Consolidated Review of
Analysis of the HTTPS Certificate Ecosystem

1. Strengths

Very comprehensive dataset, larger than anything reported before, with interesting observations regarding the deployment of X.509 certificates in the Internet. They use a comprehensive dataset to shed light on several alarming practices within the certificate authority community.

Thorough analysis of the current state of HTTPS certificates and certificate authorities. They analyze the data long several dimensions; including top browser trusted authorities, diversity in the organizations signing certificates, validity of browser certificates, trends in the root key size and the change in the market share of the authority over the year. It documents the current status of the certificate ecosystem and highlights several practices within the community that could be improved. The authors provide some lessons learned (Section 9) that could help guide future deployment of HTTPS.

2. Weaknesses

The paper's main strength is in the size of the dataset, while the key takeaways are somewhat incremental to previous work. While the authors provide some lessons learned, the papers focuses primarily on analysis, and there is less in the way of a plan of action to secure the HTTPS certificate ecosystem. Some results are a bit overblown, or not put in the appropriate context.

The collection and analysis methodology is pretty straightforward. The paper reports the current state of the certificate ecosystem based on single port probe data. While the paper discussed several worrisome practices within the community, it does not present any cases of known abuse. It maybe primarily because there is limited insight that can be gained from just the probe data and other techniques would need to be used to detect such abuse.

Leans somewhat heavily on prior work.

Several typos throughout the paper indicated somewhat sloppy execution.

3. Comments

This is a very interesting paper about the certificate ecosystem. The description of the measurement methodology is clear and I particularly appreciate the discussion in 4.2 that describes measures taken to reduce the active scanning and how you addressed concerned network operators. The results teach the reader a lot about the current state of the certificates in the Internet, and raises real concerns that should be addressed by CAs. While some of the results were reported in previous works (e.g., [16]), I think there is value for this paper, both due to the large dataset it uses and for the additional insights it provides.

This was a fun paper to read and seems like a great fit for IMC. It makes an important contribution to further our understanding of HTTPS usage and certificate ecosystem. The authors collect a comprehensive data set and perform a through analysis. This would be a great paper to have in the program.

While the authors look at data from scans over 11 months, it would be interesting to see how the data has changed over the past several years.

The authors identified some security weaknesses: is it known whether these weaknesses have been taken advantage of? How frequently? Are more security breaches occurring over time?

Can authors say anything about the changes that should be made to the HTTPS certificates that would have the biggest impact in practice? What can users of HTTPS websites do to move these changes forward?

I do have some comments that may help improve the overall presentation and positioning:

- The authors do a good job of discussing prior work on this topic, and there has been quite a bit of recent prior work. It seems to me that there are two points of overlap: methodology and results. In terms of methodology, the paper's techniques do not seem new, but the dataset if bigger and more comprehensive it seems. It would be good to identify what aspects of the methodology are new if any, and which aspects are borrowed/improved upon. In terms of results, the paper states in the related work section that prior results have focused on similar issues except using less comprehensive data sets. What is missing from the rest of the paper though is which of your quantitative results are new/improvements and which re-confirm prior results. I would have liked to see the new results clearly highlighted in each section. In particular, you say that while [5] focuses on just ALexa top 1M, you do a more comprehensive analysis of the address space. Yet, you don't comment on how the sites outside the top 1M differ from those in the list, assuming there are any differences at all.
- There are numerous typographical errors (too many to list). Case in point is reference [5], which is quite crucial to your paper as you build right off it. Unfortunately, this reference is broken - conference and year info is missing! There are many other such issues with writing.
- You definitely must perform another pass on the writing of the paper -- there are many typos and sentences that are badly written throughout the paper that need to be fixed. Also, please spend more time explaining tables and plots before discussing the takeaways. Some examples -- Table 5 is not clear and many of the plots are cumulative (and you should state that). Also, in terms of style - there are too many whitespaces in the paper that surround the plots, and these should be removed.
- I found some of the discussion confusing and other parts overblown: e.g., a very basic question: why did you need multiple scans? What new information were you hoping to learn from each scan? How often do the scans need to happen to glean the info you want to glean? In section 5.2 you say that 61 root authorities directly sign 41,000 leaf certificates. This does not seem significant given the total number of leaf certificates. At many places you don't put your results in context even though you clearly have the data to do so; e.g., are the anomalous sites mostly in some specific AS or geographical region?

The good news is that these issues are fixable through careful writing and minimal additional analysis.

External Reviewer: The authors include discussion regarding whether they missed certificates due to the same IP address hosting multiple TLS endpoints. But they do not discuss the issues of NAT, managed TLS services and such. The authors might perform some experiments with known examples to see what, if anything, their approach misses.

4. Summary from PC Discussion

This paper was accepted without discussion.

5. Authors' Response

In this work, we present the results of a large-scale study of the HTTPS ecosystem based on data we collected over the past year using 110 Internet-wide scans of hosts listening on port 443. Our aims are to map the trust relationships within the CA ecosystem, identify practices that are putting the ecosystem at risk, provide up-to-date metrics on the deployment of HTTPS, and draw lessons and recommendations for future practice. We are also releasing our dataset for use by the research community. It includes the full results of our scans, including 42 million distinct TLS certificates that were served from 108 million IP addresses between June 2012 and August 2013.

We thank the anonymous reviewers and our shepherd, Udi Weinsberg, for their feedback and suggestions. In response, we have expanded on the lessons from our study in a new discussion section, which additionally explores options for improving the security of the HTTPS ecosystem. We better explain how the methodology we use presents a different, more comprehensive perspective on this ecosystem than techniques used in previous work. We have also added context for several of the experiments, and we have rewritten much of the text for clarity and polish.