Securing IoT Devices using Geographic and Continuous Login Blocking: A Honeypot Study

Fredrik Heiding[a], Mohammad-Ali Omer, Andreas Wallström and Robert Lagerström[b]
Division of Network and Systems Engineering, KTH Royal Institute of Technology, Stockholm, Sweden

Keywords: IoT, GeoIP, Fail2ban, Honeypot, Cowrie, p0f, Conpot, Snort, Suricata, Geographic Blocking.

Abstract: IoT (Internet of Things) devices have grown exponentially in the last years, both in the sheer number of devices and concerning areas of applications being introduced. Together with this rapid development we are faced with an increased need for IoT Security. Devices that have previously been analogue, such as refrigerators, door locks, and cars are now turning digital and are exposed to the threats posed by an Internet connection. This paper investigates how two existing security features (geographic IP Blocking with GeoIP and rate-limited connections with fail2ban) can be used to enhance the security of IoT devices. We analyze the success of each method by comparing units with and without the security features, collecting and comparing data about the received attacks for both kinds. The result shows that the GeoIP security feature can reduce attacks by roughly 93% and fail2ban by up to 99%. Further work in the field is encouraged to validate our findings, create better GeoIP tools, and to better understand the potential of the security techniques at a larger scale. The security features are implemented in aws instances made to simulate IoT devices, and measured with honeypots and IDSs (Intrusion Detection Systems) that collect data from the received attacks. The research is made as a fundamental work to later be extended by implementing the security features in more devices, such as single board computers that will simulate IoT devies even more accurately.

1 INTRODUCTION

Internet of Things is an exploding field. The Mozilla Foundation estimates the number of IoT devices will be 30 billion in 2020[1]. As they increase in numbers, IoT devices are likely to become an integral part of our everyday lives. The number of innovations and updates come at a striking pace and concepts that recently seemed futuristic are becoming everyday technology. Smart devices, smart homes, and even smart cities, are gaining momentum, such as the recently proclaimed Alphabet powered smart city in Toronto[2]. The increase of IoT devices requires an equal increase of IoT Security. Devices such as door locks, baby monitors and even IIoT (Industrial IoT) present even more pressing security risks (Atzori et al., 2010; Sha et al., 2018).

Unfortunately, new security vulnerabilities are found at a rapid pace. In 2018 almost 1400 new vulnerabilities were reported every month and in 2016 one of the world's largest DDoS (Distributed Denial-of-Service) attacks was executed from an IoT botnet (Kolias et al., 2017). The number of found security breaches are likely to keep growing with the increase of IoT.

Brute-force attacks (systematically trying different login/password combinations to enter a system) are the most common type of attacks towards IoT devices. 87% of all compromised IoT devices are attacked by a brute-force attack targeting the Telnet or SSH password[3]. Furthermore, according to a study by Kaspersky lab, 60% of all attacks targeting IoT devices originate from only five countries: Brazil, China, Japan, Russia, and the US[4].

This article investigates two security features: one to protect them against brute-force attacks (fail2ban) and one to protect them against attacks originating

[a] https://orcid.org/0000-0001-7884-966X

[b] https://orcid.org/0000-0003-3089-3885

[1] https://Internethealthreport.org/2018/ accessed 2019-07-30.

[2] https://www.ft.com/content/bf5949f2-96ab-11e9-8cfb-30c211dcd229 accessed 2019-07-30.

[3] Source: https://www.cvedetails.com/browse-by-date.php accessed 2019-08-02.

[4] Source: Kaspersky lab, https://securelist.com/new-trends-in-the-world-of-iot-threats/87991/ accessed 2019-08-02.

Heiding, F., Omer, M., Wallström, A. and Lagerström, R.
Securing IoT Devices using Geographic and Continuous Login Blocking: A Honeypot Study.
DOI: 10.5220/0008954704240431
In *Proceedings of the 6th International Conference on Information Systems Security and Privacy (ICISSP 2020)*, pages 424-431
ISBN: 978-989-758-399-5; ISSN: 2184-4356

form foreign countries (GeoIP). The security features are implemented in aws instances made to simulate IoT devices, as a first step to analyze the potential of the features. Since the results proved successful to enhance our defences against the given attacks, and the given attacks are the most common towards IoT devices, further work will aim to implement the techniques in single board computers and real IoT devices to continue the work

In total 16 honeypots were launched, eight with the security features and eight without. The honeypots with security features received a large reduction in received attacks. For the GeoIP module, the reduction was up to 93% and for fail2ban the reduction was up to 99%.

