Why Eve and Mallory (Also) Love Webmasters

A Study on the Root Causes of SSL Misconfigurations

Sascha Fahl, Yasemin Acar, Henning Perl, Matthew Smith
FKIE, Fraunhofer
Bonn, Germany
{sascha.fahl,yasemin.acar,henning.perl,matthew.smith}@fkie.fraunhofer.de

ABSTRACT

Previous research showed that the SSL infrastructure is a fragile system: X.509 certificate validation fails for a non-trivial number of HTTPS-enabled websites resulting in SSL warning messages presented to users. Studies revealed that warning messages do not provide easy-to-understand information or are ignored by webbrowser users. SSL warning messages are a critical component in the HTTPS infrastructure and many attempts have been made to improve these warning messages. However, an important question has not received sufficient attention yet: Why do webmasters (deliberately) deploy non-validating, security-critical X.509 certificates on publicly available websites? In this paper, we conduct the first study with webmasters operating non-validating X.509 certificates to understand their motives behind deploying those certificates. We extracted the non-validating certificates from Google's webcrawler body of X.509 certificates, informed webmasters about the problem with the X.509 certificate configuration on their website and invited a random sample of the respective webmasters to participate in our study. 755 webmasters participated, allowing us insight into their motives. While one third of them admitted to having misconfigured their webserver accidentally, two thirds of them gave reasons for deliberately using a non-validating X.509 certificate.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous

General Terms

SSL, Webmasters, Usable Security, User Study

1. INTRODUCTION

For the authentication of a server during an SSL handshake, clients perform multiple validation steps to check whether the server's X.509 certificate is trustworthy or should better be rejected. Self-signed and expired certificates, certificates that were signed by an unknown CA, certificates that are not delivered with a complete issuer chain as well as certificates issued for the wrong hostname result in rejection. Whenever there is a problem with X.509 certificate validation, modern webbrowsers generate warning messages so users can decide how to proceed with the (possibly) critical X.509 certificate in question.

Further research revealed that misconfigurations of HTTPS-enabled webservers are a widespread issue. In 2009, the EFF conducted an internet-wide scan of all public IPv4 addresses on port 443 and collected the respective X.509 certificates[1]. Since then, multiple projects crawled the public part of the Internet for HTTPS certificates and analyzed different aspects of their deployment in the wild. Holz et al. [5], Akhawe et al. and the ICSI Certificate Notary [1] collected X.509 certificates either actively or passively and concluded certain properties of the current CA-based SSL infrastructure: While previous research provides valuable insights into the current SSL ecosystem, their focus is solely on technical aspects of SSL configurations or on the behavior of webbrowser users confronted with SSL warning messages but leaves out the following interesting questions: (1) Why are HTTPS-enabled websites operated with non-validating X.509 certificates at all? (2) How many misconfigured websites are frequently visited with webbrowsers and hence throw SSL warning messages? (3) For how many users do SSL warning messages occur unexpectedly? Based on the knowledge that a non-negligible percentage of SSL handshakes fail and that a large percentage of warning messages is dismissed by users[2], we conducted the first qualitivate study with website administrators to investigate the root causes for X.509 misconfiguration that cause browser warning messages. We collected 755 study results to assess the motivation for the use of non-validating X.509 certificates on the web. We were interested in the reasons and motivation for webmasters to operate non-validating X.509 certificates, how these webmasters assess the operation of their non-validating certificates and the number of affected users and the protected data types. Additionally, we were interested in suggestions to improve the usability of certificate configuration.

Our findings suggest that a remarkable number of websites that operate non-validating X.509 certificates either do so intentionally or are not actively in use and hence do not trigger warning messages at all. However, we also find that many administrators misconfigure their HTTPS webservers due to the high complexity of SSL configuration options or

ASIA CCS'14, June 4–6, 2014, Kyoto, Japan.
Copyright 2014 ACM 978-1-4503-2800-5/14/06 ...$15.00.
http://dx.doi.org/10.1145/2590296.2590341.

