Table of Contents

A Abstract

Distributed Denial of Service (DDoS) attacks are some of the most effective and dangerous attacks faced by Internet services today. DDoS attacks both misuse and target core Internet infrastructures and services. One of these is the Domain Name System (DNS). The DNS performs the crucial task of translating human-readable domain names into IP addresses, but also support most Internet applications, content distribution platforms, and many security services. DDoS attacks on DNS infrastructure can thus have devastating effects. In a nutshell, attacks that compromise the DNS can severely disrupt the Internet itself. To the extent that individual networks can minimize the risks of such attacks, and mitigate their impacts when they do occur, we believe rigorous and systematic measurement and analysis are essential to quantifying these risks as well as the benefits of proposed solutions.

The goal of this proposal is to comprehensively analyze the DDoS ecosystem targeting the DNS – attack sources, targets, and characteristics observed in DDoS attack traffic data – and to assess vulnerabilities and single points of failure that threaten the resilience of the DNS under such DDoS attacks. Combining these two perspectives will yield a clear view on the threat landscape facing the DNS, and generates actionable intelligence enabling real-world improvements to the resilience of the DNS against attacks. The intelligence generated by the MADDVIPR project will aid *protection* of the DNS as well as facilitate *prevention* of attacks against the DNS.

In more detail, this proposal has two pillars:

1. **Identifying DNS single points of failure and vulnerabilities** – The first pillar of MADDVIPR is to map single points of failure in the global DNS and vulnerabilities in DNS deployment that DDoS attacks can exploit. Examples include: mismatches between DNS parent- and child-zones; concentration of DNS operations with single organizations or facilities; and cross-domain vulnerabilities, where attacks on one domain affect many other domains.

2. **Mapping the DNS DDoS ecosystem** – While DDoS attacks against the DNS have clearly proven their efficacy (consider, for example, the Nov-Dec 2015 attacks against the DNS root, or the more recent attack against DNS operator Dyn), our knowledge of these attacks relies mostly on anecdotal evidence. In the MADDVIPR project, our first goal is to systematically investigate the DDoS ecosystem for attacks against the DNS. This implies gaining a structural understanding about attack sources, attacks and targets, identifying trends and, ultimately, mapping the risk zones of the Internet.

We will leverage the strengths of both project partners to collaboratively synthesize our two perspectives into a coherent unified view, enabling: a) immediate actionable intelligence on ongoing attacks against the DNS, b) insight into high-priority risks for DNS operators, based on historical information on attacks and analysis of vulnerabilities and single-points-of-failure, and c) clear guidance for improving DNS resilience to prevent future attacks.

B Performance Goals

This proposal intends to address the challenges raised in the second technical topic area of the call, *Distributed Denial of Service (DDoS) Attacks and the Domain Name System (DNS)*. In particular, this proposal focuses on tackling the first subtopic in that area, *a. Protection of the DNS against DDoS attacks*.

As the call states: *"If you can stop the DNS, you effectively stop most Internet communication"*. This strong statement underlines the critical role that the DNS plays in almost all Internet services, from the web to e-mail to voice and video streaming and conversation. Our dependency on the DNS makes it imperative (as the call states) to protect the DNS against DDoS attacks. But before we can protect the DNS, we first need to know *what we are protecting against* and *what we are protecting*. The main objective of this proposal is to develop and extend tools that generate *actionable intelligence* to support *protecting* the DNS against DDoS attacks, including, wherever possible, to *prevent* attacks before they occur.

We propose to pursue this objective by setting the following goals: (a) perform a systematic analysis of single points of failure and vulnerabilities in the current DNS using comprehensive active DNS measurement data from the OpenINTEL project that covers 60% of the global DNS name space; (b) analyze the DNS DDoS ecosystem, identifying attack sources, targets, and characteristics; and (c) synthesize the results of goals (a) and (b) into a coherent unified view that considers both ongoing attacks and current vulnerabilities. In the process we will develop and operate a measurement and analysis system that generates actionable intelligence that can protect the DNS against DDoS attacks, and also provides recommendations for prophylactic measures that minimize the risk and impact of such attacks.

