Understanding the Impacts of Google's CT Policy Update on the WebPKI

HARRY BRAAM, University of Twente, The Netherlands

This research investigates the impacts on the Public Key Infrastructure (PKI) ecosystem of a recent Google policy update [11], which removed the explicit dependency on the Google Certificate Transparency (CT) logs on CT compliance. Digital certificates were introduced in the 1990s to verify the identity of parties involved in digital communications mitigating man-in-the-middle attacks. Third-trusted parties, known as Certificate Authorities (CA), were introduced to achieve this. CT was introduced in the 2010s to remove the need to trust CAs blindly. CT forces CAs to log their issued certificates in a public append-only log run by various parties such as Google, Cloudflare, and others. By requiring registration in one or more logs, this measure prevents the existence of unidentified certificates for a given domain. In the early days of CT, Google required certificates to be registered in at least one Google log for the certificate to be accepted. This requirement was recently loosened, and logging certificates to a Google CT log is no longer required. We compared the logging policies of the top 5 CAs before and after the policy update. We also looked at where domains were logged before and after the policy update. Finally, we investigated how other browser vendors have updated their CT policies in response to a changing CT ecosystem. We found little to no impact on the logging behavior of the major CAs due to this policy update. However, some smaller CAs have reduced their reliance on Google CT. Next, based on the results obtained during the analyses we propose two future research directions.

Additional Key Words and Phrases: Certificate Authority, Certificate Transparency, Public Key Infrastructure, Certificate Logging, PKI

1 INTRODUCTION

The Internet has become a crucial part of our lives, allowing us to communicate quickly and easily with people and institutions worldwide. However, as more sensitive information is shared online, the risks of cybercrime have also increased [4, 6]. Establishing secure and trustworthy online communications is of the utmost importance for security, which is where the X.509 standard comes into play. It provides a standardized framework for issuing and managing certificates used for authenticating stakeholders [14]. It succeeds by making use of asymmetric keys which allows for the creation of a public and private key. The unique aspect of these keys is that one can be used for encryption and the other one for decryption without revealing each other [19]. Asymmetric keys can be used to create digital certificates using a three-cornered trust model. By adding a trusted third party to an interaction, the user can be sure of the service's identity without verifying it personally. These third-trusted parties are called Certificate Authorities (CAs) [14].

The Internet Engineering Task Force (IETF) uses the X.509 standard to authenticate secure connections for its Transport Layer Security (TLS) protocol [8]. When a client connects to a web server using TLS, the server sends its digital certificate to the client. Digital certificates contain the server's public key, which the client can use to send messages to the server, some identifiable information about the server, and certificate details. The certificate is signed using a CA's private key. On the client side, it is then checked against a predetermined list of trusted CAs, and the client can use the server's public key contained in the certificate to send secure messages to the server. That way, man-in-the-middle attacks that attempt to intercept and modify the server's public key contained in the certificate would be prevented as re-signing the certificate with the CA's private key would be needed [8].

For digital certificates to be trusted, it is imperative that CAs can be trusted and are not compromised. Having a second certificate for the same domain in the hands of bad actors would allow for man-in-the-middle attacks [10]. Such an attack could be executed by intercepting the server's certificate and passing on the fraudulently acquired certificate with the attacker's public key. The attacker would be able to listen in on the server-client communication and even modify client-server communication [10]. Placing such a safety-critical system in the hands of a handful of companies and governments is bound to cause issues [20]. The lack of trust in CAs is not purely theoretical. For instance, in 2011 the Dutch DigiNotar CA was compromised [18], resulting in fraudulent certificates for websites such as Google, Facebook, and Skype. These certificates were used to intercept traffic and steal login credentials from users highlighting the need to update the CA system.

To address these risks, Certificate Transparency (CT) was introduced by engineers at Google as a system for publicly logging TLS certificates issued by a CA in an append-only log. In exchange for logging certificates, a Signed Certificate Timestamp (SCT) will be issued as proof of inclusion. CT logs provide transparency and accountability for the issuance of certificates [12]. The integrity of these logs could be verified by monitors. Domain admins could monitor the logs to get informed when new certificates for their domain get issued, allowing for detection by domain owners of fraudulently issued digital certificates. The domain owner can quickly trigger the mechanism for revoking the fraudulently acquired certificates and take appropriate countermeasures.

Due to the limited number of CT logs available when they were first introduced, Google made it mandatory for a certificate to be logged in at least one CT log operated by Google to be accepted in Chrome. In March 2022, Google made the decision to update its CT policy as the CT landscape evolved and matured. Going forward, any two distinct logs acknowledged by Google will suffice, except for certificates with a validity period exceeding 180 days. For such certificates, a total of three different logs is required [11].

The goal of this research: As mentioned earlier, CT plays a crucial role in ensuring security in the modern WebPKI, offering transparency and accountability in certificate issuance. CT logs serve as a critical security system that safeguards a significant amount of personal data. However, Google made an announcement on the 10th of February 2022 that it will be eliminating the explicit dependency

TScIT 39, July 7, 2023, Enschede, The Netherlands

for CT compliance with its CT logs [11]. To the best of our knowledge, the impact of the policy changes announced by Google in 2021 has not been investigated yet. Therefore, the goal of this research will be to provide insights into the impact of CT policy changes and inform future CT policy updates to enhance web security.

2 RELATED WORK

CT is, by now, a moderately researched topic. Much research has been focused on the TLS landscape, with CT being only mentioned [12]. Most CT-focused research can be split into two tracks, one proposing extensions or improved implementations [13, 15, 22, 23] and a track of research extracting data from the current implementation to verify its correct working [3, 12, 17, 21] or extract general information about the web [2, 12]. Amann et al. [3] examined the use of CT in the context of general improvements to the TLS ecosystem since 2011. The paper focuses on the deployment and use of these new systems. Scheitle et al. [21] analyze the evolution of CT over time and the implications of exposing Domain Name System (DNS) names for security and privacy. They showed that a large part of connections support CT. Additionally, they demonstrated that CT logs are the target of scanning campaigns to find hidden domains.

Gustafsson et al. [12] presents a comprehensive analysis of the CT landscape, characterizing eleven CT logs through both active measurements and passive observations within a university network. The study reveals significant variations in root store selection and the methodology employed for adding new certificates. Google-operated logs, which employ crawling, exhibit larger root stores and include a broader range of certificates, mirroring web traffic browsers encounter in various contexts. In contrast, CA-operated logs show a smaller diversity of certificates and tend to comply with Chrome's Extended Validation (EV) certificate policy. The research further examines cross-log submissions and highlights differences in submission rates of Domain Validated (DV) certificates across various CAs over time.

The authors mentioned the availability of certificates in a series of Google CT logs. The study found a high percentage of certificates available in Google CT logs, typically in the high nineties. Additionally, the paper sees the results of another policy shift concerning the mandatory registration of EV certificates in CT logs, which became effective in January 2015. This policy change marked the first instance of required CT logging. The analysis observed bulk registrations in various CT logs, indicating compliance with the new requirement. The study is partially focused on distinguishing between EV, DV, and Organization Validated (OV) certificates. However, it is worth noting that this distinction has lost its significance, as major browsers have ceased to significantly differentiate between these different types of certificates.

Korzhitskii et al. [16] analyze the evolving root store landscape. The authors examine CT logs and compare their root stores to those of major software vendors. They observe that CT technology is highly established and widely used by CAs and internet clients. The study reveals that root stores have grown larger over time, but there are issues with certain roots being included by only a few log operators and the presence of compromised certificates in some logs. The authors emphasize the importance of careful and transparent management of CT policies and root stores.

Apple and Google rely on a few logging operators, and CT logs do not sufficiently cover some WebPKI roots trusted by major software vendors. The authors identify problems such as duplicate entries in root lists, anomalies in root store presentation, and logs violating policies while being considered trusted. They call for improved management practices to ensure the integrity and effectiveness of CT in securing the internet ecosystem. How the recent policy change has impacted the WebPKI has to the best of our knowledge, not been investigated yet.

3 METHODOLOGY

3.1 Research Questions

We aim to focus on the biggest and most influential CAs as they are the vanguard of WebPKI. Next to this, they hold most of the market, with the top 5 CAs holding 90% of the issued certificates in our research period. Our study will address the following research question (RQ) and sub-research questions (SRQ):

RQ: What is the impact of removing the explicit dependency on Google CT logs on CA compliance with CT requirements?

- **SRQ1**: What logs did the top 5 most-used CAs rely on before the policy update to achieve CT compliance, and how did their logging practices change after the policy update?
- **SRQ2**: What proportion of CAs have significantly reduced their reliance on Google CT logs after the policy update?
- **SRQ3**: How have other browser vendors (Mozilla, Microsoft, Apple) updated their CT Policies in response to changes in the CT ecosystem, and what impact has this had on the WebPKI?

