Prevalence of IoT Protocols in Telescope and Honeypot Measurements

Lionel Metongnon[1,2,*] and Ramin Sadre[1]

[1] *Université catholique de Louvain, Belgium*
[2] *Université d'Abomey-Calavi, Bénin*
E-mail: lionel.metongnon@uclouvain.be
[*] *Corresponding Author*

Received 15 December 2018; Accepted 19 December 2018;
Publication 02 April 2019

Abstract

With the arrival of the Internet of Things (IoT), more devices appear online with default credentials or lacking proper security protocols. Consequently, we have seen a rise of powerful DDoS attacks originating from IoT devices in the last years. In most cases the devices were infected by bot malware through the telnet protocol. This has lead to several honeypot studies on telnet-based attacks.

However, IoT installations also involve other protocols, for example for Machine-to-Machine communication. Those protocols often provide by default only little security. In this paper, we present a measurement study on attacks against or based on those protocols. To this end, we use data obtained from a /15 network telescope and three honey-pots with 15 IPv4 addresses. We find that telnet-based malware is still widely used and that infected devices are employed not only for DDoS attacks but also for crypto-currency mining. We also see, although at a much lesser frequency, that attackers are looking for IoT-specific services using MQTT, CoAP, UPnP, and HNAP, and that they target vulnerabilities of routers and cameras with HTTP.

Keywords: Internet measurement, IoT, IoT attacks, IoT protocols.

1 Introduction

With the arrival of the Internet of Things (IoT), many devices appear online with default credentials or lack proper security protocols [7, 20].

Journal of Cyber Security and Mobility, Vol. 8_3, 321–340.
doi: 10.13052/jcsm2245-1439.832

Furthermore, many IoT devices are difficult to update and are, once installed, forgotten by their users. Consequently, we observe more large-scale attacks DDoS attacks from these devices [2]. A recent example is the Mirai malware [16] and its variants that targets unsecured IoT devices and infects them with botnet software.

The most successful attacks misusing IoT devices started with an infection through telnet. Therefore, the majority of recent publications in this area have focused on this protocol. However, IoT installations often also involve other protocols, such as the Constrained Application Protocol (CoAP), the Message Queuing Telemetry Transport protocol (MQTT), and the protocols used by Universal Plug and Play (UPnP) for configuration and service discovery. Those protocols are specialized lightweight protocols and provide by default only little security. Due to the enormous "success" of telnet-based attacks, the prevalence of those protocols in attack traffic has not been widely studied.

The goal of this paper is to close this gap. We present a measurement study on attacks against common protocols used by IoT devices. To this end, we use data obtained from a large /15 network telescope operated by SurfNet[1] and three honeypots with 15 IPv4 addresses. We find, not unexpectedly, that telnet-based malware is still widely used. Furthermore, we observe that the infected devices are employed not only for DDoS attacks but also for crypto-currency mining. However, we also see that attackers are looking for IoT installations using the above mentioned other protocols. Such attacks are much rarer than telnet-based ones, but it can be expected that their frequency will increase in the near future in light of the increasing popularity of IoT and M2M communication.

The rest of this paper is organized as follows: We discuss related work in Section 2. We briefly present the studied protocols in Section 3. We describe the experimental setup in Section 4 and discuss results in Section 5. We conclude the paper in Section 6.

This paper is an extension of prior work published in [14]. While the methodology and the conclusions have stayed the same, our honeypot experiments are now based on a dataset that is more than twice as large as the original one. Consequently, all related figures and statistics in Section 5.2 have been replaced and the discussion updated. The longer measurement period allows us to observe the evolution of incoming honeypot traffic on a larger timescale. We have also added a new Section 5.3 with results from the Shodan database on the IP addresses contacting the honeypots.

[1] https://www.surf.nl/en/about-surf/subsidiaries/surfnet

2 Related Work

The usage of network telescopes and honeypots for security research has a long tradition. By design, nearly[2] all traffic reaching them is malicious (e.g. a network scan), unintended (e.g. caused by misconfiguration), or the result of attacks with spoofed IP addresses, the so-called *backscatter* traffic. A general discussion on telescopes and honeypots is out of the scope of this paper and we refer the reader to seminal papers such as [15]. Instead, we will focus on IoT in the following.

The main intention of attacks against IoT devices is to turn them into bots for botnets. Kolias *et al.* [11] point out many reasons why IoT devices are interesting for botnets, such as the fact that they are online 24/7, their lack of security features, and their poor maintenance. In addition, their capability to launch powerful DDoS attacks has been underestimated for a long time and therefore their protection has not been given high priority by network administrators. The sheer number of available vulnerable devices allows to perform highly distributed DDoS attacks. This is a major difference to most non-IoT based DDoS attacks; in fact, Krämer *et al.* reported in [12] that major amplification-based DDoS attacks were short-lived and 96% came from single sources.

IP-based cameras are attractive victims for botnet operators because they often have good network connectivity and are poorly secured. Several powerful DDoS attacks performed by such cameras have been reported [3, 4, 8]. Pa Pa *et al.* [16] implemented a honeypot for telnet-based attacks and identified at least four distinct DDoS malware families supporting as many as nine different CPU architectures.

