CAs are all trusted equally in the SSL public key infrastructure, a problem amplified by the fact that the major web browsers trust hundreds of different firms to issue certificates for any site. Each of these firms can be compelled by their national government to issue a certificate for any particular website that all web browsers will trust without warning. Thus, users around the world are put in a position where their browser entrusts their private data, indirectly, to a large number of governments (both foreign and domestic) whom these individuals may not ordinarily trust.

In this paper, we introduce a new attack, the *compelled certificate creation attack*, in which government agencies compel (via a court order or some other legal process) a CA to issue false certificates that are then used by law enforcement and intelligence agencies to covertly intercept and hijack individuals' secure communications.

We also show how currently available surveillance products are advertised in a way that suggests that this attack is more than a theoretical concern, but is likely in active use; at least one private company is supplying government customers with specialized covert network appliances specifically designed to intercept SSL communications using deceptively created certificates.

In order to protect users from these powerful government adversaries, we introduce a lightweight defensive browser add-on that detects and thwarts such attacks. Finally, we use reductive analysis of governments' legal capabilities to perform an adversarial threat model analysis of the attack and our proposed defensive technology. We believe that this form of legal threat model analysis is itself new to the computer security literature.

In section 2 we provide a brief introduction to CAs, web

browsers and the man-in-the-middle attacks against them. In section 3 we discuss the presence of government-controlled CAs in the browsers. In section 4, we describe the compelled certificate creation attack and then in section 5, we present evidence that suggests it is being used. In section 6 we introduce our browser based add-on, and in section 7, we analyze its effectiveness via a threat model based analysis. Finally, we present related work in section 8 and conclude in section 9.

2. CERTIFICATE AUTHORITIES AND THE BROWSER VENDORS

In this section, we provide a brief overview of the roles played by the Certificate Authorities in the public key infrastructure, the browser vendors in picking the certificate authorities that they include in the browsers, and existing man-in-the-middle-attack techniques that circumvent SSL based security.

2.1 Certificate Authorities

> *"[Browser vendors] and users must be careful when deciding which certificates and certificate authorities are acceptable; a dishonest certificate authority can do tremendous damage."*
> — RFC 2246, The TLS Protocol 1.0 [15]

CAs play a vital role in the SSL *public key infrastructure* (PKI). Each CA's main responsibility is to verify the identity of the entity to which it issues a certificate.[1] Thus, when a user visits https://www.bankofamerica.com, her browser will inform her that the bank's certificate is valid, was issued by VeriSign, and that the website is run by Bank of America. It is because of the authenticity and confidentiality guaranteed by SSL that the user can continue with her transaction without having to worry that she is being phished by cyber-criminals.

CAs generally fall into one of three categories: Those trusted by the browsers ("root CAs"), those trusted by one of the root CAs ("intermediate CAs" or "subordinate CAs"), and those neither trusted by the browsers nor any intermediate CA ("untrusted CAs"). Furthermore, intermediate CAs do not necessarily have to be directly verified by a root CA — but can be verified by another intermediate CA, as long as the *chain of trust* eventually ends with a root CA.[2]

From the end users' perspective, root CAs and intermediate CAs are functionally equivalent. A website that presents a certificate signed by either form of CA will cause the users' browser to display a lock icon and to change the color of the location bar. Whereas certificates verified by an untrusted CA and those self-signed by the website owner will result in the display of a security warning, which for many non-technical users can be scary [40], confusing, and difficult to bypass in order to continue navigating the site [56].

As the CA system was originally designed and is currently implemented, all root CAs are equally trusted by the browsers. That is, each of the 264 root CAs trusted by Microsoft, the 166 root CAs trusted by Apple, and the 144 root CAs trusted by Firefox are capable of issuing certificates for any website, in any country or top level domain [18]. For example, even though Bank of America obtained its current SSL certificate from VeriSign, there is no technical reason why another CA, such as GoDaddy, cannot issue another certificate for the same site to someone else. Should a malicious third party somehow obtain a certificate for Bank of America's site and then trick a user into visiting their fake web server (for example, by using DNS or ARP spoofing), there is no practical, easy way for the user to determine that something bad has happened, as the browser interface will signal that a valid SSL session has been established.[3]

Of course, GoDaddy is extremely unlikely to knowingly provide such a certificate to a malicious third party. Doing so would almost certainly lead to significant damage to its reputation, a number of lawsuits, as well as the ultimate threat of having its trusted status revoked by the major web browsers.[4] Therefore, it is in each CAs' self-interest to ensure that malicious parties are not able to obtain a certificate for a site not under their own control.

It is important to note that there are no technical restrictions in place that prohibit a CA from issuing a certificate to a malicious third party. Thus, both the integrity of the CA based public key infrastructure and the security users' communications depend upon hundreds of CAs around the world choosing to do the right thing. Unfortunately, as will soon be clear, any one of those CAs can become the weakest link in the chain.

2.2 Web Browsers

There is no technical standard that specifies how web browsers should select their list of trusted CAs. As a result, each browser vendor has created their own set of policies to evaluate and approve CAs [3, 2, 1]. Since there is no evidence to suggest that any browser has knowingly or incompetently approved a rogue CA, we do not discuss each particular vendors' policies in depth.

---

[1] The level of verification performed by the CA depends upon the type of certificate purchased. A domain registration certificate can be obtained for less than $15, and will typically only require that the requester be able to reply to an email sent to the administrative address listed in the WHOIS database. Extended Validation (EV) certificates require a greater de of verification.

[2] Dan Kaminsky describes this aspect of the CA chain of trust as: "You can just walk up to a certificate authority and say, 'Yeah, so I spent a lot of money on my CA and it doesn't work with anyone outside my company. Um, here's a pile of money and I promise to be good.' No really, you can just buy a root certificate, effectively. It's not expensive, it's not that difficult, and there's an unknown number of companies out there – not just the certificate authorities but all of the companies that have intermediate certificates – they can all issue certificates for your domain [30]."

[3] Even if the user examines the more complex security information listed in the browser's SSL interface, she will still lack the information necessary to make an informed trust decision. Since GoDaddy is a valid certificate authority and has issued millions of other valid certificates, there is no way for the user to determine that any one particular certificate was improperly issued to a malicious third party.

[4] The browser vendors wield considerable theoretical power over each CA. Any CA no longer trusted by the major browsers will have an impossible time attracting or retaining clients, as visitors to those clients' websites will be greeted by a scary browser warning each time they attempt to establish a secure connection. Nevertheless, the browser vendors appear loathe to actually drop CAs that engage in inappropriate behavior — a rather lengthy list of bad CA practices that have not resulted in the CAs being dropped by one browser vendor can be seen in [38].

What does merit further attention is the method by which the browser vendors deliver and update their list of root CAs and the in-browser user interface provided to end-users to view and manage them.

