Superspreaders: Quantifying the Role of IoT Manufacturers in Device Infections

Elsa Rodríguez[1] ✉, Arman Noroozian[1,2], Michel van Eeten[1], *and* Carlos Gañán[1]
(1) Delft University of Technology, (2) University of Amsterdam
{e.r.turciosrodriguez, a.noroozian, m.j.g.vaneeten, c.hernandezganan}@tudelft.nl

Abstract

The influx of insecure IoT devices into the consumer market can only be stemmed if manufacturers adopt more secure practices. It is unlikely that this will happen without government involvement. Developing effective regulation takes years. In the meantime, governments have an urgent need to engage manufacturers directly to stop the damage from getting worse. The problem is that there are many thousands of companies that produce IoT devices. Where to start? In this paper, we focus on identifying the most urgent class: the manufacturers of IoT devices that get compromised in the wild. To identify the manufacturers of infected IoT, we conducted active scanning of Mirai-infected devices. Over a period of 2 months, we collected Web-UI images and banners to identify device types and manufacturers. We identified 31,950 infected IoT devices in 68 countries produced by 70 unique manufacturers. We found that 9 vendors share almost 50% of the infections. This pattern is remarkably consistent across countries, notwithstanding the enormous variety of devices across markets. In terms of supporting customers, 53% of the 70 identified manufacturers offer firmware or software downloads on their websites, 43% provide some password changing procedure, and 26% of the manufacturers offer some advice to protect devices from attacks. Our findings suggest that targeting a small number of manufacturers can have a major impact on overall IoT security and that governments can join forces in these efforts, as they are often confronted with the same manufacturers.

1 Introduction

Insecure Internet-of-Things (IoT) devices are still flooding the market, even though the damage that these devices can cause has been evident for years. In response, many governments have issued baseline security recommendations and guidelines for security by design for IoT [1–4]. While useful, such guidelines do not address the underlying root cause: many manufacturers lack the incentives to adequately secure their devices. Similarly, engaging with certain other actors, for instance ISPs who are in a position to mitigate part of this problem on a short-term basis, still leaves much to be addressed [5]. A consensus is emerging that governmental interventions are required to overcome the incentive problem [6]. While governments are debating long-term solutions regulatory strategies like apportioning liability and setting minimum security standards, and some liability frameworks are being proposed [7], there is a short-term need to engage manufacturers to reduce the current influx of insecure devices. To illustrate: in 2016, the U.S. Federal Trade Commission (FTC) lodged a complaint against ASUS because the company "failed to take reasonable steps to secure the software on its routers". Through a consent order, the FTC got ASUS to "establish and maintain a comprehensive security program subject to independent audits for the next 20 years" [8].

For governments, the process to engage manufacturers directly is resource intensive and can only be applied to a limited set. Where to start? The question of which manufacturers to engage is complicated by the enormous complexity of the IoT ecosystem. There are markets around many different product types, each with different populations of manufacturers. Kumar *et al.* [9] found a long tail of 14,000 companies, though just 100 of them were responsible for 90% of devices in their observations. There are geographical factors at play also. Product types and manufacturers will vary across different countries and continents. Last, but not least, governments lack reliable data on the security practices of these manufacturers.

As part of a collaboration with the Dutch government, this paper presents an empirical approach to identify the priority targets for governmental intervention: the manufacturers of IoT devices that get compromised in the wild. For those manufacturers, the evidence for the lack of adequate security of their devices is compelling, as is the fact that this lack is causing harm. To identify device types and manufacturers, we build on recent advances in large-scale device discovery and identification. As a basis for governmental action, though, these studies have certain drawbacks that we need to overcome. Some studies rely on privileged access to internal

network data from home [9, 10] or ISP networks [11, 12]. For our purpose, this approach would create selection bias towards the manufacturers in the limited set of networks where such access could be obtained. Other studies use Internet-wide scans, typically focused on developing scalable methods for identification [13]. For scanning, most studies use device fingerprints that were developed for a specific set of devices that the researchers had access to. In other words, these scans can only detect a sample of devices that the researchers knew about and could test beforehand. It is unknown how these samples relate to the population of devices in the wild. This means that all other IoT devices are simply out of scope. In our case, however, we need to identify manufacturers for a specific population of compromised devices in the wild, not for a set of devices that was predefined. The population in the wild will at best partially overlap with those discovered in the large fingerprint-based studies. Finally, most studies focus on IoT devices in general, not on compromised devices. The few exceptions have serious limitations; one is based on an observation period of a single day [14], one only identifies high-level device categories [15], and one relies on third parties such as Shodan [16] to identify manufacturers [17].

Our approach starts with two months of real-time observations of the IP addresses of compromised devices that were scanning a /16 darknet with the Mirai fingerprint. As most compromised IoT resides in consumer networks [18], we focused our analysis on devices in 355 ISP networks that together have the bulk of the market share in 68 countries. Each real-time observation was immediately followed up with an active scan of that IP address that collected banners and Web-UI pages for the device. We then manually labelled the unique device fingerprints in an attempt to identify as many manufacturers and devices as possible. We opted for manual labelling because our goal is to provide data that is as accurate, explainable and complete as possible, since it will provide the basis for regulatory interventions. The goal was not to improve on existing scalable identification techniques. Our approach was able to identify 31,950 compromised devices attributed to 70 manufacturers. We aim to answer these questions: (*i*) Which manufacturers are associated with compromised IoT across 68 countries? (*ii*) How variable is the set of manufacturers across different countries? (*iii*) What are these manufacturers doing to remediate the insecurity of their devices? In sum, we make the following contributions:

- We present the first systematic analysis of which manufacturers share attributed infections for infected IoT devices in 68 countries.
- We develop a transparent and reproducible approach to identify manufacturers of infected devices that can be applied across jurisdictions and that does not rely on privileged access to network data.
- Our results demonstrate a strong pattern of concentration: while we find 70 manufacturers in total, just 9 of them share around 50% of all infections. This pattern is quite consistent across multiple jurisdictions, thus supporting international regulatory collaboration in engaging these manufacturers.
- Notwithstanding the variety across markets, geographical areas and legal frameworks, the set of manufacturers associated with infected devices is remarkably consistent across countries. The manufacturers related to around half of the share attributed to infections were present in at least 47 (69%) of the 68 countries.
- We analyze what, if any, firmware or software was provided to download by the manufacturer, and we found that out of the 70 manufactures 53% offer firmware or software to download on their websites. We checked if the manufacturers provide any password changing procedure, and 43% of them do. Finally, we checked whether or not there was some advice to protect the devices from attacks, and 26% of the manufacturers offer advice to protect the devices.

2 Context

IoT manufacturers continue to bring devices into the market at an incredible pace [19]. Many governments want to unleash the potential of this technology—e.g., the European Union (EU) highlighted IoT in its vision of the digital single market [20].

In light of the security issues associated with IoT, the engineering community keeps working on defining security standards. In 2019, the IETF published `RFC8520` (Manufacturer Usage Description (MUD) [21]) aiming at providing a white list of their devices' traffic so third parties such as ISPs could identify anomalous traffic flows that do not match the MUD profile [22]. Governments have also acknowledged the need to intervene and define common guidelines to secure IoT devices.

On the side of governments, various countries are trying to change the behavior of firms in the IoT markets. The UK government released guidelines of what they consider a secure IoT product [4]. At the European level, the EU Agency for Cybersecurity (ENISA) has released good practices for secure IoT software development [3]. In the United States, the National Telecommunications and Information Administration (NTIA) created a bill to increase the transparency of the whole supply chain of IoT devices by encouraging the "Software Bill of Materials" (SBOM) [23]. In addition, the National Institute of Standards and Technology (NIST) created a inter-agency report (NISTIR 8259) to help manufacturers incorporate security into their IoT devices. NISTIR 8259 renders guidance on how manufacturers could provide post-sale security of IoT devices and on how to communicate security to customers [1]. NISTIR 8259A [2] asserts a baseline of security that an IoT device needs to provide through technical

means. These government efforts could, in the long run, result in more secure IoT devices. Similar to the E.U. General Data Protection Regulation, which is increasingly considered the default global standard for privacy [24], these IoT security policies might get manufacturers to follow them in all the countries where they have presence, rather than differentiate devices per jurisdiction.

In the Netherlands, there are discussions about an update-obligation law for 2021, which would make sellers of IoT devices responsible for supplying updates, rather than directly imposing this obligation on manufacturers [25]. Ahead of European and national regulation, the Dutch government—more precisely: the Ministry of Economic Affairs and Climate—wants to start conversations with manufacturers of poorly-secured devices to improve IoT security [26]. Our study is conducted in collaboration with the ministry and meant to provide the basis for the selection of manufacturers that the government will engage with.