The *Mirai* malware was observed first in 2016. Infected hosts scan the Internet for other targets (leaving out some address ranges belonging to Hewlett-Packard, General Electric, and the US Department of Defense) and use dictionary attacks to gain access to the devices on telnet ports 23 and 2323. Many victims are cameras and routers that are still using their factory user name and password. A distinct feature of the original Mirai malware is the fact that it sends very efficiently "stateless" TCP SYN packets with the initial sequence number (ISN) identical to the target IP address [5].

Hajime is a Mirai-like botnet using BitTorrent's DHT protocol for peer discovery and the uTorrent Transport Protocol (uTP) for data exchange [6]. Its purpose is still unknown.

[2] An exception is for example the traffic from benign network scans used by researchers for Internet wide studies.

BrickerBot is a more exotic malware because it makes infected devices completely unusable. In that way, the authors wanted, so they claim, to protect the Internet from unsecure devices. BrickerBot contains attack vectors for telnet, SSH, HTTP, HNAP and SOAP [10].

Our findings confirm several of the observations made in the above publications, respectively update them. As an example for the latter, we find that only 10% of the Mirai-like telnet traffic respects the ISN pattern, which indicates that attackers have adapted the original source code of Mirai [13]. More important, there are, to our knowledge, no existing studies on the prevalence of attacks against the other IoT-related protocols discussed in this paper.

A completely different usage of network telescopes can be found in [17]. Shaikh et al. use a combination of passive and active (scanning) measurement to discover and characterize unsolicited IoT devices.

3 Background

In this section, we give a short introduction to the protocols studied in this paper.

Telnet is a text-oriented interaction protocol based on TCP providing a bidirectional communication to a device. It is typically used to send commands to a terminal service. While not used anymore on servers and workstations, it is still very popular on embedded systems because of its simplicity. The access to the command line is granted after authentication, however owners of IoT devices often do not change the default credentials.

CoAP is a protocol for machine-to-machine (M2M) communication, running on top of UDP [18]. Encryption with DTLS is optional. CoAP servers, i.e. the IoT devices, provide a REST interface that can be used by applications to retrieve sensor data or to change configuration parameters. Like in HTTP, resources are addressed by URIs. Clients can obtain a list of the available resources of a CoAP server by requesting the resource `/.well-known/core`.

MQTT [1] is another protocol for M2M communication. In contrast to CoAP's client-server model, MQTT uses a publisher-subscriber model where a *broker* server forwards the data received from the publishers (i.e. the IoT devices) to the subscribers. MQTT uses TCP as transport protocol with optional usage of SSL. Most noteworthy, connections to the broker are always initiated by the publishers and subscribers. In theory, this means that IoT devices using MQTT do not need to accept incoming connections and only the broker is of interest for MQTT-based attacks.

UPnP is a set of network protocols used in home automation to discover the presence of devices in a network and to establish services, such as data sharing and streaming [9]. UPnP is activated by default on many home network devices like home routers, web cams, printers, and cameras. UPnP's service discovery protocol SSDP uses UDP.

Finally, *HTTP* is used to exchange hypertext messages. Many devices, such as routers, provide a web interface for configuration that can be accessed through HTTP. HTTP is also the underlying protocol of many other application protocols. For example, the Home Network Administration Protocol (HNAP) by Cisco is used for the management of network devices. A device supporting HNAP will reply with a valid response to HTTP requests for the resource `/HNAP1` [19].

4 Experimental Setup

In this section, we describe our measurement setup to investigate the presence of IoT-related attacks in the Internet.

4.1 Telescope

SURFnet is a company of SURF, the collaborative ICT organization for Dutch education and research. Its mission is to support innovation through their infrastructures. We use the tool *eemo*[3] to access their /15 telescope. We capture 10 minutes of packet data every hour from 2017-09-25 to 2018-02-28 with an interruption in October. Since the telescope is passive and therefore does not respond to incoming traffic, we use it only to record connection attempts as well as to identify the most used protocols and targeted ports. Compared to our honeypots with 15 IP addresses (see below), the telescope with its large address range gives a better insight into what to expect in number of attacks on a large scale, while we use the honeypots for the detailed analysis of the individual attacks.

4.2 Honeypots

A honeypot is software mimicking specific services or devices and is used to attract malicious traffic. In our experiments, we are interest in IoT-related service and protocols. SURFnet provides us 15 IP addresses that we use to deploy the three honeypots *Cowrie*, *dionaea* and *HoneyPy*.

[3] https://github.com/SURFnet/eemo

Cowrie[4] is a medium interaction honeypot impersonating a server with SSH and telnet service and easily cracked login credentials. Attackers who connect to the honeypot see a fake filesystem. All TTY message exchanges and attempts to download and install malware are recorded. We use version 1.1 (we started with commit 3d12c8c54 and an upgrade to 4f0fc85e02) of the honeypot.

Dionaea[5] is also a medium interaction honeypot. Similar to Cowrie, it records all message exchanges. We use version 0.6.0 (we started with commit ac971c3ab and upgrade to 02492e2b973) with the following services: EPMAP, FTP, HTTP, Memcache, MQTT, MSSQL, MySQL, PPTP, SIP, SMB, and UPnP.

