Almazarqi, H. A., Marnerides, A., Mursch, T., Woodyard, M. and Pezaros, D. (2022) Profiling IoT Botnet Activity in the Wild. In: 2021 IEEE Global Communications Conference (GLOBECOM), Madrid, Spain, 07-11 Dec 2021, ISBN 9781728181042

(doi: 10.1109/GLOBECOM46510.2021.9686012)

This is the Author Accepted Manuscript.

There may be differences between this version and the published version. You are advised to consult the publisher's version if you wish to cite from it.

http://eprints.gla.ac.uk/257918/

Deposited on: 15 November 2021

Enlighten – Research publications by members of the University of Glasgow
http://eprints.gla.ac.uk

Profiling IoT Botnet Activity in the Wild

Hatem A. Almazarqi*, Angelos K. Marnerides*, Troy Mursch†, Mathew Woodyard†, and Dimitrios Pezaros*
*School of Computing Science, University of Glasgow, Glasgow, Scotland, UK
h.almazarqi.1@research.gla.ac.uk, angelos.marnerides, dimitrios.pezaros@glasgow.ac.uk
† Bad Packets LLC, Chicago IL, USA
troy,mat@badpackets.net

***Abstract*—Undoubtedly, the Internet of Things (IoT) contributes significantly to daily mission-critical processes underpinning a number of socio-technical systems. Conversely, its rapid adoption has extensively broadened the cyber-threat landscape by virtue of low-cost IoT devices that are manufactured and deployed with minimal security. Evidently, vulnerable IoT devices are utilised by attackers to participate into Internet-wide botnets in order to instrument large-scale cyber-attacks and disrupt critical Internet services. Since the 2016 outbreak of the first IoT Mirai botnet there has been a continuous evolution of Mirai-like variants. Tracking these botnets is challenging due to their varying structural characteristics, and also due to the fact that malicious actors continuously adopt new evasion and propagation strategies. This work provides a new measurement study highlighting specific behavioural properties of Mirai-like botnets in terms of their propagation. We provide a comprehensive analysis conducted on real Cyber Threat Intelligence (CTI) feeds gathered for a period of 7 months from globally distributed attack honeypots and pinpoint the evolutionary port scanning patterns, targeted vulnerabilities and preferred services pursued by Mirai-like botnets. We identify the most frequently active Mirai-like malware binaries and we are the first to report the evolution of a new, P2P-based variant. In parallel, we provide evidence related to the lack of vendor-specific patching through highlighting unpatched vulnerabilities. Moreover, we pinpoint the inadequacy of widely used IP blacklisting databases to timely list malicious IP addresses. Thus, arguing in fair of integrating honeypot information from diverse Internet vantage points within the design of next generation botnet defence mechanisms.**

***Index Terms*—IoT, botnets, Internet measurements, Mirai, malware, attack honeypots, cyber threat intelligence.**

I. Introduction

Recently and post-2016, IoT botnets have been one of the most common means of instrumenting large-scale cyber attacks and organised cyber terrorism activities. Such activities are diverse and could be in the form of Distributed Denial of Service (DDoS) attacks as well as Advanced Persistent Threats (APTs) with a number of spamming, phishing and ransomware campaigns [1]. IoT botnets can be defined as a group of compromised IoT devices ('bots') which are infected with malware and controlled via a single entity ('a malicious actor') or organised groups of 'hacktivists'. Such devices include, but are not limited to, Internet-enabled DVRs, smart meters, programmable logic controllers, wearables and home routers. Evidently, the rush of deploying IoT-oriented services has led manufacturers to take minimal security considerations particularly for low-cost IoT devices. In addition, policy-makers are unable to catch up with the consumer-oriented IoT market and thus challenging for them to enforce adequate policies on manufacturers explicit to security [1], [2].