The major browsers (Internet Explorer, Firefox, Chrome and Safari) have all adopted slightly different policies for managing and displaying the list of trusted CAs: Firefox is the only major browser to maintain its own database of trusted CAs, while the other three browsers instead rely upon a list of CAs provided by the operating system. However, since two of these three browser vendors are also major players in the computer operating system business, the line between browser and operating system tends to be rather blurry.

In years past, Microsoft, like the other vendors, included hundreds of CAs in its Windows operating system *Trusted Root Store*. Users who discovered the relevant user interface were able to view and manage the full list of CAs. However, in response to criticism from large enterprise customers, Microsoft reduced the number of certificates in the trusted store in subsequent OS versions down to just a handful.[5]

It would be easy for a naive user (or security researcher) comparing the various CA databases through the user interfaces provided by Microsoft, Apple and Mozilla to conclude that Microsoft has adopted a far more cautious approach in trusting CAs than its competitors, since the user interface of a fresh installation of Windows Vista or Windows 7 will list less than 15 CAs in the operating system's Trusted Root Store. Unfortunately, this interface is extremely misleading as it does not reveal the fact that Microsoft has opted to trust 264 different CAs. The company's own documentation reveals that:

> "Root certificates are updated on Windows Vista [and Windows 7] automatically. When a user visits a secure Web site (by using HTTPS SSL) [...] and encounters a new root certificate, the Windows certificate chain verification software **checks the appropriate Microsoft Update location for the root certificate**. If it finds it, it downloads it to the system. To the user, the experience is seamless. **The user does not see any security dialog boxes or warnings. The download happens automatically, behind the scenes** [3]."

Thus, any web browser that depends upon Microsoft's Trusted Root Store (such as Internet Explorer, Chrome and Safari for Windows) ultimately trusts 264 different CAs to issue certificates without warning, although only a handful of them are listed in the operating system's user interface. While Microsoft clearly describes this in its online developer documentation [3], no mention of this rather important design decision is made in the browser or the operating system certificate management user interface, where interested users are most likely to look.

---

[5]The former product manager for Internet Explorer told the authors that "a very few enterprises who chose to control their own trust decisions raised concerns regarding a trusted store pre-loaded with 70–100 root CAs as a potential for abuse. For this and several other reason Microsoft has since reduced the number of root certificates in the trusted store [53]."

2.3 Man in The Middle

> "Any website secured using TLS can be impersonated using a rogue certificate issued by a rogue CA. This is irrespective of which CA issued the website's true certificate and of any property of that certificate."
> — Marc Stevens *et al.* [55]

While an exhaustive explanation of man in the middle attacks against SSL is beyond the scope of this article, we at least provide a brief introduction to the subject. Over the past few years, the SSL protocol has been subject to a series of successful attacks by security researchers, some exploiting flaws in deployed systems while others made use of social engineering and other forms of deception [32, 52, 34, 43, 45].

It is because SSL protected web connections flow over a number of other insecure protocols that it is possible for attackers to intercept and hijack a connection to a SSL protected server (these are known as *man in the middle attacks*). It is only once the browser has received and verified a site's SSL certificate that the user can be sure that her connection is safe.

However, this step alone is often not enough to protect users. Sites that supply self-signed certificates, or that exploit unpatched vulnerabilities in the certificate handling code in the browsers can still trigger the display of the SSL lock icon, yet without providing the user with the associated security protections that they would normally expect.

Security researcher Moxie Marlinspike has repeatedly attacked the SSL based chain of trust, revealing exploits that leverage both browser design flaws, as well as social engineering attacks against end-users. His *sslsniff* [35] and *sslstrip* [36] tools automate the task of performing a man-in-the-middle attacks, and when supplied with a valid SSL certificate (obtained via a rogue CA for example), can be used to intercept users' communications without triggering any browser warnings.

3. BIG BROTHER IN THE BROWSER

Microsoft, Apple and Mozilla all include a number of national government CAs certificates in their respective CA databases.[6] These government CAs, like all other root CAs included by the browsers, must satisfy the requirements detailed in each browser vendor's CA policies, and are included for legitimate reasons: Many governments embed cryptographic public keys in their national ID cards, or do not wish to outsource their own internal certificate issuing responsibilities to private companies.

While it may be quite useful for Estonian users of Internet Explorer to trust their government's CA by default (thus enabling them to easily engage in secure online tasks that leverage their own national ID card), the average resident of Lebanon or Peru has far less to gain by trusting the Estonian government with the blanket power to issue SSL certificates for any website. Thus, users around the world are put in

---

[6]For example, Microsoft's Root Certificate Program includes the governments of Austria, Brazil, Finland, France, Hong Kong, India, Japan, Korea, Latvia, Macao, Mexico, Portugal, Serbia, Slovenia, Spain, Switzerland, Taiwan, The Netherlands, Tunisia, Turkey, United States and Uruguay [4].

a position where their browser entrusts their private data, indirectly, to a number of foreign governments whom those individuals may not ordinarily trust.

As an illustrative and hypothetical example of what is currently possible the Korean Information Security Agency is able to create a valid SSL certificate for the Industrial and Commercial Bank of China (whose actual certificate is issued by VeriSign, USA), that can hypothetically be used to perform an effective man-in-the-middle attack against users of Internet Explorer.

While this might at first seem like an extremely powerful attack, there are several reasons why governments are unlikely to use their own CAs to perform man in the middle attacks.

First, while *some* governments have succesfully petitioned the browser vendors to include their CA certificates, not all governments have done so. Thus, for example, the governments of Singapore, the United Kingdom and Israel (among many others) do not have state-run CAs that are included by any of the major browsers. These governments are therefore unable to create their own fake certificates for use in intelligence and other law enforcement investigations where snooping on a SSL session might be useful.

Second, due to the fact that the SSL chain of trust is *non-repudiable*, any government using its own CA to issue fake certificates in order to try and spy on someone else's communications will leave behind absolute proof of its involvement. That is, if the Spanish government opts to issue a fake certificate for Google Mail, and the surveillance is somehow discovered, anyone with a copy of the fake certificate and a web browser can independently trace the operation back to the Spanish government.

4. COMPELLED ASSISTANCE

Many governments routinely compel companies to assist them with surveillance. Telecommunications carriers and Internet service providers are frequently required to violate their customers' privacy — providing the government with email communications, telephone calls, search engine records, financial transactions and geo-location information.

In the United States, the legal statutes defining the range of entities that can be compelled to assist in electronic surveillance by law enforcement[7] and foreign intelligence investigators[8] are remarkably broad.[9] Examples of compelled assistance using these statutes include a secure email provider that was required to place a covert back door in its product in order to steal users' encryption keys [48], and a consumer electronics company that was forced to remotely enable the microphones in a suspect's auto-mobile dashboard GPS navigation unit in order to covertly record their conversations [37].

Outside of the United States, and other democratic countries, specific statutory authority may be even less important. The Chinese government, for example, has repeatedly compelled the assistance of telecommunications and technology companies in assisting it with its surveillance efforts [33, 29].