HoneyPy[6] is a low-interaction honeypot. We use its services TFTP, TR-069.1 and TR-069.2 (a management protocol on port 5555 and 7574 used by some ISPs to configure the modems of their subscribers), and TelnetIoT (telnet on port 2323). In addition, we use a module developed by Dany Yarakou[7] that emulates a CoAP server providing several resources. Clients can access a resource containing a simple counter number, a resource returning the current time, resources sending large data in blocks, a separate large resource that sends intermediate ACKs to indicate that the processing of the request is delayed, and finally the `/.well-known/core` resource that gives a list of all the existing resources on the server.

We did not operate the honeypots continuously over the measurement period. This was partially caused by technical difficulties that forced us to shut them down for maintenance. For example, the interruption in January 2018 was caused by an update to mitigate Meltdown/Spectre.

To get more information about the hosts accessing the honeypots, we used Shodan[8], a well-known search engine for Internet-connected devices. Shodan periodically scans IP addresses in the Internet and collects information about the host behind the addresses, including information on their location, available services, and operating system. The Shodan databases can be queried programmatically through an API or manually through a web user interface.

[4] https://github.com/micheloosterhof/cowrie
[5] https://github.com/DinoTools/dionaea
[6] https://github.com/foospidy/HoneyPy
[7] https://github.com/DanMistyk/HoneyPy
[8] https://www.shodan.io

5 Results

We present the results of our analysis of the data obtained from the telescope (Section 5.1) and the three honeypots (Section 5.2).

5.1 Telescope Results

Figure 1 shows the number of packets that the telescope received per day. Due to technical issues, no data was collected in the period from 2017-10-05 to 2017-11-01, on 2018-01-20, 2018-02-14, and 2018-02-19. Apart from that, the setup was not changed during our experiments.

With 96.1% of all observed packets, TCP is the dominant protocol, followed by UDP (2.8%) and ICMP (1.1%).

Figure 2 indicate the disappearance of the ICMP traffic after 2017-12-08 without any specific reason. These ICMP traffic were already very shallow on the telescope and are present only on 47 days during the measurement. We identify 29 different universities probing with 137,760 packets send covering only 3.53% of the total probes received (Table 1). However, University of Michigan is the 11 top AS in probes quantity with 74,061 on the period, That

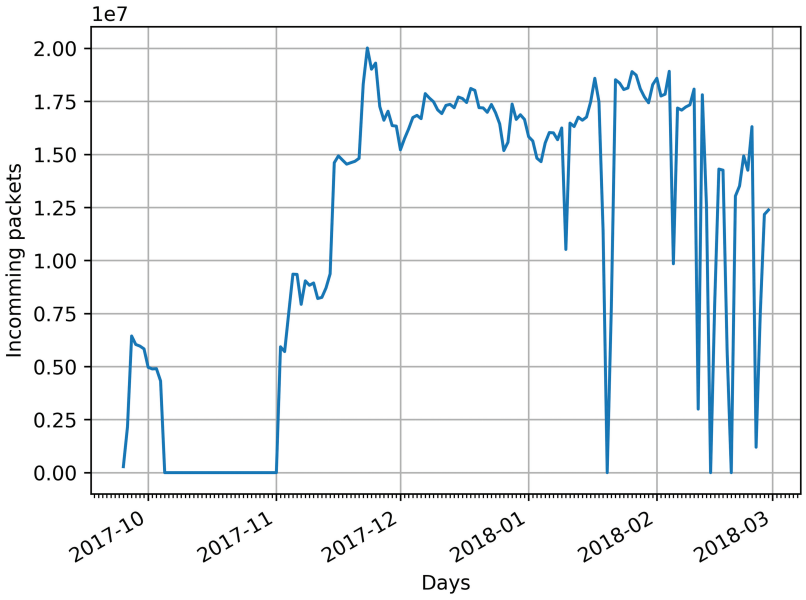

Figure 1 Number of packets per day reaching the telescope. Note the scaling factor of 10^7 for the y-axis.

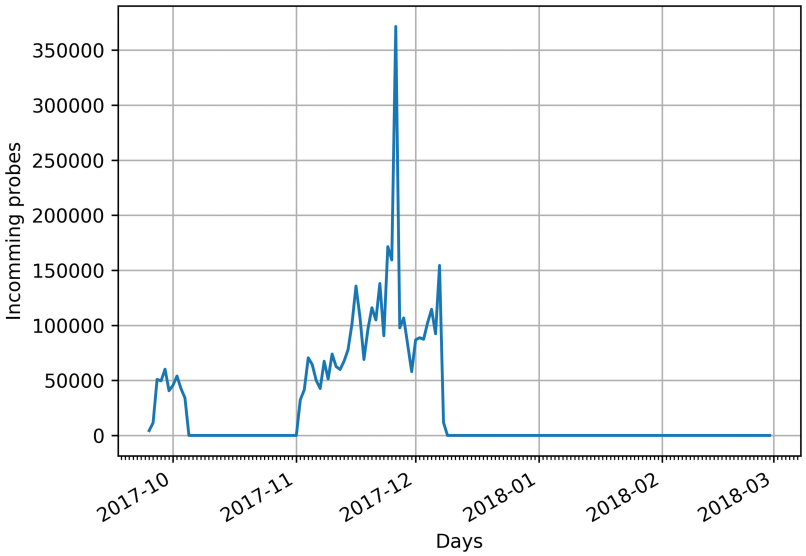

Figure 2 Number of packets per day reaching the telescope. Note the scaling factor of 10^7 for the y-axis.