2 RELATED WORK

Gupta et al. (2017) created a firewall for IoT devices and A. K. Simpson (2017) looked at simplifying IoT Security for household appliances by creating a home manager . Altolini et al. (2013) researched encryption for IoT devices and S. Raza (2013) created a new way of intrusion detection for IoT systems, focusing on information spoofing attacks.

Watson (2015) presents potentially negative side effects of IoT devices in smart homes, discussing downsides such as isolation and incapacity if a disconnection occurs. He also proposes an increased risk of crime such as burglary as a consequence of monitoring user experiences. However, little focus is made on the prevention of attacks or hardening of IoT devices.

Implementation of honeypots to test IoT Security has been done in some studies. For example by modeling and clustering hostile activities towards IoT machines using the Multivariate Hawkes Process (Sun et al., 2018). Oliveri and Lauria (2018) used honeypots to identify IoT botnets and Sohal et al. (2018) focus on securing IoT devices using both honeypots and a Hidden Markow model to identify harmful devices within the network and Cloud-of-Things environments.

The analyzed articles often fail to address reasonable countermeasures to harden IoT devices, but rather focus on exposing a vulnerability or developing a way to detect vulnerabilities. The purpose of our article is to confirm the potential of security techniques that can protect against the attacks most common towards IoT devices, while at the same time simulating devices with a similar nature as IoT devices, to gain an idea of how these would be hardened by the implementation.

Some research has been done regarding specific networking services. Pa et al. (2015) found a large increase in Telnet attacks and analyzed them using a custom-built honeypot for IoT devices, the IoTPOT. They claim previous honeypots are insufficient at analyzing Telnet based attacks towards IoT devices due to an incapacity of handling different incoming commands. Once again, focus is on capturing the attacks rather than preventing them. Geo-based blocking or rate-limited connections are not used.

Similar studies were made for SSH attacks against IoT devices, for example by Dowling et al. (2017) to assess SSH attacks against ZigBee honeypots and Valli et al. (2013) who assess SSH attacks using Kippo honeypots.

Some studies use or briefly mention the fail2ban software but they do not focus on its implementation towards IoT devices, such as Yu (2016) and Florin B. Manolache (2014).

GeopIP has been investigated or briefly mentioned in some general studies relating to IoT security, for example Shrivastava et al. (2019) investigating attacks towards IoT devices with a Cowrie honeypot. However, the mention of geo-blocking is brief and the article focus on classification of attack types rather than prevention of them.

3 BACKGROUND

In this section we clarify the terms used in the article, such as honeypots, GeoIP, and fail2ban.

3.1 Rate Limiting and Fail2ban

Rate limiting connections restrict the amount of incoming traffic based on a criterion. In this article we use the Failban software to block IP addresses that make three failed login attempts within 24 hours[5]. Rate limiting can increase protection against brute-force attacks that attempts to break a password by systematically generating and trying a large amount of potential combinations (characters, letters and numbers)[6].

3.2 GeoIP and Iptables

Geographic blocking is to exclude traffic from certain geographic areas. In this article we use GeoIP to im-

[5] https://www.fail2ban.org/ accessed 2019-07-30.

[6] Source: kaspersky lab, https://securelist.com/new-trends-in-the-world-of-iot-threats/87991/ accessed 2019-08-02.

plement a geographic blocking where we reject all IP connections from other countries[7].

Iptables is a utility program to set up, maintain and inspect IP filtering rules in the Linux Kernel. Using iptables GeoIP add-on, we can add rules to block IP addresses based on geographic location[8].

3.3 Honeypot

A honeypot is a resource set up to monitor unauthorized activity. When an attacker enters the honeypot, the honeypot collects information about security holes in its own system as well as information about the intruder. Honeypots used in this article are Cowrie, Conpot, p0f, Snort, and Suricata.

The Cowrie honeypot logs all SSH and Telnet connection attempts to the server[9]. Cowrie does not log attacks such as trying to exploit a known vulnerability on the server.

