DNS in Botnets and Advanced Persistent Threats

Peter McLaren, Gordon Russell, Bill Buchanan
School of Computing
Edinburgh Napier University
Edinburgh, UK
p.mclaren@napier.ac.uk, g.russell@napier.ac.uk, b.buchanan@napier.ac.uk

***Abstract*—Botnet and Advanced Persistent Threat malware remains a major challenge for cyber security. Communications between hosts infected by the malware and controllers for directing attacks, extruding data or updating the malware depend on reliable communications channels. As controllers need to change IP address in order to evade detection, DNS is commonly used to resolve the IP address of the controller domains. This paper is a preparatory study for a deeper analysis of features which can reliably identify the establishment of botnet and advanced persistent threat command and control channels. This is achieved by a review of major research, analysis of feature classes and identification of publicly available sources of benign and malware network traffic.**

Keywords—Domain Name System; botnet; advanced persistent threat; Domain Generation Algorithms; fast-flux

I. INTRODUCTION

Infected hosts in botnets and Advanced Persistent Threat (APTs) are most effective when able to receive commands to perform malicious acts. Bot commands are delivered to the hosts from controller(s), managed by a botmaster. A key element for this delivery is a reliable communication and control channel between bot and controller(s). As APTs are also a network of infected hosts with external controllers, APTs may be detected through discovery of the command and control channel. Although the commands can be requested by the host ('pull') or distributed by the controller(s) ('push') the creation of the channel is initiated by the infected host.

Where the controller IP address is unknown, the DNS protocol is used to resolve its domain name. P2P bots which do not have a command and control channel often DNS to resolve an address from which their binaries can be updated. The use of DNS by malware for establishing a command and control channel can be categorised as one of the following:

1) Simple DNS request/response.

Early bots used DNS to resolve the IP address of the controller. Once investigators had discovered the IP address further communication could be blocked by blacklisting.

2) Fast-flux

Fast-flux (FFSN) occurs where many IP addresses are associated with a single fully qualified domain name (FQDN). For botnets the IP addresses are changed very frequently by changing the FQDN DNS A records. The weakness of fast-flux is the use of a single domain name. This enables sink-holing - where traffic for the domain is routed to managed IP addresses – to be a successful defence mechanism. Examples of bots that have used fast-flux are Waledac[1] and Storm[2].

3) Domain generation algorithm (DGA)

DGA is a mechanism where random domain names are algorithmically generated and the bot then uses DNS to resolve a subset. To discover the domain names, the bot software needs to be reverse engineered to find the algorithm. A judicious use of seeds for the algorithm makes defending against it more difficult. Bots that have used DGA include Kraken[3] and Conficker[4].

This paper is a preparatory study prior to undertaking a deeper analysis of bot and APT DNS traffic. Common bots using fast-flux or DGA are reviewed in Section II. In Section III significant prior research is identified. Features which might distinguish malware DNS traffic including those established by prior research are reviewed in Section IV. The datasets which could be used for analysis are presented in Section V. Finally, Section VI is a summary of the paper and proposed further research.

II. BOTS

For DGA botnets, attention has focused on two, Conficker (C and D variations) and Kraken, which came into existence over 8 years ago although is it believed that at least 25 new DGAs have been discovered since 2013[5]. For FFSN botnets, it is likely that the largest botnet was Storm. A brief discussion of these three bots follows.

A. *Conficker*

The first Conficker attacks were seen in 2008. Although the Conficker-C bot uses peer-to-peer protocols for command distribution, it uses a centralized model for updating its binaries. For this purpose, a bot uses a pseudo-random generator to generate a 50,000 possible domains per day, a substantial increase from the 250 for Conficker-B. However, it only chooses 500 of the 50,000 domains and it attempts to resolve each of these once during the day. Furthermore, these requests are also randomized between 10 and 50 seconds so as to reduce the possibility of detection. The domain generating algorithm uses the current date as a seed [6].

B. Kraken

Although Kraken was believed around 2009 to be the biggest botnet in the world controlling as many as 500,000 bots and had a brief resurgence in 2014, it is now considered to be defunct[3]. The earliest versions of the Kraken DGA used the same seed so that the same domains were generated and take-down was relatively easy. Later versions use a time-dependent seed using a commonly used pseudo-random generator known as a linear congruential generator to create domain names[3].