C Detailed Technical Approach

C.1 Introduction

Distributed Denial of Service attacks are one of the most critical cyber-threats on the modern-day Internet. They are cheap, effective, and they keep growing in intensity. Widely reported examples of attacks include: the attack on Spamhaus in 2013 [3] (300Gbps), or on a Cloudflare customer in 2014 [13], or more recently, the attacks against the hosting company OVH (2016, 1Tbps) [12] and the attack leveraging the firepower of the Mirai botnet against the service and DNS provider Dyn [5] (2016, allegedly 1 Tbps).

The Domain Name System forms part of the core of the Internet. DNS provides the vital function of translating human readable domain names into IP addresses, thus acting as the phone book of the Internet. It also serves as a support infrastructure for most applications, commercial content distribution platforms, and many security services [7]. DDoS attacks against the DNS can therefore have devastating effects (consider, for example, the fallout of the attack on Dyn [5]). To the extent that individual networks can minimize the risks of such attacks, and mitigate their impacts when they do occur, we believe rigorous and systematic measurement and analysis are essential to quantifying these risks as well as the benefits of proposed solutions.

Securing the DNS against DDoS attacks is not trivial. First of all, the DNS is a highly structured and distributed infrastructure, which makes the DNS inherently more robust against attacks and

incidents. But it also means that the DNS probably has one of the largest attack surfaces among the core Internet infrastructures.

Second, as a consequence of its distributed nature, the DNS is managed in widely different ways, and local policies often shape management and security decisions. For example, the adoption of DNSSEC, although improving, is still as low as 3.5% of domains globally [15].

Third, if we look at how DDoS attacks on the DNS are performed, such attacks are often indistinguishable from regular DNS traffic, making mitigation hard. Even worse, current mitigation solutions may dampen the attack, but introduce new problems. For example, one mitigation approach is *ingress traffic rate limiting*, which limits the amount of DNS traffic reaching a DNS server. However, since it is often impossible to distinguish DNS attack traffic from legitimate traffic, rate limiting will not stop only malicious traffic, but also potentially penalize regular users.

The rise of DDoS attacks has paved the way to a new market for DDoS protection systems, i.e. appliances and services aimed at stopping malicious traffic from hindering a certain service [7]. In the Netherlands, a notable initiative in this respect is NaWas *(de nationale anti-DDoS Wasstraat)*, a *clearing house* that mitigates DDoS attacks by serving as a sink for DDoS traffic, preventing traffic overload (and thus combating the denial-of-service attack) of a victim network. However, clearing houses are not always a feasible solution for all customers, both for economic and for privacy reasons. Moreover, they introduce a single point of failure by aggregating traffic toward a single entity, such as the attack on Dyn illustrates.

In this proposal, we advocate taking a step back to reconsider: in order to protect the DNS against DDoS attacks, we first need to analyze *what we are protecting* and *what we are protecting against*.

To do this, we propose to map vulnerabilities and single points of failure that threaten the resilience of the DNS under attack, and to comprehensively map the DDoS ecosystem targeting the DNS (attack sources, attacks and targets). Combining these two perspectives will create a clearer view of the threat landscape facing the DNS, and generates actionable intelligence enabling real-world improvements to resilience against attacks of the DNS itself, and advancing the state-of-the-art beyond the need to urgently resort to clearing houses for protection. The main goal of the MADDVIPR project is to generate intelligence that will offer *protection* of the DNS as well as facilitate *prevention* of attacks against the DNS.

Our approach is based on the following pillars, which we will discuss in detail in the following sections: (i) *Identifying DNS single points of failure and vulnerabilities*, which aims at identifying how the DNS can become the target of future DDoS attacks; (ii) *Analyzing the DNS DDoS ecosystem*, which will provide a comprehensive overview of current DDoS attacks against the DNS by studying the attackers, the attacks and targets; and (iii) developing the MADDVIPR framework: a coherent, unified view of the DNS DDoS ecosystem and the DNS single points of failure and vulnerabilities, that will yield actionable information for operators on how to improve their own DNS resilience against DDoS attacks.

A core part of our approach is that we will use the unique datasets that both project partners have created, maintain, and use. The University of Twente will contribute data from the OpenINTEL project [16][1], which collects daily active measurements of 60% of the global DNS name space, including the main top-level domains such as `.com`, `.net` and `.org`, and a growing number of country-code top-level domains, including the `.nl` ccTLD. OpenINTEL has been collecting

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

data since February 2015, and thus provides a unique longitudinal view of the evolution of large parts of the DNS. CAIDA will contribute data from the UCSD Network Telescope[2], a large (/8) globally routed block of IPv4 address space that CAIDA monitors for incoming traffic. This instrumentation provides visibility into, among other things, backscatter from ongoing DoS attacks [2, 11].

