Towards an Independent and Resilient DNS

Moritz Müller and Roland van Rijswijk-Deij

Stichting Internet Domeinregistratie Nederland (SIDN)
and University of Twente, The Netherlands
{m.c.muller, r.m.vanrijswijk}@utwente.nl

Abstract. The Domain Name System (DNS) is the crucial naming system of the Internet. Before clients can establish a connection with a service they look up its address in the DNS. Therefore, DNS operators have a large responsibility and a lot of power: if the DNS is not available, clients effectively cannot establish a connection to a service. Also, DNS operators can track and manipulate requests by clients. This can leak personal information, lead to censorship and cause security issues.
DNS is designed to be in the hands of many. This makes it more resilient against outages and independent from parties that misuse their power. In practice, however, the DNS becomes concentrated at a few providers – often with negative consequences: e.g., when the DNS provider Dyn suffered an outage in 2016, many services on the Internet went offline.
In this PhD research we aim at studying and improving the resilience and independence of the DNS infrastructure. We analyze where on the Internet the DNS is concentrated and discuss the impact of these points of concentration on the resilience and independence. Further, we propose measures to mitigate these impacts.

1 Introduction

The Domain Name System (DNS) is *the* naming system of the Internet. It is critical because a lookup in the DNS precedes almost every connection setup.

By design, the DNS is a distributed system. Its name-space is hierarchical, where each zone (the root, .com, and example.com) is served by authoritative name servers. These are queried by recursive resolvers, which clients use to lookup information in the DNS.

Figure 1 demonstrates such a lookup. Assume the client wants to visit the website *example.com*. To do so, it first needs to look up the IP address of *example.com* in the DNS. It employs a recursive resolver that walks through the DNS hierarchy until it reaches the name server that is authoritative for the requested domain. The server responds with the IP address and the resolver returns it to the client. Finally, the client uses the address to connect to the website.

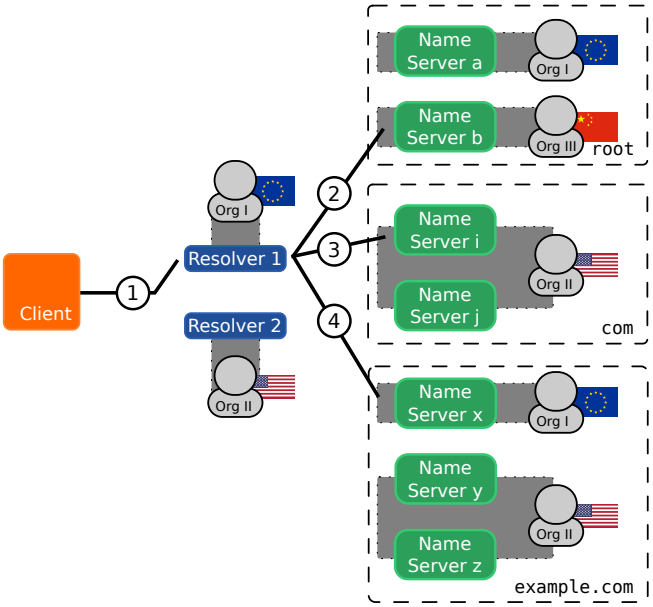

Fig. 1. An example of a DNS query and its involved components and stakeholders.

This architecture gives DNS operators a lot of responsibility but also a lot of power: on one side, operators of zones like the root, .com, or example.com have the responsibility to keep their name servers *available* all the time. If in Figure 1, none of the name servers of example.com are reachable, the services of example.com become effectively unavailable. This responsibility becomes even larger if the operator manages zones with many domains, like the root or a TLD.

On the other side, operators of recursive resolvers have a lot of power over their clients. First, operators can track every DNS request, which can reveal sensitive and possibly confidential information [1]. Second, because a DNS request precedes virtually every connection setup on the Internet, resolver operators are gatekeepers. They can block access to certain domains or serve their clients a wrong answer. Last, they also need to keep their resolvers available all the time. If clients cannot reach a resolver to serve their queries, they become effectively offline.