C. Storm

Storm was known to be around before 2007 and is believed to have controlled between 1 and 5 million bots and been responsible for a majority of spam at the time. It was an early implementer of measures to evade detection including fast-flux, peer-to-peer (P2P) protocol, encrypting communications, frequent update of the bot binaries and a denial of service mechanism to inhibit reverse engineering[7].

III. Related work

Researchers have investigated features of DNS which might be used to detect bot command and control channels although to date little analysis is available for advanced persistent threats. Selected research papers, to which reference will be made later, are shown in Table 1. They are also classified by whether the analysis applies to fast-flux, DGA or to all types of DNS command and control traffic. Although it is not an exhaustive list, papers are concentrated in the period from 2008 to 2012 suggesting decreased research in bot and APT detection. Particularly, FFSN research seems to have reduced, possibly reflecting less use of FFSN for DNS.

Table 1 Significant DNS Bot Research

Authors	All	FFSN	DGA	Cited
Antonakakis et al 2010 [8]	Y			193
Antonakakis et al 2011 [9]	Y			119
Antonakakis et al 2012 [10]			Y	117
Bilge et al 2012 [11]	Y			242
Frosch et al 2013 [12]	Y			2
Holz et al 2008 [13]		Y		272
Nazario et al 2008 [2]		Y		144
Perdisci et al 2012 [14]		Y		33
Schiavoni et al 2014 [15]			Y	14
Stalmans et al2011 [16]		Y		33
Yadav et al 2010 [17]			Y	143
Yadav et al 2011 [18]			Y	40

IV. Features

A. Overview

The detection of botnets using DNS has used active probing or passive monitoring. As suggested by their respective names 'active probing' communicates directly with the domains in the flux network while 'passive monitoring' analyses DNS network traffic together with other quasi-static sources of information. Although passive monitoring has become popular since its introduction by Weimer[19], there are disadvantages to either approach and we think both should be considered.

A number of features have been used to detect botnet command and control channels in DNS traffic. These features can be assigned to one of the following categories:

- Network features. Features which can be associated with a domain indirectly by an analysis of the IP addresses that are returned in DNS responses.
- Temporal features. Features such as the number of DNS requests/responses to specific domains over a pre-determined period.
- Lexical features. Benign services tend to choose readable domain names as these are likely to be more memorable and therefore user-friendly. By contrast botmasters, particularly those using domain generating algorithms, are not concerned with memorable names so lexical differences may exist.
- Zone features. Zone features are derived from SOA resource records and cover the minimum and maximum time to live (TTL), the difference between those values as well as the number of serial number changes for the SOA record, and the minimal values for expiry time, retry time, and refresh time.
- WHOIS features. Static data for domains can be obtained from registries tools including online WHOIS sites and the WHOIS tool available in various flavors of Unix. The Ubuntu WHOIS tool has been used in this paper.

As shown in Table 2, research papers have often covered features across multiple categories.

Table 2 DNS Bot Research by Category

Author/Class	Net-work	Tem-poral	Lexi-cal	Zone	Whois
Antonakakis et al 2010[8]	Y				Y
Antonakakis et al 2011 [9]	Y				
Antonakakis et al 2012 [10]	Y		Y		
Bilge et al 2012 [11]	Y	Y	Y		
Frosch et al 2013 [12]	Y		Y	Y	Y
Holz et al 2008 [13]	Y				
Nazario et al 2008 [2]	Y			Y	
Perdisci et al 2012 [14]	Y				
Schiavoni et al 2014 [15]			Y		
Stalmans et al 2011 [16]	Y		Y		Y
Yadav et al 2010 [17]			Y		
Yadav et al 2011[18]	Y		Y		

The contribution made by these researchers for the different categories of features is reviewed in the following sections.

B. Network Features

Network features can be seen as direct or indirect attributes of the IP addresses obtained from DNS records. The explanation of this distinction follows. Direct attributes can be inferred directly from the DNS response fields: Address (A), Name servers (NS), Mail Servers (MX). For example, when a response to a DNS A record request has one or more IP addresses, the address range is a direct attribute. Indirect attributes require additional information sources to be derived. For example, the number and locations of the autonomous system number (ASN) associated with the IP addresses are indirect attributes.