The Conpot honeypot logs all attacks directed at Industrial Control Systems (ICS). ICS is an umbrella term used to describe different control and monitor systems in industrial complexes. Conpot uses common industrial control protocols to emulate large and intricate infrastructures, aimed to convince a hacker they entered a real industrial complex[10]. The Conpot honeypot is interesting since Industrial IoT (IIoT) is gaining momentum. Smart devices are being implemented in both ICS's and other parts of the industrial infrastructure, some of which is critical and require extra security like nuclear and electrical power plants Ungurean et al. (2014).

The p0f honeypot logs all types of connections to the server[11]. p0f will therefore also log what the other honeypots in this paper logs. However, p0f does not log detailed information such as what user name and passwords attackers use to try logging in to the system. Neither does p0f log the type of attack.

Snort and Suricata are similar as they both log events based on specific rules. They are both technically defined as IDSs (Intrusion Detection Systems) rather then Honeypots, meaning they are more oriented to monitor a system for intrusion than trying to lure attackers in and pretending to be a real system. They can be used as honeypots however, as done in this article. The rules used in our paper are the default rules, checking for attackers trying to exploit a wide range of different security vulnerabilities, such as web specific exploits, malware, trojans, DNS and SQL attacks [12][13].

With the data from Snort and Suricata it is possible to see the number of hacking attempts targeted at the machine. One problem with Snort and Suricata is that they only log attacks that match a rule. An attacker using a new type of exploit not listed as one of the rules would not get logged. Because of this Snort and Suricata might report fewer attacks than were received.

4 METHOD

Here we describe how the honeypots with and without security features were launched, the security settings of each honeypot and how they were configured.

4.1 Launching Honeypots without Security Features

We launched eight honeypots of five different types, each type logging different attacks. Table 1 shows which types and how many implementations of each type we used. To launch and manage the honeypots Modern Honey Network was used[14]. The honeypots were active for five days before taken down, more specifically between 2019-04-08 6pm and 2019-04-13 5pm UTC. The honeypots were launched in AWS' data center "US East (N. Virginia)" and each one was assigned a public IPv4 address from that region. The machines were not assigned IPv6 addresses. The firewall was configured to allow connections from everywhere and to all ports. The different types of honeypots used were Cowrie, Conpot, p0f, Snort, and Suricata, all launched on VPS's with a clean Ubuntu 16.04 LTS installation. For all the stated honeypots the allowed ports where derived from the default rulesets, as found in each honeypot's documentation.

4.2 Launching Honeypots with Security Features

Another version of each honeypot listed in table 1 was launched with added security features. GeoIP was set to block all incoming connections from outside

[7] https://www.maxmind.com/en/geoip2-services-and-databases accessed 2019-07-30.

[8] https://linux.die.net/man/8/iptable accessed 2019-07-30.

[9] https://www.Cowrie.org/posts/2015-07-05-Cowrie/ accessed 2019-07-30.

[10] http://Conpot.org accessed 2019-07-30.

[11] http://lcamtuf.coredump.cx/p0f3/ accessed 2019-07-30.

[12] https://www.Snort.org accessed 2019-07-30.

[13] https://suricata-ids.org accessed 2019-07-30.

[14] https://github.com/threatstream/mhn accessed 2019-07-30.

Table 1: Honeypots launched without security features.

Honeypot type	Number
Conpot	2
Cowrie	2
p0f	2
Snort	1
Suricata	1

the United States. The Maxmind GeoLite2 Country database was used to translate IP addresses to country codes[15]. Fail2ban was used to block all IP addresses that unsuccessfully tried to connect via SSH three times within 24 hours[16].

Fail2ban was installed on the second Cowrie honeypot. Cowrie is the only honeypot we launched that monitors SSH login attempts specifically and therefore, fail2ban was only installed on the Cowrie honeypot. Table 2 shows which honeypots were given which security features.

5 RESULTS

In this section, we compare the number of received attacks for the honeypots with and without security features. The different country of origin for the attacks are presented, as well as a difference in the frequency of attacks towards certain honeypots. Lastly, the average attacks per IP address are displayed.

5.1 Honeypots without Security Features

The honeypots without security features received a combined total of 100,274 attacks (between 2019-04-08 6pm and 2019-04-13 5pm UTC). A breakdown of the attacks can be seen in table 3.

