Cross-domain Tracking with TLS Session Resumption

Erik Sy
University of Hamburg
sy@informatik.uni-hamburg.de

1. DESCRIPTION

TLS is arguably the most important encryption standard in the Internet and can provide strong protection against network-based attackers and mass surveillance. To address these attacks, TLS needs to preserve the integrity and confidentiality of the payload and as far as practicable also metadata of a communication.

TLS session resumption makes use of an abbreviated handshake and leaks, by design, metadata such as an identifier for the previous session or the knowledge of a cryptographic secret. Thus, the comparison of these metadata between the initial and a resumed TLS connection allows an attacker to correlate these sessions with each other. As a result, TLS session resumption supports the creation of user profiles by online tracking services and mass surveillance. Tracking with TLS session resumption is independent from tracking mechanisms such as those based on HTTP Cookies or IP addresses. Hence, tracking with TLS session resumption can be used to improve existing tracking mechanisms by for example identifying a user after a change of his/ her IP address or the HTTP Cookies are deleted.

TLS session resumption contributes to performance enhancements during the resumed handshake, which allow a reduced number of round trips to complete the handshake or savings on computational expensive public key operations. These performance gains on the one hand, and the leakage of informations about previous sessions of a user on the other hand, present a tradeoff situation. To balance both of them, each TLS client and server can independently restrict the maximal time period between the initial and the resumed handshake. This time period is described in the following of this text as session resumption lifetime. While the draft of TLS 1.3 [4] proposes as an upper limit of the session resumption lifetime seven days, an evaluation of real-world session resumption lifetimes for day-to-day web browsing is still missing. In our talk, we provide an analysis of the server- and browser-side configuration of TLS session resumption. We noticed, that around 80% of the Alexa Top Million domains [3] with support for TLS session tickets, make use of resumption lifetimes of less than or equal to 10 minutes, while the remaining domains use considerably longer resumption lifetimes.

In regard of countermeasures, a simple method to clear the TLS cache of the web browser and thus prevent session resumption, is to shut down the application or device. However, especially smartphones are rarely shut down in order to stay connected to the mobile network, while their usage accounts for approximately the half of all web browsing sessions [5]. Furthermore, the web browser with its TLS cache can remain active for multiple days in the background of the smartphone's operating system. This leads to a situation, where even session resumption lifetimes of seven days can be achieved by common users, because the runtime duration of the application is even longer. In this talk, we report results of our analysis of the TLS session resumption behaviour of popular web browser, where we observed maximal session resumption lifetimes of two days.

Furthermore, we evaluate the impact of the session resumption lifetime on the capability of an online service to track its user. For this purpose, we use a pseudonymized DNS data set [2], to approximate the probability for online services with which they can recognise a user using TLS session resumption. Based on evaluations of this data set, with DNS traffic of 3859 users over a period of two months, a session resumption lifetime of seven days would lead to an average session resumption rate of 95,6% for all revisits of online services. However, popular online service achieve even higher session resumption rates, which allows them to almost continuously track their user.

Finally, we look into third party tracking, which refers to a practise, where a tracker, which is not the website directly visited, can track the user's visit to a website. Previous research [1] has shown that Google can track users across nearly 80% of the Alexa Top Million domains [3] by utilising its third party domains. Our results on the session resumption behaviour of popular web browser show, that third party tracking services can often track users across multiple domains due to the current browser configuration practise. We propose, that web browser apply besides the temporal limitation of TLS session resumption also a contextual restriction, such that session resumptions for third parties are only allowed in the context of the same visited website.

Overall, our results show cause for concern, but also encouraging signs. While TLS session resumption is an effective tracking mechanism on the session layer of the OSI model, and can be used by third parties to track users across multiple domains of their browsing sessions, it is a straightforward task for browser vendors to restrict this usage of TLS session resumption.

2. REFERENCES

[1] S. Englehardt and A. Narayanan. Online tracking: A 1-million-site measurement and analysis. In *Proceedings of the 2016 ACM CCS*, pages 1388–1401. ACM, 2016.

[2] D. Herrmann, C. Banse, and H. Federrath. Behavior-based tracking: Exploiting characteristic patterns in DNS traffic. *Computers & Security*, 39:17–33, 2013.

[3] Alexa Internet Inc. Alexa top 1,000,000 sites, 2018. URL: `s3.amazonaws.com/alexa-static/top-1m.csv.zip`.

[4] E. Rescorla. The Transport Layer Security (TLS) Protocol Version 1.3, 2018. URL: `https://tools.ietf.org/pdf/draft-ietf-tls-tls13-24.pdf`.

[5] StatCounter. Desktop vs Mobile vs Tablet Market Share Worldwide, 2018. URL: `gs.statcounter.com/platform-market-share/desktop-mobile-tablet/worldwide`.