[1] https://www.eff.org/observatory

due to a misunderstanding of the security features of SSL. Our contributions can be summarized as follows: (1) We conduct the first user-study with webmasters of HTTPS-enabled websites to identify the root causes for SSL warning messages in modern webbrowsers. (2) We find that a large amount of non-validating certificates is meant to be that way and clicking through them can be classified as deliberate. (3) We find that mainly websites with a manageable user count throw certificate validation errors. In many of these cases the users were previously informed, have probably been helped with the installation of the respective CA or trust the certificate and thus are not shown a warning message when browsing the site. (4) We report that a remarkable amount of websites employing non-validating certificates as can be found by crawlers are not meant to be actively used, are only ever accessed by crawlers and thus do not trigger real world warning messages for users. (5) We find that a substantial number of webmasters are overwhelmed by the complexity of SSL and the configuration parameters offered by HTTPS webservers. (6) We provide a list of suggestions to improve the usability of X.509 certificate configuration on webservers given by the webmasters.

2. BACKGROUND

The SSL protocol provides authentication based on the X.509 public key infrastructure[2], protects data confidentiality using symmetric encryption, and ensures data integrity with cryptographic message digests. SSL is commonly used to secure websites and mail servers, prevent network attackers from eavesdropping or replaying the client's messages, and is generally considered security best practice for websites using the HTTPS protocol[3].

2.1 Browser Certificate Validation

Basically X.509 certificate validation in browser software consists of the following steps: (1) check if the certificate was digitally signed by a trusted issuer, (2) check if the certificate is not already expired, (3) check if the certificate was issued for the requested hostname.
(1) This step verifies whether a server's certificate was issued by a trusted CA of which modern browsers usually include more than 100. Common reasons for failures in practice are X.509 certificates that were self-signed or signed by unknown CAs, e. g. CAs used in an enterprise context. (2) An X.509 certificate has a validity period of typically 12 or 24 months. In this step browsers check if the certificate was delivered within its validity period and otherwise reject it. (3) Finally, browsers check if the given certificate's common name or subject alternative names match the server's hostname. A widespread real world reason for hostname mismatches are certificates used for hostnames without the `www.` prefix although the certificate was only issued for the prefix or vice versa.

2.1.1 *SSL Warning Messages*

In case X.509 certificate validation fails, modern browsers show their users warning messages. These warning messages might imply that a Man-In-The-Middle attack occurred, i. e. that an attacker exchanged the server's original certificate with their own, to eavesdrop or alter data sent between the client and server. Another reason for failures are misconfigured servers, e. g. a webmaster did not update an already expired certificate or operates a certificate for an invalid hostname. In case browsers are certain that an attack is occuring, the presented warning message cannot be bypassed. If there is a chance that the warning message is a false positive (i. e. the website's administrator (deliberately) misconfigured the X.509 certificate), browsers will show a bypassable warning message, discouraging users from clicking through.

3. RELATED WORK

To the best of our knowledge, we present the first study with webmasters on the root causes of misconfigured X.509 certificates for HTTPS webservers. Although no other studies focused on the same topic, the motivation for our work is built on related work which we will discuss in this section. As mentioned above, several projects either actively or passively measured – or, respectively, are still measuring – the condition of the HTTPS ecosystem. Akhawe et al. [1] passively collected SSL handshakes of multiple US universities and 300,000 users over a period of nine months in 2012 and 2013, concentrating on the frequency of X.509 certificate validation errors in SSL handshakes. Overall, they found that 98.46% of the 3.9 billion SSL handshakes they monitored validated correctly, while 1.54% failed for different reasons: 70.51% used an unknown issuer, 2.99% a self-signed certificate, in 7.65% of all handshakes the certificate was expired and 18.82% of all handshakes generated hostname validation errors. Due to the unlikeliness of an actual Man-In-The-Middle attack, they assume all validation errors to be false positives. Durumeric et al. [3] presented ZMap – a fast internet-wide scanner – and conducted 110 scans of the world-wide HTTPS infrastructure over one year, collecting more than 42 million unique certificates of which 6.9 million were browser trusted. In 2009, Sunshine et al. [6] conducted lab studies with over 400 internet users to evaluate the effectiveness of browser SSL warning messages, as well as their human understandability, finding that participants made unsafe choices when confronted with warning messages. They suggest reducing the number of warning messages altogether, taking the decision whether to trust an unsafe connection or not out of the users' hands. In 2013, Akhawe and Porter Felt [2] used Firefox and Chrome's telemetry feature to measure click-through rates for SSL (and other) warning messages for different browsers in situ. Over a period of two months, they collected 16,704,666 SSL warning impressions for Chrome and 10,976 for Firefox. However, they were not able to see the respective handshakes or certificates that led to the warnings, thus they assume that almost all warning messages they saw were false positives.