It is expected that benign domains will change IP addresses less frequently than malicious domains that use fast-flux as the latter change to evade detection. Furthermore, indirect

attributes for benign domains, such as ASNs, are expected to be concentrated across fewer values than malicious domains. However, content delivery networks (CDNs) which aim to provide the most effective path to content for hosted clients present a potential exception to the above rules. Nevertheless, researchers have been able to use network features to distinguish CDNs from malicious domains [12].

Holz et al [4] investigated the number of A records and NS servers returned in an A response as a higher number would be expected to uncover domains that were fast-flux. Nazario et al [2] assessed the number of IP addresses in A record responses, whether IP addresses are more than a /16 subnet apart, the number of associated ASNs, and for nameservers, the number of entries, whether they are a /16 subnet apart , and the number of associated ASNs to create an indicator of the possible presence of a FFSN bot. Perdisci et al [14] considered a number of features including six network features principally concerning differences in the number of new distinct IPs, IP prefixes, domains compared to previous periods. Bilge et al [20] introduced an additional network feature, the number of distinct domains that share the returned IP addresses, in combination with other factors to distinguish malware from benign domains. Antonakakis et al [8] used the number of IPs, ASNs associated historically with a domain to build a reputation system. Antonakakis et al [9] believing that their earlier work had weaknesses produced a new system based on requester characteristics as bot requests are expected to come from a variety of IP addresses but in a consistent manner unlike requests from benign domains.

We are not aware of any analysis of mail server network features or the distribution of the IP addresses in A, NS and MX records for the same domain so this may be an area for further investigation.

C. Temporal Features

When fast-flux or DGA is used for DNS, it has been proposed that malicious domains will often show a sudden increase followed by a sudden decrease in the number of requests. Furthermore, when the domains are used for scams, for example, it is expected that there will be a sudden increase in the number of requests [11].

A key challenge with temporal analysis is establishing the appropriate window size. Success in selecting a value might be short-lived as bot-writers can amend the software so the host beacons to the controller outside the window.

D. Lexical Features

A variety of lexical features have been analyzed. Frosch et al [12] considered whether differences in proportions could be found between benign and malicious domains for three features: the number of digits in the domain name compared to length of domain, the number of consonants compared to length of domain, and the number of consonants compared to amount of vowels (called 'vocals' in the study). Bilge et al [11] considered two features: the ratio of the numerical characters to the length of the domain name and the ratio of the length of the longest meaningful substring (i.e., in a dictionary) to the length of the domain name. Yadav et al [17] considered three features to distinguish benign and malicious domains: 1) differences in the distribution of alphanumeric characters in groups of domains 2) differences in the distribution of two consecutive character sets in groups of domains using two different measures and 3) edit distance defined as the number of transformations required to change one string to another. Subsequently, Yadav et al [18] used edit distance as a measurement of entropy to analyze benign and malicious domain names. Alternative entropy measures will be considered in future work.

E. Zone Features

Zone features are obtained from SOA records. The fields obtained in a standard response are Zone-serial, Zone-refresh, Zone-retry, Zone-expire, Zone-minimum and TTL. Nazario et al [2] analyzed SOA responses to demonstrate that short TTLs, and minimum retry values distinguish bot from benign DNS requests. Frosch et al [12] found that although there were overlaps, malware and benign domains had significantly different values for the other zone features. The definition of zone features used by Antonakakis et al [8] is different to ours and is not included in this section.

F. WHOIS Features

The results from issuing the WHOIS command depends on the registry where the domain is registered. Based on the currently reviewed set, the fields in the left hand column of Table 3 should be available. Aside from where the field values have been distinguished as malicious and therefore blacklisted, e.g. domain names known to host malware, other WHOIS features have been analyzed by researchers.

Lexical features are considered elsewhere and blacklisting is well-defined. For the remaining possible distinguishing features there have been a number of studies. Antonakakis et al [8] analyzed the following WHOIS features: distinct registrars and diversity of registration dates for top-level, second-level and third-level domains. Frosch et al [12] analyzed the age of domains i.e. the delta between the current date and the registration date, and found that the mean for benign domains in the sample was 2276 days whereas the mean for malicious domains was 416 days. Nazario et al [2] analyzed the time between registration of a domain name and the first appearance of it in collected fast-flux data logs by accessing WHOIS data for all the domains.