Table 1 Top 10 of AS owners by involved IP addresses

Autonomous System Owner	Number of Distinct IP Addresses
No. 31, Jin-rong Street	9,373
PT Telekomunikasi Indonesia	4,030
CHINA UNICOM China169 Backbone	2,652
VNPT Corp	1,509
China Unicom Beijing Province Network	903
National Internet Backbone	796
TELEFNICA BRASIL S.A	692
PT Excelcomindo Pratama (Network Access Provider)	635
PJSC Rostelecom	618
TOT Public Company Limited	597

means the quantity is not negligible despise the percentage. Only No. 31, Jin-rong Street and CHINA UNICOM China169 Backbone appear in Table 2 sending many packets matching the quantities of their IP addresses.

Figure 3 shows the number of packets received by the telescope per destination port. We can then see that port 23 (telnet) is the most active port with 29.37% of all packets followed by port 22 (SSH) and and port 2323.

Table 2 Top 10 of AS owners by incoming traffic

Autonomous System Owner	Number of Incoming Packets
Hangzhou Alibaba Advertising Co.,Ltd.	574,699
No. 31, Jin-rong Street	462,393
SuperNetwork s.r.o.	410,297
Linode, LLC	199,525
HLL LLC	193,630
CHINA UNICOM China169 Backbone	156,772
Los Nettos	92,063
QuadraNet, Inc	89,633
Digital Ocean, Inc.	83,659
Guangdong Mobile Communication Co.Ltd.	81,609

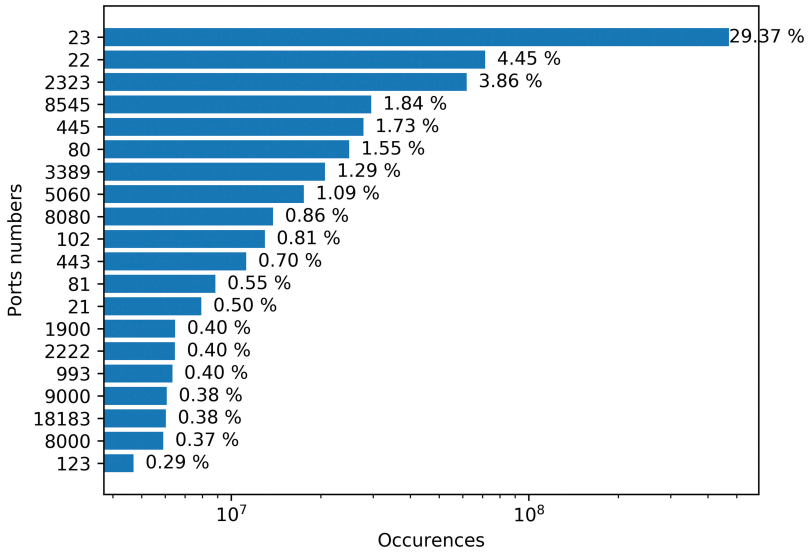

Figure 3 Incoming packets per port on the telescope (top 20). Note the scaling factor of 10^8 for the x-axis.

Officially, port 2323 has been assigned by the IANA to the *3d-nfsd* service, but we know from Mirai that it also probes this port for telnet. Port 8545 is the default port of the JSON RPC protocol. Its popularity could be explained by the fact that it is used by Ethereum clients. The ports associated with HTTP and HTTPS (port 80, 81, 8080, 8081, 8090, 443) and services such as SMB (port 445), FTP (port 21), and IMAP (port 993) also appear frequently.

Among IoT-specific protocols, UPnP's Simple Service Discovery Protocol (SSDP) on port 1900 is the most popular one (0.4%). In contrast, only 0.017% of the packets target port 1883 (MQTT), 0.018% target port 8883 (MQTT over SSL), and 0.005% target port 5683 (CoAP). Visibly, these protocols are not (yet) popular among attackers. Possible reasons for this could be that (a) only few devices use these protocols, (b) effective attacks against the corresponding services are not known, or (c) most affected devices are hidden behind NAT and firewalls or use IPv6 instead of IPv4. Points (a) and (c) are partially confirmed by the Shodan search engine: a search for port 1900 returns 4,096,473 results, compared to only 48,190 for port 1883.

5.2 Honeypot Results

5.2.1 General overview

Figure 4 shows the number of packets received by the honeypots per day. It should be noted that these (and the following) numbers are not directly comparable to those from the telescope since the honeypots are not passive and encourage connecting hosts to exchange more packets. The periods where the honeypots were not active are visualized as a zero baseline, for example in October 2011 and August 2018. Despite this gaps in the data, we can see that the