To distribute the responsibility and power, the DNS is designed to be in the hands of many [2]: as shown in Figure 1, clients can select from multiple resolvers. Also, each domain can have multiple name servers and each server can be operated by a different provider.

This design should make the DNS resilient against outages and avoids that every component is controlled by the same organization. If the authoritative name server x in Figure 1 of example.com fails, then there are still two other servers available. Also, if the client does not trust resolver 1 to serve the correct answers it can choose resolver 2 instead.

Incidents in the past, however, indicate that the actual implementation of the DNS does not follow these principles anymore. A major Distributed Denial of Service attack on the DNS provider Dyn in 2016 rendered many popular web services unavailable for hours [3] and affected domains like twitter.com, linkedin.com and airbnb.com. All of which had multiple name servers configured, but every one of them was hosted at Dyn. When Dyn became unavailable, so did every name server.

Similar trends have been observed at recursive resolvers as well. Internet service providers (ISPs) are forced by governments to block DNS requests, e.g. to adult content [4]. At the same time they make it harder for their customers to choose other resolvers than the one provided by the ISP [5]. If the ISP has to block certain queries at their resolvers a (less tech-savvy) user has no other choice than to accept this censorship.

This indicates, the DNS in fact becomes more concentrated at a few organizations. This apparently has negative implications for the resilience and independence of the DNS and thereby, the Internet as a whole.

Several proposals have tried to address at least parts of these issues but usually involved a complete re-design of the DNS and did not see any wide spread deployment[6].

In this PhD research we carry out a comprehensive study of all aspects that influence the resilience and independence of the DNS. Our goal is to analyze (*RQ1*) how concentrated the DNS is, (*RQ2*) study its implications on the resilience and independence and discuss (*RQ3*) countermeasures that can be applied to the existing DNS architecture.

Previous research shows that name servers of popular domain names are largely concentrated at a few providers [7]. We choose a broader view: we carry out a comprehensive study of the domain name space and take all aspects that influence the resilience and availability of the DNS into account.

2 Goal, Research Questions, and Approach

The ultimate goal of this research is to:

> **Goal** Study and (further) improve the resilience and independence of the DNS infrastructure.

To achieve this goal we have to answer three research questions: in *RQ1* we lay the foundation for *RQ2* and *RQ3*.

> **RQ1** How are components of the DNS distributed across networks, organizations and states?

The purpose of this research question is to measure if and where on the Internet the DNS is concentrated and examine its drivers. Concentration can occur on a technical, organizational and national level and can be motivated by technical, economical or social drivers.

We will study the following components: (a) the recursive resolvers (boxes in the middle column in Figure 1), (b) the authoritative name servers (right column), (c) the network links that connect the components (bold lines), and (d) the managing organizations and countries.

We, among others, use the active DNS measurement platform OpenINTEL to map which name servers are authoritative for the domain space[1]. We want to understand in which networks these servers are located, who is responsible for running the servers, and under which jurisdiction those organizations fall.

Similar research is necessary for recursive resolvers. In order to study their concentration, we need to understand, which resolvers are "important" for the DNS eco-system. Based on passive DNS query data, collected at TLDs and the root, as well as query data of recursive resolvers, we will develop and evaluate a methodology to identify resolvers that serve large parts of the Internet users [8,9].

As soon as we have identified these resolvers, we measure where they are located in the network infrastructure and who is responsible for them. This helps us to answer the question:

> **RQ2** Is the current DNS infrastructure a threat to the resilience and independence of the DNS?

We develop threat scenarios and apply them to the points of concentration, identified in *RQ1*. In these scenarios the availability and integrity of the DNS, but also the privacy of its users are threaten. We base them on realistic threats observed in the wild, for example: 3rd-parties that want to take down critical components of the DNS, but also the operators of these components. They manipulate the DNS, e.g. for their own interest or because they are forced by a nation state, and can cause unintentional outages or changes in the DNS through mis-configuration.