C.2 Pillar (1) – Identifying DNS single points of failure and vulnerabilities

Protecting the DNS against DDoS attacks starts with correctly selecting and configuring DNS servers. Conversely, studying existing DNS configurations can be remarkably revealing of DNS vulnerabilities. For this project we consider two vulnerabilities in particular: (i) *single points of failure* and (ii) *vulnerabilities due to misconfigurations or suboptimal configurations*. We define a single point of failure as the situation where authoritative information for a domain is available from only a single DNS operator, i.e., there is no redundant source in case this single source is unreachable. Potential misconfigurations, or suboptimal configurations are, for example, mismatches between DNS delegations in the parent and child zone or cross-domain vulnerabilities (if one domain is attacked, others are also affected as collateral damage), etc.

Our approach to identifying these vulnerabilities consists of two steps:

1. *Identifying single points of failure* – Using data collected by the OpenINTEL project, we can identify single points of failure by mapping authoritative name servers back to operators. The simple but incomplete approach is to use the DNS hostnames of these nameservers to perform this mapping, but we propose to more exhaustively search this space by considering other topological information, e.g. the autonomous systems or IP prefixes that host name servers. The main challenge is to compose a set of views of OpenINTEL data that combines these different ways to identify operators, but still yields a consistent and coherent view of single points of failure.

2. *Identifying misconfigurations and suboptimal configurations* – We will first perform a systematic analysis of both good practices in terms of configuring DNS for domains, and of common configuration errors. Based on this analysis, we will use longitudinal data collected by the OpenINTEL project to quantify the occurrence of misconfigurations, and to analyse whether we can observe trends in the frequency at which these occur (investigating the question: does the resilience of the DNS improve, degrade or remain stable over time? Do trends differ by network types or size?). The main challenge here will be defining suitable signatures of such misconfigurations to look for in the sizable datasets from OpenINTEL.

We will go one step further to consider the potential impact of proposed approaches to minimize or mitigate DNS vulnerabilities. In particular, there are currently several IETF drafts at various stages of maturity that aim to improve the resilience of the DNS under attack. For example, some have proposed running local copies of zones [9] or serving stale data from DNS caches under certain conditions [10]. Once we have identified common vulnerabilities, we can survey the state-of-the-art in terms of standardised or proposed approaches to improve DNS resilience, and analyze to what extent these approaches actually address the vulnerabilities we have discovered. This

[2] http://www.caida.org/projects/network_telescope/

analysis process can guide future standardisation of such solutions and help prioritise ongoing work in the IETF to improve DNS resilience.

Finally, in order to facilitate reproducibility of this part of our research, it is important that the research community has access to OpenINTEL data. However, much of the data collected by OpenINTEL is subject to contracts with TLD registries (such as, e.g., VeriSign for the `.com` TLD), that prohibit sharing of this data in bulk. DHS has invested in research infrastructure to support navigation of privacy and legal issues in data-sharing: the *Information Marketplace for Policy and Analysis of Cyber-risk & Trust (IMPACT)*[3]. CAIDA is a long-standing contributor of data to IMPACT, and has proposed the HI-CUBE platform to provide the IMPACT research community with analytics, query tools and visualization of sensitive data sources available through IMPACT. The Netherlands is a trusted partner of the IMPACT project. A final goal of this pillar is to assess how we can make OpenINTEL data available through IMPACT, potentially via CAIDA's proposed HI-CUBE platform.

C.3 Pillar (2) – Mapping the DNS DDoS ecosystem

A thorough understanding of the characteristics of (Distributed) Denial-of-Service attacks on the DNS plays an essential role in both attack prevention and effective protection. For this reason, in MADDVIPR we will devise a methodology that maps the DNS (D)DoS attack ecosystem. The mapping will involve a macroscopic analysis of trace data on past and present attacks. We will base our method on two data sources:

- The UCSD Network Telescope offers an excellent vantage point to capture trace data of DoS attacks in which attackers try to disguise the source of malicious network traffic by applying (uniformly) random IP spoofing [11]. CAIDA has analyzed the observable "backscatter" traffic reaching the Telscope as a result of such DoS attacks for many years, allowing us to assess historical trends.