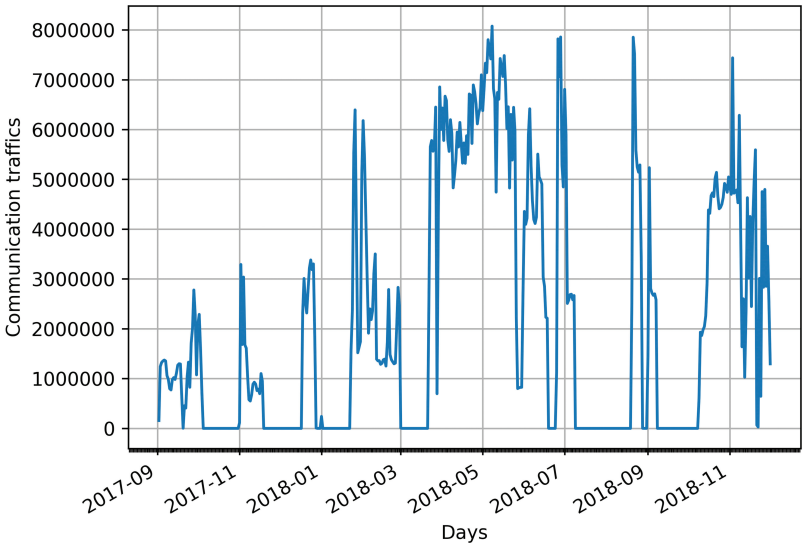

Figure 4 Incoming packets per day on the honeypots.

traffic intensity in 2018 is higher than in 2017. While receiving more than $3{\cdot}10^6$ per day was a rare event until March 2018, it has become the normal situation in the last nine months of the measurement period. The violent fluctuations appearing in the time series in the form of distinct peaks are caused by sharp and short raises in the number of long telnet connections. We will discuss this in Section 5.2.3.

In total, around 74.5% of the received packets were targeting UDP and TCP services. The remaining 25.5% are discovery attempts with ICMP. We can see a huge disparity regarding the ICMP probes received on the honeypots compared to the telescope. Figure 5 shows the distribution of the ICMP packets on the honeypots. A total of 269,367,609 probes were sent from only 6,539 different source addresses to the 15 addresses of the honeypots. We find that only 31.38% of these addresses also sent TCP traffic and only 9.13% also sent UDP traffic. We find 87 sources sending more than a million probes, respectively. Table 3 shows the top 10 most active probers. We have an average of 962,027.175 probes per measurement day on the honeypots. The number-one prober is active 279 days out of 280 measurement days with a peak of 41,191 probes on 2017-12-22. This source IP did not sent any other types of packets.

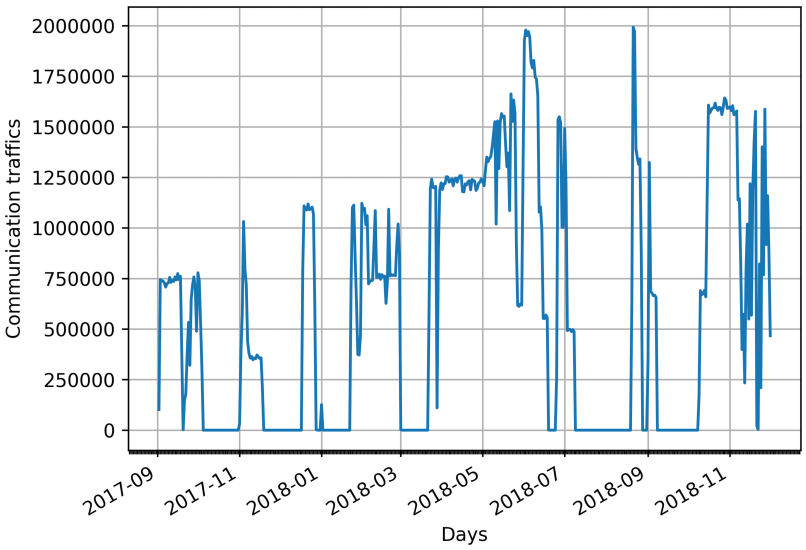

Figure 5 Incoming ICMP probes per day on the honeypots.

Table 3 Top 10 sources sending ICMP probes. For privacy reasons, only the first 24 bits of the addresses are shown

IP Address	Number of Probes
203.205.144.*	6,773,751
49.51.128.*	3,625,689
119.28.96.*	3,625,612
184.105.66.*	1,221,235
183.60.64.*	3,622,302
182.254.4.*	3,622,035
203.195.225.*	3,621,781
113.99.136.*	3,620,991
115.159.138.*	3,618,200
183.232.164.*	3,618,094

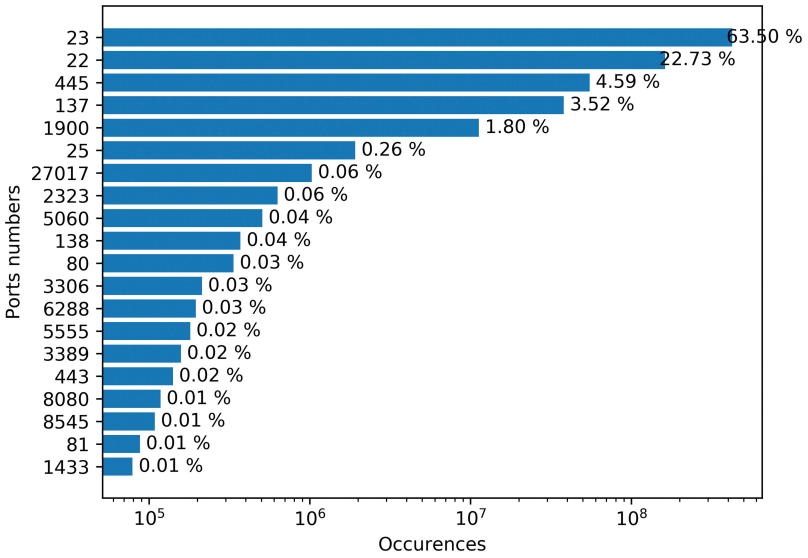

Figure 6 Incoming packets per port on the honeypots (top 20). Note the logarithmic x-axis.