Of the three honeypot types that had two instances running (Conpot, p0f and Cowrie), Conpot and p0f received a similar amount of attacks for both instances but the Cowrie honeypots had a slight difference, one instance receiving roughly 25% fewer attacks than the other. One of the Cowrie instances had a similar amount of attacks (20,156) as the p0f instances (19,134 and 18,986) even though Cowrie specifically targets SSH attacks and p0f detects all connections as attacks. The other Cowrie sensor had a significantly higher amount of attacks than all other honeypots.

5.2 Honeypots with Security Features

The honeypots with security features received a combined total of 6,184 attacks (between 2019-04-24 6pm and 2019-04-29 5pm UTC). This means a reduction of attacks by 94 %. A breakdown of the attacks can be seen in table 4.

5.3 Comparison

A direct comparison of the results with and without security features is presented in table 5 and table 6. In table 5 Conpot and p0f display the average received attacks between the types' two instances with and without security features. Cowrie display the average received attacks for the two instances without security features and the specific amount of attacks for the Cowrie honeypot with GeoIP and the Cowrie honeypot with fail2ban. For all honeypots, the received attacks decreased by at least 87%, sometimes up to 99%. Note that the Cowrie honeypot with the enabled GeoIP security feature received significantly fewer attacks than any other honeypot.

We can see that the honeypots with the most attacks per IP address also has the biggest reduction in attacks per IP address.

5.4 Origin of Attacks

As seen in figure 1 most attacks originate from Ireland (23.5%), followed by the Netherlands (14.7%) and Germany (13.8%). In total, these countries contributed to more than 50 % of the attacks. However, most of these attacks were targeting the Cowrie honeypots. If we analyze the origin of attack for all honeypots but the Cowrie the results differ, as seen in figure 2. Excluding Cowrie most attacks come from the US (23.6%) followed by Russia (17.9%) and China (17.6%). In total there were 53,048 attacks against all honeypots except Cowrie. For the industrial system honeypot Conpot more than 50 % of the attacks originated from China, as seen in figure 3. The Conpot honeypots received 6,719 attacks in total.

6 DISCUSSION

6.1 Geographic Blocking

The results show that GeoIP has a potential to heavily reduce the number of received attacks, in our case between 87.0% - 99.8%. p0f received the lowest reduction, 87.0%, and Cowrie received the highest at

[15] https://dev.maxmind.com/geoip/geoip2/geolite2/ accessed 2019-07-30.

[16] https://www.fail2ban.org/ accessed 2019-07-30.

Table 2: Honeypots launched with security features.

Honeypot Type	Number	Security Feature
Conpot	2	GeoIP
Cowrie	1	GeoIP
Cowrie	1	Fail2ban
p0f	2	GeoIP
Snort	1	GeoIP
Suricata	1	GeoIP

Table 3: Attacks against honeypots without security features.

Honeypot	Attacks	Attacks / hour
Conpot-1	3,073	25.8
Conpot-2	3,509	29.5
Cowrie-1	27,060	227.4
Cowrie-2	20,156	169.4
p0f-1	19,134	160.8
p0f-2	18,986	159.5
Snort-1	3,905	32.8
suricata-1	4,235	35.6

Table 4: Attacks against honeypots with security features.

Honeypot	Security Feature	Attacks	Attacks / hour
Conpot-1	GeoIP	269	2.3
Conpot-2	GeoIP	152	1.3
Cowrie-1	GeoIP	36	0.3
Cowrie-2	fail2ban	189	1.6
p0f-1	GeoIP	2,301	19.3
p0f-2	GeoIP	2,647	22.3
Snort-1	GeoIP	301	2.5
suricata-1	GeoIP	289	2.4

Table 5: Comparison of attacks for honeypots with and without security features.

Honeypot	Attacks without Security Features	Attacks with Security Features	Attack Reduction
Conpot (GeoIP)	3,291	210	93.6%
Cowrie (GeoIP)	23,608	36	99.8%
Cowrie (fail2ban)	23,608	189	99.2%
p0f (GeoIP)	19,060	2,474	87.0%
Snort (GeoIP)	3,905	301	92.2%
suricata (GeoIP)	4,235	289	93.2%

Table 6: Average number of attacks per IP against each honeypot type.