There are fields such as renewal and updated date which do not appear to have been considered in research to date.

Table 3 Standard WHOIS fields

Field	Possible features	Benign Field Examples
Domain	Blacklisted, Lexical Features	napier.ac.uk
Registered For:	Blacklisted, Frequency	Napier University
Domain Owner	Blacklisted, Frequency	Edinburgh Napier University
Registered By:	Blacklisted, Frequency	
Servers:	Blacklisted name, address range, location	ns3.napier.ac.uk 146.176.7.1
Registrant Contact:	Blacklisted	C & IT Help Desk
Registrant	Blacklisted	Edinburgh Napier etc.

Address:		
Renewal date:	Delta between field & created or updated	Thursday 21st Jul 2016
Entry updated:	Delta between field & created	Thursday 16th April 2015
Entry created:	Delta between field & current date	Friday 7th November 2003

V. DATASETS

Recently, network analysis research has been improved by the availability of a number of publicly available packet capture traces and other network datasets. These datasets, listed in Table 4, have been classified by the suppliers as benign (G), malware (M) or blended (B). For reasons of privacy most benign datasets have anonymized data which may impact analysis. Other risks with the use of such datasets are whether the classification is correct (e.g. does data classified as benign actually contain malware) and whether the traffic is sufficiently recent to be of relevance.

Table 4 Benign & Malware Datasets

Dataset	Class	Description
Contagio [21]	M	Small network traces for specific malware instances e.g. Sality, AlienSpy
CAIDA [22]	B	Historical internet traces. Data is anonymised.
CTU-13 [23]	M, B	13 datasets including blended benign and malware as well as the malware used for blending. Benign data is anonymised
IMPACT [24]	G	5 datasets: 1 history of IP addresses used in DNS traffic, 4 CDN traffic captures
ISCX [25]	G, M, B	2 benign datasets and 5 blended datasets labelled ISCX 2012, 1 training and 1 testing dataset comprising a subset of the ISOT dataset, a subset of the ISCX 2012 dataset and all of the CTU malware datasets
ISOT[26]	B	1 dataset including Zeus, Storm and Waledac traffic. Benign traffic has truncated DNS fields
UPC [27], [28]	G	3 datasets for Linux, Windows XP and Windows 7 hosts. Traffic includes HTTP, DNS and other protocols. Data not anonymised but simulated.

VI. SUMMARY

Research into fast-flux and domain generating algorithms has been progressing for almost a decade. However, new datasets, bots and additional possible features for distinguishing traffic provide an opportunity for further exploration.

The next phases of this project will include the following: (1) a revisit of previous results using publicly available datasets (2) analysis additional DNS-related fields e.g. using WHOIS fields and MX values and (3) application of machine learning algorithms to datasets discover new feature combinations.

REFERENCES

[1] A. Caglayan, M. Toothaker, D. Drapaeau, D. Burke, and G. Eaton, "Behavioral patterns of fast flux service networks," in *System Sciences (HICSS), 2010 43rd Hawaii International Conference on*, 2010, pp. 1–9.

[2] J. Nazario and T. Holz, "As the net churns: Fast-flux botnet observations," in *Malicious and Unwanted Software, 2008. MALWARE 2008. 3rd International Conference on*, 2008, pp. 24–31.

[3] J. Bader, "Kraken's two domain generation algorithms," 2015. .

[4] K. Burton, "The Conficker Worm," 2012.

[5] Plohmann, Daniel, "DGArchive: A deep dive into domain generating malware," Dec-2015.

[6] S. Shin, G. Gu, N. Reddy, and C. P. Lee, "A large-scale empirical study of conficker," *IEEE Trans. Inf. Forensics Secur.*, vol. 7, no. 2, pp. 676–690, Apr. 2012.

[7] Porras, Phillip, Saıdi, Hassan, and Yegneswaran, Vinod, "SRITechnical-Report-10-01-Storm-Analysis.pdf," SRI, Oct. 2007.

[8] M. Antonakakis, R. Perdisci, D. Dagon, W. Lee, and N. Feamster, "Building a dynamic reputation system for DNS," in *USENIX security symposium*, 2010, pp. 273–290.