---

[7]See: 18 U.S.C. §2518(4).

[8]See: 50 U.S.C. §1805(c)(2)(B).

[9]A thorough survey of the ways in which technology firms can and have been compelled to violate their customers' privacy can be found in [51].

Just as phone companies and email providers can be forced to assist governments in their surveillance efforts, so too can SSL certificate authorities. The *compelled certificate creation attack* is thus one in which a government agency requires a domestic certificate authority to provide it with false SSL certificates for use in surveillance.

The technical details of this attack are simple, and do not require extensive explanation.[10] Each CA already has an infrastructure in place with which it is able to issue SSL certificates. In this compelled assistance scenario, the CA is merely required to skip the identity verification step in its own SSL certificate issuance process.

For the purposes of our analysis, we assume that a CA cannot refuse to comply with a lawful court order. However, it may be possible, via a *warrant canary* or a similar technique, for a CA to communicate the existence of a secret court order to the Internet community [44]. For example, a representitive from one CA has informed us that his organization's disaster contigency plans include court orders, and that his technical infrastructure includes a "kill switch" that enables him to move to a new physical location, and nullify data at the data center [39]. We do not evaluate the effectiveness of such measures in this paper.

When compelling the assistance of a CA, the government agency can either require the CA to issue it a specific certificate for each website to be spoofed, or, more likely, the CA can be forced to issue a intermediate CA certificate that can then be re-used an infinite number of times by that government agency, without the knowledge or further assistance of the CA.

In one hypothetical example of this attack, the US National Security Agency (NSA) can compel VeriSign to produce a valid certificate for the Commercial Bank of Dubai (whose actual certificate is issued by Etisalat, UAE), that can be used to perform an effective man-in-the-middle attack against users of all modern browsers.

5. SURVEILLANCE APPLIANCES

In October 2009, one of the authors of this paper attended an invitation only conference for the surveillance and lawful interception industry in Washington, DC.[11] Among the many vendor booths on the trade show floor was Packet Forensics, an Arizona based company that sells extremely small, covert surveillance devices for networks.

The marketing materials for the company's 5-series device reveal that it is a 4 square inch "turnkey intercept so-

---

[10]The legal issues relating to this kind of compelled assistance are far more complex. Any US government agencies compelling such CA assistance would almost certainly rely on the assistance provisions highlighted earlier. However, it is unclear if such compelled assistance would be lawful, due to the fact that it would interfere with the CA's ability to provide identity verification services. Such compelled assistance would also raise serious First Amendment concerns, due to to the fact that the government would be ordering the CA to affirmatively lie about the identity of a certificate recepient.

[11]The author caused national headlines in December of 2009, when he released an audio recording of one of the panel discussions at the same conference in which telecommunications company employees bragged about the extent of their cooperation with government agencies, including the extent to which they provide consumers' GPS location information [50, 61].

lution," designed for "defense and (counter) intelligence applications," capable of "packet modification, injection and replay capabilities" at Gb/sec throughput levels. The company proudly boasts that the surveillance device is perfect for the "Internet cafe problem." Most alarming is the device's ability to engage in active man-in-the-middle attacks:

> "Packet Forensics' devices are designed to be inserted-into and removed-from busy networks without causing any noticeable interruption [. . .] This allows you to conditionally intercept web, e-mail, VoIP and other traffic at-will, even while it remains protected inside an encrypted tunnel on the wire. Using 'man-in-the-middle' to intercept TLS or SSL is essentially an attack against the underlying Diffie-Hellman cryptographic key agreement protocol [. . .] To use our product in this scenario, [government] users have the ability to *import a copy of any legitimate key they obtain* **(potentially by court order)** or they can generate 'look-alike' keys designed to give the subject a false sense of confidence in its authenticity [42]."

The company has essentially packaged software equivalent to *sslstrip* into a 4 square inch appliance, ready for government customers to drop onto networks, at a price that is "so cost effective, they're disposable."

When contacted by a journalist from Wired News in March 2010, Packet Forensics spokesman Ray Saulino initially denied the product performed as advertised in its sales materials, or that anyone used it. But in a follow-up call the next day, Saulino changed his stance, telling the journalist that:

> "The technology we are using in our products has been generally discussed in internet forums and there is nothing special or unique about it [. . .] Our target community is the law enforcement community [49]."

Furthermore, while Packet Forensics has not disclosed a list of its customers, the firm's website reveals that the 5-series device was authorized for export to foreign firms and governments by the United States Bureau of Industry and Security on July 7, 2009 [41].

6. PROTECTING USERS

The major web browsers are currently vulnerable to the compelled certificate creation attack, and we do not believe that any of the existing privacy enhancing browser add-ons sufficiently protect users without significantly impacting browser usability.

In an effort to significantly reduce the impact of this attack upon end-users, we have created *Certlock*, a lightweight add-on for the Firefox browser. Our solution employs a Trust-On-First-Use (TOFU) policy (this is also known as 'leap-of-faith' authentication) [54, 10], reinforced with a policy that the country of origin for certificate issuing does not change in the future. Specifically, our solution relies upon caching CA information, that is then used to empower users to leverage country-level information in order to make common-sense trust evaluations.

In this section, we will outline the motivations that impacted the design of our solution, discuss our belief in the potential for users to make wise country-level trust decisions, and then explore the technical implementation details of our prototype add-on.

6.1 Design Motivations

The compelled certificate creation attack is a classic example of a low probability, high impact event [12]. The vast majority of users are extremely unlikely to experience it, but for those who do, very bad things are afoot. As such, it is vital that any defensive technique have an extremely low false positive rate, yet be able to get the attention of users when an attempted SSL session hijacking is detected.

Most users are unlikely to know that this threat even exists, and so it is important that any protective system not require configuration, maintenance, nor introduce any noticeable latency to users' connections. Given the low likelihood of falling victim to this attack, most rational users will avoid any protective technology that requires configuration or slows down their Web browsing [25].

Furthermore, to achieve widespread adoption (even moreso if the browser vendors are to add similar functionality to their own products), any protective technology must not sacrifice user privacy for security. Information regarding users' web browsing habits should not be leaked to any third party, even if that party is 'trusted' or if it is done so anonymously. The solution must therefore be self-contained, and capable of protecting the user without contacting any remote servers.

We believe that most consumers are unaware of how SSL functions, what a CA is, the role it performs, and how many companies are trusted by their browser to issue certificates. Expecting consumers to learn about this process, or to spend their time evaluating the business practices and trustworthiness of these hundreds of firms is unreasonable. Nevertheless, the security of the current system requires each user to make trust decisions that that they are ill equipped (nor willing) to perform.

We also believe that consumers do not directly trust CAs. Aside from the biggest CAs such as VeriSign and large telecommunications firms local to their country,[12] it is unlikely that consumers have ever heard of the vast majority of the hundreds of companies entrusted by their web browser to issue certificates. Thus, it is just as unreasonable to expect an American consumer to make a trust decision regarding a certificate issued by Polish technology firm Unizeto Technologies as it is to expect a Japanese consumer to evaluate a certificate issued by Bermuda based QuoVadis. However, both of these CAs are trusted by the major browsers, by default.