3 Ethical considerations

To answer our research questions, we deployed active scanning of IP addresses where we detected a device infected with Mirai malware. Since active scanning has ethical implications, especially when conducted in consumer broadband networks, we got the approval of the board of ethics of our university to start with this research (Application #993). Our Data Protection Impact Assessment (DPIA) and a Data Management Plan (DMP) were also reviewed and approved. we briefly discuss relevant ethical considerations organized around the principles laid out in the Menlo Report [27].

First, *Respect for Persons*. Since we cannot identify the owners of the devices located at the IP address that we scan, let alone being able to contact them, we cannot get their prior consent. On the IP address of the server conducting the active scans, we set up a web page with information about the project and an opt-out mechanism. We received four opt-out requests during the scanning period, and we removed these IP addresses from our dataset and from further scans.

Second, *Beneficence*. An unintended harm is that in a rare number of cases the screenshots captured from the scans would contain sensitive data, such as a customized NAS access login page that contained a personal picture. The data of all scans were stored on a secure server with access limited to the researcher team. The raw data was removed after the analysis was completed. The benefit of this research to the owners of the devices that we scanned back is that our findings regarding the manufacturers are part of a governmental project that aim to get manufacturers to better support these – and all other – users with insecure devices, a longer-term benefit that is underlined by the presence of a compromised device on the home network of the users involved in the scans.

Third, *Justice*. The selection of IP addresses to be scanned was driven completely by the observation of Mirai scanning traffic originating from these addresses towards the darknet. Within this set, additional selection was made by focusing on IP addresses from broadband consumer networks. This process does not bias against specific user groups within the consumer population.

Fourth, *Respect for Law and Public Interest*. This study is co-funded by the central government and designed in partnership with them. It is part of the government 'Roadmap for Secure Hardware and Software' [26].

4 Methodology

To identify which manufacturers share attributed infections of the bulk of the compromised IoT devices in each country, we setup a (near) real-time data collection pipeline to gather information on infected IoT devices observed in the wild. This pipeline ran for a period of two months (July to September 2020). Subsequent steps were executed to process the pipeline data and arrive at a labeled data set of compromised devices and their manufactures. Figure 1 illustrates a high-level overview of our methodology. Steps 1 and 2 capture the real-time data collection pipeline. Steps 3 and 4 consist of the subsequent data processing and labeling components which were executed offline at a deferred time. Below, we provide more details on each step.

Data collection

To collect data on infected IoT devices, we implemented a data pipeline tied to a `/16` darknet through which we gather real-time observations on IPv4 addresses scanning the darknet with a Mirai malware fingerprint. We match all incoming darknet packets against the fingerprint developed in prior work on detecting Mirai [15] to filter and extract Mirai scan traffic from our darknet. All matching packets are buffered over 1 minute intervals and stored as PCAP network packet capture files which are queued for further processing in our pipeline.

Next – in step 2 of the pipeline – we extract source IP addresses from the queued PCAP files, and scan back all source IPs in (near) real-time to gather additional information on each entry. The data gathered here includes the set of responsive TCP ports at each IP, protocol banners for a set of pre-selected TCP services common to IoT devices (FTP, Telnet, SSH, HTTP(s), SSL/TLS), as well as screenshots of Web-UI content if publicly reachable through any of the exposed ports. This additional data helps us determine whether we are scanning back and potentially talking to a single device or multiple devices, and is simultaneously used to identify the IoT device(s) behind each IP and their manufacturers in later steps.

To gather this information, we first use `Masscan` [28] – a highly scalable TCP port scanner – to detect all open and responsive ports on the IP addresses in our data pipeline. We

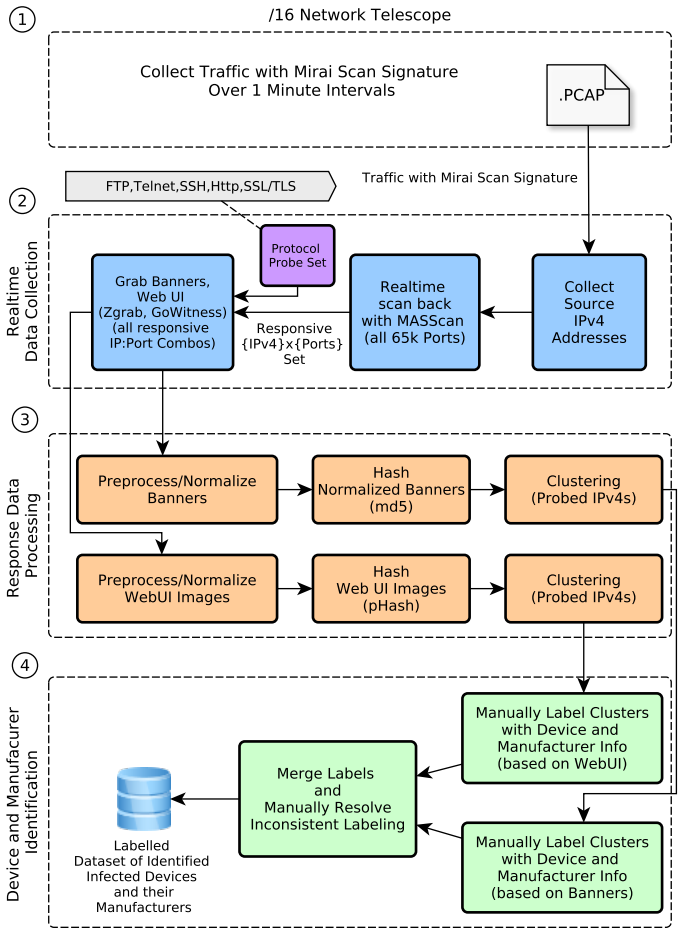

Figure 1: Data collection and processing pipeline

then feed its output to zgrab [29] to collect protocol banners for the previously mentioned set of common TCP services. We also feed the Masscan output to custom scripts to rapidly detect HTTP content on any responsive port whose output is then, in turn, used to collect screenshots of Web-UI content using `gowitenss` [30] a scalable Web-UI content collector implemented in the Go language. Note that we also probe for services on non-standard ports. Prior work has already demonstrated that the number of services running on non-standard ports are far more substantial than commonly assumed [31].

As a result of the large number of possible non-standard port and service combinations that need to be scanned, our pipeline has been tuned and highly optimized for collecting data expeditiously with all non-critical processing (and analysis) of the collected data deferred to subsequent steps. To further complicate matters, it is also crucial to maintain a near real-time scan back throughput in our pipeline due to potential IP churn. As IP addresses churn over time, the correspondences between IPv4 addresses and the IoT devices associated with each IP address will also change. An IP that previously corresponded to a network camera, now points to a home router for instance.

To maintain the necessary high scan back throughput we have implemented two main optimizations within the pipeline: First, we designed our pipeline to avoid scanning back an IP address that has already been scanned within the past 24 hours. We are assuming here that most IP addresses will churn at a rate slower than 24 hours. A secondary reason for this optimization is ethical as we want to avoid directing unnecessary scan traffic to IP addresses and devices that have already been scanned recently. Our second optimization is due to observing a handful of IP addresses in our data pipeline that had all 65k ports exposed. We suspect these IP addresses to have pointed at improperly configured honeypots rather than actual infected IoT devices. We also observed a handful of IP addresses with an unusually high number of exposed ports. As a result, we optimized our pipeline to only grab banners and screenshots from the standard ports "21", "23", "80", "8080", "8081", "443" in combination with “FTP”, “telnet”, “HTTP”, and/or “SSH” services when running into any IP address with more than 1,000 exposed ports after having scanned them via Masscan.

Note that a limitation of our approach – as well as all prior studies that employ comparable techniques to detect infected IoT devices from outside networks – is that an IP address does not have a one-to-one correspondence with a uniquely identifiable IoT device. With respect to the cardinality of the correspondence two corner case scenarios are possible in our case: (**i**) that multiple Mirai infected devices appear as having a single IP address due to Network Address Translation (for instance when multiple infected devices are sitting behind a router) (**ii**) that the infected IoT device is itself a router hosting an arbitrary number of other clean or infected IoT devices behind its NAT. With respect to these cases, we have adopted the following procedure: if the only device that we see accessible through our collected scan back data was a router and that router is known to be vulnerable to Mirai infection vectors, then we consider the router as the infected device. On the other hand if multiple devices have been detected, as long as they have known vulnerabilities to Mirai, all are considered infected.

