

Lab Session 1

The Internet of Names: Big Data Analysis for DNS

Anna Sperotto¹, Mattijs Jonker¹, Christian Dietz²

¹University of Twente, The Netherlands

²Universität der Bundeswehr München, Germany

a.sperotto@utwente.nl, m.jonker@utwente.nl, christian.dietz@unibw.de

Abstract. The Domain Name System (DNS) is part of the core infrastructure of the Internet. Tracking changes in the DNS, therefore, provide valuable information about the evolution of the Internet. Think about adoption of protocols (e.g., IPv6 and DNSSEC) and applications (e.g., cloud e-mail providers), distribution of content (Web domains), and network security (e.g., botnets). Since February 2015, the University of Twente, SURFnet, and SIDN run a largescale active measurement of the DNS, which cover the domain names in the .com, .net, and .org zones. Since February 2016, the .nl zone has also been added. In total, our measurement currently queries over 50 % of the DNS name space on a daily basis. The measurement results are stored in an Hadoop cluster for later analysis [1].

The goal of this hands-on tutorial is to familiarize the participants with DNS, DNS measurements, and possible research application. The session will start with a general introduction to the measurement including a few example use cases. Then, we will briefly introduce the participants to a virtualized lab environment, in which they can experiment with the data themselves. The remainder of the session is then spent “hackathon”-style, in groups, each of which will present their experiences and possible findings from the data at the end of the session in a short presentation. The lab environment will contain real data for the Alexa Top 1 Million domains.

Reference

1. van Rijswijk-Deij, R., Jonker, M., Sperotto, A., Pras, A.: The internet of names: a DNS big dataset. In: SIGCOMM 2015 Poster Paper, ACM, London, UK, August 17–21 (2015)