3.2 Dataset

The dataset used for analysis in this study is scraped from CT logs hosted on various platforms and is stored on a cluster maintained by the University of Twente. The dataset includes the following CT servers: Cloudflare Nimbus, DigiCert Nessie, DigiCert Yeti, Google Argon, Google Xenon, Let's Encrypt Oak, Sectigo Mammoth, and Sectigo Sabre. These CT logs have been scraped from 2021 until 2023 to provide a diverse and comprehensive coverage of certificate issuance. The dataset comprises a substantial amount of information, with over 15 billion rows. It encompasses a diverse range of certificates issued by various CAs, totaling 854 in number[1]. As mentioned before, the certificates are not evenly distributed over CAs, with the top 5 authorities commanding a staggering 90% of the certificates issued in 2021-2023, which can be seen in Table 1. Additionally, the dataset encompasses many domains, representing over 970 million distinct domains. The period analyzed in this study spans from January 2021 to June 2023, allowing for an extensive examination of trends before and after the policy update.

3.3 Analysis of CT Data

We will employ two distinct methods to examine the data: firstly, by evaluating the total count of certificates issued by each CA per CT

[1]Based on unique values of the organization field.

CAs	Percentage of Total
Let's Encrypt	72.42%
Cloudflare, Inc.	5.32%
cPanel, Inc.	5.31%
Sectigo Limited	3.83%
DigiCert Inc	3.32%

Table 1. Percentage of total certificates 2021-2023

log before and after the policy update, and secondly, by analyzing specific combinations of domains and CAs in a period before and after the policy update. Looking at domain CA combinations aims to determine whether domains previously present in a Google CT log are no longer present after the policy update. By employing these two analysis approaches, we believe they will complement each other effectively. Assessing the absolute number of certificates can potentially reveal alterations in CT logging behavior. However, it is important to note that this correlation is not guaranteed, as there could be various explanations for such changes. Therefore, the second analysis method, which focuses on specific domains, assures that the observed changes are genuine.

We conducted an analysis of the absolute number of certificates published by each CA and CT server. To distinguish between various CAs, we focused on the organization field of the issuer. This field was preferred over the common name field as the common name field differentiates between departments within a CA, which adds complexity to the analysis. Additionally, it can be assumed that a company generally adheres to a single CA policy. The dataset has over 800 distinct organization fields to focus our analysis. Only the top 5 CAs were chosen. These were selected by looking at the total number of certificates published from 2021 to 2023, leading to the selection of the CA's Let's Encrypt, Cloudflare, Cpanel, Sectigo, and DigiCert. To establish a suitable timeframe for our study, we selected the period from 2021-01-01 to 2023-06-01. This duration spans 29 months, starting 13 months from the announcement of the policy update and ending 13 months after its implementation. We chose this timeframe because it aligns with the maximum validity period of certificates accepted by major browser vendors and determined on CA/Browser conference [9], which is 398 days. To organize the certificates based on their date, the *not_after* field was selected. This choice was made because temporally sharded logs use this field to partition certificates across different iterations of the logs, allowing for easy identification of the specific log iterations in which the certificates were logged.

Looking at a specific domain (CA combination before and after the policy update), we divided the data into two time periods: one that occurred 398 days before the policy announcement and another that occurred 398 days after the policy was implemented. For domain CA combinations, a list was made of CT log servers where the certificate was logged. This was done for both time frames, after which the CA domain combinations in both sets were selected, leaving us with a list of CT log servers before and after the policy update for a domain. CAs with fewer than a thousand domains in this dataset were excluded, as the natural variability could disproportionately affect the results. Determining the threshold for significance is challenging. However, we have opted for a 10% reduction in domains that were previously recorded in the Google CT log but are no longer observable following the implementation of the policy modification. We are confident that this method is sensitive enough to identify any changes in the logging policy.

3.4 Policy Analysis

The methodology for examining the changes in criteria to gain CT compliance by browser vendors (Google, Mozilla, Microsoft, Apple) and their impact on WebPKI involves the following steps. Firstly, a literature review was conducted to gather existing knowledge on the topic. Official sources where the vendors publish their CT policies were identified, including websites and documentation. Next, the data from WayBackMachine and GitHub were utilized to analyze historical snapshots and previous commits to track policy changes. The analysis was focused on the period from 2018 to the present (June 2023) and on embedded SCTs. SCTs delivered through TLS or the Online Certificate Status Protocol (OCSP) are very uncommon [3], and there are stricter requirements for embedded SCTs. This has led CAs to adhere to the embedded SCT policy more closely, as it is the most commonly followed and strictest policy. The gathered data were analyzed to identify patterns and trends. Finally, the findings were interpreted to draw conclusions on the evolution of CT policies and their impact on WebPKI.

4 RESULTS

In this section, we present the results obtained from the analysis of the CT data and policy data. The section starts with the analysis of the absolute number of registered certificates by the top 5 CAs. Next, we examine the number of certificates issued by each CA which were present in a Google CT log before and after the policy update. Finally, we look at changes in CT policy by major browser vendors. These results provide insights into changes in CT logging behavior by CAs and the impact on the WebPKI.

4.1 Timed Data

In this analysis, we will examine the total count of certificates recorded in each CT log per CA. It is important to note that fluctuations in the data, such as spikes or drops, when transitioning between CT servers, can be explained by the ordering of certificates based on their expiration dates. If a CA offers certificates with a validity of 30 and 90 days and chooses to stop logging to a CA server at date X, it is expected that certain certificates show up at X + 30 and a different group at X + 90. This explains the lack of clean breaks and starts when switching CT servers. Next to this for multiple CAs, there is a gap in logging certificates in Digicert Yeti for 2022. This is possibly related to the bit flip Digicert yeti experience in that year [5]. CAs might not have wanted to deal with the replacement log and abandoned Yeti for that year instead.

4.1.1 cPanel. Figure 5 in Appendix shows the logging behavior of cPanel. Notably, in the middle of 2021, cPanel made a transition by switching its preferred CT log server from Let's Encrypt Oak to

Cloudflare Nimbus. This shift to Cloudflare Nimbus is also observed in Figure 1.

Furthermore, cPanel's interaction with Google Argon is worth mentioning. The data suggest that cPanel initially registered a small number of certificates or conducted tests by registering certificates in Google Argon. However, this logging activity ceased in 2022, implying a change in their approach or reduced reliance on Google Argon for logging purposes.

On the other hand, cPanel consistently logged certificates in Google Xenon, as depicted in Figure 1, both before and after the policy update. This finding suggests that cPanel maintained a steady logging behavior in Google Xenon and did not exhibit any significant alterations in response to the policy update.

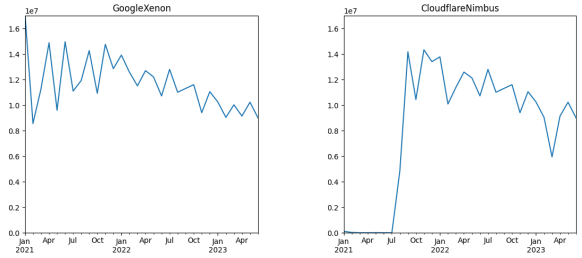

Fig. 1. cPanel, Registered certificates in Xenon and Nimbus

4.1.2 Cloudflare. Figure 6 in Appendix shows Cloudflare's logging behavior across various CT servers. Cloudflare has used a majority of CT servers available. Notably, in 2023, Cloudflare stopped logging in Cloudflare Nimbus and instead started logging in Let's Encrypt Oak, as can be seen for Cloudflare Nimbus in Figure 2. This is an interesting choice as Cloudflare Nimbus is hosted by Cloudflare itself.

Regarding Digicert Yeti, Cloudflare ceased logging into this CT log in 2022 but resumed logging in 2023. This might be related to the Digicert Yeti bit flip [5].

When considering the policy update, no significant changes in logging behavior can be observed. Both before and after the policy update, Cloudflare continues to make use of Google CT logs, as indicated by their logging activity in Google Xenon, as depicted in Figure 2. This finding suggests that Cloudflare's logging behavior remained consistent and unaffected by the policy update.

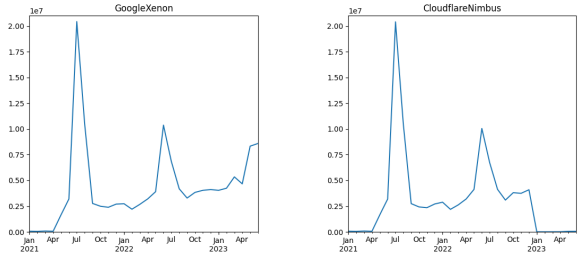

Fig. 2. Cloudflare, Registered certificates in Xenon and Nimbus

4.1.3 DigiCert. Figure 7 in Appendix shows DigiCert's logging behavior across various CT servers. Similar to Cloudflare, DigiCert ceased logging in Cloudflare Nimbus in 2023 and transitioned to logging in Let's Encrypt Oak. Additionally, DigiCert did not log certificates in Digicert Yeti in 2022, which might be related to the Yeti bit flip [5]. However, it is worth noting that DigiCert did continuously use its own CT server Nessie, as depicted in Figure 3.