In total, we scanned back 4,873,430 IP addresses using our pipeline. From this set, we selected the subset located in broadband ISPs for analysis. For these networks, we can have the highest confidence that the devices that scanned the darknet are actual consumer IoT devices, rather than scanners or other systems. Prior work also found that the overwhelming majority of compromised devices are located in broadband ISPs [18]. We used a reliable dataset of the Autonomous Systems (ASes) operated by broadband providers in 68 countries, developed in prior work [5, 32–34], to filter and select the aforementioned subset.

Selecting for IP addresses in these networks, we had a set of 61,154 unique IP addresses. After removing results that consisted only of errors, such as 404, 401, we had a dataset

of banners and Web-UI screenshots for 59,657 unique IP addresses.

For some devices, we could collect only banner data, but no screenshots of Web-UIs. For others, it was the reverse.

Processing and labelling

With the data obtained from the pipeline, we later performed two processes in parallel: (i) classification of banners, and (ii) classification of Web-UI images as depicted in steps 3 and 4 in Figure 1.

We first normalized the collected Web-UI images and banners. Similar to [35], we created heuristics with regular expressions to replace values such as date and time information, content-length field of HTTP response until the banners of the same devices were aggregated. We then hashed the images and banners. In the case of the images, we used perceptual hashing [36] to allow for minor variations in the Web-UIs. The banners were hashed via the MD5 algorithm. We then clustered the results based on the hash values.

Finally, we manually labeled the resulting unique 3,547 Web-UI image hashes and 566 banner hashes in our dataset and applied the labels to our clusters of data. These labels included manufacturer and device type. We identified manufacturers and device types based on the logos of the Web-UI and the text present in the banner. Hence, this labeling approach does not include original equipment manufacturer (OEM) because we observe the name of the brand that is marketing the devices. Regarding the device type, this approach sometimes allows to determine the category of the device (i.e. IP camera) and in some cases, it allows to get the specific model of the device. Finally, we resolved inconsistencies between the banners and Web-UI labels and we obtained our final labeled data set.

For about half of the devices, the collected responses did not contain any information from which we were able to identify the manufacturer. These results were labelled as 'unknown'. From the 59,657 IP addresses, we managed to apply an informative label for 31,231 (52.3%) of them, corresponding to 31,950 devices. In total, we labelled devices for 70 unique manufacturers. As shown in Figure 2, we could identify 49 manufacturers via the banner data and 21 via the Web-UI images. There was an overlap of only 13 manufacturers. This underlines that any identification method would need to combine various types of data.

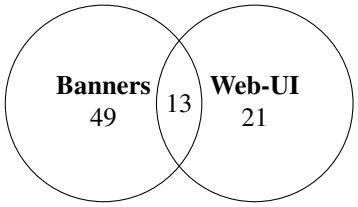

Figure 2: Number of unique manufacturers identified per data type

Data Collection on Firmware available and Manufacturer Security Advice

After we compiled a set of manufacturers and devices during steps 1-4, we also investigated what remediation options or security advice was being offered by the manufacturers. More specifically, we collected data for three categorical variables: (i) whether or not there was a software or firmware to download for the device model or device category; (ii) whether or not there was information provided on how to change the password for the device model or device category; and (iii) whether or not there was any security related information to protect the device model or device category from attacks. We followed an approach similar to [37], where researchers analyzed how security features and advice were presented to users in the manuals and support pages for 220 IoT devices. First, we identified the manufacturer's website. Since our approach sometimes allows to determine the category of the device (e.g. IP camera) and in some cases, it allow us to determine the model of the device (e.g. RT-AC5300), to accomplish this, we used Google's search engine with the following terms: "Device category" AND "Manufacturer" (e.g. IP camera Avtech) or "Device model" AND "Manufacturer" (e.g. RT-AC5300 ASUS). From the Google results we identified the manufacturer's website, which typically contains the manufacturer name in the domain name. Next, within the website, we manually inspect for "Device category" AND "manual" or "guide" or "quick start" or "Device model" AND "manual" or "guide" or "quick start" (depending on whether we had obtained the device category or the model) to check if the device model or category of the device had a user manual available. In cases where the search in the manufacturer's website was not fruitful, we used Google's search engine with the following terms to find the manuals: "Device category" AND "manual" or "guide" or "quick start" AND "Manufacturer" (e.g. IP camera manual Avtech) or "Device model" AND "manual" or "guide" or "quick start" AND "Manufacturer" (e.g. RT-AC5300 manual ASUS), depending on whether we had obtained the device model or the device category. In cases, where we had only the "Device category" (e.g. "IP camera"), we picked one random device of the category.

Next, we manually inspected for the "Device category" or "Device model", and we checked if there was a firmware (FW) or software (SW) to download available in the website. In the manual, we checked if there was any information on how to enable automatic "firmware upgrade" or "firmware update". The outcome was coded as yes or no depending on whether or not we found any FW/SW to download available either in the website or if we found any way to do automatic firmware upgrade or update in the manual. Next, within the documentation related to the "Device category" or "Device model" on the website or in the manual, we searched for the word 'password" to find whether the material contained any password change procedure for the user of the device. The

outcome was coded as yes or no depending on if a password procedure was found or not. Finally, we searched for 'security' as a keyword to inspect whether there was any information related to how to protect the device from attacks or make it more secure. The outcome once again was coded as yes or no.

To code our data, two researchers independently visited each manufacturer website and the manuals. Once they coded the three outcomes, they resolved inconsistencies by double checking the website and the manuals together. Figure 3 summarizes the method to check manuals and websites.

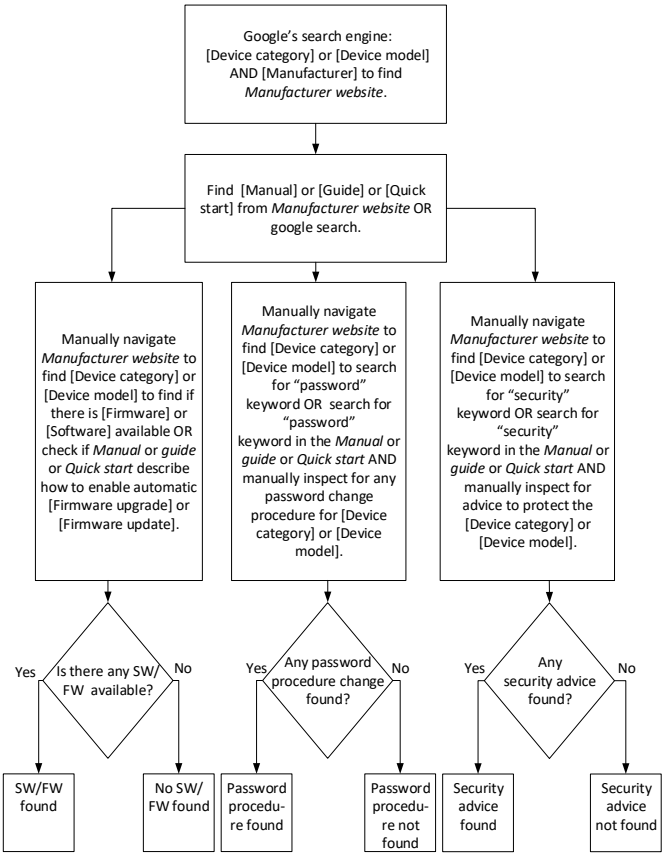

Figure 3: Method to check manuals and websites

5 Findings

Our final dataset contains data on infected devices located in 68 countries and attributed to 70 unique manufacturers—or labelled as 'unknown', where we could not identify the manufacturer. The number of devices seen in each country is highly variable. Table 1 depicts data for the top 20 countries with most infected devices.

Similar to [14] and [15], we find countries such as China, Vietnam, Brazil and the United States leading the number of infections. This suggests that numnber of infections is correlated to the number of broadband connections. This makes intuitive sense: more broadband subscribers means more devices connected to their networks, thus a higher risk of infections. To get a sense of the relative size of the number of devices in relation to the number of broadband subscribers, we have also included those statistics. We used Teleography data [38] of the first quarter of 2018 to calculate the total number of subscribers of the Internet Service Providers in each country. The last column contains the number of infected devices per 100,000 subscribers. There we see that countries with a large consumer broadband base have lower infection rates compared to many smaller countries.

5.1 Manufacturers

Which manufacturers are responsible for the largest share of IoT infections in each country? Figure 4 shows that around 42% of the infections can be attributed to just nine manufacturers. Around 9% is attributed to all 61 other manufacturers combined ('Others (61 Manuf)'). We decided to group the 61 manufacturers because they were a long tail with a share lower than 2% of the infections. The remainder consists of devices we could not attribute ('unknown').