4. WEBMASTER STUDY

While previous research either focused on a technological analysis of the deployed X.509 certificates in the wild or evaluated the users' behaviour when faced with an SSL warning message, our work incorporates the third important pillar in the SSL infrastructure: the webmasters of HTTPS-enabled websites. Knowing the technical reasons why SSL handshakes fail and produce warning messages and how users react to those warning messages are important aspects. However, to achieve a better understanding of the whole picture,

[2] https://www.ietf.org/rfc/rfc5280.txt

[3] https://tools.ietf.org/html/rfc2818

we conduct the first study with webmasters to assess the root causes why webmasters operate non-validating X.509 certificates.

4.1 Methodology

To find websites that operate non-validating X.509 certificates, we gathered certificates deployed in the wild in a first step. We applied a technique different from previous work to collect X.509 certificates from websites: We used the body of certificates Google's webcrawler collected over a period of 12 months. The webcrawler collected X.509 certificates for 55,675,334 ($\sim$ 55.7 million) different hosts (identified as different domain names), resulting in a body of 4,487,463 X.509 certificates and their corresponding hostnames. This certificate body overcomes two essential problems common to other approaches reported in literature: (1) Actively crawling X.509 certificates for the complete IPv4 space such as Holz et al.[5] and the EFF SSL observatory resulted in a comprehensive map of IPv4 addresses and corresponding X.509 certificates - in this case one cannot deduce for which hostname the certificate was configured, hence post-validation does not allow for hostname verification. (2) Passively recording X.509 certificates similar to Akhawe et al. [1] only collects X.509 certificates for the websites their users visit - although they collected both X.509 certificates and SNI[4] values for the corresponding SSL handshakes, they might have missed an essential part of the HTTPS-enabled part of the Internet. The X.509 certificate body of Google's webcrawler provides both, an encompassing list of X.509 certificates and their corresponding hostnames of the publicly available part of the Internet and the possibility to perform all three steps of X.509 certificate validation in postprocessing steps: (1) CA signature validation, (2) expiration checks and (3) hostname verification. We used the webcrawler's certificate body and performed the following steps to select candidates for our study: Firstly, we re-validated all X.509 certificates, using the NSS library as proposed by Akhawe et al. [1] which gave us the following results:

Error Type	#Certificates	
Valid	3,876,497	(86.38%)
Self-Signed	89,981	(2.0%)
Expired	309,350	(6.89%)
Hostname Mismatch	146,941	(3.27%)
Unknown Issuer	64,694	(1.44%)

Altogether, our re-validation left us with 610,966 X.509 certificates that generate warning messages when users visit the corresponding websites. We picked a random sample of 50,000 of all failed X.509 certificates [5] and subsequently revisited all websites of our 50,000 certificate sample to learn the current SSL configuration status of the corresponding webserver. This left us with 46,934 X.509 certificates and their corresponding webservers.

We decided to get in contact with all of the affected webmasters. Therefore we started by extracting email addresses from the collected X.509 certificates. Whenever we found an email address pointing to a Certificate Authority's or a webhosting provider's info address, we ignored it. For all other email addresses, we ran DNS queries for MX entries for the email's domain. In case of a positive response, we stored the email address for contacting the webmaster later. In order to contact the webmasters for whom we did not find email addresses embedded in the certificate, we decided to send an email to `webmaster@domain.com`. Altogether, we sent 46,145 emails to email addresses either embedded in an X.509 certificate or to the domain's corresponding webmaster email address. We sent 40,480 emails to webmaster@domain.com and 5,664 to embedded addresses. 37,596 of those could not be delivered, leaving us with 8,549 successfully delivered emails. We received 755 complete responses to our survey, a response rate of 8.83%. We decided on a set of questions that would take only 5-7 minutes to answer, including two free text questions why the webmasters were using exactly this X.509 certificate on their website and the free text prompt to report problems with the configuration and wishes to make configurations for HTTPS more usable. We were mainly interested in the following aspects:
(1) **Reasons and Use Cases for employing HTTPS:** We were interested in how the website was primarily accessed, how many users were visiting it and in which context (e.g. commercial, private etc) the website was mostly used.
(2) **Technical Knowledge concerning SSL:** We asked several questions to asses how much the webmasters knew about SSL and if they had set it up themselves; we asked for estimations for the pricing of X.509 certificates and problems they had with SSL.
(3) **Risk Assessment Concerning Misconfigured SSL:** We asked how important SSL was for their website and how strong the risk for users was due to the non-validating certificate.
(4) **Complaints, Wishes and Suggestions for SSL:** In the end, we asked them to fill in a free text about if they had problems with configuring the certificate for their webserver, also asking them for complaints and ideas to "make things better".

