MADDVIPR

MAPPING DNS DDOS VULNERABILITIES TO IMPROVE PROTECTION AND PREVENTION

SUMMARY

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 supports 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 project 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.

HIGH-LEVEL ARCHITECTURAL DIAGRAM OF THE PROPOSED SYSTEM

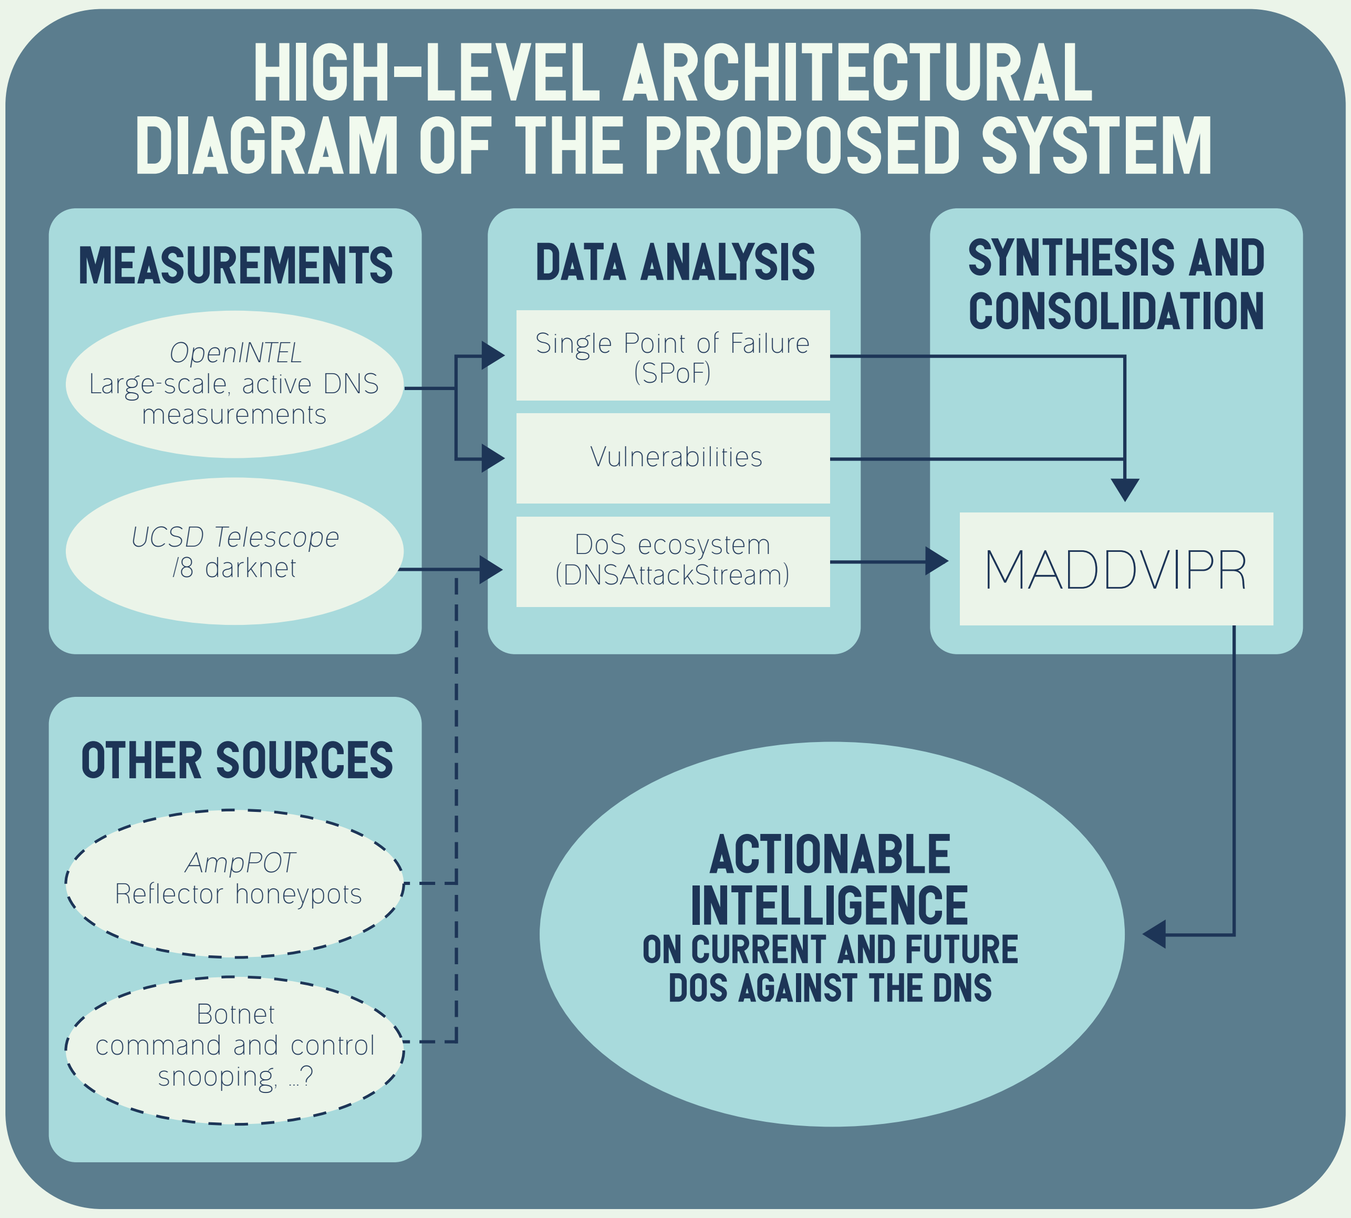

IDENTIFYING DNS SINGLE POINTS OF FAILURE AND VULNERABILITIES

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, mis- matches 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.

1. IDENTIFYING SINGLE POINTS OF FAILURE

Using data collected by the OpenINTELproject,wecan 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 sys- tematic 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.

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