- We will also use other previously proven data sources to account for attack types that do not involve uniformly random spoofing. One example is data from the AmpPot project [8], which leverages honeypots to capture traces of DoS attacks that involve the abuse of reflectors. The operators of AmpPot have committed to sharing data in a letter of support.

The data sources listed above provide trace data of targeted hosts and the characteristics of malicious network traffic sent toward these hosts. Where opportunities arise, we will augment our analysis with other sources of trace data such as evidence from botnet command and control servers. Our methodology will build on previous work that uses the same trace data to analyze the ecosystem of DoS attacks on Web infrastructure [6]. We will integrate the resulting methodology into a new software platform – *DNSAttackStream* – that automatically analyzes the DNS (D)DoS attack ecosystem. *DNSAttackStream* will generate (near) real-time intelligence on ongoing attacks as they show up in the Network Telescope and other data sources, and will identify attack sources, targets, and characteristics.

[3] `https://impactcybertrust.org/`

C.4 Synthesizing a unified view

Pillars (1) and (2) provide two complementary views of the DNS DDoS problem. From the one side (Pillar (1)) we are now able to identify SPoF and vulnerabilities that can be exploited in case of DDoS attack against the DNS; from the other (Pillar (2)), we have gained an overview of what is attacked in practice, based on continuous network measurements. The next step is therefore to **synthesize a unified view** of the DNS DDoS ecosystem, SPoFs and vulnerabilities in order to create **actionable intelligence for DNS protection and attack prevention**.

A motivational example: the 2016 attack on Dyn On October 21, 2016, DNS hosting provider Dyn was attacked. The attack targeted Dyn's authoritative name server infrastructure, and hit Dyn's infrastructure on the East Coast of the U.S. in particular. This resulted in a number of high-profile Internet companies, e.g., Twitter and PayPal, becoming unreachable for several hours.

While the direct effects of the attack have been covered extensively in both the mainstream and tech media, there is a secondary story that got far less coverage. The attack was such a success because the Internet companies that were most affected **exclusively used Dyn's DNS platform**. This meant that when that platform became unreachable, effectively these companies also became unreachable. Dyn had in practice become a **single point of failure**.

Using OpenINTEL data, we are able to identify Dyn as a potential point of failure by analyzing the exclusive use of the Dyn's DNS platform by domains in a certain zone (in this example, we focused on the `.com` zone, as this is the largest TLD). Analysis of DNS configurations for `.com` domains before and after the attack sheds light on the impact of such a SPoF on the DNS under attack.

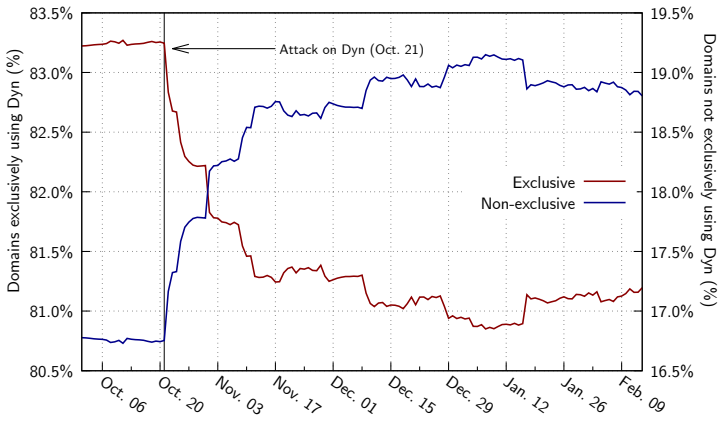

Figure 1: Exclusive and non-exclusive use of Dyn's DNS service at the time of the Oct. 2016 attack (source [15])

Figure 1 shows the fraction of domains that used Dyn's services exclusively versus those that use both Dyn's service and another DNS operator, for the four months following the October 2016 attack. There is a steep drop in exclusive use of Dyn after the attack: over 4,000 domains changed from exclusive use to non-exclusive use. Such behavior was not observed before the attack, which means that the aftermath of the attack marks a notable trend change. Analysis of the domains who changed their configuration after the attack shows a wide range of domain owners changing their behavior, including large Internet service providers, media ventures and gaming platforms.