4.2 Ethics

Our University has no formal IRB process, but the Privacy Officer also consults on ethical matters. The purpose of making contact with the affected webmasters was two-fold: (1) we intended to inform webmasters of the misconfiguration of their website and (2) kindly asked them to support our research. However, webmasters would benefit from our email without participating in the study.
We were aware that sending emails to all candidates at once could cause resentment in the recipients. To reduce negative side-effects, we specifically contacted the webmaster of the website by sending an email to the webmaster@domain address that is specifically intended for questions and comments concerning technical problems, as stated in RFC2142 [6]. However, we felt we were offering the administrators valuable information. Our results confirm our estimation of the situation: most webmasters reacted in a grateful or at least friendly way, some nicely explained why we saw the invalid certificate or thanked us for alarming them to the non-validating certificate and some of them said they wanted to fix the SSL configuration of their websites immediately.

[4] Server Name Indication (Cf. `http://www.ietf.org/rfc/rfc3546.txt`

[5] we conservatively estimated a success rate of reaching 10% of the administrators, and another 10% response rate to the study, which would have provided us with 500 answers

[6] https://www.ietf.org/rfc/rfc2142.txt

4.3 Study Results

Of the 755 webmasters, 154 (20.4%) operated websites with an expired certificate, for 250 (33.1%) websites hostname validation failed (13 were also expired), 160 (21,2%) websites used an X.509 certificate issued by a CA not included in the Mozilla truststore and 191 (25.3%) websites operated self-signed certificates.

4.3.1 Reasons and Use Cases for employing HTTPS

The primary access method was said to be via browsers in 681 cases, 15 by apps, 15 by embedded systems, and 44 stated they did not know. 319 webmasters estimated they had less than one hundred visitors per month, 165 estimated between one hundred and thousand, 95 between thousand an ten thousand, 66 between ten thousand and a hundred thousand, 19 between a hundred thousand and a million, 5 more than a million. Of the 134 webmasters who offered information about their websites' users, 84 (62.7%) said it was used only by themselves (primarily for administrative purposes), 11 (8,2%) said it was mostly used by friends and 39 (29,1%) said it was used by their company and colleagues. An important question of interest was why exactly the websites operated a non-validating X.509 certificate. Of the 495 webmasters who gave information as to why the certificate was configured in a way that would throw a warning message, 330 said they had configured it in such a way on purpose. W713, who operated a self-signed certificate, stated: *"The site is a development system not accessed by customers or the public and the warning message "issue" is known internally."*, while W49 stated: *"The X.509 certificate is used for access to sensitive parts of the site. It is only being used by skilled operators, i. e. people who are able to check the fingerprint of the certificate to determine its authenticity and then store it for subsequent uses."* W23 on the other hand mentioned: *"Using SSL with a commercial CA issued certificate that is not under total control by myself is inherently insecure, since every CA owner can hijack the security and all providers that acquired an intermediate CA certificate can do so. And in the last years we have seen how weak some CAs are protected against cybercriminals. So it's much more secure for users to accept a certificate that was signed by my own CA once and get cautious when it changes."* Another prominent statement came from W31: *"Our users are explicitly required to provision the CACert.org root CA before visiting the website. The website generates no warning then."* While 495 webmasters stated they use the questionable certificate deliberately, 165 webmasters said they had accidentally misconfigured their X.509 certificate. Again, the reasons for employing a non-validating certificate are manifold. W218 for example stated, *"I am the administrator of a website in the medical domain that must be HIPAA compliant and the HIPAA guidelines require HTTPS for websites. Since we did not want to spend money on a commercial X.509 certificate we decided to use a self-signed certificate."*, while W284 stated *"Actually you see the warning message because the certificate was issued for xxx.com and not www.xxx.com. If you click through the warning message you will be redirected to the correct website and will see no warning message at all."*. W98 stated *"I'm using one certificate for many sites (my server did not support SNI until a recent update), so I had to list every one of them in the alternative domain name. Since it is tiresome to add an additional entry for every one of them to account for the the correct subdomain, I did not bother [...] But your survey brought up to my attention this cases and I will fix the issue immediately"*
The *deliberate setup* group self-reported a mean SSL technical knowledge of 4.08, while the *misconfiguration* group reported a mean SSL technical knowledge of 3.80. The deliberate group rated their data sensitivity to be a mean of 2.48 out of 5, while the misconfiguration group rated it to be a 2.43. 101 of the participants stated that their websites were not actively in use or said that they were sure that there were no hyperlinks pointing to their website and hence no browser warning messages would ever be thrown.

