Combining Website Fingerprinting Attacks against Tor users with DNS sniffing attacks

Benjamin Greschbach, KTH Stockholm
joint work with Tobias Pulls (Karlstad) and Philipp Winter,
Laura M. Roberts, Nick Feamster (Princeton)

May 15, 2016

I will present our ongoing work on analyzing the combination of Website Fingerprinting Attacks with DNS sniffing attacks on Tor exit relays. This is a relevant threat model for Tor users because recent measurements show that some DNS resolvers, such as Google's 8.8.8.8, are used by a substantial fraction of all Tor exit relays. We analyze the advantage that a website fingerprinting attacker would get if she controls or observes such a popular DNS resolver. We aim to develop recommendations for Tor exit relay operators and users how to protect against this new type of attack.

Tor, a low-latency, decentralized anonymization network can be used to conceal which websites a user is browsing. Tor users redirect their connections through three nodes of the Tor network: an entry guard, a middle relay and an exit relay, before they connect to the intended destination. For each hop the user adds another layer of encryption so that no node in the network sees both the original source and final destination of a connection.

When using Tor in a proper way, even DNS requests are redirected through the Tor network to avoid leakage of the information they contain to an adversary. When browsing a website through Tor, the exit relay that the client chose for this connection both resolves all necessary DNS names and fetches the website content. Exit relays can either resolve DNS names themselves or make use of recursive resolvers such as their ISP's DNS resolver or well known public resolvers such as Google's 8.8.8.8. Measurements suggest that Google's DNS resolver might see almost one third of all DNS requests issued by Tor exits.

Website Fingerprinting (WF) attacks aim to identify which website a user is browsing to by only looking at encrypted and anonymized traffic. A passive attacker that is located close to a Tor user, for example by sniffing the user's connection to a wireless hotspot, can use website fingerprinting attacks to try to identify which website the user is visiting.

Usually website fingerprinting attacks are seen as a classification problem: training data is collected by visiting a set of possible websites over Tor and collecting the resulting encrypted traffic traces between the user and the entry guard. Then a machine learning algorithm is trained on these labeled traces to build a classifier. When the target user visits a website, the adversary observes the encrypted traffic trace originating from the user's device and uses the classifier to determine which website the user most likely visited.

There are different opinions in the research community on how successful this attack can be. Some authors report very high accuracies for WF attacks and see them as a relevant threat for Tor users. Others doubt the practical feasibility of the attack under realistic assumptions. Several defenses against WF attacks have been proposed, some of them are already implemented in Tor. The practical feasibility and effectiveness of the defenses are also debated in the research community.

Combining WF attacks with DNS data: In our ongoing work, we are looking at the threat posed by an attacker that combines WF attacks on the Tor ingress side (close to the target user) with bulk-measurements of DNS requests on the Tor egress side (DNS requests issued by Tor exit relays). Even if the observed DNS requests cannot be linked to Tor users directly, they might improve the website fingerprinting attack accuracy when being used as additional feature or to reduce the candidate set of possible websites visited by the targeted Tor user.