Consumers are simply told to look for the lock icon. What happens in the browser to produce that lock icon, is assumed by users to be reliable. We believe that it is our responsibility as security technologists to make sure that what happens behind the scenes does in fact protect the average users' privacy and security.

This is not to say that we think that users are clueless — merely that browsers currently provide them with little to no useful contextual information without which such complex decisions are extremely difficult.

---

[12]For example, Verizon in the United States, Deutsche Telekom in Germany or Swisscom in Switzerland.

6.2 Country-Based Trust

We believe that many consumers are quite capable of making basic trust decisions based on country-level information. We are not alone in this belief. Since March 2010, Google has been providing country-level warnings to users of its Google Mail service when it detects that their account has been accessed from a potentially suspect IP address in a different country [16].

Thus, a consumer whose banking sessions are normally encrypted by a server presenting a certificates signed by a US based CA might become suspicious if told that her US based bank is now using a certificate signed by a Tunisian, Latvian or Serbian CA.

To make this trust evaluation, she doesn't have to study the detailed business policies of the foreign CA, she can instead rely on common sense, and ask herself why her Iowa based bank is suddenly doing business in Eastern Europe. In order to empower users to make such country-level evaluations of trust, CertLock leverages the wealth of historical browsing data kept by the browser.

Individuals living in countries with laws that protect their privacy from unreasonable invasion have good reason to avoid trusting foreign governments (or foreign companies) to protect their private data. This is because individuals often receive the greatest legal protection from their own governments, and little to none from other countries. For example, US law strictly regulates the ability of the US government to collect information on US persons. However, the government can freely spy on foreigners around the world, as long as the surveillance is performed outside the US. Thus, Canadians, Swedes and Russians located outside the United States have absolutely no reason to trust the US government to protect their privacy.

Likewise, individuals located in countries with oppressive governments may wish to know if their communications with servers located in foreign democracies are suddenly being facilitated by a domestic (or state controlled) CA.

6.3 Avoiding False Positives

A simplistic defensive add-on aimed at protecting users from compelled certificate creation attacks could simply cache all certificates encountered during browsing sessions, and then warn the user any time they encounter a certificate that has changed. In fact, such an add-on, Certificate Patrol, already exists [5].

The problem with such an approach is that it is likely to suffer from an extremely high false positive rate. Each time a website intentionally changes its certificate, the browser displays a warning that will needlessly scare and soon desensitize users. There are many legitimate scenarios where certificates change. For example: Old certificates expire; certificates are abandoned and or revoked after a data breach that exposed the server private key; and many large enterprises that have multiple SSL accelerator appliances serving content for the same domain use a different certificate for each device [31].

By adopting a Trust-On-First-Use policy, we assume that if a website starts using a different certificate issued by the same CA that issued its previous certificate, there is no reason to warn the user. This approach enables us to significantly reduce the false positive rate, while having little impact on our ability to protect users from a variety of threats.

We also believe that there is little reason to warn users if a website switches CAs within the same country. As our threat model is focused on a government adversary with the power to compel any domestic CA into issuing certificates at will, we consider CAs within a country to be equals. That is, a government agency able to compel a new CA into issuing a certificate could just as easily compel the original CA into issuing a new certificate for the same site. Since we have already opted to not warn users in that scenario (described above), there is no need to warn users in the event of a same-country CA change.

By limiting the trigger of the warnings to country-level changes, we believe that we have struck a balance that will work in most situations.

6.4 Implementation Details

Our Certlock solution is currently implemented as an add-on to the Firefox browser.

The Firefox browser already retains history data for all visited websites. We have simply modified the browser to cause it to retain slightly more information. Thus, for each new SSL protected website that the user visits, a Certlock enabled browser also caches the following additional certificate information:

> A hash of the certificate.
> The country of the issuing CA.
> The name of the CA.
> The country of the website.
> The name of the website.
> The entire chain of trust up to the root CA.

When a user re-visits a SSL protected website, Certlock first calculates the hash of the site's certificate and compares it to the stored hash from previous visits. If it hasn't changed, the page is loaded without warning. If the certificate has changed, the CAs that issued the old and new certificates are compared. If the CAs are the same, or from the same country, the page is loaded without any warning. If, on the other hand, the CAs' countries differ, then the user will see a warning (See Figure 1).

At a high level, this algorithm is quite simple. However, there are a few subtle areas where some complexity is required.

Because governments can compel CAs to create both regular site certificates as well as intermediate CA certificates, any evaluation of a changed site certificate must consider the type of CA that issued it.

While the web browser vendors do not vouch for the trustworthiness of any of the root CAs that they include, we believe it is reasonable to assume that the browser vendors do at least verify the country information listed in each of their root CAs. Therefore, we are able to trust this information as we evaluate changed certificates.

When Certlock detects a changed certificate, it must also determine the type of CA that issued the new certificate. If the new certificate was issued by a root CA, then Certlock can easily compare the country of the old certificate's CA to the country of the new root CA. However, if the new certificate was issued by an intermediate CA, then we have no way of verifying that the issuing CA's country information is accurate.

As an illustrative and hypothetical example of what is currently possible, the Spanish government could compel a Spanish CA to issue an intermediate CA certificate that
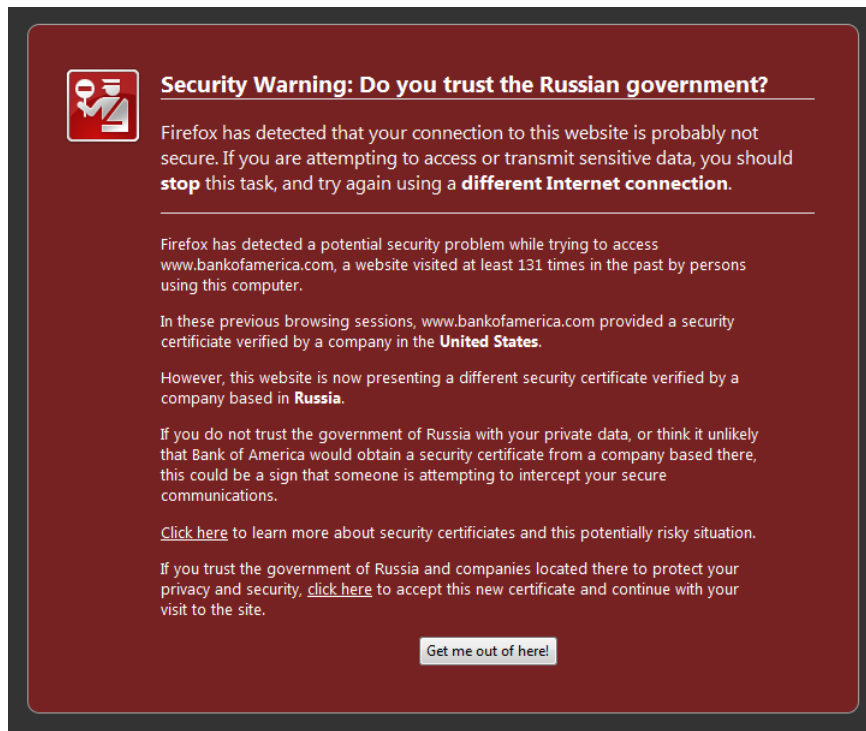
Figure 1: The warning displayed to users of Certlock.