It is interesting to observe that DigiCert logged a small number of certificates in Sectigo Sabre. The reason behind this logging activity is unclear, but it could be attributed to testing purposes or specific requests from buyers. Further investigation would be required to determine the exact motive behind this logging behavior.

DigiCert did not demonstrate any changes in logging behavior after the policy update. Both before and after the update, DigiCert continued to register certificates in Google Xenon and Argon, as illustrated for Xenon in Figure 3.

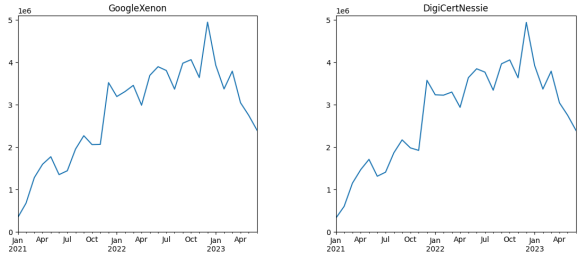

Fig. 3. DigiCert, Registered certificates in Xenon and Nessie

4.1.4 Let's Encrypt. Figure 8 in Appendix provides insights into Let's Encrypt's logging behavior, the largest CA in terms of the number of certificates issued. They log a significant number of certificates in Cloudflare Nimbus and Let's Encrypt Oak, but the volume is much higher in Google Xenon and Argon, as seen in Figure 4. The reasons behind these varying volumes are unknown. A theory could be that certificates with a validity of greater than 180 days are logged in these servers since they require three or more distinct CT logs. However, this would be false as the average certificate duration of Let's Encrypt certificates is 90 days in all CT servers.

Let's Encrypt also logged certificates in DigiCert Yeti, but after 2022 no further certificates from Let's Encrypt were found there. This is unusual as it is expected that a few certificates are present, as seen in Digicert Nessie and the Sectigo CT servers. The exact reasons for this change in logging behavior in DigiCert Yeti are not known and would require further investigation or additional information. Let's Encrypt also logged a small number of certificates in Sectigo Mammoth. Overall, Let's Encrypt logs the highest number of certificates in the Google CT servers, and no visible changes in logging behavior were observed after the policy update.

4.1.5 Sectigo. Figure 9 in Appendix shows the logging behavior of Sectigo. One observation is that Sectigo now logs certificates in Google Xenon and Cloudflare Nimbus. Additionally, the figure shows that Sectigo previously relied on Let's Encrypt Oak but phased out its usage in the middle of 2021. Instead, Sectigo started

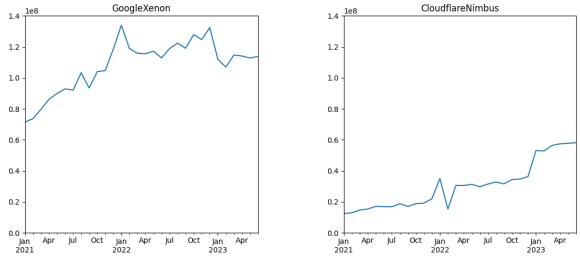

Fig. 4. Let's Encrypt, Registered certificates in Xenon and Nimbus

relying on Cloudflare Nimbus as its chosen CT log server. Next to this, it appears that Sectigo does not make use of its own CT servers, Sectigo Mammoth and Sabre.

Furthermore, Figure 9 reveals that Sectigo registers some certificates in Google Argon. This is likely done to comply with the requirement of having certificates with a validity longer than 180 days logged in three or more distinct CT logs. By logging into Google Argon, Sectigo ensures adherence to this requirement, which is also reinforced by the average certificate duration, which is 367-375 days in Google Argon compared to 126-190 days in Nimbus and Xenon. There is no clear update in the logging policy after the policy update. Sectigo kept logging its certificate to Google CT servers.

4.2 CA Domain Data

Figure 6 in Appendix shows how much Google CT logs were used before and after the policy update. Only CAs with over 1000 certificates in the dataset are shown. *Total* shows the number of domains registered at each CA before and after the policy update. *Before* shows the number of domains present in either Google Argon or Xenon before the policy update. *After* shows the number of domains in Google CT logs after the policy update. This data further confirms the results of the timed data, with the top 5 CAs having more than 99.9% of their certificates in a Google CT server before and after the policy update.

If a CA consistently continues to publish certificates to Google CT logs, the percentage of domains registered after the policy change is expected to be high. However, slight variations can occur due to test certificates, publishing delays to CT logs, or other data-related issues. Out of 73 CAs, 23 registered at least 10% fewer domains in Google CT logs after the policy update. This represents a significant increase compared to before the policy change, where all CAs registered 99.2% of their domains in Google CT logs.

The 5 biggest CAs with less than 90% of their domains in Google CT logs have been selected to investigate their case further. These CAs are *Actalis* (63%), *Cybertrust Japan* (69%), *GlobalSign* (81%), *home.pl* (76%), and *Unizeto Technologies* (77%). These are all relatively small CAs. Their timed data has also been gathered, presented in Figures 11 to 15 in Appendix.

Looking at these charts, it is clear that each CA still actively publishes Google CT logs, so it is not possible to conclude that they completely dropped their reliance on Google CT logs. However, a significant portion of their certificates are no longer present in Google CT logs. This might be attributed to a slow move away from Google CT logs, oversight of the CA, a small-scale test, or a lag in publishing to Google CT logs.

4.3 Policy Changes

This section presents the findings regarding the CT policies of Mozilla and Microsoft, Apple, and Google.

Regarding Mozilla, it was observed that they do not actively check the CT status of certificates. This implies that Mozilla does not enforce the requirement for certificates to be logged in a CT log to be trusted by their browsers. This issue has been open on their forum since 2016 [1]. On the other hand, Microsoft's CT policies were not found to be explicitly published. However, it is worth noting that since the Edge browser is built upon the Chromium code base [7], there is a possibility that Microsoft follows the predefined CT policies established by Google.

In order to achieve CT compliance, Google implemented different requirements for the number of CT logs depending on the lifetime of the certificate, as depicted in Table 2. Furthermore, Google mandated that at least one SCT from a Google-operated log and another SCT from a different log operator must be present.

Following suit, Apple also enforced CT compliance for all certificates starting October 15, 2018. The compliance criteria set by Apple were similar to those of Google and were also based on the certificate lifetime as shown in Table 2. However, in light of the decision to disallow certificates with a lifetime greater than 398 starting September 1, 2020, Apple updated its CT policy on April 21, 2021. According to the updated policy, certificates with a validity of 180 days or less required two SCTs from separate logs, while certificates with a validity period of 181 to 398 days required three SCTs from separate logs. Additionally, Apple mandated the use of at least two logs from different operators.

Subsequently, Google modified its CT policy on April 15, 2022. This update eliminated the requirement for at least one Google CT log and aligned the required number of SCTs per time frame with Apple's policy. Furthermore, Google adopted the requirement of utilizing logs from different operators. This CT policy change by Apple and the slow policy update by Google resulted in a significant variation in CT policies for more than a year.

Certificate Lifetime	Number of SCTs from distinct CT Logs
< 15 months	2
>= 15 and <= 27 months	3
> 27 and <= 39 months	4
> 39 months	5

Table 2. Number of distinct CT log required

5 DISCUSSION

While the conducted analysis provides valuable insights into the logging behavior of specific CAs and their CT policies, it is crucial to acknowledge some potential issues that may impact the generalizability and interpretation of the results.

Firstly, the methodology focuses only on a selected group of CAs. This limited scope means the findings may not accurately

represent the entire WebPKI ecosystem. Other CAs, which may have different logging practices or responses to policy changes, are not included in the analysis. Therefore, caution should be exercised when extrapolating the results to the broader landscape of CAs.

Another potential issue is related to the analysis of CA domain data. The proposed approach assumes that CAs promptly update their logging policies in response to the policy update. However, it is possible that there might be a time lag between the policy change and the implementation of the updated logging behavior by CAs. As a result, in the analyzed dataset, a certificate may not be registered in a Google CT log. However, a previous certificate for the same domain might still be present as all the CT servers are aggregated, creating a misleading impression that the CA continues to use Google CT logs. This discrepancy could affect the accuracy of the analysis and the conclusions drawn.