In the case of Dyn, we are only able to perform a post-facto analysis of the effect of a SPoF. The MADDVIPR project will seek to operationalize this capability, in order to facilitate proactive identification of SPoF and exploitable DDoS vulnerabilities in the DNS, which, combined with

knowledge about the DDoS ecosystem, will yield actionable information to enables networks to minimize the risk and mitigate the impact of future DoS attacks.

This step of the project will concentrate on the following activities:

- **Identification of the impact of a possible attack** – By combining knowledge of attack targets, attack trends and SPoF, we will ascertain the impact an attack could potentially have on the DNS infrastructure and collateral damage on other services.
- **A view of future attacks** – By combining the DNS DDoS ecosystem with knowledge of the vulnerabilities identified in Pillar (1), we will identify weak points in the practical use of the DNS and its configuration that might lead to future attacks.
- **Prioritization of risks** – Finally, we will study how we can use such combined knowledge to create a clear prioritization and ranking of SPoF and vulnerabilities that are a major risk for the DNS. Such knowledge will be of direct use to operators and security experts in attack mitigation and prevention.

C.5 Ethical considerations

This proposal tackles, with a strongly measurement-based methodology, the identification of weak points in the DNS that have the potential to be exploited in DDoS attacks. As such, the topic at hand is highly sensitive, and therefore needs to be addressed in a ethically-conscious manner.

Both applicants have a proven track record in including ethical considerations in their measurement studies. The Faculty of Electrical Engineering, Mathematics and Computer Science at the University of Twente, to which the UT applicants are affiliated, has an established and recognized Ethics Committee whose role is to assess the ethical issues related to research projects conducted within the faculty. In addition to this, the applicants have personally been involved in ethics-related studies in the field of Internet Security. For example, prof. dr. ir. Aiko Pras has co-organized the *Dagstuhl seminar "Ethics in Data Sharing" (Dagstuhl Seminar 14052)*[4], which led to a publication describing best practices in sharing data [4] and has inspired the SURFnet[5] *Data Sharing Policy* containing legal and ethical guidelines relating to data sharing for research purposes [14]. On the U.S. side, CAIDA was a driving force behind the Menlo Report [1], which offered a model for ethical impact assessment of networking and cybersecurity research. Ethics of network measurement research is also a key topic addressed in CAIDA's annual Active Internet Measurement Systems (AIMS) workshop series.

Specific to this proposal, the OpenINTEL project partners have take particular care to create a methodologically responsible design of the OpenINTEL active DNS measurement platform. A review of the ethical considerations regarding the OpenINTEL measurement is available at [15] (Appendix C.3). OpenINTEL is therefore compliant with guidelines such as having a clearly identifiable intent and contact point, providing an opt-out mechanism and being designed such as to minimize the overall impact on the DNS infrastructure. Similarly, CAIDA carefully curates the

[4] `http://www.dagstuhl.de/en/program/calendar/semhp/?semnr=14052`

[5] SURFnet is the Dutch National Research and Education Network, to which dr. ir. Roland van Rijswijk-Deij is currently affiliated.

Network Telescope data and ensures scientific integrity of the data by vetting the personnel that has access to it.

Finally, both applicants strongly believe in open access to research results and are committed to releasing high-quality curated datasets. At the same time, both applicants are aware that the research results for this proposal may prove highly sensitive. Actionable intelligence on vulnerabilities in the DNS can also be misused by attackers. Therefore, our data disclosure approach will include a thorough ethical and privacy impact assessments, in consultation with other principal investigators of DHS' IMPACT project, to weigh the benefits of publicly releasing data against the potential harm of empowering attackers. Where appropriate to safeguard the security of networks ("first do no harm"), we will restrict access to research and analysis results. If both applicants agree, we can still share such results with fellow researchers and operators. We plan to use IMPACT as the framework to share data from this project wherever possible.

C.6 References

[1] M. Bailey, D. Dittrich, E. Kenneally, and D. Maughan. The Menlo Report. *IEEE Security & Privacy*, 10(2), 2012.

