A Measurement-based Investigation of DNS Hijacking (Abstract)

Authors: Rebekah Houser, Shuai Hao, Zhou Li, Daiping Liu, Chase Cotton, and Haining Wang
Presenter: Zhou Li, UC Irvine

DNS hijacking attacks have garnered substantial attention over the past few years, leading to a renewed drive to create and strengthen defenses against this type of event. The Sea Turtle campaign has been somewhat of a catalyst in this area. It has prompted several groups involved in security research to publish details of the attacks, recommendations for defenses, and warnings that this incident could be a forerunner of new and increasingly serious DNS-focused attacks. We expect to see many new or improved methods of DNS hijacking detection and prevention researched and implemented over the next few years. In this talk, I will describe a recent measurement study [1] we have done in measuring the real-world DNS hijacking attacks at large scale, through which we hope to provide some guidance on the development of new defense mechanisms.

There are two unique challenges that we try to address in this study. Firstly, there is a lack of clarity in threat models related to DNS hijacking, which has caused some confusion about which type of attack should be studied under this theme. Secondly, there is a lack of in-depth research into certain types of DNS hijacking attacks. Although there are some, such as MITM attacks, that have been studied closely, others, such as domain hijacking are unpredictable, and short-lived, on which no work has examined in-depth.

I will talk about how we address these two challenges. Firstly, I will present a taxonomy of DNS hijacking that aligns different attack vectors with the DNS infrastructure. Then, I will show the methodology and result of our new measurement study guided by the taxonomy. Specifically, by skimming over security incidents from 2008 to 2020, we identified 34 relevant incidents from news stories, retrieved over 27,000 Indicators of Compromises (IOCs), and augmented them with passive DNS data from Farsight's DNSDB, which informs us when an attack happens, how records were changed by the attacker, etc. By conducting quantitative and qualitative analysis on this dataset, we identified features that help to detect the presence of DNS hijacking, whose effects were examined with the real-world attacks logged by the dataset.

Yet, there are some limitations of this work at its current form and we are working on the improvements. First, the detection systems we built depend heavily on the appearance of new nameservers and AS owners in a domain’s DNS records, which exclude the situation that an attacker may be able to gain access to nameservers under the same apex domain. Second, the effectiveness of our detector is not very consistent across different testing periods (e.g., 0.57 for 2016-2020, compared to 0.73 for 2010-2013), which might be caused by the shift of attacker’s strategies. These observations suggest new datasets (with ground-truth), new features or machine-learning techniques might be needed for better results.

References:

[1] Rebekah Houser, Shuai Hao, Zhou Li, Daiping Liu, Chase Cotton, and Haining Wang. A Comprehensive Measurement-based Investigation of DNS Hijacking. In the proceeding of the 40th International Symposium on Reliable Distributed Systems (SRDS), virtual, September, 2021.