DNS Transparency Logs (Abstract)

Raffaele Sommese - Mattijs Jonker
{r.sommese, m.jonker}@utwente.nl

Introduction

Historically, access to zone files and domain information has been subject to confidential requirements by TLDs operators. The reason behind these choices was related to privacy issues regarding domain names and potential marketing implications.
Although it was possible to obtain access to zone files, the legal requirements in place raise a barrier for the research community, given the large number of (time-consuming) legal agreements required.
Moreover, in the past few years, with the explosion of new gTLDs, the number of zone files has increased tenfold.
To address this issue, ICANN created the Centralized Zone Data Service (CZDS), a service for requesting and obtaining access to zone files of different TLDs in a quick and standardized way. This service enables many researchers to perform studies on the DNS ecosystem.
Unfortunately, the zone files are published in CZDS once per day.
This granularity is sufficient for long-term trend analysis, but has limits when it comes to intraday events.
Our proposal is to create a system, analogous to Certificate Transparency (CT) logs, to facilitate access to changes inside the DNS (and especially TLD zones) with more fine-grained timing information.
Despite the primary goal of auditing possible fraudulent certificate issuance, CT logs have proven to be a fundamental resource for operators and researchers, enabling numerous studies on the X.509 certificate ecosystem.

Motivation, Requirements and Use Cases

The “DNS Transparency Logs” would provide near real-time insights into changes inside a zone file (e.g. insert of a new domain, nameserver, or glue records). Moreover, the log could be enriched with the information provided via RDAP (or WHOIS), such as registrar details and registration, renewal, and expiring date and times.
Several ccTLDs, such as .ch and .se, already allowed users to perform Incremental Zone Transfer Protocol (IXFR) queries to obtain zone updates. However, this possibility is strictly TLD-dependent.
Being able to detect changes inside the DNS ecosystem in real-time, can enable researchers to perform reactive measurements and/or correlate the changes with other events (e.g. DDoS attacks).
We identify several use cases:

1. **Detection of DNS hijacking events**: By having the real-time knowledge of DNS name server change in the infrastructure, the detection and the study of the impact of NS Hijacking events will be simplified.

2. **Behavioral analysis of newly registered domains**: Getting insights on newly registered domains as soon as possible can help researchers and operators analyze their behavior. This will allow them to quickly and proactively identify possible malicious usage.
3. **Operators' behavior for infrastructure under attack:** Monitoring live changes of DNS infrastructure during attacks can help researchers to understand and assess operators' practices.

The infrastructure required to enable the “DNS Transparency Logs” is already partly deployed for the communication between registries and registrars as Extensible Provisioning Protocol (EPP).

The main challenge to address is how to create a bridge to share the “non-sensitive” content with the third parties interested.

The EPP communications contain, in fact, information that could be subject to privacy concerns, such as registrant personal data. Also, business concerns should be taken into account: registries may want to hide information regarding the market share of registered domain names across different registrars. Finally, trademark protection is another concern, in order to avoid commercial exploitation of newly registered trademarks.

What are we hoping to discuss at DINR 2021

During the DINR 2021, we want to discuss the technical and legal challenges of the proposed platform and hear from the community about other possible use-case studies for this platform.