Fundamentally and similarly with conventional botnets, IoT botnet operation revolves around a single or a number of command and control (C&C) servers that are instrumented by a malicious actor or 'hacktivist' groups. Depending on the malware variant and also the botnet's scanning and propagation strategy, C&C servers interact with Loader and Report servers as well as with devices that are simply infected (i.e., bots) [3]. The communication channel amongst the aforementioned entities varies and it defines the architecture of a given botnet to act under a centralised or a distributed fashion. Commonly, centralised botnets are underpinned by protocols such as IRC and HTTP/HTTPS whereas P2P-based protocols form the basis for distributed botnets [1], [3], [4]. The evolution of botnet development by organised APT groups (e.g., APT41 group [1]) has demonstrated that modern IoT botnets are resilient to detection by ISP policies and intrusion detection systems (IDS) due to advanced evasion techniques such as protocol obfuscation, Fast-Flux and DNS-oriented Domain Generation Algorithms (DGA) [3].

Since the 2016 Mirai outbreak where more than $550K$ IP-enabled DVRs were instrumented for a series of DDoS attacks and disrupted services almost across half of the global Internet [5], new Mirai-variants have emerged [6]. Efforts to analyse the dynamic properties of new Mirai-like variants are challenging and require a systematic and macroscopic view of the Internet in order to [1]–[3], [6]: (i) adequately inform corresponding governmental or industrial bodies and organisations in charge of infrastructure defense and (ii) equip next generation automated cyber threat intelligence (CTI), IDS and anomaly detection mechanisms with enriched information regarding evolving scanning, establishment and propagation strategies of new botnet variants.

In this work, we provide a 7-month measurement study of Mirai-like variants. With the correlation of Internet-wide CTI feeds gathered from globally distributed honeypots and global IP blacklist databases we point intrinsic characteristics during their propagation phase. Thus, the main contributions of this paper are:

[1]FireEye report on APT41:https://content.fireeye.com/apt-41/rpt-apt41/

1) Novel views on the macroscopic nature of evolving Mirai-like botnet activity with respect to current scanning and infection strategies.
2) A first insight on the activities of Mozi; a new Mirai-like P2P-based botnet first seen in late 2020.
3) Assessment of IP blacklist efficiency on capturing the evolving behaviour of IoT botnets.

The remainder of this paper is structured as follows: Section II provides background information on basic botnet entities whereas Section III provides an overview of related work. Section IV describes the datasets and methodology used in this work. Section V is dedicated on presenting our findings. Finally, Section VI summarises and concludes this work.

II. Background

A. Botnet structure

As discussed over recent studies assessing the publicly released Mirai and BASHLITE source code (e.g., [1], [5], [7]), the typical Mirai-like botnet architecture is commonly centralised and not distributed (i.e., P2P-based communication), and encapsulates a number components as shown in Figure 1.

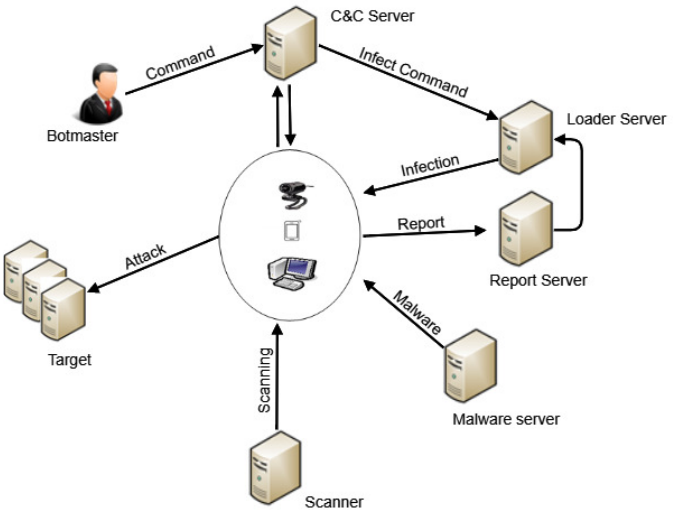

Fig. 1: IoT botnet components.

Mirai-like botnets are controlled by a botmaster(s) instrumenting the following entities:

1) Bot(s): IoT device(s) infected and compromised by Mirai-like malware acting on behalf of the Botmaster.
2) Command and Control Center (C&C): server in charge of botnet control issuing commands to bots for launching different types of attacks (e.g., spamming, DDoS).
3) Scanners: IoT devices or servers used to identify other vulnerable IoT devices scanning the IPv4 address space for open TCP/UDP ports.
4) Report Server(s): server keeping record of scan results, active bots and stolen credentials.
5) Loader(s): servers obtaining scan results and credentials from report server(s) in order to login to vulnerable IoT devices, and instruct them to download botnet malware.
6) Malware server(s): servers that can also act as Loader servers hosting malware code that will be downloaded by compromised IoT devices.