Honeypot	Avg number of attacks per IP (without)	Avg number of attacks per IP (with)
Conpot	17.4	12.8
Cowrie (GeoIP)	101.1	1.7
Cowrie (fail2ban)	101.1	4.6
p0f	6.5	6.5
Snort	6.5	1.8
suricata	7.3	1.9

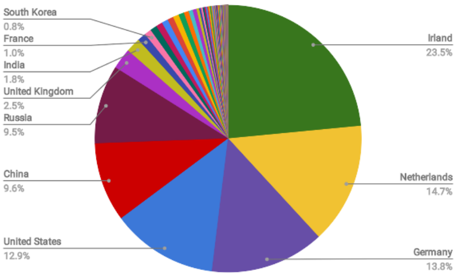

Figure 1: Origin of attacks, all honeypots.

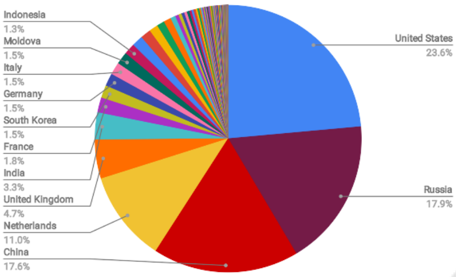

Figure 2: Origin of attacks, all honeypots but Cowrie.

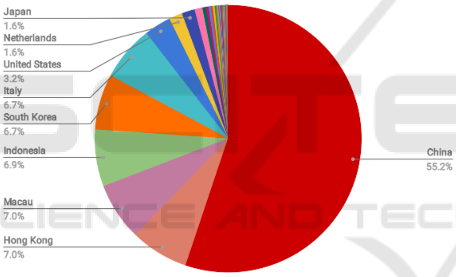

Figure 3: Origin of attacks targeting the Conpot honeypots.

99.8%. This makes sense since p0f logs all incoming connections to the machine and all incoming connections are not hostile. Some could for example be search engines trying to index the Internet. However, note that web search engines such as Google and Bing would not crawl our honeypots since our server does not host a web server, neither are the IP's to our honeypots listed anywhere on the web.

Geographic blocking is not a novel technique in itself, but the interesting thing to note is that it currently has a very limited implementation in IoT devices. Since the results of its implementation in this study where highly successful, and its application in IoT devices makes sense as stated below, it is a good indicator that more work should be continued in this area.

The implementation of geographic blocking makes sense for IoT devices since they often work inside a specific context, with the end-user nearby. For example, smart lights, smart TVs, and smart refrigerators are all located within the home of a user. Most of the users will be relatively close to, or at least in the same country as these devices. When the user is abroad, a VPN (Virtual Private Network) can be used to access the devices when needed Shrivastava et al. (2019).

It is also worth to mention, once more, that geographic blocking should be used as a complementary security feature, not the sole means of protection. A VPN can be used to spoof an attackers location from a foreign country to bypass the geo-blocking, making it vulnerable to specific, targeted attacks and since the geographical blocking allows connection from the same country (or county, city etc.), we must implement security features to protect against attacks from within.

6.2 Fail2ban and Cowrie

The average attack per IP address for Cowrie was notably higher than for any other honeypot. This makes sense since Cowrie logs SSH brute-force attacks where the attacker repeatedly tries to log in to a device, resulting in a large number of connection attempts. The fail2ban rules blocking IP addresses after three unsuccessful login attempts within 24-hours was successful. It reduced attacks by 99.2% and reduced the average attacks per IP address from 101.1 to 4.6. The average number of attacks after the reduction, 4.6, is higher than the ban threshold of three. This can happen since there is a delay between an SSH connection is made and the actual ban happens. This delay can allow some IP addresses to make more than three login-attempts before the ban takes place.

6.3 Whitelisting

Given the high reduction in attacks obtained by implementing the security mechanisms, we could also consider a complement to fail2ban and GeoIP, namely to whitelist allowed IP addresses and strictly ban all others. This could make sense for IoT devices since they are often used by a limited amount of IP addresses. The downside of this approach is that it requires the user to fill in the allowed IP addresses manually, which may be tedious for a non technical user. But given a good graphical interface it could be a solid option for increased security.

