Threat Identification Using Active DNS Measurements

Olivier van der Toorn (✉) and Anna Sperotto

University of Twente, Enschede, The Netherlands
{o.i.vandertoorn,a.sperotto}@utwente.nl

Abstract. The DNS is a core service for the Internet. Most uses of the DNS are benign, but some are malicious. Attackers often use a DNS domain to enable an attack (e.g. DDoS attacks). Detection of these attacks often happens *passively*, which leads to a reactive detection of attacks. However, registering and configuring a domain takes time. We want to *pro-actively* identify malicious domains during this time. Identifying malicious domains before they are used allows to pre-emptively stop an attack. We aim to accomplish this goal by analysing active DNS measurements. Through the analysis of active DNS measurements there is a window of opportunity between the time of registration and the time of an attack to identify a threat before it becomes an attack. Active DNS measurements allows us to analyse the configuration of a domain. Using the configuration of a domain we can predict if it will be used for malicious intent. Machine Learning (ML) is often used to process large datasets, because it is efficient and dynamic. This is the reason we want to use ML for the detection of malicious domains. Because our results are predictive in nature, methodology for validation of our results need to be developed. At the time of the detection ground truth is not (yet) available.

1 Introduction

A core part of the Internet infrastructure is the Domain Name System (DNS). This system performs the translation from a domain name to an IP address. Most uses of the DNS are benign, but there are malicious uses of the DNS. This malicious activity may have a devastating impact on the Internet at large. A chief example is the Distributed Denial of Service (DDoS) attacks against the service and DNS provider Dyn [1], which made a large portion of Web sites inaccessible. However, examples like the 2016 DDoS attack against an unnamed European media organization which peaked at 363Gbps [2] demonstrate what kind of traffic abused DNS servers can generate.

Many of these attacks require a domain name. For example, as a (controlled) means of amplification (in a DDoS attack), or as a dynamic way of hosting a Command and Control (C&C) server for botnets. Research in the field of detection and mitigation has already indicated that the DNS plays a central role as source of data for security. The work in [3, 4], for example, analyzes user-generated DNS traffic to identify botnet command-and-control traffic. In practice, there exists a plethora of mechanisms to protect us from cyberthreats (e.g., firewalls,

blacklists, ACL, IDS, DDoS protection services etc.). What all these solutions have in common is that they are *reactive*. Threat intelligence based on DNS is frequently based on *passive* DNS measurements. These systems require client DNS activity before they are able to detect attacks.

We want to change the approach from *reactive* to *pro-active* threat identification. There is a window of opportunity between the registration of a domain and when it is first used in an attack. We want to identify malicious domains during this time. We aim to do so by combining active DNS measurements with Machine Learning.

2 Goal, Research Questions and Approach

The goal of this research is to develop methodology for *pro-active* threat identification of malicious domains through *active* DNS measurements. We want to detect malicious domains between registration time of a domain, and the first attack. Rather than when the attack is on-going. Because we want to predict if a domain is malicious or not, we also need to develop a method of validation.

We hypothesize that active DNS measurements allow us to do pro-active detection of malicious domains. Since attackers frequently register a DNS domain and configure it in a specific manner before an attack begins. Active DNS measurements allows us to analyse the configuration of domains, and make a prediction about the nature of the domain. To achieve our goal we have defined the following research question.

> ***RQ*$_M$: How can we use active DNS measurements to pro-actively identify malicious domains, and what are the benefits of such an approach?**

Our approach in answering RQ_M will be part measurement-based and part experiment based. Figure 1 shows the envisioned path from start to finish. Each node represents a sub-research question, which is discussed in the following sections.

2.1 Survey the landscape

There exists a plethora of cyber-attacks, for example botnets, DDoS attacks, and worms. While we would like to detect all of them, if an attack makes no use of the DNS, our approach will not be feasible in the detection this attack. We need to research which attacks make use of the DNS, and how. To this end we have defined the following research question.