B. Botnet propagation

Port scanning is the main method instrumented by a botmaster for identifying susceptible hosts. Evidently, different botnets have their own carefully crafted scanning methods that may look identical to routine scans performed by network operators for aspects of service management [8]. The sole purpose of an adversary during the scanning process is to obtain a better view of devices that operate over vulnerable services attached on open TCP/UDP ports that are also responsive to scanning probes. Port scanning is stratified into two major categories; (i) vertical, and, (ii) horizontal. In vertical scans, multiple ports are scanned on the same target [9]. Vertical scans are useful for gathering information to attack a particular victim host or when a targeted attack is planned to be instrumented over particular web services. On the contrary, horizontal scans are considered when the same port is scanned over multiple targets [9]. As reported in [10], modern botnets may demonstrate hybrid scanning properties involving both vertical as well as horizontal scans.

III. Related Work

Over the last decade, several studies were conducted in order to profile and understand the behavioural properties of general IoT botnets (e.g., [1], [11]) whereas other studies focused particularly on understanding the properties of botnet scan traffic (e.g., [4], [8], [9]). The work in [12] was one of the major studies to demonstrate the importance of utilising improved honeypots for composing realistic samples from malicious IoT traffic. However, the proposed scheme was focused on a quite restricted set of honeypots whereas in this work we operate with Internet-wide feeds from globally distributed honeypots. Evidently, most of the aforementioned studies focused mainly on the outcomes resulted by IoT botnets and didn't assess the relationship between particular IoT vulnerabilities and the malware binaries exploiting them as we do in this work.

The seminal study in [5] was the first to profile the original 2016 Mirai botnet outbreak and contributed towards the Internet-wide, AS-level analysis as well as the impact of shared source code which led to the proliferation of Mirai variants. Moreover, the work in [3] and [6] assessed the impact of Mirai-like variants on industrial control systems (ICS) and how Mirai-like variants exploit DNS records respectively. Nonetheless, all aforementioned studies didn't capture the modern scanning characteristics of new variants and in parallel did not provide a recent overview of the Mirai-like structural properties (i.e., centralised or P2P) as conducted in this work. In addition, there were no clear indications of the types of instruction set architectures targeted by Mirai-like variants as also shown herein.

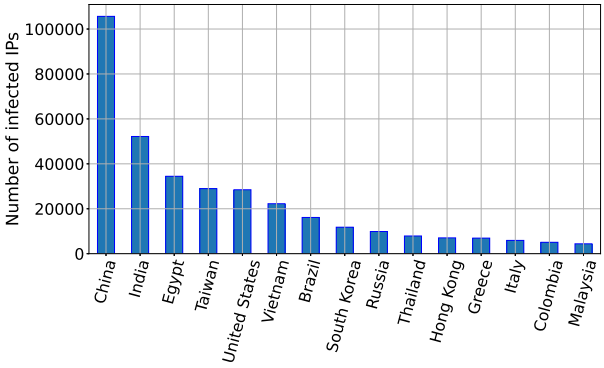

Fig. 2: Top fifteen countries ranked based on the number of Mirai-like infected IP addresses.

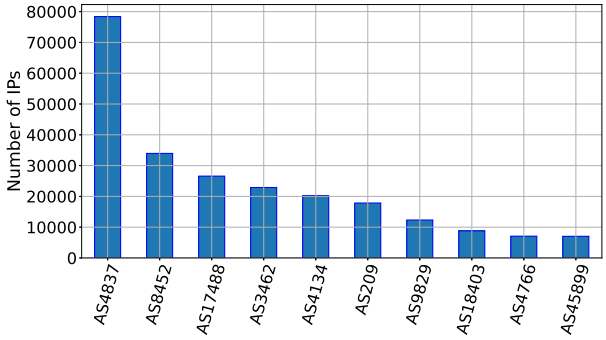

Fig. 3: Top ten ASes ranked based on the number of Mirai-like infected IP addresses.