It is essential to consider potential confounding factors that may influence the observed logging behavior of CAs. The analysis assumes that the policy update solely influences changes in logging behavior. However, other external factors, such as evolving industry practices or operational considerations, can also shape CA practices and CT logging behavior. These factors should be considered when interpreting the findings and drawing conclusions.

5.1 Data Artefacts and Peculiarities

The dataset reveals several instances of missing values in fields where such occurrences should not be common. Notably, a significant number of certificates lack information in crucial fields such as the organization and domain fields. Of particular concern are the CT logs hosted by Let's Encrypt, which exhibit a higher prevalence of data issues than other CT servers. For example, the number of certificates with missing organization fields in Let's Encrypt logs reaches hundreds of thousands, whereas in other servers, this count remains in the tens or low hundreds. It is worth noting that this observation holds true even when accounting for the fact that Let's Encrypt does not register more certificates than other servers.

Let's Encrypt CT logs also display instances of missing date information. This highlights the need for CT log hosts to implement more robust safeguards to prevent the acceptance of incomplete or broken certificates. By doing so, the overall safety and reliability of WebPKI can be enhanced.

Table 3 shows the timestamps in Coordinated Universal Time (UTC) standard of the first and last certificate present in the dataset. For a comprehensive view, please refer to the complete version in Appendix (Table 4). The table highlights a lack of consistency among CT log operators regarding the start and end dates of their logs, even across different iterations.

It is common to find overlapping active logs or variations in time zones used by CT operators in different years. These inconsistencies can potentially lead to misleading interpretations. For instance, assuming that a 2021 log contains all the certificates issued in that year would be false, as there could be a gap of up to a week in certificates in the case of Let's Encrypt Oak.

To address these issues and ensure clarity, CT log operators should strive for uniformity in their practices. It is essential to use the UTC standard consistently and establish clear and consistent start and end dates for log periods. By adhering to these standards, the results derived from CT log analysis will be more accurate, reliable, and easier to interpret.

CT server	First	Last
DigiCertYeti-21	2021-01-01 00:00:00	2022-01-01 00:00:00
DigiCertYeti-22	2021-12-31 23:13:04	2022-12-31 22:06:56
DigiCertYeti-23	2022-12-31 23:00:16	2024-01-01 00:00:00
LetsEncryptOak-21	2021-01-01 00:00:00	2022-01-07 00:00:00
LetsEncryptOak-22	2021-12-31 23:00:16	2023-01-06 23:00:16
LetsEncryptOak-23	2022-12-31 23:00:16	2024-01-06 23:00:16

Table 3. Date of first and last certificate in each CT log in UTC standard

6 FUTURE WORK

In this section, we propose two possible directions for future research. First, we were often left guessing the reasons for shifts in logging behavior or choices in CT servers. Future work could address this issue and involve closer collaboration with CAs and log operators to gain access to information not publicly available during this research. Understanding the reasons behind logging policy shifts by CAs would provide valuable insights into the decision-making process. By establishing partnerships with CAs, researchers could understand the factors that influence changes in CT policies, such as emerging security threats, industry trends, or regulatory requirements.

Furthermore, conducting interviews or surveys with representatives from CAs and log operators can help shed light on their perspectives and motivations behind policy shifts. This approach would provide a more comprehensive understanding of the decision-making process and enable researchers to explore potential correlations between policy changes and external factors.

Second, as mentioned in the results, future work could look into the absence of Let's Encrypt certificates in Digicert Yeti starting in 2023. The lack of Let's Encrypt certificates is interesting, given their widespread usage and the open submission process. This anomaly contrasts other CT servers, where a consistent presence of Let's Encrypt certificates can be observed.

To explain this anomaly, further research could explore potential reasons behind the absence of Let's Encrypt certificates in Digicert Yeti. One possible speculation is that the increased bandwidth requirements resulting from Let's Encrypt's logging activities in Digicert Yeti in 2022 led to Let's Encrypt being blacklisted for subsequent log iterations. Understanding the technical or policy factors contributing to this possible exclusion would require closer collaboration with Digicert and Let's Encrypt, obtaining insights into the log operator's decision-making process and any constraints imposed on certificate submission. If Let's Encrypt were to be excluded against their will, it would go against the precedent and spirit of the CT ecosystem.

7 CONCLUSION

In this section, we summarize the findings of the analysis and answer the research questions.

SRQ1 *What logs did the top 5 most-used CAs rely on before the policy update to achieve CT compliance, and how did their logging practices change after the policy update?*
The top 5 most-used CAs did not exhibit significant changes in their logging practices after the policy update. This suggests a consistent adherence to their existing CT logging strategies, indicating that they already had robust systems in place that aligned with CT compliance requirements. Their logging practices did not show a move away from Google CT logs throughout the observed period.

SRQ 2: *What proportion of CAs have significantly reduced their reliance on Google CT logs after the policy update?*
The analysis revealed that approximately 30% of CAs have significantly reduced their reliance on Google CT logs after the policy update. However, it is important to note that none of the CAs have completely stopped publishing certificates in a Google CT log. This raises ambiguity about whether these reductions are temporary due to testing purposes or represent a definitive shift away from Google CT logs.

SRQ 3: *How have other browser vendors (Mozilla, Microsoft, Apple) updated their CT Policies in response to changes in the CT ecosystem, and what impact has this had on the WebPKI?*
Mozilla does not actively check CT status, while Microsoft's CT policies were not explicitly found but may align with Google's predefined policies. Google and Apple have implemented CT compliance requirements, with Google initially mandating SCTs from Google and different log operators. Over time, Google and Apple diverged in 2021, creating a more unclear CT landscape. However, with the recent Google policy update, the requirements match again, creating a more clear CT landscape.

RQ: *What is the impact of removing the explicit dependency on Google CT logs on CA compliance with CT requirements?*
In conclusion, the observed data suggests that the policy change had a minimal immediate impact on the WebPKI. The major CAs represent the bulk of newly issued certificates and have not notably updated their logging policy. Some smaller CAs have removed a small selection of certificates from the Google CT logs. However, The Google CT logs remain the preferred option for most CAs, indicating its popularity and stability. While the policy change has opened the door for potential shifts in logging behavior, the current analysis does not indicate any significant shifts in the short term. However, it is worth noting that the dynamic nature of the WebPKI ecosystem and evolving industry practices may lead to changes in the future.

REFERENCES

[1] 2016. 1281469 - Implement Certificate Transparency support (RFC 6962). https://bugzilla.mozilla.org/show_bug.cgi?id=1281469

[2] Maarten Aertsen, Maciej Korczyński, Giovane C. M. Moura, Samaneh Tajalizadehkhoob, and Jan van den Berg. 2017. No domain left behind: is Let's Encrypt democratizing encryption?. In *Proceedings of the Applied Networking Research Workshop*. 48–54. https://doi.org/10.1145/3106328.3106338 arXiv:1612.03005 [cs].

[3] Johanna Amann, Oliver Gasser, Quirin Scheitle, Lexi Brent, Georg Carle, and Ralph Holz. 2017. Mission accomplished?: HTTPS security after diginotar. In *Proceedings of the 2017 Internet Measurement Conference*. ACM, London United Kingdom, 325–340. https://doi.org/10.1145/3131365.3131401

[4] Ross Anderson, Chris Barton, Rainer Boehme, Richard Clayton, Carlos Ganan, Tom Grasso, Michael Levi, Tyler Moore, Stefan Savage, and Marie Vasek. [n. d.]. Measuring the Changing Cost of Cybercrime. ([n. d.]).

[5] Andrew Ayer. 2021. Yeti 2022 not furnishing entries for STH 65569149. https://groups.google.com/a/chromium.org/g/ct-policy/c/PCkKU357M2Q/

[6] Andreea Bendovschi. 2015. Cyber-Attacks – Trends, Patterns and Security Countermeasures. *Procedia Economics and Finance* 28 (Jan. 2015), 24–31. https://doi.org/10.1016/S2212-5671(15)01077-1

[7] Windows Experience Blog and Joe Belfiore. 2018. Microsoft Edge: Making the web better through more open source collaboration. https://blogs.windows.com/windowsexperience/2018/12/06/microsoft-edge-making-the-web-better-through-more-open-source-collaboration/

[8] Sharon Boeyen, Stefan Santesson, Tim Polk, Russ Housley, Stephen Farrell, and David Cooper. 2008. *Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile.* Request for Comments RFC 5280. Internet Engineering Task Force. https://doi.org/10.17487/RFC5280 Num Pages: 151.

[9] CA/BROWSER FORUM. 2021. Baseline Requirements for the Issuance and Management of Publicly-Trusted Certificates. 1.7.4 (May 2021).

[10] Franco Callegati, Walter Cerroni, and Marco Ramilli. 2009. Man-in-the-Middle Attack to the HTTPS Protocol. *IEEE Security & Privacy* 7, 1 (Jan. 2009), 78–81. https://doi.org/10.1109/MSP.2009.12 Conference Name: IEEE Security & Privacy.

