A longitudinal study of DNS traffic: Understanding current DNS practice and abuse

A thesis submitted in fulfilment of the requirements for the degree of MASTERS IN COMPUTER SCIENCE of RHODES UNIVERSITY

By Ignus van Zyl

December 2015

ABSTRACT

This thesis examines a dataset spanning 21 months, containing 3,5 billion DNS packets. Traffic on TCP and UDP port 53, was captured on a production /24 IP block. The purpose of this thesis is twofold. The first is to create an understanding of current practice and behavior within the DNS infrastructure, the second to explore current threats faced by the DNS and the various systems that implement it. This is achieved by drawing on analysis and observations from the captured data. Aspects of the operation of DNS on the greater Internet are considered in this research with reference to the observed trends in the dataset, A thorough analysis of current DNS TTL implementation is made with respect to all response traffic, as well as sections looking at observed DNS TTL values for ,za domain replies and NX DOMAIN flagged replies. This thesis found that TTL values implemented are much lower than has been recommended in previous years, and that the TTL decrease is prevalent in most, but not all EE TTL implementation. With respect to the nature of DNS operations, this thesis also concerns itself with an analysis of the geoloeation of authoritative servers for local (,za) domains, and offers further observations towards the latency generated by the choice of authoritative server location for a given ,za domain. It was found that the majority of ,za domain authoritative servers are international, which results in latency generation that is multiple times greater than observed latencies for local authoritative servers. Further analysis is done with respect to NX DOM AIN behavior captured across the dataset. These findings outlined the cost of DNS miseonfiguration as well as highlighting instances of NXDOMAIN generation through malicious practice. With respect to DNS abuses, original research with respect to long-term scanning generated as a result of amplification attack activity on the greater Internet is presented. Many instances of amplification domain scans were captured during the packet capture, and an attempt is made to correlate that activity temporally with known amplification attack reports. The final area that this thesis deals with is the relatively new field of Bitflipping and Bitsquatting, delivering results on bitflip detection and evaluation over the course of the entire dataset. The detection methodology is outlined, and the final results are compared to findings given in recent bitflip literature.