Figure 6 shows the number of received packets per port. Telnet is by far the most popular protocol, followed by SSH. We will discuss this in the next section.

Table 4 shows the AS owners of the IP addresses contacting the honeypots. Those are not necessarily the most active IP addresses in terms of interactions with the honeypots because many of them only probe ports without further activities. Indeed, Table 5 shows a completely different set of top AS (with

Table 4 Top 10 of AS owners by involved IP addresses

Autonomous System Owner	Number of Distinct IP Addresses
TELEFONICA BRASIL S.A	147,865
CHINA UNICOM China169 Backbone	70,181
No. 31, Jin-rong Street	67,014
Magyar Telekom plc.	52126
PJSC Rostelecom	45,987
Korea Telecom	40,536
Uninet S.A. de C.V.	35,066
OVH SAS	22,268
Turk Telekom	13,994
Data Communication Business Group	13,688

Table 5 Top 10 of AS owners by incoming traffic

Autonomous System Owner	Number of Incoming Packets
Global Layer B.V.	107,121,191
FranTech Solutions	71,071,903
Digital Ocean, Inc.	69,130,300
Shenzhen Tencent Computer Systems Company Limited	52,137,841
Tencent Building, Kejizhongyi Avenue	43,822,755
Regionalnaya Kompaniya Svyazi Ltd.	38,828,910
Aruba S.p.A.	35,167,720
Hostio Solutions B.V.	23,124,288
SoftLayer Technologies Inc.	18,656,831
No. 31, Jin-rong Street	18,093,637

the exception of No. 31, Jin-rong Street) if we sort them by the number of sent packets. Of course, it could be that attackers from those AS are behind NATs, in this way resulting in a smaller number of visible addresses.

5.2.2 SSH

Before we present the results for the various IoT-related protocols, we want to briefly discuss the traffic received by Cowrie on port 22. In number of packets, SSH is the second most active protocol (Figure 6). One could expect that SSH sessions are used in a similar way as telnet. However, a manual inspection of the traffic reveals that it mostly consists of connection attempts without further communication or of reflected traffic. In both cases, we don't have enough information to decide whether they are related to IoT services. 50.68%

(99,379,132 packets) of the ssh traffic come from the Autonomous System Owner `Global Layer B.V.` We also spotted 249 packets coming from well-known scanning projects like Shodan but they don't have a significant impact on our analysis due to their small number.

5.2.3 Telnet

All 85,788,366 login attempts on telnet follow the same pattern using a list of default credentials with the two most frequent ones being `root/vizxv` (8.1% of all attempts) and root/12345 (3.7%). On average, attackers send 10 packets per telnet connection. This low number is mostly due to the presence of many short scans on ports 23 and 2323.

After a successful login, we see two different procedures: Some attackers write directly commands (or shellcode) on the terminal, the others first download a script file containing the commands. Attacks typically start with a series of `enable` commands to enable built-in functionalities of bash. Then, they create files with random content to check the existence of standard directories like /, /sys, /proc, /dev, /dev/pts, /run/bin, /dev/shm, /run/lock, /proc/sys/fs /proc/sys/fs/binfmt_misc/, /boot, /home. In this way, they try to identify honeypots. Finally, they download malware using tftp, curl or wget, that gives them root access to the system. They also download modified versions of services like ssh, telnet and ntpd to avoid future intrusions.

As mentioned in Section 5.2.1, the traffic on the honeypots is far from being evenly distributed over the measurement days. We saw two sharp rises on January 25 and February 1 where the average number of sent packets per telnet connection went up to 41 packets. On these days, the honeypots were mostly under attacks from the same IP address sending series of commands over the telnet connection before downloading binaries, whereas on the other days the script-based approach was dominant.

We see mostly Mirai and Mirai-like attacks. All these attacks show the typical sequence of commands targeting embedded systems with busybox[9] that gave Mirai its name [5]. However, Mirai's network signature, i.e. SYN packets with the ISN identical to the destination IP, is visible in only 10% of the attacks. We see variants that do not show this signature and that also include other changes (like a modified set of the IPs hardcoded in the original Mirai). Several of these attacks also try to install miner software (namely minerd, xmrminer, and cpuminer-multi) for crypto-currencies, we therefore conclude that attackers try to use IoT devices to harvest crypto-currencies. A deeper

[9] https://busybox.net/

analysis of the different binaries shows miner software for bitcoin, ethereum and litecoin mostly.