[2] K. Benson, A. Dainotti, K. Claffy, A. Snoeren, and M. Kallitsis. Leveraging Internet Background Radiation for Opportunistic Network Analysis. In *Proc. of ACM IMC 2015*.

[3] P. Bright. Spamhaus DDoS grows to Internet-threatening size. http://arstechnica.com/security/2013/03/spamhaus-ddos-grows-to-internet-threatening-size/.

[4] S. Dietrich, J. van der Ham, A. Pras, R. van Rijswijk-Deij, D. Shou, A. Sperotto, A. van Wynsberghe, and L. Zuck. Ethics in Data Sharing: developing a model for best practice. In *2nd Cyber-security Research Ethics Dialogs & Strategy (CREDS II)*, 2014.

[5] S. Hilton. Dyn Analysis Summary Of Friday October 21 Attack. http://dyn.com/blog/dyn-analysis-summary-of-friday-october-21-attack/, Oct. 2016.

[6] M. Jonker, A. King, J. Krupp, C. Rossow, A. Dainotti, and A. Sperotto. A Third of the Internet is Under Attack: a Macroscopic Characterization of the DoS Ecosystem. In *Proc. of ACM IMC 2017*.

[7] M. Jonker, A. Sperotto, R. van Rijswijk-Deij, R. Sadre, and A. Pras. Measuring the Adoption of DDoS Protection Services. In *Proc. of ACM IMC 2016*.

[8] L. Krämer, J. Krupp, D. Makita, T. Nishizoe, T. Koide, K. Yoshioka, and C. Rossow. AmpPot: Monitoring and Defending Against Amplification DDoS Attacks. In *Proc. of RAID 2015*.

[9] W. Kumari and P. Hoffman. Decreasing Access Time to Root Servers by Running One on Loopback. RFC 7706 (Informational), https://tools.ietf.org/html/rfc7706, November 2015.

[10] D. Lawrence and W. Kumari. Serving Stale Data to Improve DNS Resiliency. https://tools.ietf.org/html/draft-tale-dnsop-serve-stale-01, June 2017.

[11] D. Moore, C. Shannon, D.J. Brown, G.M. Voelker, and S. Savage. Inferring Internet Denial-of-service Activity. *ACM Transactions on Computer Systems*, 24(2):115–139, 2006.

[12] OVH. The DDoS that didn't break the camel's VAC. https://www.ovh.com/us/news/articles/a2367.the-ddos-that-didnt-break-the-camels-vac.

[13] M. Prince. Technical Details Behind a 400Gbps NTP Amplification DDoS Attack. https://blog.cloudflare.com/technical-details-behind-a-400gbps-ntp-amplification-ddos-attack/.

[14] R. van Rijswijk-Deij. Ethics in Data Sharing: a best practice for NRENs. In *Proc. of TNC 2015*.

[15] R. van Rijswijk-Deij. *Improving DNS Security: A Measurement-Based Approach*. PhD thesis, 2017. CTIT Ph.D. thesis series no. 17-430.

[16] R. van Rijswijk-Deij, M. Jonker, A. Sperotto, and A. Pras. A High-Performance, Scalable Infrastructure for Large-Scale Active DNS Measurements. *IEEE Journal on Selected Areas in Communications*, 34(7), 2016.

D Testing and Evaluation

Both the University of Twente and CAIDA have access to the type of high-performance computing environments required to analyse and process the large scale datasets (OpenINTEL, Telescope, potentially others) used in the proposal. In addition to this, we intend to use the extensive social networks that both collaborators have in the DNS operator community to solicit feedback on the results that emerge from the MADDVIPR project. A prime example of a forum in which we expect to get substantial feedback is the Domain Name System Operations Analysis and Research Center (DNS-OARC). CAIDA is one of the founding members of DNS-OARC and the University of Twente is an academic member. DNS-OARC organizes two annual meetings where members can present work to and solicit feedback from fellow members. This membership includes all of the big names in DNS operations, ranging from registry operators for top-level domains to large Internet brands. We intend to leverage this particular channel to get early feedback from organizations that will benefit directly from MADDVIPR's results. Finally, the major global bank ING, the Dutch National Cyber Security Center NCSC, and the Center for IT-Security, Privacy and Accountability (CISPA, Saarland University) expressed a particular interest in MADDVIPR. These organizations will be part of an external review board to be set up by the University of Twente.