Collection Period			
01/07/2020 - 22/02/2021			
IP addresses	**Origin Countries**	**ASes**	**Connection attempts**
422,003	203	7,265	1,599,647

TABLE II. Summary of Mirai-like CTI feeds gathered from the Bad Packets honeypots used in this work.

IV. Dataset Description & Methodology

A. Dataset Description

In order to identify structural patterns of Mirai-like botnet activity we correlate cyber threat intelligence (CTI) feeds, Internet geolocation data provided by MaxMind[2] and four IP address blacklist databases over a period of 7 months. The CTI feeds are collected by Bad Packet's 11 globally distributed honeypots located in: the United States (3 in Las Vegas, Nevada, 1 in Minden, Nevada, 3 in Los Angeles, California), Russia (2 in Moscow), and Brazil (2 in Sao Paulo). Every honeypot emulates a number of IoT-based services based on profiles of IoT devices, consumer-grade routers, and even enterprise VPN endpoints. Incoming traffic form malicious actors targeting our honeypots is captured and further indexed using Splunk. In order to monitor and capture the best bad packets, we employ sinkhole domains previously used by DDoS botnet threat actors. Mirai-like activity is determined via a fingerprinting method comparing the TCP SYN sequence values with the IP address value. Evidently, any Mirai-like activity initiated by a given infected host has the sequence value of the first TCP SYN packet to be equal with the senders IP address [3]. As summarised in Table II, we extracted 1,599,647 connection and control attempts as well as malware payload uploaders. These observation were stemmed from 422,003 distinct IP addresses located across 7,265 Autonomous Systems (ASes) spanning 203 countries between July 2020 and February 2021.

As already mentioned, we utilise MaxMind's GeoLite 2 database to explore the geographic distribution of infected IPs in our dataset. Fig. 2 shows the distribution of infected IPs among the top fifteen countries. Although Mirai scans originate from 203 countries, 60% of scan traffic originates only from five countries: (i) China: 25%, (ii) India: 12.4%, (iii) Egypt: 8.2%, (iv) Taiwan: 6.7% and, (v) USA: 6.8%.

Mirai-like botnet herders constantly scan the IP address space in order to locate vulnerable IoT devices and further expand a given botnet. As shown in Fig. 3, we have mapped the infected IPs in our datasets to their corresponding Autonomous Systems (ASes). It is revealed that 56% of the infected IPs in our collected dataset only reside in the top ten ASes across the top ten 10 countries from Fig. 2. Evidently, the remaining 44% of Mirai-related traffic is spread sporadically over random ASes spanning 193 countries.

B. Methodology

Since the traffic dynamics imposed by Mirai-like botnets hold a high level of randomness in both scanning and instrumentation, we exploit the properties of Shannon entropy as used in other studies [8]. Hence, we measure the amount of information obtained by observing CTI feed logs through the Shannon entropy formulation given by:

$$H(X) = -\sum_{i=1}^{n} p_i \; log_2 \; p_i$$

In practice, we compute the distribution of targeted ports denoted by p_i in order to identify their dispersity or concentration with respect with their information entropy $H(X)$.

The range of values obtained by sample entropy relies on N, i.e., the number of unique values observed in the sampled set of packets which in our case is the port number. The value of sample entropy could be in the range $(0,\ log_2\ N)$ with a 0 value indicating that the distribution is maximally concentrated having all observations be the same. When the distribution is maximally dispersed, i.e., $n_1 \ = n_2 \ = ... = n_N$, sample entropy takes on the value $log_2\ N$. In general, we conduct exploratory normalised entropy overviews of timeseries observations related to the frequency we observe IP addresses in our honeypots and the corresponding destination ports they interact such as to profile their scanning behaviour.

[2]MaxMind: https://www.maxmind.com/en/home

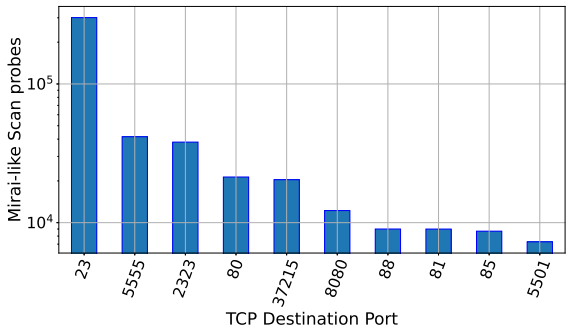

Fig. 4: Top 10 targeted ports determined by the number of connection attempts received by each port.