There is a significant percentage of unknown manufacturers in our data. Could this potentially change the pattern of concentration that we found? In other words, could other major manufacturers be present in that set of unknown devices? To explore this issue, we looked at the frequency of the hashes. If hashes are seen only rarely, then it is very unlikely that these devices—and thus their manufacturers—make up a significant share of the population. (They could, of course, belong to one of manufacturers that we already identified. This would not change the overall picture, though, because of the small numbers involved). A high frequency for specific hashes, on the other hand, could point to the presence of a large share for a manufacturer.

We found that 57 unique hashes corresponding to banners and mainly for that part of the data we could not obtain Web-UI that provides information to allow us labeling the manufacturer either. We had 4 hashes corresponding to Web-UI. We plotted the cumulative probability of the number of hashes and less than 5 hashes have around 85% probability of showing more often (see Figure 5).

Some banners provided information about the device, mainly "DVR" and "NAS", but no manufacturer information. There were 17 (30%) hashes that correspond to DVRs and 6 (10%) to NASes. For instance, we got responses like `220 NAS FTP server ready`. We are confident this is an IoT device, but it is not possible to determine with this information the manufacturer name. The rest was a long tail that included FTP servers and Bftpd servers probably used by NASes, set-top boxes, and routers, however, it was not possible to determine the manufacturer either. We tried checking with different

Table 1: Average number of infected IoT devices seen per day for the top 20 countries (July–September 2020)

Country	**Infected Devices** (unique daily)	**Subscribers**	**Infected devices** (per 100k subs.)
Vietnam	10856	11959829	91
Taiwan	6627	4417500	150
China	6363	352767000	2
Rusia	2573	24125823	11
Brazil	2388	23529853	10
Indonesia	2240	7100350	32
Thailand	2183	8463797	26
United States	2136	94085580	2
Korea	1592	19073673	8
Turkey	1442	12159767	12
Mexico	1247	17432549	7
Italy	1227	16201874	7
Malaysia	1209	2492325	49
Iran	1092	10230000	11
Greece	1035	3912680	26
Egypt	1006	5133000	20
Romania	962	4513930	21
Germany	900	30868800	3
France	855	27292191	3
India	598	16410909	4
Others	13292	226277100	4

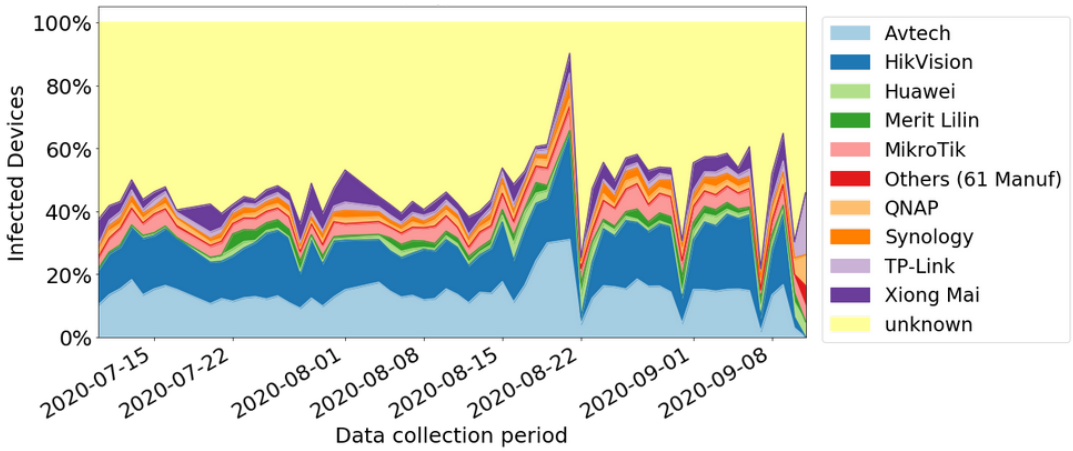

Figure 4: Top 10 Manufacturers over the period of observation

sources to determine if specific banner texts are unique to some manufacturers, but we did not succeed. Although this is a limitation of this method, which we discuss more in section section 8, the frequency of the hashes gives us some confidence that the remainder of manufacturers that we could not identify will not change the overall picture.

The devices of nine manufacturers were responsible for a large share of the global set of infections. How dominant is this pattern at the level of countries? In other words, does each government have to engage with a different set of manufacturers or does the pattern of concentration hold across countries? Figure 6 shows that, overall, the same manufacturers are responsible for a high share of the infections in most of the top 20 countries with the most infections. We aggregate the data of the other 48 countries codes under 'Others (48 CC)' as well as the data of the rest of the 61 manufacturers that were not on the top 9 under 'Others (61 Manuf) '. To calculate the ratio we divided the total number of infected devices in a country by the total number of infections attributed to a manufacturer. There is some variability, of course. In some countries, the share of 'unknown' is very high. Furthermore, in some countries the share of 'others' manufacturers is larger than that of the nine manufacturers. Still, in many countries the same manufacturers are present in the population of infected

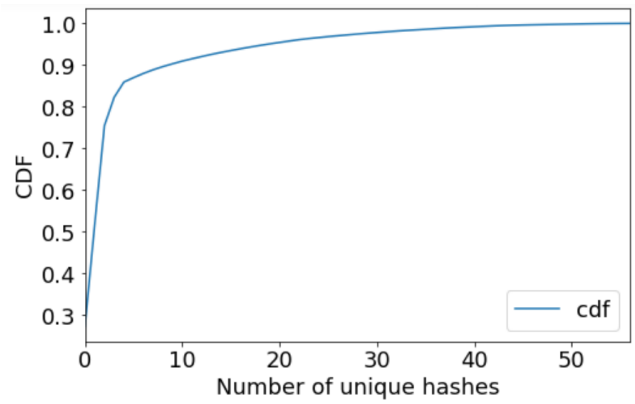

Figure 5: Frequency of hashes of non-identified devices

Manufacturer	HQ	Presence (%)	Share attributed infections (%)
HikVision	China	54 (79%)	28%
Avtech	Taiwan	47 (69%)	25%
MikroTik	Latvia	40 (59%)	7%
Xiong Mai	China	50 (74%)	7%
Synology	Taiwan	28 (41%)	3%
Merit Lilin	Taiwan	26 (38%)	3%
TP-Link	China	36 (53%)	3%
QNAP	Taiwan	34 (50%)	3%
Huawei	China	28 (41%)	3%

Table 2: Manufacturer presence across countries

devices.

Table 2 quantifies this more clearly. We checked which manufacturers are present in the population of infected devices in each country. Meaning that the manufacturer at least appeared once in the data of that country. We can see that HikVision devices are in the infected population in 54 (79%) of the 68 countries in our measurements. Avtech devices show up in 47 (69%) of all countries. Those two together are present in most countries and they represent over half of all infections that we could attribute to a manufacturer.

This suggests that international collaboration among regulators in various countries is a feasible path. This would not only bundle scarce resources on the side of governments, but is also more likely to influence manufacturer behavior through collective action. An obvious starting point would be coordination at the level of the European Union. When we look at the distribution in the E.U. countries in Figure 7, we also observe the same nine manufacturers associated with most of the infections. We aggregate the data of the other 40 countries under the label 'Others (40 CC)' as well as the data of the rest of the 61 manufacturers that were not on the top 9 under 'Others (61 Manuf)'. As before, to calculate the ratio we divided the total number of infected devices in a country by the total number of infections attributed to a manufacturer.

Table 2 also demonstrates that the locations of the manufacturers' headquarters (HQ) are highly concentrated in China and Taiwan. This suggests another path for coordination, where the governments of those countries could help facilitate improved security practices in the manufacturing processes, in order to safeguard access to overseas markets thus this can give some leverage to governments to discuss with them their security postures since their IoT products are being imported to their countries.

In sum, the dataset gives a clear answer to our first two research questions. First, which manufacturers are associated with the compromised IoT across 68 countries? It turns out that just nine manufacturers are associated with about half of all infections. Second, how variable is the set of manufacturers across different countries? We find that—notwithstanding regional and country-level differences in consumer preferences, regulatory regimes, and market access—this pattern is remarkably stable across countries.

5.2 Devices

Although our main focus is on device manufacturers within this study, it is informative to get a sense of which types of devices dominate the infected population. Figure 8 shows that, where we were able to ascertain the device type, almost 80% of the infected devices are Digital Video Records (DVRs) and IP cameras. As described in the method, as long as the devices were vulnerable to Mirai, they were considered infected. When checking the manuals, most of these devices had weak hard-code credentials. This is line with [9] work, which describe that guessable passwords vulnerable to attacks are used by the manufacturers in some of these device categories.