falsely listed the country of the intermediate CA as the United States. This rogue intermediate CA would then be used to issue site certificates for subsequent surveillance activities. In this hypothetical scenario, let us imagine that the rogue CA issued a certificate for Bank Of America, whose actual certificate was issued by VeriSign in the United States. Were CertLock to simply evaluate the issuing CA's country of the previously seen Bank of America certificate, and compare it to the issuing country of the rogue intermediate CA (falsely listed as the United States), CertLock would not detect the hijacking attempt. In order to detect such rogue intermediate CAs, a more thorough comparison must be conducted.

Thus, in the event that a new certificate has been issued by an intermediate CA, Certlock follows the chain of trust up to the root CA, noting the country of every CA along the path. If any one of these intermediate CAs (or the root CA itself) has a different country than the CA that issued the original certificate, then the user is warned.

7. THREAT MODEL ANALYSIS

In this section, we outline several *hypothetical* scenarios in which a man-in-the-middle attack may be desired. In each example scenario, we examine the government's available surveillance options, consider the suitability of the compelled certificate creation attack, and evaluate the ability of CertLock to detect and thwart the attack. A condensed summary of the threats that CertLock defends against is also presented in Figure 2.

7.0.1 Scenario A

Actual CA	VeriSign (USA)
Compelled CA	VeriSign (USA)
Website	Citibank (USA)
Location of Suspect	USA
Spying Government	USA

In this scenario, the United States government compels VeriSign to issue a certificate for use by a law enforcement agency wishing to spy on communications between a suspect located in the United States and Citibank, her United States based bank.

This attack is impossible for CertLock to detect, because the CA issuing the fake certificate is also the same that issued the legitimate certificate. However, we believe that this scenario is extremely unlikely to occur in the investigations of end users. This is because if a government adversary is able to obtain a court order compelling VeriSign's cooperation, it can just as easily obtain a court order compelling Citibank to disclose the suspect's account information.

While there are perhaps a few volunteer run Internet providers that will do anything possible to avoid delivering user data to government agents, we believe that the vast majority of corporations will eventually comply. Outright refusal could potentially result in seizure of corporate assets, and the jailing of executives—consequences that profit focused shareholders would likely wish to avoid. As a related example, in 2006, Google very publicly fought a subpoena from the US Deparment of Justice requesting aggregate search request records. However, once a court ruled on the matter, the company complied and provided the government with 50,000 URLs from the Google search index [59]. As such,

Spying Government	Country of Actual CA	CertLock Protects?
X	X	No
X	Y	Yes

Figure 2: A trust matrix evaluating CertLock. In short, the tool only protects users from compelled certificate creation attacks when the Spying Government and the Country of the Actual CA are not the same.

our threat model specifically excludes the rare category of ISPs willing to say no to government requests at all costs, and instead focuses on typical, law-abiding corporations that provide services to most users.

7.0.2 Scenario B

Actual CA	VeriSign (USA)
Compelled CA	GoDaddy (USA)
Website	Citibank (USA)
Location of Suspect	USA
Spying Government	USA

In this scenario, the United States government compels GoDaddy, a CA located in the United States to issue a certificate for an intelligence agency wishing to spy on communications between a suspect located in the United States and a bank also located in the United States (CitiBank), which obtained its legitimate SSL certificate from VeriSign.

Just as with Scenario A, this attack is extremely unlikely to occur. This is because any government agency able to compel GoDaddy is also capable of obtaining a court order to compel VeriSign or Bank of America. By simple reduction, any attacker capable of Scenario B is also capable of Scenario A. CertLock does not detect attacks of this type.

7.0.3 Scenario C

Actual CA	VeriSign (USA)
Compelled CA	VeriSign (USA)
Website	*Poker.com* (USA)
Location of Suspect	USA
Spying Government	USA

In this scenario, US law enforcement agents are investigating a US-based online gambling website and the US-based users of the service. The agents wish to first obtain evidence that illegal activity is occuring, by monitoring the bets as they are placed via SSL encrypted sessions, before they later raid the offices of the company and seize their servers. In order to surveil the communications between users and the gambling website, law enforcement officials compel VeriSign to issue an additional certificate for the site, which is then used to intercept all communications to and from the website.

In this scenario, where both ends of the SSL connection are under investigation by the government, the compelled certificate attack is a highly effective method for covertly gathering evidence. However, because the issuing CA does not change, CertLock is unable to detect this attack and warn users.

In general, attack scenarios in which both the end-user and the website are under surveillance are beyond the scope of our threat model.

7.0.4 Scenario D

Actual CA	VeriSign (USA)
Compelled CA	TeliaSonera (Finland)
Website	Aktia Bank (Finland)
Location of Suspect	Finland
Spying Government	Finland

In this scenario, a resident of Finland is accessing her Aktia Savings Bank online account, which obtained its legitimate SSL certificate from VeriSign, a US firm. The Finnish intelligence services are interested in getting access to the suspect's online transaction data, and thus seek to compel TeliaSonera, a domestic CA to issue a certificate for the surveillance operation.

This scenario is not identical to scenario A, however it is quite similar. Again, if the Finnish government is able to compel a domestic CA into assisting it, we assume that it could just as easily compel the Finnish bank into providing the suspect's account details. While we believe that this attack scenario is unlikely, should it occur, CertLock will detect it.

7.0.5 Scenario E

Actual CA	VeriSign (USA)
Compelled CA	TeliaSonera (Finland)
Website	Google Mail (USA)
Location of Suspect	Finland
Spying Government	Finland

In this scenario, a US executive is travelling in Finland for business, and is attempting to access her secure, US-based webmail account using the Internet connection in her hotel room. Finnish authorities wish to intercept her communications, but due to Google's use of SSL by default for all webmail communications [46], the government must employ a man-in-the-middle attack. This scenario is thus an ideal candidate for a compelled certificate creation attack, since the Finnish authorities have no leverage to compel the assistance of Google or VeriSign. This scenario is also one that is easily detected by CertLock.

7.0.6 Scenario F

Actual CA	VeriSign (USA)
Compelled CA	VeriSign (USA)
Website	CCB (China)
Location of Suspect	USA
Surveilling Government	USA

In this scenario, a Chinese executive is travelling in the United States for business, and is attempting to acccess her China Construction Bank account using the Internet connection in her hotel room. US Government authorities wish to get access to her financial records, but are unwilling to

let the Chinese government know that one of their citizens is under investigation, and so have not requested her records via official law enforcement channels.