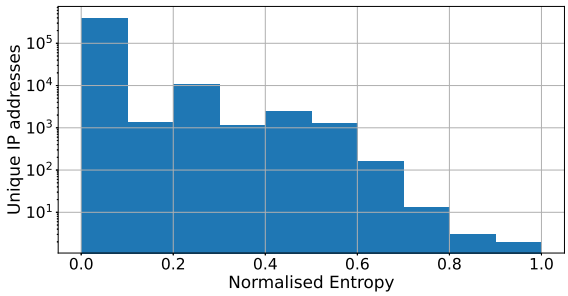

Fig. 5: Entropy distribution of destination ports scans from unique IP addresses: i) Low entropy values: IP addresses scan a small number of TCP ports, ii) High entropy: IP addresses scan random and multiple TCP ports.

V. Results

A. Scanning phase

Based on the publicly released Mirai source code, any botnets underpinned by Mirai-like malware would aggressively and randomly scan on TCP ports 23 and 2323 [6]. Both ports are dedicated for the Telnet protocol, enabling remote connection to a given IoT device. Nonetheless, as shown by Fig. 4 there are also other TCP ports targeted. It is evident, that new Mirai variants have expanded their target range on TCP ports by including vulnerabilities that are likely to persist on applications running HTTP/HTTPS services (e.g., web servers on TCP port 8080) and also HTTP-based protocols over TCP port 5555 enabling auto-configuration and remote management of home routers, modems, and other customer premises equipment (CPE). Moreover, in contrast with the original Mirai, we witness Mirai variants to scan for Universal Plug and Play (UPnP) services running on TCP port 37215. Thus targeting networked devices such as printers, mobile devices and Wi-Fi access points.

Fig. 5 represents the entropy distribution for the frequency in which TCP destination source ports are scanned by IP addresses. Based on the resulted distribution, it is evident that a large proportion of infected IPs has relatively low entropy. Thus, dictating that their scanning strategy is focused on specific TCP ports and their corresponding protocol-related vulnerabilities. IP addresses with higher entropy values seem to be more flexible and include more TCP ports in their scanning phase. Nonetheless, a much smaller portion of around 2000 IP addresses demonstrated random scanning properties over multiple TCP ports.

In order to determine the range of ports that are targeted by each Mirai-like scanner we assessed the number of unique ports related to each scan attempt. As demonstrated by Fig. 6, each scanner may scans a maximum of 15 ports with a minimum of 2 in every scanning session. Hence, in contrast with discussions (e.g., [6], [8]) on the full randomness of scanning strategies, we identify that even new Mirai variants have a carefully crafted and strategic scanning procedure.

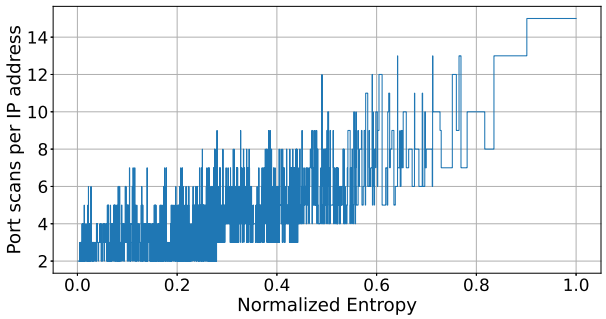

Fig. 6: Entropy distribution relating the frequency on the number of individual ports targeted by each Mirai-like scanner.

B. Infection phase

Through manual inspection over our CTI feeds we verify that Mirai-like botnets rely on brute force attacks on responsive IoT devices that operate over a vulnerable protocol. As demonstrated by Table III we identify, that the greatest majority of exploit attempts was related with vulnerabilities underpinning Remote Code Execution (RCE) on IP DVR, and TCT CCTV cameras of various vendors and also devices utilising the Android Debug Bridge (ADB). A much smaller fraction targeted home network access points (i.e., HNAP) and in particular NetGear routers. Evidently, we observe most of the exploits being related to more than one Common Vulnerability Exposure (CVE) tags indicating that IoT devices operating vendor-specific services have been unpatched for more than 7 years (e.g., AVTECH exploit). Therefore dictating the inadequacy of vendors on providing patching updates.