Because we have identified the important components of the DNS in *RQ1* we can measure the impact of the identified threats on the users of the DNS. These threats lead us to the third and last research question:

> **RQ3** Which countermeasures are suitable to (further) increase the resilience and independence?

We focus on countermeasures that address the most harmful and probable threats, identified in *RQ2* and discuss countermeasures on a technical, operational and organizational level. In our approach, we assess the effect and practicality of existing measures using testbeds and production systems (e.g. at the Dutch ccTLD operator). Further, our goal is to design novel approaches as well that have not been proposed at standardization organizations or implemented before. Countermeasures include mechanisms to make resolvers more resilient against outages of name servers, methodologies to reduce the risk of operational failures, or regulatory measures that split up points of concentration.

[1] https://www.openintel.nl

We demonstrate that we have achieved our research goal by using qualitative and quantitative metrics. We deploy our proposed countermeasures at resolvers and name servers in production and can rely on anecdotal evidence from operators to prove their efficiency.

3 Preliminary Results

In the previous months, we took first steps to answer *RQ2* and *RQ3*.

To identify threats in *RQ2* we need to understand how resolvers interact with name servers. We measured their behavior and observed that more than 5% of the resolvers in the wild send every query to only one name server [10]. If this server becomes unavailable the resolvers are, temporarily, unable to receive an answer to their query. These resolvers threaten the availability of the DNS.

One reason for unavailable name servers are mis-configurations of the DNS Security Extensions (DNSSEC). DNSSEC is designed to protect the integrity of the DNS but can also cause outages if not configured correctly. We developed and tested a methodology that reduces the risk of outages for DNSSEC operators [11]. Thereby, we increase the resilience of the DNS and contribute to *RQ3*.

Acknowledgment This PhD research is carried out in part time with the support of SIDN (https://www.sidn.nl) and is supported by SIDN Labs and NLnet Labs through the project Self-managing Anycast Networks for the DNS (SAND) project, phase 3 (http://www.sand-project.nl).

References

1. Bortzmeyer, S.: DNS Privacy Considerations. RFC 7626 (Informational) (2015)
2. Mockapetris, P.: Domain names - concepts and facilities. RFC 1034 (November 1987)
3. Dyn Inc.: Incident Report: DDoS Attack Against Dyn Managed DNS. https://www.dynstatus.com/incidents/nlr4yrr162t8 (2016)
4. BBC: Q&A: UK filters on legal pornography. http://www.bbc.com/news/technology-23403068 (2013)
5. Edmonds, Robert: Disappearing Choice of Recursive DNS Services in Home Networks. https://indico.dns-oarc.net/event/28/session/11/contribution/58/material/slides/0.pdf (2018)
6. Kalodner, H.A., Carlsten, M., Ellenbogen, P., Bonneau, J., Narayanan, A.: An empirical study of namecoin and lessons for decentralized namespace design. In: WEIS. (2015)
7. Bates, S., Bowers, J., Greenstein, S., Weinstock, J., Zittrain, J.: Evidence of Decreasing Internet Entropy: The Lack of Redundancy in DNS Resolution by Major Websites and Services. Working Paper 24317, National Bureau of Economic Research (2018)
8. Wullink, M., Moura, G.C., Müller, M., Hesselman, C.: ENTRADA: A high-performance network traffic data streaming warehouse. In: Network Operations and Management Symposium (NOMS), 2016 IEEE/IFIP, IEEE (2016) 913–918

9. DNS OARC: DITL Traces and Analysis. https://www.dns-oarc.net/oarc/data/ditl/2017 (2017)
10. Müller, M., Moura, G.C.M., de O. Schmidt, R., Heidemann, J.: Recursives in the Wild: Engineering Authoritative DNS Servers. In: Proceedings of the 2017 Internet Measurement Conference. IMC '17, New York, NY, USA, ACM (2017) 489–495
11. Müller, M., Chung, T., van Rijswijk-Deij, R.: Rolling with Confidence: Managing the Complexity of DNSSEC Operations (Poster). In: ICT-Open Proceedings. (2018)