4.3.2 Technical Knowledge Concerning SSL

We asked the webmasters who had set up the X.509 certificate for their HTTPS server. 613 stated they had set it up themselves, 63 certificates were set up by a coworker, 12 by a retired coworker, 68 by their service provider and 11 did not know who set up the certificate. We asked the participants to self-report their technical knowledge of SSL on a 5-point likert-scale between *very low* and *very high*. 12 self-reported their SSL knowledge as *very low*, while 236 said their technical knowledge of SSL was *very high*. In the mean they rated it as a 3.96. Many administrators were uninformed about the pricing of X.509 certificates and strongly overestimated the actual costs. At total of 87 (11.5%) webmasters did not know their configuration could lead to browser warning messages prior to our survey (those who did know reported a mean technical SSL knowledge to be 4.02, those who did not reported their mean technical SSL knowledge to be 3.50). One interesting finding was that six survey participants stated they do not need the features provided by CA-issued X.509 certificates. They argued that for their use cases they did not need the authenticity features CA-issued certificates provide but only rely on "strong encryption" to transport sensitive data. These statements demonstrate a lack of understanding of the SSL security features. Although they rely on strong encryption, they oversee the fact that without properly verifying the identity of the server, no secure communication channel can be established in a reliable way since a Man-In-The-Middle attacker could easily exchange the original with a malicious certificate.

4.3.3 Risk Assessment Concerning Misconfigured SSL

We asked the participants to rate the importance of HTTPS for the operation of their website on a 5-point likert-scale from *not important at all* to *very important*. 217 said that HTTPS is not important at all, while 190 stated that it is very important for their websites' users. In the mean they rated it as a 3.00. We also asked them to rate the sensitivity of the data their website serves via HTTPS on a 5-point likert-scale ranking from 1 as *not sensitive at all* to 5 as *very sensitive*. 252 rated the sensitivity of their data as not sensitive at all, while 81 rated it as very sensitive. In the mean they rated it as a 2.53. We also had them rate the risk the non-validating X.509 certificate they are using poses to their users. On a 5-point likert-scale they could choose from *very low* to *very high*. 524 rated the risk as very low, while 23 said the risk for their users was very high. In the mean they rated it as a 1.55 Of the 755 respondents, 612 stated their users never complained about the occurring warning message, 7 administrators reported they receive complaints at least once a week, 12 received complaints on a monthly ba-

sis, 26 on a yearly basis, 77 receive complaints less often and 21 could not remember how often they receive complaints.
We had been looking for interesting correlations between certificate error types and certain self-reported values. However, except for the results described above, we did not find a statistically significant correlation that would have helped us predict what kind of error would occur because of which characteristic. This underlines the importance of increasing the usability of X.509 certificate configuration and deployment in general, as well as building in more failsafe mechanisms and generally taking the weight of correct and secure X.509 certificate configuration and deployment away from webmasters.
At the end of the study, we asked the webmasters to describe problems they encountered with setting up SSL for their website and suggestions they had to *make X.509 certificate configuration more usable*. We present their concerns and suggestions in what we call the *Admins' Wishlist*.

4.4 Admins' Wishlist

We asked our participants to describe improvements they would like to add to make X.509 certificate configuration for HTTPS webservers easier and what they think is missing in the current system. In the following section we analyze their statements and describe what most participants find lacking in the current system. Of the 755 responding administrators, 87 offered suggestions, some of them more than one. Their suggestions can be categorized into six different groups:

Lowering The Price: 13 of the participants mentioned that the current price range for X.509 certificates that do not throw warning messages in browsers is not adequate. They find that paying a high amount of money for such a low cost task such as digitally signing an X.509 certificate is not fair and they would like to see a change in the current pricing policy of commercial Certificate Authorities. They criticize that the current CA infrastructure *"is a money printing machine without providing strong security for both service providers and their users"* (W29). Nine of them asked for a CA that issues free certificates that are accepted by popular browsers. Four participants complained that they have to configure X.509 certificates for multiple subdomains and that current wildcard certificates are too expensive. They wished to get access to cheaper wildcard certificates to reduce the number of false positive warnings on their websites.