Assuming a response to a given Mirai-related scan from a vulnerable device, a handshake between the IoT device and a Report server is conducted. Our investigation revealed that the Report server redirects the vulnerable IoT device to a Loader or Malware server through a URL encoded in the payload of the first session packet. The encoded URL contains the location of the Mirai-like malware binary that is present on the Loader server having as a result the vulnerable device to download the actual binary.

In general, we identify 569 IP addresses mapped to one or

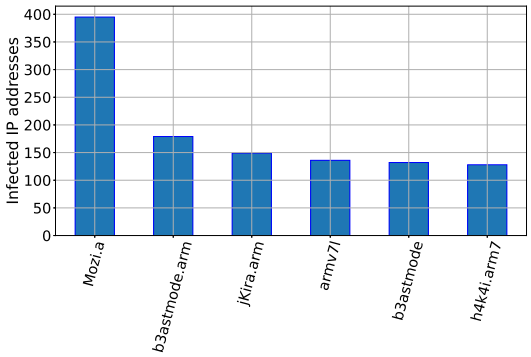

Fig. 7: Successful number of exploits of Mirai-variants on compromised IP addresses. Mozi.a, a new P2P Mirai variant overtakes by far most of exploits.

Vulnerability	Compromise Attempts	CVE Tag
AVTECH Exploit	85560	CVE-2013-4981 CVE-2013-4980
MVPower DVR RCE	66969	CVE-2018-10562 CVE-2018-10561 CVE-2017-17215
Android Debug Bridge (ADB)	47823	CVE-2019-6005
HNAP	4280	CVE-2015-2051 CVE-2020-10173 CVE-2020-9054 CVE-2018-17173
TVT (Generic OEM) DVR Targeted	2042	CVE-2017-8225 CVE-2017-5174 CVE-2017-7927'

TABLE III. The five most frequentl exploits across all Mirai-like variants indicating their mapping with multiple Common Vulnerability Exposure (CVE) tags that were unpatched on the infected devices.

more IoT devices that were successfully exploited with Mirai-like malware. As evident by Fig. 7, we observe around 400 of the exploits to be resulted by the propagation of a 2020 Mirai variant, Mozi.a. Through backtracking the properties of the Mozi.a binary, it was revealed that this particular Mirai-like botnet operates purely on P2P protocols. Hence, its expansion has progressed much more aggressively than the rest of the Mirai variant counterparts that relied on more centralised structural properties (e.g., Kira.arm). In addition, the Mozi.a variant is able to infect devices running on either ARM or x86 processor architectures, whereas the majority of the rest of the variants are purely focusing on ARM. Thus, centralised Mirai-like botnets are likely to compromise low-cost IoT devices running dedicated ARM architectures, whereas distributed variants such as Mozi.a are far more inclusive on more general-purpose IoT devices.

C. Botnet activity duration

In this section, we assess the duration of botnet activity by individual IP addresses as observed in our honeypots. Nonetheless we re-emphasise that an IP address may represent more than one IoT device and we utilise IP addresses as identifiers since devices may often use randomly or privately assigned IP addresses issued via DHCP or NAT.

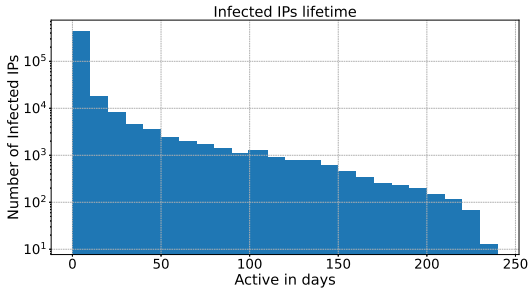

Fig. 8: Activity duration for infected IP addresses participating in Mirai-like botnets.