***RQ*$_{\mathbf{M.1}}$**: Which cyber-attacks make use of DNS and how do they use it?

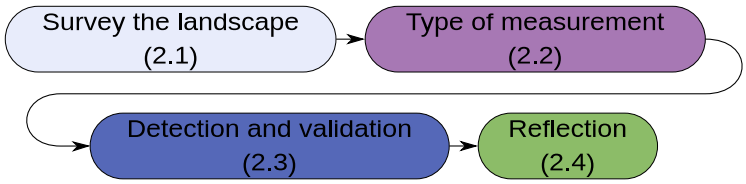

Fig. 1. Overview of the TIDE project

To answer this research question we want to perform an extensive literature survey. Additionally, we wish to interview DNS operators to obtain a complete picture of the landscape. From this research question we are able to define a set of use-cases, attacks, which will help confine the further research.

2.2 Passive Versus Active DNS Measurements

Figure 2 shows a typical DNS setup. On the left are clients querying a resolver. This resolver, in turn, queries the authoritative name servers to answer the queries of the clients. Two types of DNS measurements are visible in this Figure. Passive DNS measurements typically measure the traffic between resolver and authoritative name servers. Whereas active DNS measurements essentially emulate the clients by performing queries themselves.

State of the art (malicious) domain identification is typically based on passive DNS measurements [5, 6]. While effective they lead to a *reactive* detection.

To investigate the benefits, and drawbacks, of an active approach we have defined to following research question.

> ***RQ*$_{M.2}$: What are the strengths and weaknesses of both types of DNS measurements with respect to the attacks?**

To approach $RQ_{M.2}$, we will study both facets, active and passive, in detail. As a starting point for this study we aim to use Entrada [7], a passive DNS measurement system, and OpenINTEL [8], an active DNS measurement system, to evaluate the differences between the two types of measurements.

We want to study how well both approaches fare in the detection of attacks. The attacks we will evaluate stem from research question $RQ_{M.1}$.

2.3 Detection and validation

We want to be able to do detections of the entire DNS namespace. Therefore the detection process needs to be efficient in the processing of large amounts of data (e.g. OpenINTEL collects 2.2 billion data points from 207 million domains every day). However, since attacks are dynamic in nature [9] the detection process should also be flexible. Machine Learning (ML) is often used in (large-scale) spam detection [10, 11] which is dynamic in nature. For these reasons we aim to use ML for the detections. There are many different ML algorithms [12] we need to investigate which algorithm is suited to our detection problem.

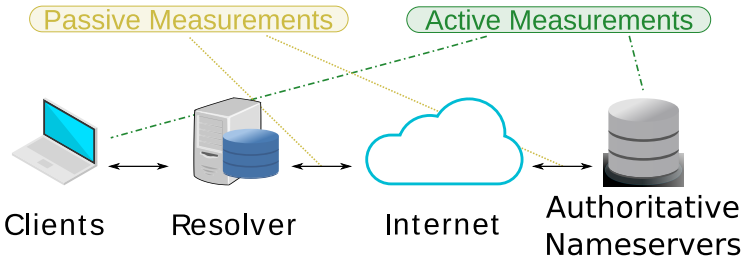

Fig. 2. Overview of the DNS system

Additionally, we face the problem of validation. Because we want to do proactive detection there is, at the time of the detection, no ground truth available to validate our detections against.

To reach these goals we have defined the following sub-question.

> ***$RQ_{M.3}$*: How can we perform efficient, large-scale, detections using Machine Learning and how do we validate these detections?**

To approach this question we want to evaluate different classifier algorithms. Subsequently, to validate the detections we need to compare our results with well-known blacklists over a long period of time. That way we are able to assess the quality of the results and evaluate the time advantage of our method.

2.4 Clustering

Finally, we want to explore if we are able to infer information about the parties behind the domains by clustering domains together. For example, in [13] the authors were able to cluster a CEO fraud campaign together, since the perpetrators used the same configuration to target multiple victims. To investigate what the benefits are of clustering approaches, such as the identification of the parties behind the act, we have defined the following sub-question.

> ***$RQ_{M.4}$*: What additional information can be obtained by clustering similar domain-configurations together?**

The approach we will take is clustering the domains we have detected together. And investigate the differences, and similarities, between the clustered domains.

3 Ethical considerations

Ethical considerations play a role throughout this project. For both types of measurements we have to take ethical best practices [14] into account. Especially with passive measurements, since these are in closer contact with user behaviour. The same counts when dealing with requests to be removed from the blacklist, these may be motivated by personal gain.