1. *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.

2. *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.

3. *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.

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. 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, which leverages honeypots to capture traces of DoS attacks that involve the abuse of reflectors.

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.

EXCLUSIVE AND NON-EXCLUSIVE USE

of Dyn's DNS service at the time of the Oct. 2016 attack

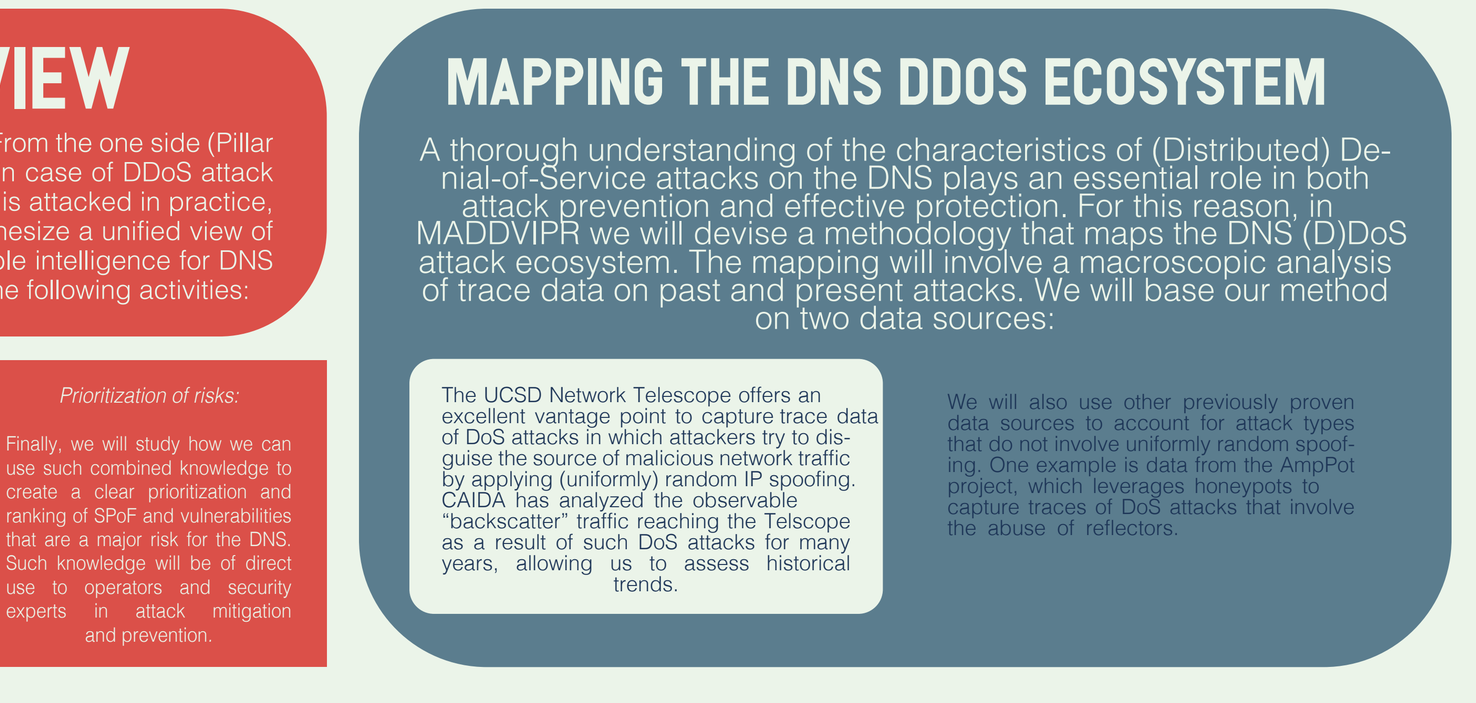

caida

U.S. DEPARTMENT OF HOMELAND SECURITY

UC San Diego

UNIVERSITY OF TWENTE.

Support for this work is provided by the U.S. Department of Homeland Security's Science and Technology Directorate under cooperative agreement FA8750-19-2-0004.