[11] Google. 2022. Chrome CT Policy Update: Removal of '1 Google Log' SCT requirement and changes to number of required embedded SCTs. https://groups.google.com/a/chromium.org/g/ct-policy/c/507lPdbbwSk

[12] Josef Gustafsson, Gustaf Overier, Martin Arlitt, and Niklas Carlsson. 2017. A first look at the CT landscape: Certificate Transparency logs in practice. In *Passive and Active Measurement: 18th International Conference, PAM 2017, Sydney, NSW, Australia, March 30-31, 2017, Proceedings 18*. Springer, 87–99.

[13] Yuncong Hu, Kian Hooshmand, Harika Kalidhindi, Seung Jin Yang, and Raluca Ada Popa. 2021. Merkle2: A Low-Latency Transparency Log System. In *2021 IEEE Symposium on Security and Privacy (SP)*. 285–303. https://doi.org/10.1109/SP40001.2021.00088 ISSN: 2375-1207.

[14] International Telecommunication Union. [n. d.]. X.509 : Information technology - Open Systems Interconnection - The Directory: Public-key and attribute certificate frameworks. https://www.itu.int/rec/T-REC-X.509-201910-I/en

[15] Salabat Khan, Liehuang Zhu, Zijian Zhang, Mussadiq Abdul Rahim, Khalid Khan, and Meng Li. 2020. Attack-Resilient TLS Certificate Transparency. *IEEE Access* 8 (2020), 98958–98973. https://doi.org/10.1109/ACCESS.2020.2996997 Conference Name: IEEE Access.

[16] Nikita Korzhitskii and Niklas Carlsson. 2020. Characterizing the root landscape of certificate transparency logs. In *2020 IFIP Networking Conference (Networking)*. IEEE, 190–198.

[17] Bingyu Li, Fengjun Li, Ziqiang Ma, and Qianhong Wu. 2020. Exploring the Security of Certificate Transparency in the Wild. In *Applied Cryptography and Network Security Workshops (Lecture Notes in Computer Science)*, Jianying Zhou, Mauro Conti, Chuadhry Mujeeb Ahmed, Man Ho Au, Lejla Batina, Zhou Li, Jingqiang Lin, Eleonora Losiouk, Bo Luo, Suryadipta Majumdar, Weizhi Meng, Martín Ochoa, Stjepan Picek, Georgios Portokalidis, Cong Wang, and Kehuan Zhang (Eds.). Springer International Publishing, Cham, 453–470. https://doi.org/10.1007/978-3-030-61638-0_25

[18] J. Ronald Prins and Business Unit Cybercrime. 2011. Diginotar certificate authority breach "operation black tulip". *Fox-IT, November* 18 (2011).

[19] R. L. Rivest, A. Shamir, and L. Adleman. 1978. A method for obtaining digital signatures and public-key cryptosystems. *Commun. ACM* 21, 2 (Feb. 1978), 120–126. https://doi.org/10.1145/359340.359342

[20] Steven B. Roosa and Stephen Schultze. 2013. Trust Darknet: Control and Compromise in the Internet's Certificate Authority Model. *IEEE Internet Computing* 17, 3 (May 2013), 18–25. https://doi.org/10.1109/MIC.2013.27 Conference Name: IEEE Internet Computing.

[21] Quirin Scheitle, Oliver Gasser, Theodor Nolte, Johanna Amann, Lexi Brent, Georg Carle, Ralph Holz, Thomas C. Schmidt, and Matthias Waehlisch. 2018. The Rise of Certificate Transparency and Its Implications on the Internet Ecosystem. In *Imc'18: Proceedings of the Internet Measurement Conference*. Assoc Computing Machinery, New York, 343–349. https://doi.org/10.1145/3278532.3278562 WOS:000511429400030.

[22] Aozhuo Sun, Bingyu Li, Huiqing Wan, and Qiongxiao Wang. 2021. PoliCT: Flexible Policy in Certificate Transparency Enabling Lightweight Self-monitor. In *Applied Cryptography and Network Security Workshops (Lecture Notes in Computer Science)*, Jianying Zhou, Chuadhry Mujeeb Ahmed, Lejla Batina, Sudipta Chattopadhyay, Olga Gadyatskaya, Chenglu Jin, Jingqiang Lin, Eleonora Losiouk, Bo Luo, Suryadipta Majumdar, Mihalis Maniatakos, Daisuke Mashima, Weizhi Meng, Stjepan Picek, Masaki Shimaoka, Chunhua Su, and Cong Wang (Eds.). Springer International Publishing, Cham, 358–377. https://doi.org/10.1007/978-3-030-81645-2_21

[23] Ze Wang, Jingqiang Lin, Quanwei Cai, Qiongxiao Wang, Daren Zha, and Jiwu Jing. 2022. Blockchain-Based Certificate Transparency and Revocation Transparency. *Ieee Transactions on Dependable and Secure Computing* 19, 1 (Jan. 2022), 681–697. https://doi.org/10.1109/TDSC.2020.2983022 Place: Los Alamitos Publisher: Ieee Computer Soc WOS:000742730400048.

APPENDIX

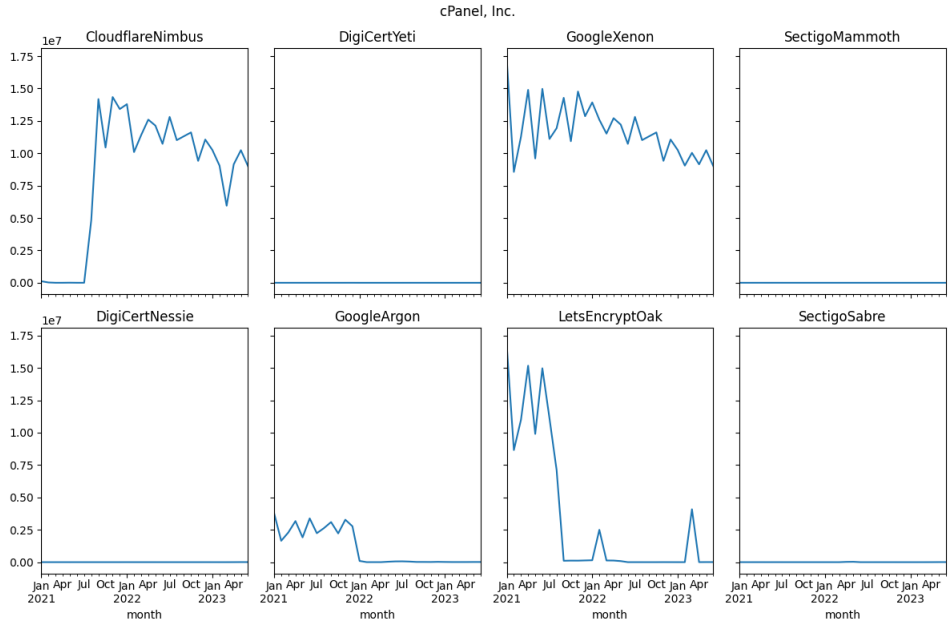

Fig. 5. For the CA *cPanel, Inc* absolute number of certificates published in each CT server

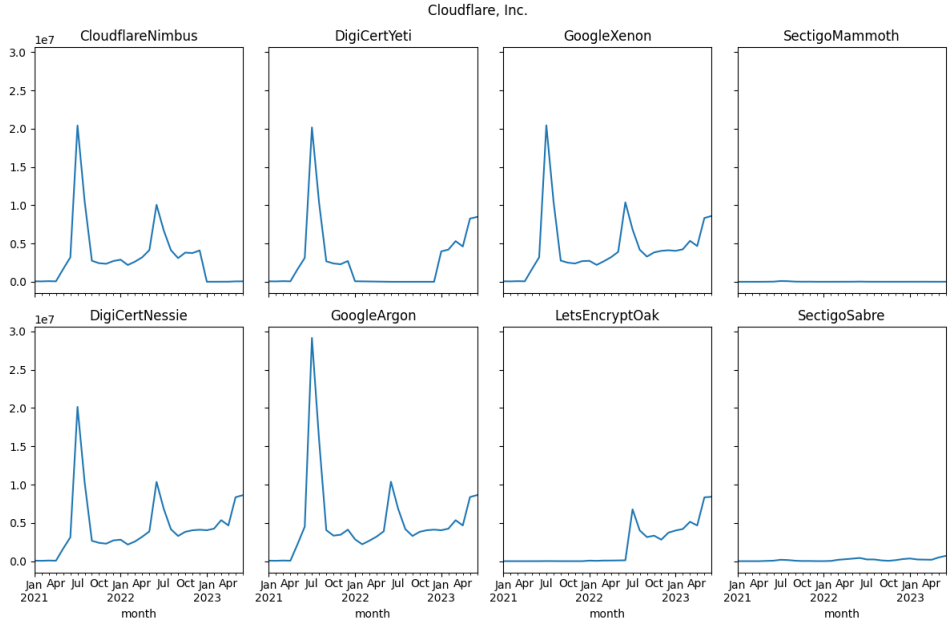

Fig. 6. For the CA *Cloudflare, Inc.* absolute number of certificates published in each CT server