Allowing CACert: 45 websites operated an X.509 certificate issued by CACert[7]. 10 of these proposed to add the CACert root CA to all popular browsers to provide an alternative to the commercial CAs issuing trusted certificates. The motivation to use a CACert certificate was two-minded: 28 of the administrators preferred CACert certificates since they did not want to support the commercial CAs and are of the opinion that basic encryption mechanisms as provided by SSL should be accessible by everyone for free. The remaining (17) did not trust the centralized trust model of commercial CAs after the breaches of DigiNotar and Commodo. They argued that the CACert's web of trust model provides more security and better protection against Certificate Authority compromise attacks.

Better Support for Non-Validating Certificates: 15 participants complained that they were forced to use certificates issued by commercial CAs to avoid SSL warning messages. They can be categorized into three different groups: Seven participants would like to change the current trust model. While two did not describe their idea of a different trust model, three would prefer a trust-on-first-use-based model such as known from the Secure Shell[8] which would allow them to use a certificate of their choice. Two other responders would prefer the TACK trust model proposed by Moxie Marlinspike[9] since they explicitly did not trust commercial CAs. Five participants would like to have an easy way to use self-signed certificates without giving concrete ideas of how such a system could work and four participants wanted to have an easier-to-use mechanism to validate certificate fingerprints to be able to securely deploy self-signed certificates for their users. Three participants were using their own CA in an enterprise environment and criticized the complicated workflow of adding their custom CAs to their users' browsers.

Better Tool Support: Six survey participants suggested to improve the tool support to generate and configure X.509 certificates for webservers. They found the command line interface for the OpenSSL tool[10] too complicated and wished for better documentation. The SSL configuration options of popular webservers were also criticized. Particularly the configuration of virtual hosts was described as very complicated and error-prone and administrators generally requested a more easy-to-use mechanism to configure X.509 certificates for multiple hostnames on a single IP address.

Auto-Update Reminder: Eight survey participants who used an already expired certificate were not aware of that fact before we contacted them. They criticized the fact that they would not receive an automatic message when their certificate expired and would like to have a service that keeps an eye on the expiration date of their certificate: *"Ideally an automatic message would be sent out to not miss the date to re-new a server's certificate"* (W643).

5. DISCUSSION

Our study reveals new findings and helps to better understand previous work in the field. While Akhawe and Felt [2] desire a 0% click-through rate for SSL warning messages, in our study 330 of 755 website administrators stated that they deliberately operate non-validating X.509 certificates and that their users are informed of the warning message beforehand. In these cases, SSL warning messages are no unexpected security warnings, but can be seen as information dialogs that users expect and to which they react by clicking through the warning because their administrator told them to. Whenever websites with non-validating certificates are re-visited by users, and the users did not add the non-validating certificate or the browser-untrusted CA to their truststore, they will repeatedly click through the warning. Since Google Chrome does not open the change to its trust store in the warning menu, it is likely that users will click-through a warning message on every visit. This is one possible explanation for the huge difference in click-through rates as reported by Akahwe et al.[2]: They count every repeated click-through in Chrome, while they count click-throughs in Firefox only at the first visit to the respective website (which may or may not have occurred during the period of their data collection). Thus, our findings sup-

[7] `http://www.cacert.org`

[8] `http://tools.ietf.org/html/rfc4253`

[9] `http://tack.io/draft.html`

[10] `http://www.openssl.org/`

port the assumptions that Chrome's click-through rate is massively influenced by re-visits of websites that operate non-validating certificates.

We found that many webmasters reported that their site was either not in use any more, or that the SSL version of the specific domain we encountered had never been meant to be accessible for users at all and had respectively never been hyperlinked anywhere on the Internet. This attests to a very important finding: Studies using datasets of SSL certificates that were accumulated by certifictae crawlers are prone to massively overreport handshake failures and hence SSL warning messages in browsers not only by assuming a possibly not applicable set of trusted certificate issuers, but also by including unused websites.

We gained valuable insights from the free texts the webmasters wrote about problems with SSL and improvement suggestions. Many of them wished for more simplicity: 165 had accidentally misconfigured SSL. Some wished for either a simpler interface to set up a webserver, others wanted an automatic renewal for expiring certificates.