5.2.4 HTTP

Only a few connections are targeting HTTP. Most of the 8846 incoming HTTP requests are not specific enough to relate them to IoT. We see 3190 requests for the `/` resource and 2417 requests with 61 different versions of attack URL against phpMyAdmin, the web-based administration tool for MySQL. We also spotted 122 requests targeting `bot.php`. Those are either attempts to exploit vulnerable chat bot scripts installed by normal users or checks to detect whether the host has been already successfully hacked by others and running a chat bot. We have filtered out those requests and manually analyzed the remaining ones for signs that the attackers were specifically targeting IoT systems.

0.76% of the HTTP requests try to access CGI resources of routers with known vulnerabilities, namely routers from Cisco (tmUnblock.cgi), Linksys (hndUnblock.cgi), and D-Link (hedwig.cgi). 0.85% of the attempts target getcfg.php to exploit vulnerabilities on D-Link DIR-6XX and DIR-8XX routers.

Much less frequent (0.12%) are attacks against Cisco's Home Network Administration Protocol (HNAP) for the management of home networks. The attackers request the resource `/HNAP1` in order to identify equipment supporting that protocol. 90% of those attacks came from only one IP address. The attacker used Firefox's user agent `Mozilla/5.0 (Windows NT 5.1; rv:9.0.1) Gecko/20100101 Firefox/9.0.1` in all attempts. Finally, 0.11% of the HTTP requests are targetting IP cameras with the Dahua backdoor (/current_config/passwd) and CVE-2017-8225 (/system.ini?loginuse&loginpas).

5.2.5 Other IoT protocols

We see 726,978 attempts to get information on UPnP devices by using UPnP's service discovery protocol. As expected, most of those service discovery requests use the URI "*", but we also see 22 requests for http://www/. We assume that the latter are from a custom implementation of the protocol since they came all from the same IP address. It is reasonable to assume that these requests are caused by a misconfiguration and are not meant to be an attack.

The 702 received CoAP messages are requests to the standard resource `/.well-known/core` with which clients can obtain the list of available resources from a server.

We received 1277 MQTT messages of mainly the four types Connect, Subscribe, Disconnect, and Ping. We observe that 87.70% of the traffic came from the same address and were concentrated on the period from 2018-01-26 to 2018-02-10. We note that the number of MQTT packets drastically decreased after February to only 81 messages exchanged in 10 days. However the number of source IP addresses is now 11.

We didn't observe any attempt to further exploit the above protocols beyond the first service-discovering interaction, although this is supported by the honeypots and many exploits of UPnP vulnerabilities are known.

5.3 Shodan Results

As explained in Section 4.2, we use Shodan to obtain more information on the hosts behind the IP addresses that contacted the honeypots. The insight that can be gained by this or similar approaches is of course limited since hosts might be behind NATs.

We find different versions of Linux (3.x, 2.6.x, 2.4–2.6) but also many Windows versions, namely (Server 2012 R2 Standard 9600, 7 or 8, Server 2008 R2 Standard 7601 Service Pack 1, Server 2012 R2 Datacenter 9600, 6.1, XP, 10 Pro 17134 etc.), Unix and FreeBSD 9.x. We even find some Playstation 4 and HP-UX 11.x. However, Shodan is only able to provide this information for a small fraction of the observed IP addresses, therefore we will not discuss this further.

71.96% of the IP addresses connecting to the honeypot listen to the typical HTTP(S) ports 80, 8080 and 443, 8.51% have telnet (port 23), and 45.03% have ssh (port 23). After checking the type and version (TP-LINK WR740N WAP http config, Linksys wireless-G WAP http config, Netgear WNR3500L WAP http config, D-Link DCS-930L_46 webcam http interface), we conclude that many hosts are IoT devices (like IP cameras) or home routers, or are behind NATs running on such routers.

6 Conclusion

In this paper, we deployed honeypots and used a telescope to observe attempts of attacks against IoT devices and infrastructures. Our theory was that since the IoT uses many different protocols with weak or no authentication, many malicious acts are performed through them. However, our study shows that telnet-based attacks against IoT devices are still the most frequent ones, even raising during our measurements, probably caused by the enormous success of

such attacks against surveillance cameras and home routers. However, while the original Mirai botnet was mostly used to perform DDoS, we see that new variants of Mirai-like malware also install miners for crypto-currencies. Furthermore, we see that those variants do not exhibit anymore Mirai's standard signature (SYN packets with the ISN identical to the destination IP address).

Attacks relying on IoT-specific protocols, such as MQTT, UPnP's SSDP and CoAP, are still much rarer. In our data, connection attempts using these protocols were limited to service discovery interactions without further attempts to find or exploit vulnerabilities. Nevertheless, we expect to see more attacks in the future using these or other IoT protocols. The Mirai attacks have drawn a lot of attention to telnet-related vulnerabilities. Once those are fixed, attackers will turn to other protocols to infect IoT devices.

Our approach has several limitations that we want to overcome in future work. First of all, we want to extend our measurements to IPv6 in order to capture attack attempts against large-scale IoT infrastructures using this protocol. Second, our measurement setup was based on the idea to study IoT protocols starting with the most widely used and standardized ones. For the future, we plan to extend the honeypots in order to support more protocols, especially proprietary ones as employed by several manufacturers of IoT devices intended for home and industrial automation.

Acknowledgements

We thank Roland van Rijswijk-Deij (SURFnet and University of Twente) for providing us access to the SURFnet infrastructure and helping us with the configuration of the measurement equipment.