6.4 Method

The honeypots were deployed using VPSs hosted by AWS. The IP-ranges for AWS is public information easily accessible on their website. This could theoretically lead to flawed results compared to an IoT device within a hidden IP address. However, in practice the results should be valid since the devices were assigned IPv4 addresses and all IPv4 addresses can be pinged within 10 hours at a 50Mbit/sec bandwidth (there are 2^{32} or roughly 4.29 billion IPv4 addresses) [17]. So even if an IPv4 address is not public it can get detected and attacked by automated bots. Furthermore, both the honeypots with and without security features were deployed in the same manner meaning the percentage attack reduction should still be accurate.

In this paper, it is assumed that the AWS VPSs are equivalent to IoT devices. By using our earlier definition IoT-devices are physical devices connected to the Internet. The used security features did not try to protect against hardware attacks but focused on attacks done via Internet. Therefore VPSs connected to Internet would face similar attacks. The operating system used on the VPSs was Ubuntu 16.04 which is common among IoT-devices together with other, similar Linux based operating systems.

6.5 Origin of Attacks

Our results aligned with the Kaspersky study in the sense that a majority of attacks came from a small number of countries. However, the most significant countries of origin differed from our study and that of Kaspersky. In our case, most attacks (74.5%) came from Ireland, the Netherlands, Germany, the US, and China, while the study by Kaspersky found most attacks (60%) came from Brazil, China, Japan, Russia, and the US. When looking at the origin of attacks for all our honeypots but Cowrie, a majority of attacks (74.8%) came from the US, Russia, China, the Netherlands, and the UK. The origin of attacks is not the focus point of this article, more than to note that a large number of attacks originate from a small number of countries. However, the specific countries of origin for the attackers is an interesting topic and could prove good material for future research, as discussed below.

6.6 Implementation

An important point of discussion is how the security mechanisms should be implemented to the IoT devices in practice. Two reasonable methods can be suggested, a central governing entity that regulate security implementation of all IoT devices, or the companies developing IoT hardware/software are responsible for implementing the security themselves. For the second alternative to be efficient, customers of IoT devices must prioritize security enough to pay an increased price for the products, since the added features are likely to come with a cost. This may speak in favor of the first alternative, to have a central governing entity, either recommending or demanding that certain security standards are met. A solid compromise could be to have a central entity conducting a list of best practices and recommendations that the manufacturers are encouraged but not enforced to follow. An example could be a scoring system that gives a higher rank if security principles are implemented, but needless to say this area is outside the scope of this article and will hopefully be treated in a future article of its own.

6.7 Future Work

Given the findings in this article we intend to continue the work by implementing the security features in a number of single board computers and actual IoT devices.

We also encourage further research to validate the findings and launch a larger set of honeypots running for a longer time. It would also be relevant to do more detailed research regarding the various areas of the honeypots, such as an in depth study of smart devices in industrial control systems, like those simulated by Conpot. A more focused study of implementing geographic blocking and rate-limited connections for IIoT devices in general would also be relevant.

In addition, it would be good to make an in depth study of the country of origin for the attackers. Future research could launch a larger number of honeypot at different times to distinguish the composition of attack origins at different dates, and/or for honeypots hosted at different locations. It would be interesting to find a pattern or better understanding of this, since there was a difference between our results and that of the Kaspersky study.

As stated in the previous section, it would also be relevant to research the best way of ensuring that security features are implemented in IoT devices. Comparing governmental legislations with company responsibilities and other methods.

In our long term research we aim to create a threat modeling and attack simulation language for IoT infrastructures (Johnson et al., 2018). This type of empirical work can feed the probability engine of such a language (Ekstedt et al., 2015).

[17] https://www.securityartwork.es/2013/01/21/how-much-does-it-take-to-ping-the-whole-Internet-12/ accessed 2019-08-06.

7 CONCLUSIONS

It seems like both geographic blocking and rate-limiting connections can increase the security of IoT devices significantly, but more work is needed to verify our findings and tailor the results towards IoT devices. In our tests, GeoIP lowered the probability of getting attacked by roughly 90% and fail2ban by up to 99%.

ACKNOWLEDGEMENTS

This work has received funding from the Swedish Centre for Smart Grids and Energy Storage (SweGRIDS).