The Open Web Application Security Project (OWASP) describes default credentials as a top threat for IoT devices [39]. Mirai's most famous attack vector is brute forcing attacks, and since Mirai's source code was released attackers can easily add credentials to the code. Authentication in IoT devices sometimes is hard-coded or manufacturers use default credentials to set up a device for the first time, and this allows attackers to perform password guessing [40]. After the initial set up, most devices do not request to change these default credentials [41]. Therefore, manufacturers of these devices can help to fix most of the infections by implementing a better password management creating unique credentials per device.

6 Updates and security advice

Our third and final question is: What are manufacturers doing to remediate the security weaknesses of their devices? To answer this, we looked at the manufacturers' websites and at manuals (see section 4). A wide-spread complaint is that many of the vulnerable devices never receive updates [42, 43]. We found that 37 (53%) of the 70 manufacturers present in

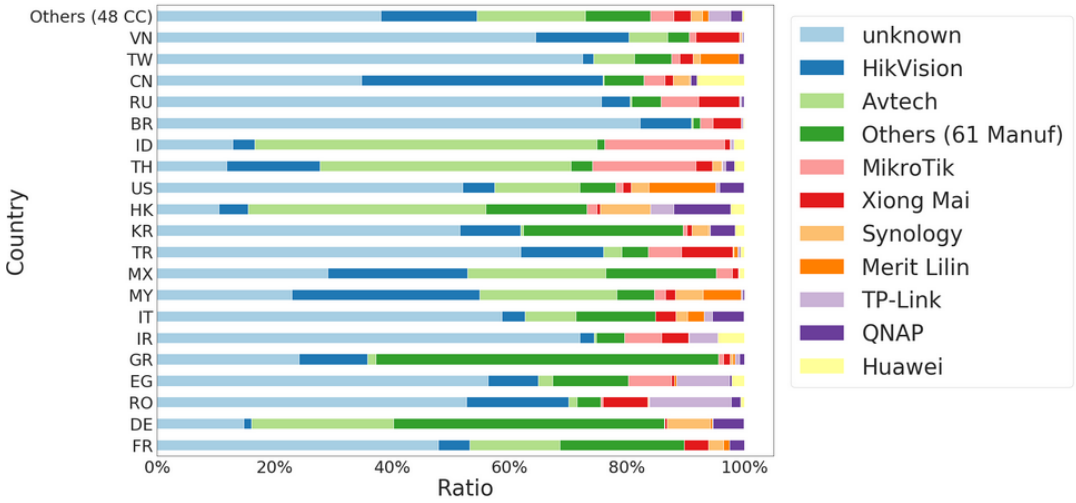

Figure 6: Share of manufacturers in top 20 countries with the most infections (countries ordered by number of infections)

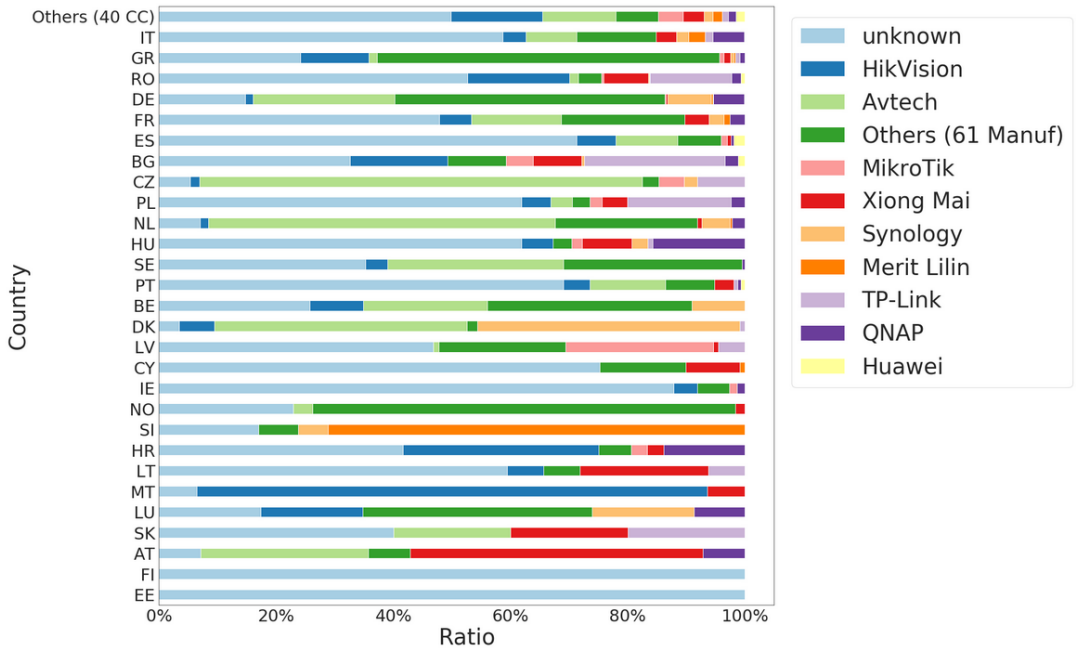

Figure 7: Share of manufacturers in E.U. countries (countries ordered by number of infections)

our dataset had either firmware or software available to download. Of the 70 manufacturers, 30 (43%) describe a password changing procedure, and 18 (26%) have some security advice on how to make the device more secure and protect it from attacks.

The picture is a bit more positive for the top 20 manufacturers associated with most infections. Of these 20, 13 (65%) had some firmware or software available to download related to their devices, and 12 (60%) describe some password changing procedures, and 8 (40%) provide some advice to protect the device from attacks or make it more secure. The two dominant manufacturers, HikVision and Avtech, had firmware and software available to download in their websites.

Table 3 presents a summary of our collected data for the top 20 manufacturers. The FW/SW update column depicts whether or not a firmware or software was found available, the password changing procedure column shows whether or not we found a password changing procedure for the device, and the last column depicts if any advice to protect the device was found or not (see more details in section 4). The data for the full set of manufacturers is provided in Appendix A. In sum, most manufacturers are making efforts to publish updates,

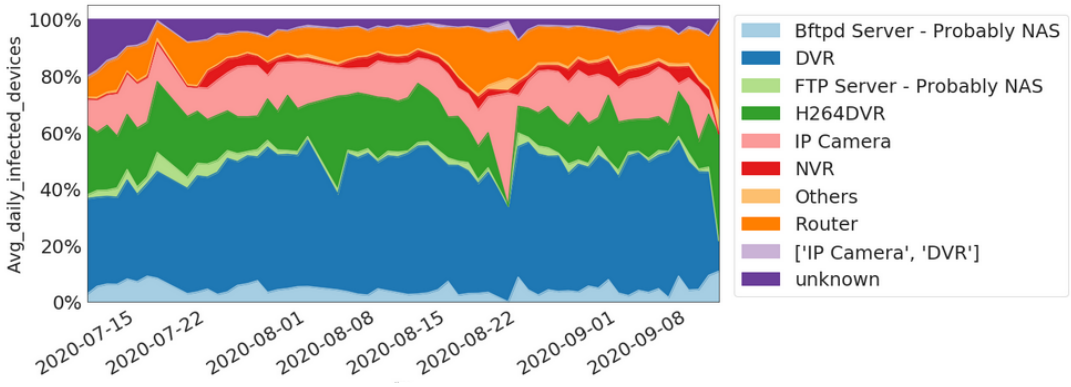

Figure 8: Top devices over the period of observation

password changing procedures, and provide security related advice. A significant group, however, does not provide one or more of these forms of support for protecting their devices. 81% lacks at least one of these three forms of support and 33% lack all three forms.

Although these findings suggest broad manufacturer support for security, it is far from complete. It is often quite difficult for consumers to find and understand the relevant information. NIST's "Foundational Cybersecurity Activities for IoT Device Manufacturers" [1] emphasizes the importance of specific IoT product information to communicate to customers and how this communication is achieved. Information such as device support, lifespan expectations, end-of life periods, how to communicate suspected vulnerabilities during and after the life span a device to the manufacturer, security capabilities of the device or manufacturer services, how to maintain security after support of the manufacturer ends, type of software updates and whom will distribute them among others things are all examples of topics that could be communicated to users according to the NIST framework.

Moreover, NIST advice states that manufacturers should provide information on whether software or firmware updates will be available, when they will be available, and how customers can verify the source and content of the update (e.g. via cryptographic hash comparison).