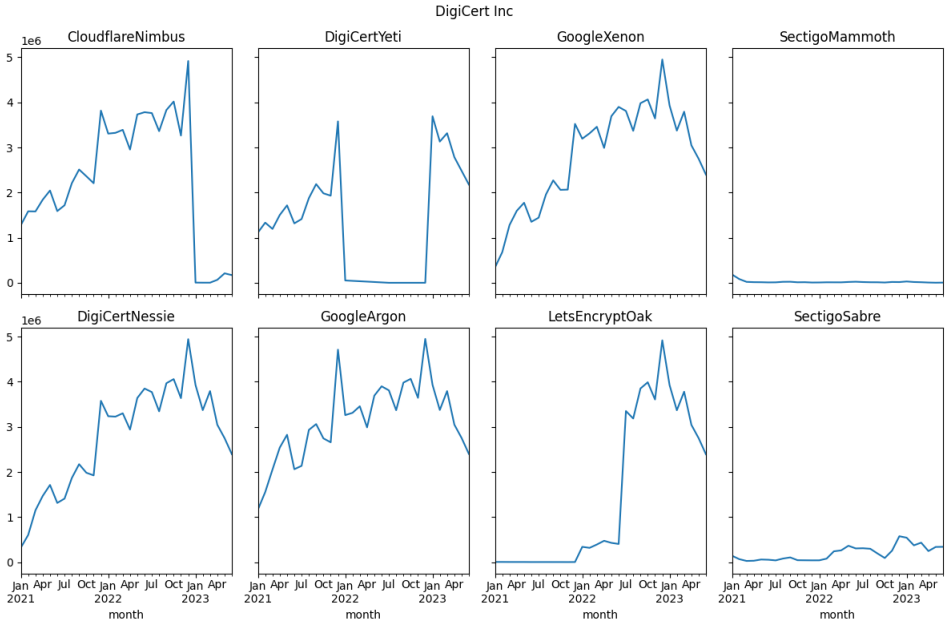

Fig. 7. For the CA *DigiCert Inc.* absolute number of certificates published in each CT server

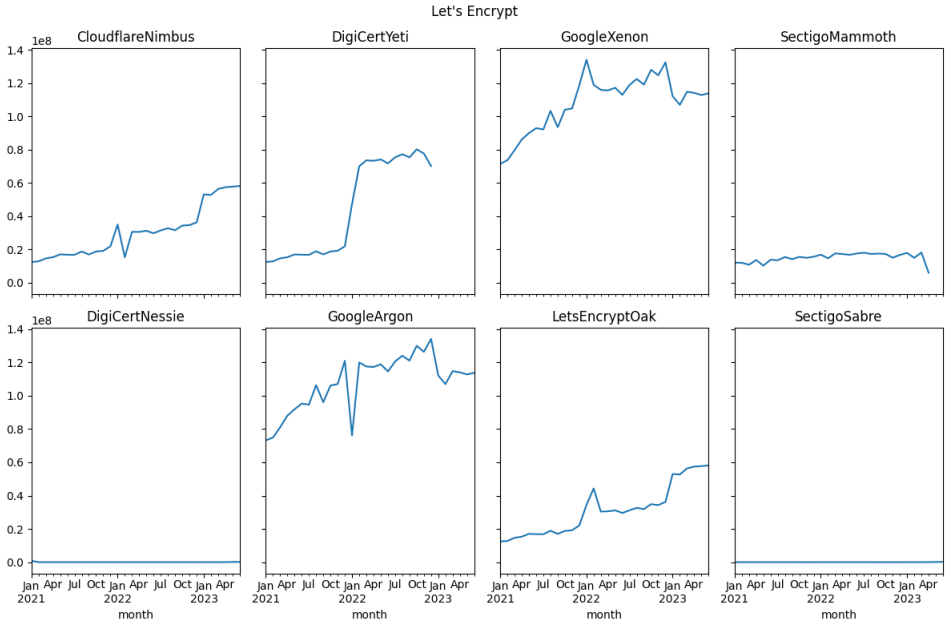

Fig. 8. For the CA *Let's Encrypt* absolute number of certificates published in each CT server

CA	Total	Before	After	Before Percentage	After Percentage
Let's Encrypt	124574763	124574514	124574274	100.0%	100.0%
cPanel, Inc.	11783099	11782773	11782764	100.0%	100.0%
Sectigo Limited	9861179	9861083	9848925	100.0%	99.88%

DigiCert Inc	4535227	4535227	4431845	100.0%	97.72%
Cloudflare, Inc.	3121651	3121651	3002122	100.0%	96.17%
Google Trust Services LLC	2692612	2692612	2692612	100.0%	100.0%
Amazon	2587572	2587568	2548832	100.0%	98.5%
GoDaddy.com, Inc.	2225948	2225948	2223726	100.0%	99.9%
ZeroSSL	1918905	1918885	1918859	100.0%	100.0%
Actalis S.p.A.	484803	484798	308503	100.0%	63.63%
GlobalSign nv-sa	446393	446393	362372	100.0%	81.18%
Microsoft Corporation	377442	377442	377290	100.0%	99.96%
TrustAsia Technologies, Inc.	269617	269617	269177	100.0%	99.84%
Starfield Technologies, Inc.	241974	241973	241846	100.0%	99.95%
Entrust, Inc.	195707	195707	151362	100.0%	77.34%
Gandi	145966	145966	145205	100.0%	99.48%
Cisco Systems, Inc.	125777	125777	123831	100.0%	98.45%
nazwa.pl sp. z o.o.	115872	115872	104723	100.0%	90.38%
Internet2	90821	90821	90479	100.0%	99.62%
GEANT Vereniging	87062	87062	86839	100.0%	99.74%
COMODO CA Limited	62165	62165	62083	100.0%	99.87%
Unizeto Technologies S.A.	60788	60788	46661	100.0%	76.76%
IdenTrust	48806	48806	45896	100.0%	94.04%
Japan Registry Services Co., Ltd.	42219	42219	42084	100.0%	99.68%
home.pl S.A.	40283	40283	30773	100.0%	76.39%
Network Solutions L.L.C.	29445	29445	29137	100.0%	98.95%
GoGetSSL	29064	29060	28758	99.99%	98.95%
QuoVadis Limited	22983	22983	18615	100.0%	80.99%
SECOM Trust Systems CO.,LTD.	22081	22081	21957	100.0%	99.44%
Verein zur Foerderung eines Deutschen Forschungsnetzes e. V.	21318	21318	18618	100.0%	87.33%
Cybertrust Japan Co., Ltd.	21097	21097	14559	100.0%	69.01%
TAIWAN-CA	18815	18813	12500	99.99%	66.44%
Buypass AS-983163327	13838	13838	13808	100.0%	99.78%
SSL Corporation	12894	12894	10174	100.0%	78.9%
Corporation Service Company	12170	12170	12030	100.0%	98.85%
Dreamcommerce S.A.	11563	11563	9200	100.0%	79.56%
The USERTRUST Network	11063	11063	11014	100.0%	99.56%
SwissSign AG	10732	10732	8446	100.0%	78.7%
Soluciones Corporativas IP, SL	9433	9433	9380	100.0%	99.44%
ATT Services Inc	9363	9363	9114	100.0%	97.34%
Trust Provider B.V.	8641	8641	8395	100.0%	97.15%
SecureCore	7326	7326	7184	100.0%	98.06%
QuoVadis Trustlink B.V.	7157	7157	7091	100.0%	99.08%
xíng zhèng yuàn (translated)	5611	5611	4678	100.0%	83.37%
T-Systems International GmbH	5218	5218	4967	100.0%	95.19%
Aetna Inc	4705	4705	4657	100.0%	98.98%
Trustwave Holdings, Inc.	4423	4423	3365	100.0%	76.08%
DHIMYOTIS	3902	3902	2993	100.0%	76.7%
DOMENY.PL sp. z o.o	2962	2962	2955	100.0%	99.76%
LH.pl Sp. z o.o.	2953	2953	2285	100.0%	77.38%
Beijing Xinchacha Credit Management Co., Ltd.	2932	2932	2264	100.0%	77.22%
WoTrus CA Limited	2779	2779	2640	100.0%	95.0%
CentralNic Luxembourg Sàrl	2650	2650	2615	100.0%	98.68%
Alpiro s.r.o.	2406	2406	2365	100.0%	98.3%
The Trustico Group Ltd	2269	2269	2246	100.0%	98.99%
Apple Inc.	2177	2177	2136	100.0%	98.12%
EUNETIC GmbH	2002	2002	1952	100.0%	97.5%
D-Trust GmbH	1960	1959	1626	99.95%	82.96%
Telia Finland Oyj	1918	1918	1367	100.0%	71.27%