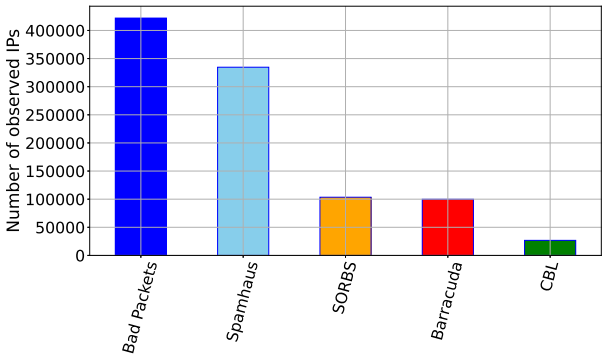

Fig. 9: Assessment on the amount of IP addresses monitored in our Bad Packets honeypots against four global IP blacklists. Spamhaus matched the greatest number.

As shown in Fig. 8, the largest number of identified IP addresses was active for less than 10 days. It was revealed that these addresses were mostly initiating scan traffic over a total of 2063 different TCP ports in the range of 0-8000. Thus, botmasters in Mirai-like variants tend to use a large number of bots for massive scans under the intention to expand their botnet but ensure that aggressive scan bots are not active for a long period. We speculate that this behaviour dictates an evasion technique from botmasters in order to stay undetected by corresponding network flooding detection mechanisms. However, observing the lifetime of IPs exceeding 10 days shows different behaviour. Evidently, only 47 unique TCP ports were scanned from IP addresses that could remain active for a much longer period reaching up to 200 days. Our analysis has also led to the conclusion that IP addresses that remained active for more than 100 days were demonstrating behaviours of Loader/Malware, Report and C&C servers.

D. Botnet IP address reputation

We investigate the reputation of all 422K IP addresses in our honeypots against four global IP blacklist databases used widely by ISPs and Internet registries; (i) Spamhaus [3], (ii) Barracuda [4], (iii) Spam Open Relay Blocking System

[3]Spamhaus:https://www.spamhaus.org/

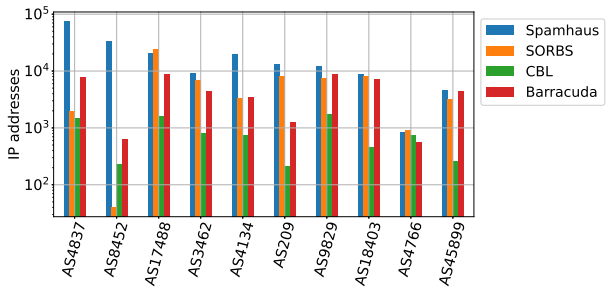

Fig. 10: AS-level distribution of correlated IP addresses captured on our honeypots against IP blacklist databases.

(SORBS) [5], and (iv) Composite Blocking List (CBL) [6]. Fig 9, demonstrates the number of IP addresses from our honeypot datasets found in each blacklist database. Evidently, not all monitored IP addresses from our honeypots were matched with blacklist entries. In particular, Spamhaus had information for 79% of the IP addresses in our datasets, SORBS 24.5% and Barracuda listed 23.6%, whereas CBL had the lowest number of hits with 6%. Moreover, through Fig. 10 we demonstrate that all blacklist databases are more effective on reporting IP addresses in particular ASes and do cover all ASes reported from our CTI honeypot feeds. Hence, they are more effective over particular regions of the Internet where potentially better inter-AS cooperation on IP blacklist reporting exists. Through this observation we therefore argue that holistic approaches on monitoring Internet-wide malicious activity should integrate input from CTI feeds and large-scale honeynets. Thus, increase the level of visibility across the Internet with additional vantage points and strengthen cyber threat assessment.

Through assessing the ratio of infected IP addresses against Spamhaus, it was identified that 77% of the correlated IP addresses were reported by ISPs for unauthenticated SMTP sessions. Thus, indicating that Mirai-variants are exploiting upper layer protocols and potentially these addresses participate into phishing campaigns. Moreover, 6% of correlated IPs listed were flagged in the past as malware hosts whereas 0.58% of them are involved into spamming activities. From both a SORBS and Barracuda as well as CBL perspective all correlated IPs were listed as addresses infected by malware.