330 webmasters had configured their webservers in a non-validating way on purpose. 15 of them wished that there was a broadly-accepted alternative to commercial CAs; in general there were complaints about the pricing of CAs. This is a very interesting finding: 20 of the 85 webmasters who suggested improvements requested a free alternative to paid CA certificates. Obviously these webmasters were not aware of the fact that there are free alternatives[11] that provide free and trusted X.509 certificates, which demonstrates that there is not only the need for a better technical education but also for a broad and basic documentation, complete with examples and links for webmasters, who understandably do not call SSL their primary field of expertise.

6. LIMITATIONS

Population: We contacted webmasters from a random sample of 50,000 websites which operated non-validating X.509 certificates without considering the popularity of the given website. Our results imply that X.509 certificate warnings occur more frequently on websites with low traffic which is often regarded as unproblematic by their webmasters: they claim that their users are aware of the presence of an SSL warning message causative certificate.

Self-Selection Bias: All our participants were self-selected. They chose to fill out the survey, which could mean that more active webmasters answered.

Bounced emails: We tried to reach webmasters either by using the contact email address in the website's X.509 certificate or the webmaster@domain.com email address. 37,596 of all emails we sent were bounced. Hence, the majority of websites with non-validating X.509 certificates does neither follow best practices and nor provide an easy-to-find contact email address. It might be possible that those webmasters have different reasons for using a non-validating X.509 certificate on their websites.

Underreporting: Some of our conclusions are drawn from answers which were given as free text. Thus we do not have data on all of our participants for several issues: Some did not report on whether their website is in use/is meant to be used at all, while others did. Not all of the users report on who the SSL connection is intended for. Therefore it is possible that we underreport the websites which are out of use, as well as the websites which exist for webmasters' use only.

7. CONCLUSIONS

We conducted the first study with webmasters who operate non-validating X.509 certificates on their HTTPS-enabled websites to understand their motives. Therefore, we used the body of 4,487,463 certificates Google's webcrawler had collected over 12 months. We identified 610,966 non-validating certificates, chose a random sample of 50,000 of these, established if they were still operating and non-validating certificates, extracted email addresses for their webmasters and emailed them. Of those emails, 8,549 were successfully delivered. Of these, 755 webmasters who operated websites with non-validating X.509 certificates responded to our study.

101 said that their website was not meant to be accessible, and that actual users would not have encountered the certificate as the webcrawler did. We found that of the 495 who reported on this issue, 330 said that their use of a non-validating certificate was deliberate, while only 165 explained it with an accidental misconfiguration. 44 of the webmasters (25% of those who had accidentally misconfigured their webserver) stated that they were confused about SSL configuration in general, strengthening the assumption already made by Fahl et al[4] that, while warning messages and user behavior are an important field of study, studies with IT professionals in general and webmasters and developers in particular are an important and often neglected issue. We confirm findings from Akahwe et al., who state that Google Chrome users tend to click through warning messages more easily: This goes well with our finding that many webmasters use non-validating certificates on purpose and inform their users about it. Clicking through these warnings will add the certificate to Firefox, while Google Chrome will show a new warning on each revisit.

8. REFERENCES

[1] D. Akhawe, B. Amann, M. Vallentin, and R. Sommer. Here's my cert, so trust me, maybe?: Understanding tls errors on the web. WWW '13, 2013.

[2] D. Akhawe and A. P. Felt. Alice in warningland: A large-scale field study of browser security warning effectiveness. In *USENIX Security Symposium*, 2013.

[3] Z. Durumeric, E. Wustrow, and J. A. Halderman. ZMap: Fast Internet-wide scanning and its security applications. In *22nd USENIX Security Symposium*, Aug. 2013.

[4] S. Fahl, M. Harbach, H. Perl, M. Koetter, and M. Smith. Rethinking ssl development in an appified world. CCS '13. ACM, 2013.

[5] R. Holz, L. Braun, N. Kammenhuber, and G. Carle. The ssl landscape: A thorough analysis of the x.509 pki using active and passive measurements. IMC '11, 2011.

[6] J. Sunshine, S. Egelman, H. Almuhimedi, N. Atri, and L. F. Cranor. Crying wolf: An empirical study of ssl warning effectiveness. SSYM'09. USENIX Association, 2009.

[11] e. g. `www.startssl.com`