To illustrate, on the Hikvision website [44], one can obtain the firmware of the device, but there is no explanation of how customers can verify the source and the content of the update. Similarly, Avtech's website [45], while providing firmware download options, does not provide visitors with information on how to verify the authenticity of the firmware content either.

Although we were not assessing if manufacturer websites comply with the NIST framework, our brief examination of their content suggested that most do not offer all prescribed information, but a more systematic analysis is necessary to comprehensively assess all manufacturer websites.

During this analysis from an end-user's view, we also found that checking a manufacturer's website or manual is quite challenging in certain cases. Most websites focus on providing commercial information about devices, features, and comparison among devices. Finding manuals, support or updates might require numerous steps to achieve. In addition, the language used can be very technical. In a handful of cases, we also ran into situation where the products were discontinued by the manufacturer and we could only find the relevant device manuals on third-party websites. This pattern is aligned with the findings of [37]. Little security is provided by the manufacturers. All of this suggests room for improvement given that all these manufacturers are producing devices that are being compromised at scale.

7 Related Work

Consumers and IoT security

An important area of IoT (in-)security research has focused on empowering consumers to consider the security and privacy implications of purchasing certain IoT devices. Vendors typically do not provide information on the security features and privacy sensitive characteristics of their products – information that may help consumers make more informed purchasing decisions – and when they do, it is often inadequate [46, 47]. Various studies have thus focused on developing security labels to better inform consumers [48, 49].

Privacy advocating organizations have also introduced valuable tools and guidelines to emphasize online safety and help consumers make more informed purchasing decisions, for instance see Mozilla's `Privacy not Included` guide [50].

The potential role of third parties in protecting consumers, for instance the role of ISPs in their capacity to mitigate IoT insecurity problems, at least as a short-term solution, has also been recently examined [5].

Nevertheless informative labels, consumer empowering tools, nor third parties like ISPs, can systematically prevent post-sale security issues in IoT products [5, 48]. They do not replace the necessity of engaging with manufacturers of

	FW/SW	Password changing procedure	Advice to protect the device
HikVision	Yes	Yes	Yes
Avtech	Yes	Yes	No*
MikroTik	Yes	Yes	Yes
Xiong Mai	Yes	No	No
Synology	Yes	Yes	Yes
Merit Lilin	No	Yes	No
TP-Link	Yes	No	Yes
QNAP	Yes	Yes	Yes
Huawei	Yes	Yes	No*
ZTE	No	No	No
Beijer Electronics	No	No	No
Zhejiang Dahua Technology Co., Ltd.	Yes	Yes	Yes
DrayTek	Yes	Yes	Yes
AVM GmbH	Yes	Yes	Yes
Domoticz	Yes	Yes	No
ASUS	Yes	No*	No*
Hichan Technology	No	No	No
ZKTeco	No	No	No
ZNDS	No	No	No
Sansco	No	Yes	No

Table 3: Manufacturers offering software/firmware and security advice

Note: The asterisk in "No" means that multiple devices of this particular manufacturer were found in our data. For some of the devices the password procedure was found, but for others not. The same holds for advice to protect the device. See Appendix A for more details for each device or category.

(compromised) devices to get them to address the security problems of already sold or newly developed IoT devices.

Regulations and standards

Leverett c.s. [51] argue that existing sectoral regulators need to determine where IoT is present in their sector and to include them into existing safety and security regulations. They also highlight the need for transparency regarding products and vendors—to which our study is contributing. The European Union Cybersecurity Act provides a voluntary certification scheme for digital products, including IoT devices, in order to increase trust and security of these products [52]. Also, product liability could lead manufacturers to comply with minimum security standards in order to reduce their exposure [7].

These long-term solutions regulatory strategies, yet do not reduce the current influx of insecure devices, and our work presents an empirical approach to identify priority targets for governmental intervention.

Internal mapping of IoT devices

Several studies use internal network scans to identify IoT devices. One study [9] used the Avast Wifi Inspector to scan 16 million home networks and found 83 million connected IoT devices. To identify the manufacturers, the researchers matched part of the device MAC address with the public IEEE Organizationally Unique Identifier (OUI) list. Another study [10] created 'IoT inspector', a tool that users can run inside their home networks to label IoT devices and their manufacturers. Similar to [9] the authors use MAC addresses to validate vendors against the OUI database. A different approach was taken in [53], which fingerprints devices using information related to the Inter Arrival Time (IAT) of packets on the local network. This method was tested with just two devices in a lab setting.

All of these methods rely on user consent and privileged access to internal network data to identify manufacturers. This limits the scalability of the approach that is needed as a basis for governmental intervention, especially when representative measurements are needed across entire countries or markets. Therefore, in our study, we build on recent work on external mapping of IoT devices instead.

External mapping of IoT devices

Numerous studies identify IoT devices in the wild based on external network scans. Most are based on developing fingerprints from known devices, e.g. in a lab setting, and then searching for these fingerprints in internet-wide scans. For example, one study builds fingerprints based on specific port configurations that are chosen by manufacturers [54]. The authors test their fingerprinting approach for 19 IoT devices and subsequently develop a hierarchical port scanning method to detect device types during external scans rather than probing whole port ranges. The approach assumes that end users will retain and not modify the specific port configurations

of their devices used for fingerprinting. In [55], the authors fingerprinted routers using the initial time to live (TTL) of two Internet Control Message Protocol (ICMP) messages to determine the brand of the routers' vendor. They highlight that the hardware distribution of different brands vary across Autonomous Systems. [11] proposed IoTFinder, which contains fingerprints for 53 devices that were developed from DNS traffic data and then compared these fingerprints to traffic from an ISP network. In [12], fingerprints are developed from a testbed setting, in this case for 96 devices belonging to 40 vendors. They then enriched their fingerprints with DNS queries, web certificates, and banners and detected IoT devices in an ISP and at an Internet Exchange Point (IXP). A different approach to generating fingerprints was presented by [56]. The authors searched the web for product descriptions of devices and then they automatically created fingerprints from these descriptions (e.g., rules to detect certain strings). This potentially scales better than generating fingerprints from analyzing the devices themselves or their firmware. However, [57] challenged the reproducibility of this method.

A common feature among these approaches is that they first develop fingerprints for a set of known devices under the control of the researchers and then conduct external scans with these fingerprints. Furthermore, some approaches—e.g., [11] and [12]—need access to ISP or IXP traffic in order to detect their fingerprints. This approach does not work for our problem of identifying a given population of devices in the wild, namely compromised devices. We cannot know which devices are in that population, let alone have them available in a lab setting for generating fingerprints.

Two other studies [14, 17] focused specifically on compromised devices. They identified the IP addresses of compromised IoT devices via attack traffic observed in darknet data. They did not develop fingerprints, however. The actual identification of the devices present at those IP addresses, was not conducted by the researchers. Instead, it relied on third-party data, most notably searching for the IP addresses in Shodan [16], a search engine that indexes a variety of internet-connected systems.

While we focus on consumer IoT, there is some overlap in approaches with the research on identifying industrial control systems (ICS) devices [58], which also relied on Shodan [16] and Censys [59]. Fingerprints were developed for individual ICS devices in order to track them over time, not for manufacturer identification. While also [60] developed a realtime ICS discovery system using ICSs protocols to discover ICS devices in the whole IPv4 space. They analyzed 17 ICSs protocols, and they did common requests that could fingerprint the ICSs devices based on the responses they obtained and that were unique to the protocols.

Like [14, 17], our study also uses attack traffic to detect the presence of compromised IoT devices, namely observing the Mirai fingerprint in darknet data. We base our analysis on a longer data collection period of two months. For our device identification, we do not rely on a black-box third party solution like Shodan. This would make it impossible to explain to manufacturers via what method their devices were identified, nor gauge how accurate this method is. Explainability and accuracy—which includes knowing the method's inaccuracy—are key requirements for providing the government with the basis to select and engage manufacturers. Rather than relying on third-party services we develop our own fingerprints, as we explained in Section 4. Different from the other studies using fingerprints, we could not start with a set of known devices to develop the fingerprints. Rather, we need a method to identify manufacturers present in a given population of compromised devices.