This scenario is almost identical to scenario E, however, there is one key difference: The legitimate certificate used by the Chinese bank was issued by a CA located in the United States and the US government has turned to the same US based CA to supply it with a false certificate. Thus, while this scenario is an ideal candidate for a compelled certificate creation attack, it is not one that can easily be detected by looking for country-level CA changes. As such, CertLock is not able to detect attacks of this type.

7.1 Why Sites Should Consider the Country of the CA They Use

Building on the information presented thus far in this paper, we can draw the following conclusions:

- Users are currently vulnerable to compelled certificate creation attacks initiated by the government of any country in which there is at least one certificate authority that is trusted (directly or indirectly) by the browser vendors.

- When users provide their private data to a company, the government of the country in which their data is located may be able to compel the provider to disclose their private data.

- When users provide their private data to a company that holds the data in country X, but uses a SSL certificate provided by a CA in country Y, users are vulnerable to both the compelled disclosure of their data by the government of country X, and interception of their private data through a compelled certificate creation attack by country Y.

- Thus, when a company that uses a certificate authority located in a country different than the one in which it holds user data, it needlessly exposes users' data to the compelled disclosure by an additional government.

It is based on this that we believe that websites best serve their users when they rely on a SSL certificate from a CA located in the same country in which their private data is stored.[13] Unfortunately, this is not a widespread practice in the industry; instead American CAs totally dominate the certificate market, and are used by many foreign organizations.

As just one example — a number of the big banks in Pakistan, Lebanon and Saudi Arabia (countries in which the US has a strong intelligence interest) all use certificates obtained from US-based CAs to secure their online banking sites.

It is because of the dominance of US CAs that CertLock is not able to equally protect users from different countries. Certlock can effectively protect users of US based services from compelled certificate disclosure attacks performed by non-US governments. Thus, it is useful for Americans travelling out of the country who may be subject to surveillance by the national government of the country in which they are travelling, and non-US persons who use US-based services and who do not wish for their own governments to get access to their data.

However, as long as companies around the world continue to rely on SSL certificates issued by American CAs, the US government will maintain the ability to perform man in the middle attacks that are practically impossible to detect with CertLock or any other country based detection mechanism.

8. RELATED WORK

Over the past decade, many people in the security community have commented on the state of the SSL public key infrastructure, and the significant trust placed in the CAs [20, 11, 19].

In 1998, James Hayes of the US National Security Agency published a paper that focused specifically on the threat of rogue insiders within a Certificiate Authority. Although the technical details of the threat outlined by Hayes are largely the same as the scenario on which we have focused (albeit with vasty different legal and policy consequences), Hayes did not address the threat of government compelled certificate creation. It is unclear if he was simply unaware of this scenario, or if the topic was too sensitive for him to discuss, given his employer. In his paper, Hayes proposed a technical solution to address the insider threat, which relied on users configuring various per-site attributes within their browser that would be used to evaluate each new site's certificate.

Crispo and Lomas also proposed a certification scheme designed to detect rogue CAs [14], while the Monkeysphere project has created a system that replaces the CA architecture with the OpenPGP web of trust [6].

Ian Grigg has repeatedly sought to draw attention to both the potential conflict of interest that some CAs have due to their involvement in other forms of surveillance, and the power of a court order to further compel these entities to assist government investigations [21, 22, 23]. In particular, in 2005, Grigg and Shostack filed a formal complaint with ICANN over the proposal to award VeriSign control of .net domain name registration, arguing that the firm's surveillance products created a conflict of interest [24].

In recent years, several browser-based tools have been created to help protect users against SSL related attacks. Kai Engert created Conspiracy, a Firefox add-on that provides country-level CA information to end-users in order to protect them from compelled certificate creation attacks. The Conspiracy tool displays the flag of the country of each CA in the chain of trust in the browser's status bar [17]. Thus, users must themselves remember the country of the CAs that issue each certificate, and detect when the countries have changed. We believe, like Herley [25], that this is an unreasonable burden to place upon end-users, considering how rarely the compelled certificate creation attack is likely to occur.

Wendlandt *et al.* created Perspectives, a Firefox add-on that improves the Trust-On-First-Use model used for websites that supply self-signed SSL certificates [58]. In their system, the user's browser securely contacts one of several notary servers, who in turn independently contact the webserver and obtain its certificate. In the event that an attacker is attempting to perform a man in the middle attack upon the user, the fact that the attacker-supplied SSL certificate, and those supplied by the Perspectives notary

---

[13]For example, all of the Hungarian banks surveyed by the authors use certificates provided by NetLock Ltd., a Hungarian CA.

servers differ will be a strong indicator that something bad has happened.

Unfortunately, the Perspectives system requires that users provide the Perspectives notaries with a real-time list of the secure sites they visit.[14] Although the scheme's designers state that "all servers adhere to a strict policy of never recording client IP addresses, period," we still don't think it is a good idea to provide users' private web browsing data to a third party, merely based on the fact that they promise not to log it.

Alicherry and Keromytis have improved upon the Perspectives design with their DoubleCheck system [8], substituting Tor exit nodes for special notary servers. Because the Tor network anonymizes the individual user's IP address, there is no way for the Tor exit nodes to know who is requesting the certificate for a particular SSL website. While the authors solved the major privacy issues that plague the Perspectives scheme, their choice of Tor carries its own cost: Latency. Their system adds an additional second of latency to every new SSL connection, and up to 15 seconds for visits to new self-signed servers. We believe that this additional latency is too much to ask most users to bear, particularly if the chance of them encountering a rogue CA is so low.

Herzberg and Jbara created TrustBar, a Firefox add-on designed to help users detect spoofed websites. The browser tool works by prominently displaying the name of the CA that provided the site's certificate, as well as allowing the user to assign a per-site name or logo, to be displayed when they revisit to each site [26].

Tyler Close created Petname Tool, a Firefox add-on that caches SSL certificates, and allows users to assign a per-site phrase that is displayed each time they revisit the site in the future. In the event that a user visits a spoofed website, or a site with the same URL that presents a certificate from a different CA, the user's specified phrase will not be displayed [13].

In May 2008, a security researcher discovered that the OpenSSL library used by several popular Linux distributions was generating weak cryptographic keys. While the two-year old flaw was soon fixed, SSL certificates created on computers running the flawed code were themselves open to attack [7, 60]. Responding to this flaw, German technology magazine Heise released the Heise SSL Guardian for the Windows operating system, which warns users of Internet Explorer and Chrome when they encounter a weak SSL certificate [57].

In December 2008, Stevens *et al.* demonstrated that flaws in the MD5 algorithm could be used to create rogue SSL certificates (without the knowledge or assistance of the CA). In response, CAs soon accelerated their planned transition to certificates using the SHA family of hash functions [55]. As an additional protective measure, Márton Anka developed an add-on for the Firefox browser to detect and warn users about certificate chains that use the MD5 algorithm for RSA signatures [9].