REFERENCES

A. K. Simpson, F. Roesner, T. K. (2017). Securing vulnerable home iot devices with an in-hub security manager. In *Pervasive Computing and Communications Workshops (PerCom Workshops), 2017 IEEE International Conference on*, pages 551–556. IEEE.

Altolini, D., Lakkundi, V., Bui, N., Tapparello, C., and Rossi, M. (2013). Low power link layer security for iot: Implementation and performance analysis. In *2013 9th International Wireless Communications and Mobile Computing Conference (IWCMC)*, pages 919–925.

Atzori, L., Iera, A., and Morabito, G. (2010). The internet of things: A survey. *Computer Networks*, 54(15):2787 – 2805.

Dowling, S., Schukat, M., and Melvin, H. (2017). A zigbee honeypot to assess iot cyberattack behaviour. In *2017 28th Irish Signals and Systems Conference (ISSC)*, pages 1–6.

Ekstedt, M., Johnson, P., Lagerström, R., Gorton, D., Nydrén, J., and Shahzad, K. (2015). Securi cad by foreseeti: A cad tool for enterprise cyber security management. In *2015 IEEE 19th International Enterprise Distributed Object Computing Workshop*, pages 152–155. IEEE.

Florin B. Manolache, Q. Hou, O. R. (2014). Analysis and prevention of network password guessing attacks in an enterprise environment. In *2014 RoEduNet Conference 13th Edition: Networking in Education and Research Joint Event RENAM 8th Conference*, pages 1–7. IEEE.

Gupta, N., Naik, V., and Sengupta, S. (2017). A firewall for internet of things. In *2017 9th International Conference on Communication Systems and Networks (COMSNETS)*, pages 411–412.

Johnson, P., Lagerström, R., and Ekstedt, M. (2018). A meta language for threat modeling and attack simulations. In *Proceedings of the 13th International Conference on Availability, Reliability and Security*, page 38. ACM.

Kolias, C., Kambourakis, G., Stavrou, A., and Voas, J. (2017). Ddos in the iot: Mirai and other botnets. *Computer*, 50(7):80–84.

Oliveri, A. and Lauria, F. (2018). Sagishi: an undercover software agent for infiltrating iot botnets. *Network Security*, 1:9–14.

Pa, Y. M. P., Suzuki, S., Yoshioka, K., Matsumoto, T., Kasama, T., and Rossow, C. (2015). Iotpot: Analysing the rise of iot compromises.

S. Raza, L. Wallgren, T. V. (2013). Svelte: Real-time intrusion detection in the internet of things. *Ad hoc networks*, 11(8):2661–2674.

Sha, L., F. Xiao, F., Chen, W., and Sun, J. (2018). Iiot-sidefender: Detecting and defense against the sensitive information leakage in industry iot. *World Wide Web*, 21(1):59–88.

Shrivastava, R. K., Bazila, B., and Hota, C. (2019). Attack detection and forensics using honeypot in iot environment. In Fahrnberger, G., Gopinathan, S., and Parida, L., editors, *Distributed Computing and Internet Technology*, pages 402–409, Cham. Springer International Publishing.

Sohal, A. S., Sandhu, R., Sood, S. K., and Chang, V. (2018). A cybersecurity framework to identify malicious edge device in fog computing and cloud-of-things environments. *Computers Security*, 74:340–354.

Sun, P., Li, J., Bhuiyan, Z. A., Wang, L., and Li, B. (2018). Modeling and clustering attacker activities in iot. *Computer*, 479:Pages 456–471.

Ungurean, I., Gaitan, N., and Gaitan, V. G. (2014). An iot architecture for things from industrial environment. In *2014 10th International Conference on Communications (COMM)*, pages 1–4. IEEE.

Valli, C., Rabadia, P., and Woodward, A. (2013). A zigbee honeypot to assess iot cyberattack behaviour. In *Patterns and patter - An investigation into SSH activity using Kippo Honeypots*.

Watson, D. L. (2015). Some security perils of smart living. *International Conference on Global Security, Safety, and Sustainability*, 534:Pages 211–227.

Yu, J. (2016). An empirical study of denial of service (dos) against voip. In *2016 15th International Conference on Ubiquitous Computing and Communications and 2016 International Symposium on Cyberspace and Security (IUCC-CSS)*, pages 54–60. IEEE.