8 Limitations and Future Work

Our approach is to scan back an IP address from which Mirai scan traffic has originated moments earlier. A core assumption behind this approach is that the scan back will actually connect with the same device from which the attack traffic was observed. In reality, there will typically be multiple devices behind the same public IP address. Some, if not most, of those devices will not be publicly reachable. In theory we might be engaging with one of those reachable devices or with the router, either of which may or may not be the infected device. While we have no certainty that the device that we scanned back is actually the infected device, we have certain indicators that increase the confidence in our approach. First, attackers behind the Mirai infections recruit devices also by scanning IPv4 addresses for publicly-reachable devices, the same logic that we apply. So if they could infect a device, that device has to be visible in an active scan. In only 1.2% of the cases did we find fingerprints for different devices at the same IP address, consistent with the fact that in most cases only a single device was accessible from the open Internet. Second, the probability that we are scanning the Internet-facing router, rather than a Mirai-infected device behind the router, is severely mitigated in light of our data. Over 60% of the devices we identified were not routers. Where we did identify routers, these models were known to be vulnerable to Mirai. This brings us to a third indicator: all devices that we identified from the banners and Web-UIs were investigated to ascertain that they were actually reported to be vulnerable for Mirai. They were, without exception.

A second limitation is that our active scanning method did not use all protocols used by consumer IoT devices. It was limited to 'ftp', 'telnet', 'http', 'SSH'. This means we might miss devices that do not operate any of these protocols. Future research might expand the set of protocols and quantify what proportion we are missing as well as try to include all IoT related protocols to better understand the infections landscape.

A third limitation is related to the fact that we only scanned devices infected with a variant from the Mirai malware family. Other malware families might bring into view additional

devices. That being said, Mirai has been the dominant malware family for years and is still being detected as a leading malware family, responsible for 21% of the IoT infected devices [61]. Furthermore, it has been reported that the different IoT malware families often compete over the same devices [62], which suggests that the Mirai population is not systematically different from other families.

A fourth limitation is related to our use of the Mirai fingerprint to identify infected devices. There is an extremely small probability that this fingerprint occurs by accident ($\frac{1}{2^{32}}$, to be precise). That still leaves open the possibility that someone sends out this fingerprint on purpose. We are not aware of any use cases for doing this. Such an activity would not be part of a honeypot design for Mirai. In any case, it is unlikely that such technical corner cases would originate in substantial numbers from consumer broadband network (as opposed to research or hosting networks).

A fifth limitation impacts our assessment of the volume of infected devices in each country. IP address reassignment (a.k.a. DHCP churn) might impact the number of infections we observed per country. To minimize the impact of churn, we assumed that IP addresses are not reassigned multiple time per day. We count infections in 24hrs long sliding windows and with each batch of scans start a new count of unique IP addresses. This significantly reduces the risk of overcounting because of churn.

Another limitation is that we could not identify the manufacturer for a significant portion—roughly about half—of all infected devices. To the best of our knowledge, no other method for identify IoT devices in the wild has achieved better rates, but this is still a limitation of our work. As we discussed in subsection 5.1, the portion of unknown devices is unlikely to impact the pattern of concentration around nine manufacturers that we uncovered.

A final limitation is related to the fact that we are not sure when the websites or manuals of the manufacturers were updated. Hence, some of the security advice could have been recently added or not up to date. Moreover, we did not check if the firmware updates were actually solving the vulnerabilities of the device, but just if there was firmware or software available to download.

9 Conclusions and Discussion

The IoT ecosystem is complex and involves many different actors. Many observers have argued that the incentives in around IoT security are misaligned. [63, 64] There is a lack of adequate information available to consumers regarding the security of the devices that they are purchasing. The costs of security failures are often borne by other stakeholders than the owners of the device or the manufacturers. So there is a market failure here that justifies government intervention. There is no single solution, of course. A recent step of the Dutch government has been a voluntary agreement with the main online electronics retailers to include in the product descriptions whether the product will receive security updates and, if so, for how long [25]. The current status is that many of these fields are still listed as 'unknown'. Many manufacturers are not supplying this information in their product description.

Any sensible strategy towards IoT security will have to change manufacturer behavior towards designing more secure devices. This is especially critical for the manufacturers associated with devices that have been getting compromised at scale in the wild. In this paper, we have investigated the manufacturers associated with the population of infected devices in 68 countries. We found that just nine manufacturers share about half of the infected devices across all countries. Notwithstanding the differences between countries in terms of consumer preferences, manufacturer presence in the market and regulatory regimes, this pattern also holds at the country level for most countries in the top 20 with most infections, as well as across European Union member states. Hence, policy makers can unite their efforts to target those to encourage them to improve their security postures. Most devices come, unsurprisingly, from China and Taiwan, the leading hardware manufacturers of the world. This concentration on the supply side of the market suggests that governments confronted with infected devices might engage their counterparts in China and Taiwan to change the behavior of the manufacturers in those countries, if only to safeguard their exports towards large markets in the U.S. and E.U.

Even though many manufacturers do provide security updates or advice, it seems that this is not enough to prevent and remediate the infections. This could be because of users' misaligned incentives [63], but it could also reflect that this support is hard to find and even harder to act on. The information on the support pages is fragmented. A user has to click different links, understand what files to download, and install them without a clear idea of what the new firmware version will or will not fix. Hence, there is room for improvement about what and how to present this information to users, as discussed in [1]. This would also reduce the cost that users have to incur to secure their devices.

The efforts that policy makers undertake can have an impact also outside their own jurisdiction. Think of how the E.U. became the *de facto* privacy regulator of the world, via the General Data Protection Regulation. Most websites adopted it globally, because it was more efficient than differentiating the setup for each jurisdiction [24]. If policy makers unify their efforts and the pattern of concentration on a handful of manufacturers holds, then a global impact is not unrealistic.

Retailers of IoT devices could also play a role, as countries such as The Netherlands are proposing [25]. If users can return these devices to retailers, then these costs would lead the retailer to exert pressure further up the supply chain and create better security incentives for manufacturers.

Government involvement is currently underway. Many countries are introducing legislation or shoring up existing

mechanisms to improve security. Our findings are a stepping stone for efforts by the Dutch government to engage the manufacturers found to be supplying most of the infected devices. Time will tell whether government pressure, in combination with empirical evidence of the problems caused by their products, is enough the start changing the security practices of these companies—and of the IoT market at large. These findings are based only on Mirai and we did not use all protocols used by consumers IoT devices, so future research could look into more IoT malware families and add additional protocols to have a more complete overview of the whole manufacturer landscape.

Acknowledgments

The authors would like to thank our anonymous reviewers for their feedback and suggestions to improve the quality of our manuscript. We would like to thank the Dutch Ministry of Economic Affairs for supporting our research. This publication is part of the MINIONS project (number 628.001.033) of the "Joint U.S.-Netherlands Cyber Security Research Programme" which is (partly) financed by the Dutch Research Council (NWO); and of the related MINIONS-TLD project, which is financed by SIDN, the .nl registry. We also wish to thank Jochem van de Laarschot for helping with the data collection of Table 3 and Appendix A.

References

[1] Michael Fagan, Katerina Megas, Karen Scarfone, and Matthew Smith. *Foundational Cybersecurity Activities for IoT Device Manufacturers*. 2020. URL: https://doi.org/10.6028/NIST.IR.8259 (visited on 03/26/2021).

[2] Michael Fagan, Katerina N Megas, Karen Scarfone, and Matthew Smith. *IoT Device Cybersecurity Capability Core Baseline*. 2020. URL: https://nvlpubs.nist.gov/nistpubs/ir/2020/NIST.IR.8259A.pdf (visited on 03/26/2021).

[3] ENISA. *Good Practices for Security of IoT - Secure Software Development Lifecycle — ENISA*. 2019. URL: https://www.enisa.europa.eu/publications/good-practices-for-security-of-iot-1 (visited on 03/24/2021).

[4] M James. *Secure by Design: Improving the cyber security of consumer Internet of Things Report*. 2017. URL: https://assets.publishing.service.gov.uk/government/uploads/system/uploads/attachment_data/file/775559/Secure_by_Design_Report_.pdf (visited on 03/26/2021).

[5] Arman Noroozian, Elsa Rodríguez, Elmer Lastdrager, Takahiro Kasama , Michel van Eeten and Carlos Gañán. "Can ISPs Help Mitigate IoT Malware? A Longitudinal Study of Broadband ISP Security Efforts (to appear)". In: *IEEE Euro Security & Privacy* (2021).

[6] Bruce Schneier. *Click here to kill everybody: Security and survival in a hyper-connected world*. WW Norton & Company, 2018.

[7] Iain Nash. "A Proposed Civil Liability Framework for Disrupting Botnets, with a particular focus on Smart Devices". Botconf. 2020.