Jackson and Barth devised the ForceHTTPS system to protect users who visit HTTPS protected websites, but who

are vulnerable to man in the middle attacks due to the fact that they do not type in the `https://` component of the URL [28]. This system has since been formalized into the Strict Transport Security (STS) standard proposal [27], to which multiple browsers are in the process of adding support. While this system is designed to enable a website to hint to the browser that future visits should always occur via a HTTPS connection, this mechanism could be extended to enable a website to lock a website to a particular CA, or CAs of a specific country.

9. CONCLUSION AND FUTURE WORK

In this paper, we introduced the compelled certificate creation attack and presented evidence that suggests that governments may be subverting the CA based public key infrastructure. In an effort to protect users from these powerful adversaries, we introduced a lightweight defensive browser based add-on that detects and thwarts such attacks. Finally, we use reductive analysis of governments' legal capabilities to perform an adversarial threat model analysis of the attack and our proposed defensive technology.

Our browser add-on is currently just a prototype, and we plan to improve it in the future. First, our currently used warning dialog text is far from ideal, and could be greatly improved with the help of usability and user experience experts. We also plan to explore the possibility of expanding the country-level trust model to regions, such as the European Union, where, for example, residents of France may be willing to trust Spanish CAs. Finally, We are considering adding a feature that will enable users to voluntarily submit potentially suspect certificates to a central server, so that they can be studied by experts. Such a feature, as long as it is opt-in, does not collect any identifiable data on the user, and only occurs when potentially rogue certificates are discovered, would have few if any privacy issues.

Ultimately, the threats posed by the compelled certificate creation attack cannot be completely eliminated via our simple browser add-on. The CA system is fundamentally broken, and must be overhauled. DNSSEC may play a significant role in solving this problem, or at least reducing the number of entities who can be compelled to violate users' trust. No matter what system eventually replaces the current one, the security community must consider compelled government assistance as a realistic threat, and ensure that any solution be resistant to such attacks.

10. ACKNOWLEDGEMENTS

11. REFERENCES

[1] Apple Root Certificate Program. `www.apple.com/certificateauthority/ca_program.html`.

[2] Mozilla CA Certificate Policy (Version 1.2). `www.mozilla.org/projects/security/certs/policy/`.

[3] Microsoft Root Certificate Program, January 15 2009. `technet.microsoft.com/en-us/library/cc751157.aspx`.

[4] Windows Root Certificate Program Members, November 24 2009. `download.microsoft.com/`

---

[14]Modern browsers already leak information about the secure web sites that users visit, as they automatically contact CAs in order to verify that the certificates have not been revoked (using the OCSP protocol). While this is currently unavoidable, we wish to avoid providing private user web browsing data to any additional parties.

download/1/4/f/14f7067b-69d3-473a-ba5e-70d04aea5929/windows\%20root\%20certificate\%20program\%20members.pdf.

[5] Certificate patrol, 2010. patrol.psyced.org/.

[6] Monkeysphere, 2010. web.monkeysphere.info/.

[7] D. Ahmad. Two Years of Broken Crypto: Debian's Dress Rehearsal for a Global PKI Compromise. *IEEE Security and Privacy*, 6:70–73, September 2008.

[8] M. Alicherry and A. D. Keromytis. Doublecheck: Multi-path verification against man-in-the-middle attacks. In *ISCC 2009: IEEE Symposium on Computers and Communications*, pages 557–563, Piscataway, NJ, USA, 2009. IEEE.

[9] M. Anka. SSL Blacklist 4.0, January 31 2010. www.codefromthe70s.org/sslblacklist.aspx.

[10] J. Arkko and P. Nikander. Weak authentication: How to authenticate unknown principals without trusted parties. In B. Christianson, B. Crispo, J. A. Malcolm, and M. Roe, editors, *Security Protocols Workshop*, volume 2845 of *Lecture Notes in Computer Science*, pages 5–19. Springer, 2002.

[11] P. S. Bance. Ssl: Whom do you trust?, April 20 2005. www.minstrel.org.uk/papers/2005.04.20-ssl-trust.pdf.

[12] M. Bussiere and M. Fratzscher. Low probability, high impact: Policy making and extreme events. *Journal of Policy Modeling*, 30(1):111–121, 2008.

[13] T. Close. Petname tool, 2005. www.waterken.com/user/PetnameTool/.

[14] B. Crispo and M. Lomas. A certification scheme for electronic commerce. In *In Security Protocols International Workshop*, page pages. Springer-Verlag, 1996.

[15] T. Dierks and C. Allen. The TLS Protocol Version 1.0. RFC 2246 (Proposed Standard), Jan. 1999. Obsoleted by RFC 4346, updated by RFCs 3546, 5746.

[16] P. Diwanji. Detecting suspicious account activity. *The Official Gmail Blog*, March 24 2010. gmailblog.blogspot.com/2010/03/detecting-suspicious-account-activity.html.

[17] K. Engert. Conspiracy — A Mozilla Firefox Extension, March 18 2010. kuix.de/conspiracy/.

[18] E. Felten. Web Certification Fail: Bad Assumptions Lead to Bad Technology. *Freedom To Tinker*, February 23 2010. www.freedom-to-tinker.com/blog/felten/web-certification-fail-bad-assumptions-lead-bad-technology.

[19] E. Gerck. First published online by the MCWG at http://mcwg.org/cert.htm (April, 1997). Invited talk at the Black Hat Briefings '99, Las Vegas, NV, July 7-8, 1999. Published by The Bell, ISSN 1530-048X, Vol. 1, No. 3, p. 8, July 2000, and available online at http://www.thebell.net/papers/certover.pdf.

[20] D. K. Gillmor. Technical Architecture shapes Social Structure: an example from the real world, February 21 2007. lair.fifthhorseman.net/~dkg/tls-centralization/.

[21] I. Grigg. VeriSign's conflict of interest creates new threat. *Financial Cryptography (blog)*, September 1 2004. financialcryptography.com/mt/archives/000206.html.

[22] I. Grigg. PKI considered harmful. October 14 2008. iang.org/ssl/pki_considered_harmful.html.

[23] I. Grigg. Why the browsers must change their old SSL security (?) model. *Financial Cryptography (blog)*, March 24 2010. financialcryptography.com/mt/archives/001232.html.

[24] I. Grigg and A. Shostack. VeriSign and Conflicts of Interest, February 2 2005. forum.icann.org/lists/net-rfp-verisign/msg00008.html.

[25] C. Herley. So long, and no thanks for the externalities: the rational rejection of security advice by users. In *NSPW '09: Proceedings of the 2009 workshop on New security paradigms workshop*, pages 133–144, September 2009.

[26] A. Herzberg and A. Jbara. Security and identification indicators for browsers against spoofing and phishing attacks. *ACM Trans. Internet Technol.*, 8(4):1–36, 2008.