CertCloud Pte. Ltd.	1752	1752	1752	100.0%	100.0%
TeliaSonera	1549	1549	1529	100.0%	98.71%
Fraunhofer	1484	1484	1459	100.0%	98.32%
Rede Nacional de Ensino e Pesquisa - RNP	1443	1443	1148	100.0%	79.56%
TBS INTERNET	1442	1442	1427	100.0%	98.96%
Chunghwa Telecom Co., Ltd.	1440	1440	1082	100.0%	75.14%
China Financial Certification Authority	1406	1406	1098	100.0%	78.09%
Deutsche Post AG	1304	1304	1274	100.0%	97.7%
TrustCor Systems S. de R.L.	1215	1215	1215	100.0%	100.0%
eMudhra Technologies Limited	1180	1179	1169	99.92%	99.07%
Open Access Technology International Inc	1158	1158	848	100.0%	73.23%
Hongkong Post	1092	1092	862	100.0%	78.94%
Wells Fargo Company	1078	1078	1046	100.0%	97.03%
NetLock Kft.	1038	1038	764	100.0%	73.6%

Table 6. Count of number of domains present in a Google CT log before and after the policy update

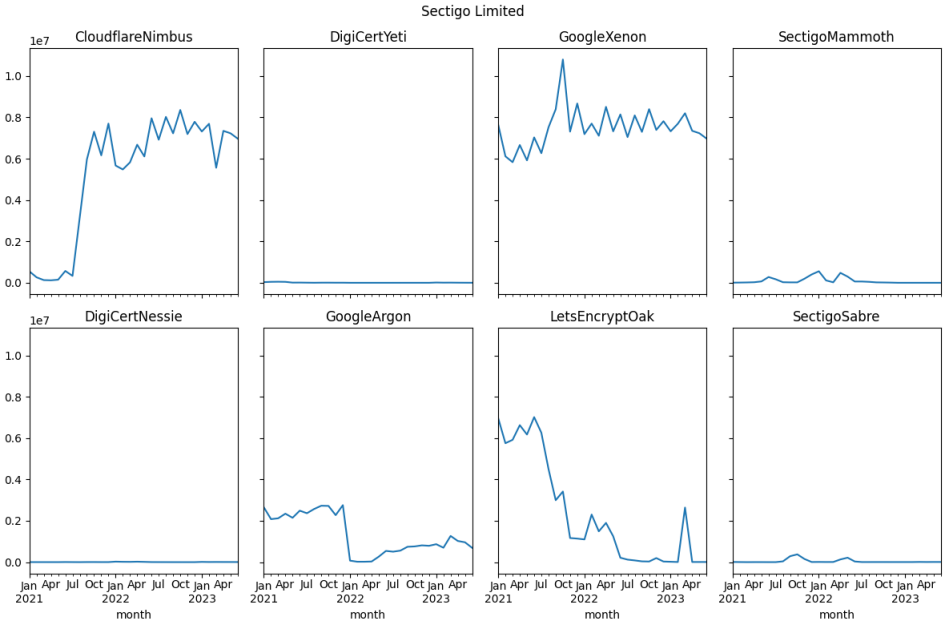

Fig. 9. For the CA *Sectigo Limited* absolute number of certificates published in each CT server

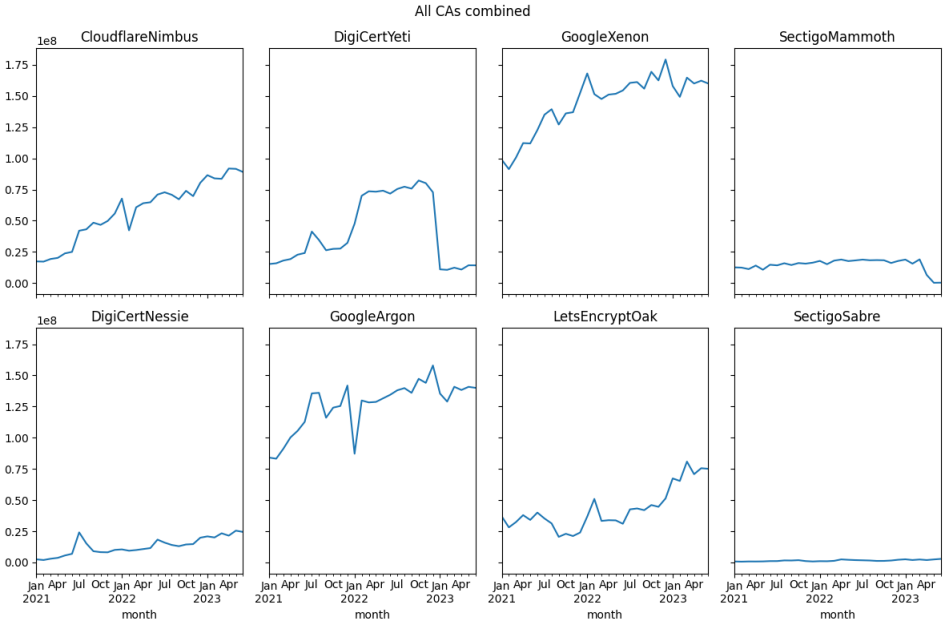

Fig. 10. Absolute number of certificates for all CAs published in each CT server

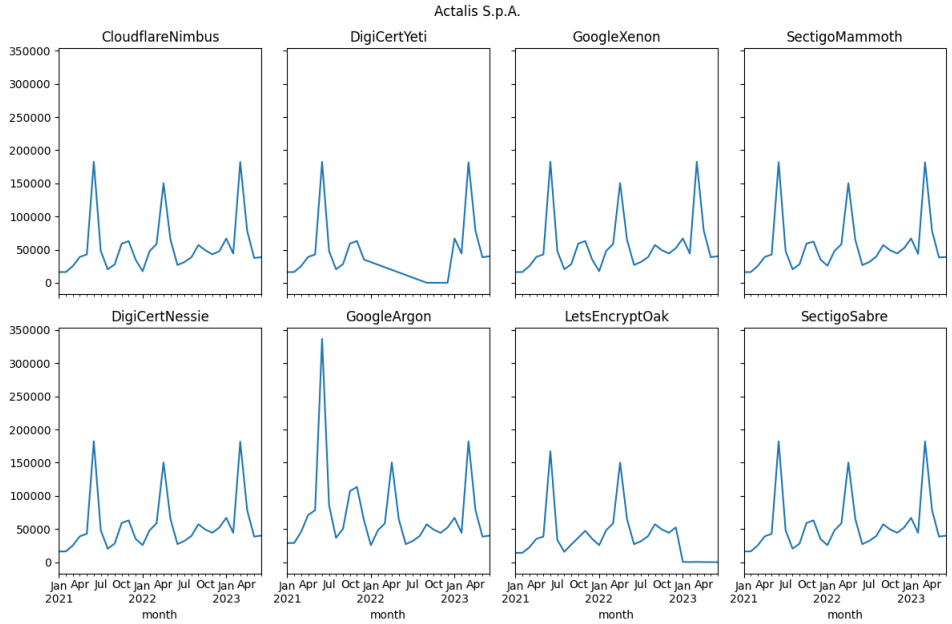

Fig. 11. For the CA *Actalis S.p.A.* absolute number of certificates published in each CT server

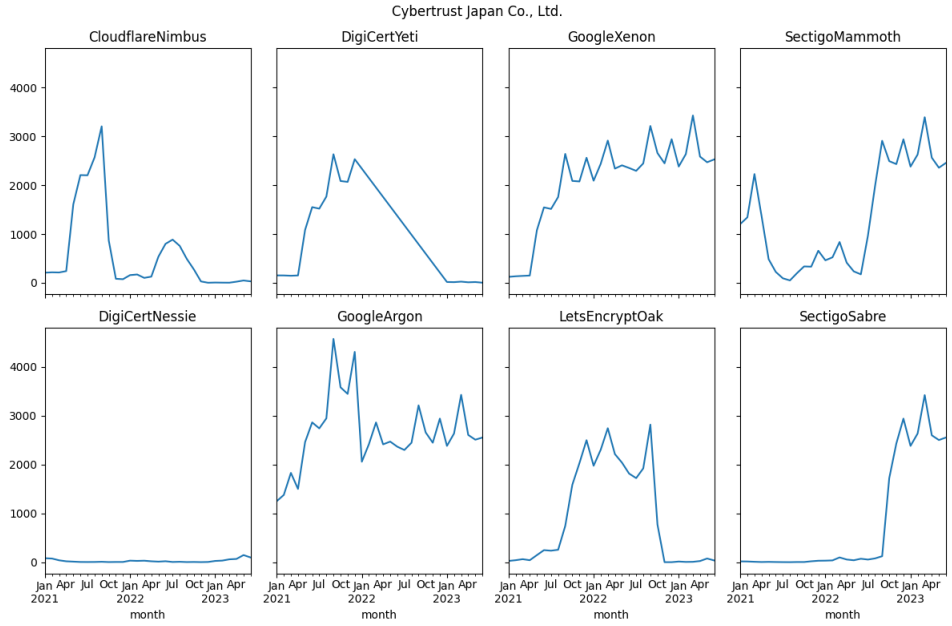

Fig. 12. For the CA *Cybertrust Japan Co., Ltd.* absolute number of certificates published in each CT server