References

[1] Andrew Banks and Rahul Gupta. Mqtt version 3.1. 1. *OASIS standard*, 29, 2014.

[2] Elisa Bertino and Nayeem Islam. Botnets and internet of things security. *Computer*, 50(2):76–79, 2017.

[3] D. Cid. Large cctv botnet leveraged in ddos attacks. https://blog.sucuri.net/2016/06/large-cctv-botnet-leveraged-ddos-attacks.html, 2016. Accessed: 2018-02-11.

[4] L. Constantin. Thousands of hacked cctv devices used in ddos attacks. http://www.pcworld.com/article/3089346/security/thousands-of-hacked-cctv-devices-used-in-ddos-attacks.html, 2016. Accessed: 2018-02-11.

[5] Alexandre Dulaunoy, Gérard Wagener, Sami Mokaddem, and Cynthia Wagner. An extended analysis of an iot malware from a blackhole network. In *TNC17*, 2017.

[6] Sam Edwards and Ioannis Profetis. Hajime: Analysis of a decentralized internet worm for iot devices. *Rapidity Networks*, 16, 2016.

[7] J. Frahim, C. Pignataro, J. Apcar, and M. Morrow. Securing the internet of things: A proposed framework. https://www.cisco.com/c/en/us/about/security-center/secure-iot-proposed-framework.html. Accessed: 2017-03-31.

[8] O. Gayer, O. Wilder, and I. Zeifman. Cctv ddos botnet in our own back yard. https://www.incapsula.com/blog/cctv-ddos-botnet-back-yard.html. Accessed: 2018-02-11.

[9] Michael Jeronimo and Jack Weast. Upnp design by example, 2003.

[10] Simon Kenin. Brickerbot mod_plaintext analysis. https://www.trustwave.com/Resources/SpiderLabs-Blog/BrickerBot-mod_plaintext-Analysis/, 2017. Accessed: 2018-03-30.

[11] Constantinos Kolias, Georgios Kambourakis, Angelos Stavrou, and Jeffrey Voas. Ddos in the iot: Mirai and other botnets. *Computer*, 50(7):80–84, 2017.

[12] Lukas Krämer, Johannes Krupp, Daisuke Makita, Tomomi Nishizoe, Takashi Koide, Katsunari Yoshioka, and Christian Rossow. Amppot: Monitoring and defending against amplification ddos attacks. In *International Workshop on Recent Advances in Intrusion Detection*, pages 615–636. Springer, 2015.

[13] Brian Krebs. Source code for iot botnet mirai released. https://krebsonsecurity.com/2016/10/source-code-for-iot-botnet-mirai-released/, 2016. Accessed: 2018-02-11.

[14] Lionel Metongnon and Ramin Sadre. Beyond telnet: Prevalence of iot protocols in telescope and honeypot measurements. In *Proceedings of the 2018 Workshop on Traffic Measurements for Cybersecurity*, pages 21–26. ACM, 2018.

[15] David Moore, Colleen Shannon, Douglas J. Brown, Geoffrey M. Voelker, and Stefan Savage. Inferring internet denial-of-service activity. *ACM Trans. Comput. Syst.*, 24(2):115–139, May 2006.

[16] Yin Minn Pa Pa, Shogo Suzuki, Katsunari Yoshioka, Tsutomu Matsumoto, Takahiro Kasama, and Christian Rossow. Iotpot: analysing the rise of iot compromises. *EMU*, 9, 2015.

[17] Farooq Shaikh, Elias Bou-Harb, Nataliia Neshenko, Andrea Patrice Wright, and Nasir Ghani. Internet of malicious things: Correlating active and passive measurements for inferring and characterizing internet-scale unsolicited iot devices. *IEEE Communications Magazine*, March 2018.

[18] Zach Shelby, Klaus Hartke, and Carsten Bormann. Rfc 7252 - the constrained application protocol (coap). 2014.

[19] Cisco Systems. Home network administration protocol (hnap) whitepaper. https://www.cisco.com/web/partners/downloads/guest/hnap_protocol_whitepaper.pdf , 2009. Accessed: 2018-03-30.

[20] Tianlong Yu, Vyas Sekar, Srinivasan Seshan, Yuvraj Agarwal, and Chenren Xu. Handling a trillion (unfixable) flaws on a billion devices: Rethinking network security for the internet-of-things. In *HotNets 2015*, 2015.

Biographies

Lionel Metongnon is a Ph.D. student at ICTEAM institute of Université catholique de Louvain at Belgium, since Spring 2015. He attended the Université d'Abomey-Calavi in Bénin where he received his B.Sc. in Electrical engineering and industrial IT in 2011 and his M.Sc. in Computer Science in 2014. His Ph.D. works focus on network monitoring and distributed Internet-scale intrusion detection for Internet of Things.

Ramin Sadre has been a professor in the ICTEAM institute of UCLouvain, Belgium, since 2014. Before that, he was an assistant professor at Aalborg University, Denmark, and a post-doctoral researcher at the University of Twente, the Netherlands. His research activities focus on performance evaluation, monitoring of networked systems, and network-based intrusion detection, targeting open Internet-wide distributed applications as well as more closed systems such as IoT and SCADA.