[9] M. Antonakakis, R. Perdisci, W. Lee, N. Vasiloglou II, and D. Dagon, "Detecting malware domains at the upper DNS hierarchy," in *USENIX security symposium*, 2011, p. 16.

[10] M. Antonakakis, R. Perdisci, Y. Nadji, N. Vasiloglou, S. Abu-Nimeh, W. Lee, and D. Dagon, "From throw-away traffic to bots: detecting the rise of DGA-based malware," in *Presented as part of the 21st USENIX Security Symposium (USENIX Security 12)*, 2012, pp. 491–506.

[11] L. Bilge, D. Balzarotti, W. Robertson, E. Kirda, and C. Kruegel, "Disclosure: detecting botnet command and control servers through large-scale netflow analysis," in *Proceedings of the 28th Annual Computer Security Applications Conference*, 2012, pp. 129–138.

[12] T. Frosch, M. Kührer, and T. Holz, "Predentifier: Detecting botnet C&C domains from passive DNS data," *Adv. IT Early Warn. Fraunhofer Verl.*, 2013.

[13] T. Holz, C. Gorecki, K. Rieck, and F. C. Freiling, "Measuring and detecting fast-flux service networks," in *NDSS*, 2008.

[14] R. Perdisci, I. Corona, and G. Giacinto, "Early detection of malicious flux networks via large-scale passive DNS traffic analysis," *IEEE Trans. Dependable Secure Comput.*, 2012.

[15] S. Schiavoni, F. Maggi, L. Cavallaro, and S. Zanero, "Phoenix: DGA-based botnet tracking and intelligence," in *Detection of intrusions and malware, and vulnerability assessment*, Springer, 2014, pp. 192–211.

[16] E. Stalmans and B. Irwin, "A framework for DNS based detection and mitigation of malware infections on a network," in *Information Security South Africa (ISSA), 2011*, 2011, pp. 1–8.

[17] S. Yadav, A. K. K. Reddy, A. L. Reddy, and S. Ranjan, "Detecting algorithmically generated malicious domain names," in *Proceedings of the 10th ACM SIGCOMM conference on Internet measurement*, 2010, pp. 48–61.

[18] S. Yadav and A. N. Reddy, "Winning with DNS failures: Strategies for faster botnet detection," in *Security and privacy in communication networks*, Springer, 2011, pp. 446–459.

[19] F. Weimer, "Passive DNS replication.pdf," presented at the 17th Annual FIRST Conference on Computer Security, 2005.

[20] L. Bilge, E. Kirda, C. Kruegel, and M. Balduzzi, "EXPOSURE: Finding malicious domains using passive DNS analysis.," in *NDSS*, 2011.

[21] M. Parkour, "Contagio," 2015. [Online]. Available: http://contagiodump.blogspot.co.uk/.

[22] CAIDA, "Center for Applied Internet Data Analysis," 2016. [Online]. Available: http://www.caida.org/home/.

[23] S. García, M. Grill, J. Stiborek, and A. Zunino, "An empirical comparison of botnet detection methods," *Comput. Secur.*, vol. 45, pp. 100–123, Sep. 2014.

[24] Impact Cyber Trust, "IMPACT," 2016. [Online]. Available: https://www.impactcybertrust.org/home#welcome.

[25] A. Shiravi, H. Shiravi, M. Tavallaee, and A. A. Ghorbani, "Toward developing a systematic approach to generate benchmark datasets for intrusion detection," *Comput. Secur.*, vol. 31, no. 3, pp. 357–374, May 2012.

[26] S. Saad, I. Traore, A. Ghorbani, B. Sayed, D. Zhao, W. Lu, J. Felix, and P. Hakimian, "Detecting P2P botnets through network behavior analysis and machine learning," in *Privacy, Security and Trust (PST), 2011 Ninth Annual International Conference on*, 2011, pp. 174–180.

[27] T. Bujlow, V. Carela-Español, and P. Barlet-Ros, "Comparison of Deep Packet Inspection (DPI) tools for traffic classification," 2013.

[28] V. Carela-Español, T. Bujlow, and P. Barlet-Ros, "Is our ground-truth for traffic classification reliable?," in *Passive and Active Measurement*, 2014, pp. 98–108.