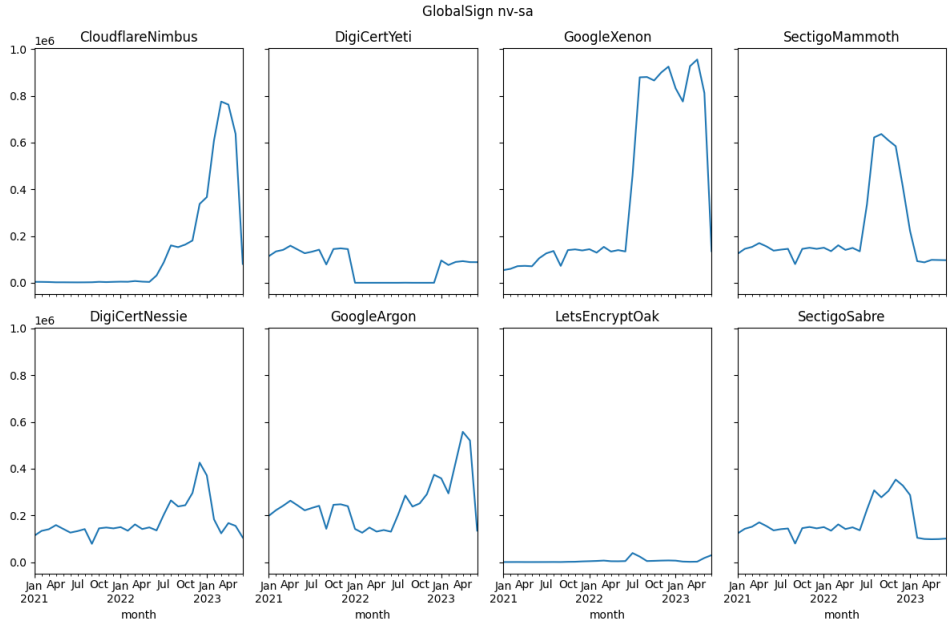

Fig. 13. For the CA *Globalsign nv-sa* absolute number of certificates published in each CT server

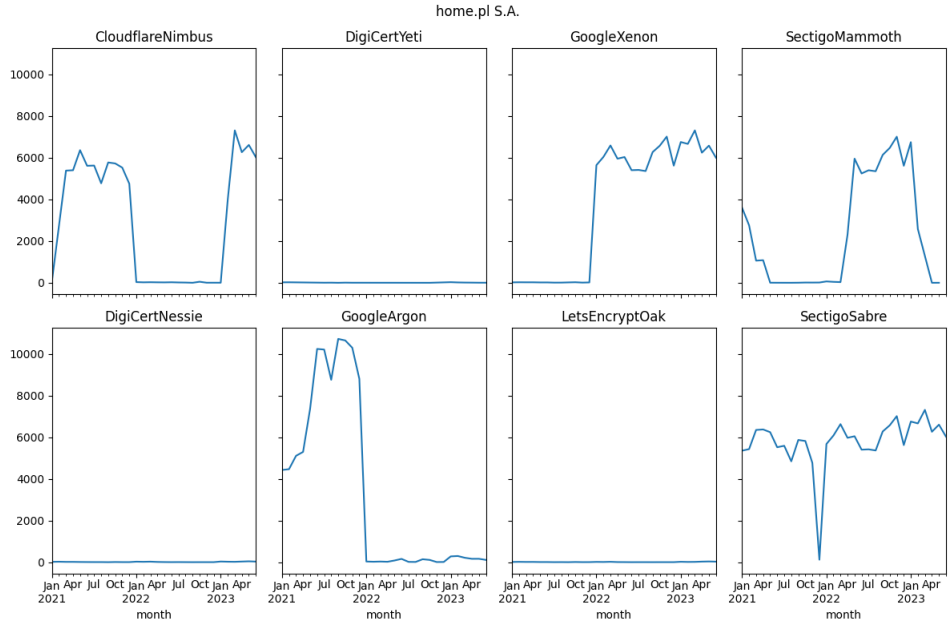

Fig. 14. For the CA *home.pl S.A.* absolute number of certificates published in each CT server

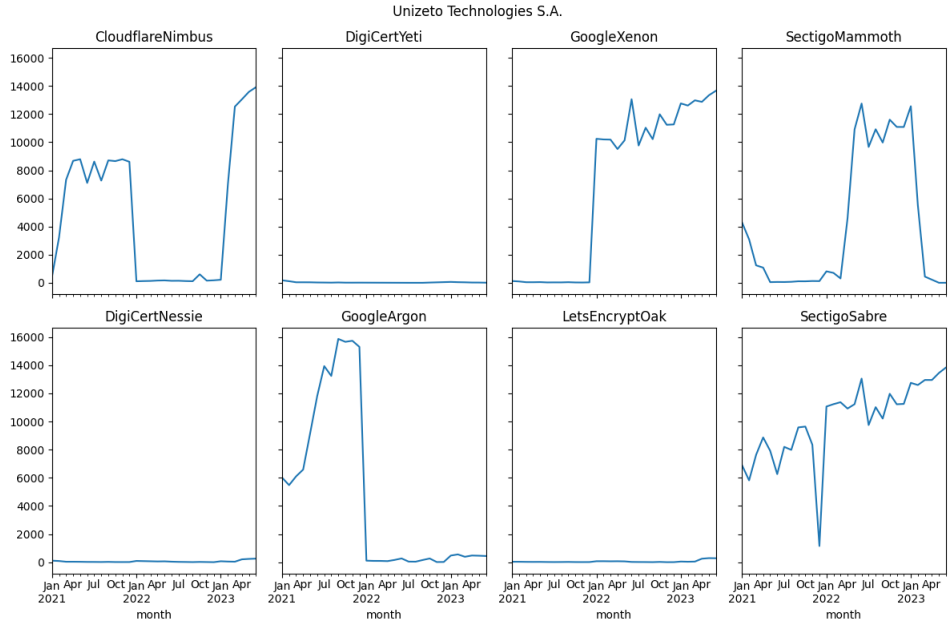

Fig. 15. For the CA *Unizeto Technologies S.A.* absolute number of certificates published in each CT server

CT server	First	Last
DigiCertNessie-21	2021-01-01 00:00:00	2022-01-01 00:00:00
DigiCertNessie-22	2021-12-31 23:00:16	2022-12-31 23:00:16
DigiCertNessie-23	2022-12-31 23:00:16	2024-01-01 00:00:00
DigiCertYeti-21	2021-01-01 00:00:00	2022-01-01 00:00:00
DigiCertYeti-22	2021-12-31 23:13:04	2022-12-31 22:06:56
DigiCertYeti-23	2022-12-31 23:00:16	2024-01-01 00:00:00
CloudflareNimbus-21	2020-12-31 23:02:24	2022-01-01 00:00:00
CloudflareNimbus-22	2021-12-31 23:00:16	2022-12-31 23:32:16
CloudflareNimbus-23	2022-12-31 23:00:16	2023-12-31 23:00:16
GoogleArgon-21	2021-01-01 00:00:00	2022-01-01 00:00:00
GoogleArgon-22	2021-12-31 23:00:16	2022-12-31 23:53:36
GoogleArgon-23	2022-12-31 23:00:16	2023-12-31 23:00:16
GoogleXenon-21	2020-12-31 23:00:16	2022-01-01 00:00:00
GoogleXenon-22	2021-12-31 23:00:16	2022-12-31 23:57:52
GoogleXenon-23	2022-12-31 23:00:16	2023-12-31 23:00:16
LetsEncryptOak-21	2021-01-01 00:00:00	2022-01-07 00:00:00
LetsEncryptOak-22	2021-12-31 23:00:16	2023-01-06 23:00:16
LetsEncryptOak-23	2022-12-31 23:00:16	2024-01-06 23:00:16

Table 4. Date of last and first certificate in each CT log in UTC standard