[27] J. Hodges, C. Jackson, and A. Barth. Strict Transport Security, December 18 2009. lists.w3.org/Archives/Public/www-archive/2009Dec/att-0048/draft-hodges-strict-transport-sec-06.plain.html.

[28] C. Jackson and A. Barth. Forcehttps: protecting high-security web sites from network attacks. In *WWW '08: Proceeding of the 17th international conference on World Wide Web*, pages 525–534, New York, NY, USA, 2008. ACM.

[29] A. Jacobs. China requires censorship software on new pcs. *The New York Times*, June 8 2009. www.nytimes.com/2009/06/09/world/asia/09china.html.

[30] D. Kaminsky. Black Ops of PKI. *26th Chaos Communication Congress*, December 29 2009. events.ccc.de/congress/2009/Fahrplan/events/3658.en.html.

[31] D. Kaminsky. Email conversation with author, February 28 2010.

[32] D. Kaminsky, M. L. Patterson, and L. Sassaman. PKI Layer Cake: New Collision Attacks Against the Global X.509 Infrastructure. In *Proceedings of Financial Cryptography and Data Security - 14th International Conference (FC 2010)*, 2010.

[33] J. Markoff. Surveillance of skype messages found in china. *The New York Times*, October 1 2008. www.nytimes.com/2008/10/02/technology/internet/02skype.html.

[34] M. Marlinspike. More Tricks For Defeating SSL In Practice. *BlackHat USA*, 2009. www.blackhat.com/presentations/bh-usa-09/MARLINSPIKE/BHUSA09-Marlinspike-DefeatSSL-SLIDES.pdf.

[35] M. Marlinspike. sslsniff, July 3 2009. www.thoughtcrime.org/software/sslsniff/.

[36] M. Marlinspike. sslsniff, December 18 2009. www.thoughtcrime.org/software/sslstrip/.

[37] D. McCullagh. Court to FBI: No spying on in-car computers. *CNET News*, November 19 2003. news.cnet.com/2100-1029_3-5109435.html.

[38] Mozilla. Potentially problematic CA practices , 2010. wiki.mozilla.org/CA:Problematic_Practices.

[39] E. Nigg. Email conversation with author, March 27 2010.

[40] J. Nightingale. SSL Question Corner. *meandering wildly (blog)*, August 5 2008. `blog.johnath.com/2008/08/05/ssl-question-corner/`.

[41] Packet Forensics. Export and Re-Export Requirements, 2009. `www.packetforensics.com/export.safe`.

[42] Packet Forensics. Marketing materials for 5-series products, 2009. `http://files.cloudprivacy.net/packet-forensics-materials.pdf`.

[43] M. Ray and S. Dispensa. Renegotiating TLS, November 4 2009. `extendedsubset.com/wp-uploads/2009/11/renegotiating_tls_20091104_pub.zip`.

[44] rsync.net. rsync.net warrant canary, 2010. `http://www.rsync.net/resources/notices/canary.txt`.

[45] S. E. Schechter, R. Dhamija, A. Ozment, and I. Fischer. The emperor's new security indicators. In *SP '07: Proceedings of the 2007 IEEE Symposium on Security and Privacy*, pages 51–65, Washington, DC, USA, 2007. IEEE Computer Society.

[46] S. Schillace. Default https access for Gmail. *The Official Gmail Blog*, January 12 2010. `gmailblog.blogspot.com/2010/01/default-https-access-for-gmail.html`.

[47] A. Shamir. Cryptography: State of the science. *ACM A. M. Turing Award Lecture*, June 8 2003. `awards.acm.org/images/awards/140/vstream/2002/S/s-pp/shamir_1files_files/800x600/Slide8.html`.

[48] R. Singel. PGP Creator Defends Hushmail. *Wired News Threat Level Blog*, November 19 2007. `www.wired.com/threatlevel/2007/11/pgp-creator-def`.

[49] R. Singel. Law Enforcement Appliance Subverts SSL. *Wired News Threat Level Blog*, March 24 2010. `www.wired.com/threatlevel/2010/03/packet-forensics/`.

[50] C. Soghoian. 8 Million Reasons for Real Surveillance Oversight. *Slight Paranoia blog*, December 1 2009. `paranoia.dubfire.net/2009/12/8-million-reasons-for-real-surveillance.html`.

[51] C. Soghoian. Caught in the cloud: Privacy, encryption, and government back doors in the web 2.0 era. In *Journal on Telecommunications and High Technology Law*, Forthcoming.

[52] A. Sotirov and M. Zusman. Breaking the Security Myths of Extended Validation SSL Certificates. *BlackHat USA*, 2009. `www.blackhat.com/presentations/bh-usa-09/SOTIROV/BHUSA09-Sotirov-AttackExtSSL-SLIDES.pdf`.

[53] C. Spiezle. Email conversation with author, February 15 2010.

[54] F. Stajano and R. J. Anderson. The resurrecting duckling: Security issues for ad-hoc wireless networks. In *Proceedings of the 7th International Workshop on Security Protocols*, pages 172–194, London, UK, 2000. Springer-Verlag.

[55] M. Stevens, A. Sotirov, J. Appelbaum, A. Lenstra, D. Molnar, D. A. Osvik, and B. Weger. Short chosen-prefix collisions for MD5 and the creation of a rogue CA certificate. In *Proceedings of the 29th Annual International Cryptology Conference on Advances in Cryptology*, pages 55–69, Berlin, Heidelberg, 2009. Springer-Verlag.

[56] J. Sunshine, S. Egelman, H. Almuhimedi, N. Atri, and L. F. Cranor. Crying wolf: An empirical study of SSL warning effectiveness. In *Proceedings of the 18th Usenix Security Symposium*, August 2009.

[57] The H Security. heise SSL Guardian: Protection against unsafe SSL certificates, July 4 2008. `www.h-online.com/security/features/Heise-SSL-Guardian-746213.html`.

[58] D. Wendlandt, D. G. Andersen, and A. Perrig. Perspectives: improving ssh-style host authentication with multi-path probing. In *ATC'08: USENIX 2008 Annual Technical Conference on Annual Technical Conference*, pages 321–334, Berkeley, CA, USA, 2008. USENIX Association.

[59] N. Wong. Judge tells DOJ "No" on search queries. *The Official Google Blog*, March 17 2006. `googleblog.blogspot.com/2006/03/judge-tells-doj-no-on-search-queries.htmll`.

[60] S. Yilek, E. Rescorla, H. Shacham, B. Enright, and S. Savage. When private keys are public: results from the 2008 Debian OpenSSL vulnerability. In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*, pages 15–27, New York, NY, USA, 2009. ACM.

[61] K. Zetter. Feds 'Pinged' Sprint GPS Data 8 Million Times Over a Year. *Wired News Threat Level Blog*, December 1 2009. `www.wired.com/threatlevel/2009/12/gps-data/`.