VI. Conclusion

The evolution of IoT botnets since the 2016 Mirai botnet outbreak has undoubtedly transformed the cyber threat landscape in the global Internet affecting numerous socio-technical systems. Given the deployment of IoT devices with minimal embedded security, the number of vulnerabilities grows and also modern IoT botnets adapt to evade detection and exploit such vulnerabilities. In parallel, the public release of the Mirai malware source code has significantly aided towards the composition of modern Mirai-like botnets instrumenting attacks at various scales over a range of services and infrastructures. In this work, we provide a comprehensive measurement study focusing on Mirai-like variants. Through the collection of CTI feeds from 10 globally distributed honeypots over a period of 7 months we provide insights related to the structural properties as well as the evolving scanning and propagation strategies initiated by such botnets. Evidently, we witness the rise of Mozi, a new P2P-based Mirai-like variant. We also identify that Mirai-like botnets tend to go beyond the typical Telnet-based scanning and target more applications to achieve propagation. In addition, we demonstrate that IP blacklist databases commonly used by ISPs and Internet registries are not sufficient for profiling IoT botnet activity. We therefore highlight the importance of integrating feeds from distributed vantage points across the Internet. Hence, contributing towards adequate insights for the development of next generation IoT botnet profiling and detection schemes.

[4]Barracuda: https://www.barracudacentral.org/
[5]SORBS: http://www.sorbs.net/
[6]CBL: https://www.abuseat.org/

Acknowledgement

The authors are grateful to Bad Packets LLC, Max Mind, Barracuda, SpamHaus, CBL and SORBS for providing their datasets. This work has received funding from the UK Engineering and Physical Sciences Research Council (EPSRC) Network Measurement as a Service project (agreement: EP/N033957/1).

References

[1] Angrishi, K., "Turning internet of things (iot) into internet of vulnerabilities (iov): Iot botnets." arXiv preprint arXiv:1702.03681 (2017).
[2] Ross, M.,Hannes, T., Jara, A. "Baseline security recommendations for IoT in the context of Critical Information Infrastructures, 2017." European Union Agency for CyberSecurity -ENISA (2017).
[3] Dwyer, O. P., Marnerides, A., K., Giotsas, V., and Mursch, T. "Profiling IoT-based Botnet Traffic using DNS." In IEEE GLOBECOM 2019.
[4] Dainotti, A., King, A., Claffy, K., Papale, F., and Pescape, A., "Analysis of a "/0" stealth scan from a botnet." IEEE/ACM Transactions on Networking 23, no. 2 (2014): 341-354.
[5] Antonakakis, M., April, T., Bailey, M., Bursztein, E., Cochran, J., Durumeric, Z., et al. "Understanding the mirai botnet." In 26th USENIX Security 2017, pp. 1093-1110. 2017.
[6] Marnerides, A. K., Giotsas, V., and Mursch, T.,. "Identifying infected energy systems in the wild." In ACM SIGEnergy e-Energy 2019
[7] Gopal, T., S., Meerolla, M., Jyostna, G., Reddy Lakshmi Eswari, P. and Magesh, E., "Mitigating Mirai Malware Spreading in IoT Environment," in ICACCI 2018 pp. 2226-2230. IEEE, 2018.
[8] Marnerides, A. K., and Mauthe, A., U.,. "Analysis and characterisation of botnet scan traffic." In IEEE ICNC 2016
[9] Bou-Harb, E., Debbabi, M., and Assi, C.,. "Cyber scanning: a comprehensive survey." IEEE Communications Surveys Tutorials 16, no. 3 (2013): 1496-1519.
[10] Zainab, A., Kaafar, M., and Jha, S.. "Early detection of in-the-wild botnet attacks by exploiting network communication uniformity: An empirical study." In IFIP Networking 2017
[11] Pa, Yin Minn Pa, Shogo Suzuki, Katsunari Yoshioka, Tsutomu Matsumoto, Takahiro Kasama, and Christian Rossow. "IoTPOT: Analysing the rise of IoT compromises." In 9th USENIX (WOOT 15). 2015.
[12] Li, Zhichun, Anup Goyal, Yan Chen, and Vern Paxson. "Automating analysis of large-scale botnet probing events." In Proceedings of the 4th International Symposium on Information, Computer, and Communications Security, pp. 11-22. 2009.