4 Preliminary Steps

We have already made a couple of preliminary steps in this project. In [15] we have applied our method to identify snowshoe spam domains. With snowshoe spam the spammer spreads the sending of spam over many hosts. Thus each individual host leaves a shallow imprint, like a snowshoe. This makes snowshoe spam hard to detect. Many spammers make use of Sender Policy Framework (SPF) to make their email appear more legitimate and increase their chance of success. SPF requires spammers to have a domain with a record for each sending host. Thus, snowshoe spam domains have many records (e.g. more than 200 `MX` records). Using active DNS measurements we were able to analyse the configuration of domains. By applying Machine Learning on this dataset we were able to detect snowshoe spam domains more than a hundred days before they appeared on regular blacklists. This research formed the basis for this Ph.D. project.

Acknowledgements

This research is funded by SIDNfonds[1]. SIDNfonds is an independent fund on the initiative of SIDN, the registrar for '.nl' domains.

References

1. Hilton, S.: Dyn Analysis Summary Of Friday October 21 Attack. https://dyn.com/blog/dyn-analysis-summary-of-friday-october-21-attack/ (October 2016)
2. Constantin, L.: Attackers use DNSSEC amplification to launch multi-vector DDoS attacks. https://www.computerworld.com/article/3097364/security/attackers-use-dnssec-amplification-to-launch-multi-vector-ddos-attacks.html (July 2016)
3. Choi, H., Lee, H., Kim, H.: BotGAD: Detecting Botnets by Capturing Group Activities in Network Traffic. In: COMSWARE. (2009)
4. Choi, H., Lee, H.: Identifying botnets by capturing group activities in DNS traffic. Computer Networks (2012)
5. Antonakakis, M., Perdisci, R., Lee, W., Vasiloglou, II, N., Dagon, D.: Detecting Malware Domains at the Upper DNS Hierarchy. In: Proc. of USENIX Security '11. (2011)
6. Bilge, Leyla and Kirda, Engin and Kruegel, Christopher and Balduzzi, Marco: EXPOSURE: Finding Malicious Domains Using Passive DNS Analysis. In: NDSS 2011. (2011)
7. Maarten Wullink, Giovane C. M. Moura, Muller, M, and Cristian Hesselman: ENTRADA: a High Performance Network Traffic Data Streaming Warehouse. In: NOMS 2016. (April 2016)
8. van Rijswijk-Deij, R., Jonker, M., Sperotto, A., Pras, A.: A High-Performance, Scalable Infrastructure for Large-Scale Active DNS Measurements. IEEE JSAC **34**(6) (2016)
9. Miller, R.: The Changing Face of Cyber-Attacks. Technical report, CA Technologies (2013) http://www3.ca.com/~/media/Files/whitepapers/the-changing-face-of-cyber-attacks.pdf.
10. Gudkova, D., Vergelis, M., Demidova, N., Shcherbakova, T.: Spam and phishing in Q1 2017. https://securelist.com/spam-and-phishing-in-q1-2017/78221/ (05 2017)
11. Bhowmick, A., Hazarika, S.M.: Machine Learning for E-mail Spam Filtering: Review, Techniques and Trends. CoRR (2016)
12. Brownlee, J.: A Tour of Machine Learning Algorithms. https://machinelearningmastery.com/a-tour-of-machine-learning-algorithms/ (November 2013)
13. Sperotto, A., van der Toorn, O., van Rijswijk-Deij, R.: TIDE: Threat Identification Using Active DNS Measurements. In: Proceedings of the SIGCOMM Posters and Demos. SIGCOMM Posters and Demos '17 (2017)
14. Dietrich, S., v. d. Ham, J., Pras, A., v. R. Deij, R., Shou, D., Sperotto, A., v. Wynsberghe, A., Zuck, L.D.: Ethics in data sharing: Developing a model for best practice. In: 2014 IEEE Security and Privacy Workshops. (2014)
15. van der Toorn, O., van Rijswijk-Deij, R., Sperotto, A.: Melting the Snow: Using Active DNS Measurements to Detect Snowshoe Spam Domains. In: NOMS 2018. (apr 2018)

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