[8] Federal Trade Commission. *ASUS Settles FTC Charges That Insecure Home Routers and "Cloud" Services Put Consumers' Privacy At Risk | Federal Trade Commission*. 2016. URL: https://www.ftc.gov/news-events/press-releases/2016/02/asus-settles-ftc-charges-insecure-home-routers-cloud-services-put?utm{_}source=govdelivery (visited on 03/20/2021).

[9] Deepak Kumar, Kelly Shen, Benton Case, Deepali Garg, Galina Alperovich, Dmitry Kuznetsov, Rajarshi Gupta, and Zakir Durumeric. "All Things Considered: An Analysis of IoT Devices on Home Networks". In: *28th USENIX Security Symposium (USENIX Security 19)*. Santa Clara, CA: USENIX Association, 2019, pp. 1169–1185.

[10] Danny Yuxing Huang, Noah Apthorpe, Gunes Acar, Frank Li, and Nick Feamster. "IoT Inspector: Crowdsourcing Labeled Network Traffic from Smart Home Devices at Scale". In: *Proceedings of the ACM on Interactive, Mobile, Wearable and Ubiquitous Technologies (IMWUT / Ubicomp* (2020).

[11] Roberto Perdisci, Thomas Papastergiou, Omar Alrawi, and Manos Antonakakis. "IoTFinder : Efficient Large-Scale Identification of IoT Devices via Passive DNS Traffic Analysis". In: *Acm Imc* (2020).

[12] Said Jawad Saidi, Anna Maria Mandalari, Roman Kolcun, Daniel J. Dubois, David Choffnes, Georgios Smaragdakis, and Anja Feldmann. "A Haystack Full of Needles: Scalable Detection of IoT Devices in the Wild". In: *Acm Imc* (2020).

[13] Kai Yang, Qiang Li, and Limin Sun. "Towards automatic fingerprinting of IoT devices in the cyberspace". In: *Computer Networks* 148 (2019), pp. 318–327.

[14] Nataliia Neshenko, Martin Husak, Elias Bou-Harb, Pavel Celeda, Sameera Al-Mulla, and Claude Fachkha. "Data-Driven Intelligence for Characterizing Internet-Scale IoT Exploitations". In: *2018 IEEE Globecom Workshops, GC Wkshps 2018 - Proceedings*. Institute of Electrical and Electronics Engineers Inc., 2019.

[15] Manos Antonakakis, Tim April, Michael Bailey, Matt Bernhard, Elie Bursztein, Jaime Cochran, Zakir Durumeric, J Alex Halderman, Luca Invernizzi, Michalis Kallitsis, and others. "Understanding the mirai botnet". In: *USENIX Security Symposium*. 2017, pp. 1092–1110.

[16] Shodan. 2020. URL: https://www.shodan.io/.

[17] Mario Galluscio, Nataliia Neshenko, Elias Bou-Harb, Yongliang Huang, Nasir Ghani, Jorge Crichigno, and Georges Kaddoum. "A first empirical look on internet-scale exploitations of IoT devices". In: *2017 IEEE 28th Annual International Symposium on Personal, Indoor, and Mobile Radio Communications (PIMRC)*. IEEE, 2017, pp. 1–7.

[18] Orçun Cetin, Carlos Gañán, Lisette Altena, Takahiro Kasama, Daisuke Inoue, Kazuki Tamiya, Ying Tie, Katsunari Yoshioka, and Michel Van Eeten. "Cleaning Up the Internet of Evil Things: Real-World Evidence on ISP and Consumer Efforts to Remove Mirai". In: *NDSS*. 2019.

[19] IoT Business News. *IoT News - The IoT in 2030: 24 billion connected things generating $1.5 trillion - IoT Business News*. 2020. URL: https://iotbusinessnews.com/2020/05/20/03177-the-iot-in-2030-24-billion-connected-things-generating-1-5-trillion/ (visited on 03/20/2021).

[20] Euroepan Commission. *The Internet of Things | Shaping Europe's digital future*. 2019. URL: https://ec.europa.eu/digital-single-market/en/internet-of-things (visited on 03/22/2021).

[21] E. Lear, R. Droms, and D. Romascanu. *Manufacturer Usage Description Specification*. Tech. rep. 2019. URL: https://www.rfc-editor.org/info/rfc8520.

[22] Michael Richardson and M Ranganathan. *Manufacturer Usuage Description for quarantined access to firmware*. Tech. rep. draft-richardson-shg-mud-quarantined-access-01. Internet Engineering Task Force, 2019. URL: https://datatracker.ietf.org/doc/html/draft-richardson-shg-mud-quarantined-access-01.

[23] *Software Bill of Materials | National Telecommunications and Information Administration*. URL: https://www.ntia.gov/SBOM (visited on 03/24/2021).

[24] Catherine Barrett. "Are the EU GDPR and the California CCPA becoming the de facto global standards for data privacy and protection?" In: *Scitech Lawyer* 15.3 (2019), pp. 24–29.

[25] NOS. *Kabinet wil verplichte updates voor 'slimme' apparaten | NOS*. 2019. URL: https://nos.nl/artikel/2315967-kabinet-wil-verplichte-updates-voor-slimme-apparaten.html (visited on 12/09/2020).

[26] Ministry of Economic Affairs and Climate Policy. *Roadmap for Digital Hard- and Software Security | Report | Government.nl*. URL: https://www.government.nl/documents/reports/2018/04/02/roadmap-for-digital-hard--and-software-security (visited on 03/25/2021).

[27] David Dittrich and Erin Kenneally. *The Menlo Report: Ethical Principles Guiding Information and Communication Technology Research*. 2012. URL: http://www.caida.org/publications/papers/2012/menlo_report_actual_formatted.

[28] Robert Graham. *GitHub - robertdavidgraham/masscan: TCP port scanner, spews SYN packets asynchronously, scanning entire Internet in under 5 minutes.* URL: https://github.com/robertdavidgraham/masscan (visited on 03/20/2021).

[29] *GitHub - zmap/zgrab2: Fast Go Application Scanner*. URL: https://github.com/zmap/zgrab2 (visited on 03/22/2021).

[30] *GitHub - sensepost/gowitness: gowitness - a golang, web screenshot utility using Chrome Headless*. URL: https://github.com/sensepost/gowitness (visited on 03/21/2021).

[31] Liz Izhikevich, Renata Teixeira, and Zakir Durumeric. "LZR: Identifying Unexpected Internet Services". In: *30th USENIX Security Symposium*. 2021.

[32] Hadi Asghari, Michel JG van Eeten, and Johannes M Bauer. "Economics of fighting botnets: Lessons from a decade of mitigation". In: *IEEE Security & Privacy* 13.5 (2015), pp. 16–23.

[33] Arman Noroozian, Michael Ciere, Maciej Korczynski, Samaneh Tajalizadehkhoob, and Michel Van Eeten. "Inferring the Security Performance of Providers from Noisy and Heterogenous Abuse Datasets". In: *16th Workshop on the Economics of Information Security. http://weis2017. econinfosec. org/wp-content/uploads/sites/3/2017/05/WEIS_2017_paper_60. pdf*. 2017.

[34] Qasim Lone, Maciej Korczyński, Carlos Gañán, and Michel van Eeten. "SAVing the Internet: Explaining the Adoption of Source Address Validation by Internet Service Providers". In: *Workshop on the Economics of Information Security*. 2020.

[35] Fumiyuki Tanemo, Mitsuhiro Osaki, Hiroaki Waki, Yutaka Ishioka, and Kazuhito Matsushita. "A Method of Creating Data for Device-information Extraction by Efficient Wide-area-network Scanning of IoT Devices". In: *2020 International Conference on Information Networking (ICOIN)*. IEEE. 2020, pp. 643–648.

[36] Buchner Johannes. *ImageHash · PyPI*. URL: https://pypi.org/project/ImageHash/ (visited on 03/30/2021).

[37] John M Blythe, Nissy Sombatruang, and Shane D Johnson. "What security features and crime prevention advice is communicated in consumer IoT device manuals and support pages?" In: *Journal of Cybersecurity* 5.1 (2019), tyz005.

[38] *TeleGeography | Home*. URL: https://www2.telegeography.com/ (visited on 03/30/2021).

[39] OWASP Fundation. *OWASP Internet of Things Project - OWASP*. 2018. URL: https://wiki.owasp.org/index.php/OWASP_Internet_of_Things_Project#tab=IoT_Top_10 (visited on 02/18/2021).

[40] BR Chandavarkar. "Hardcoded Credentials and Insecure Data Transfer in IoT: National and International Status". In: *2020 11th International Conference on Computing, Communication and Networking Technologies (ICCCNT)*. IEEE. 2020, pp. 1–7.

[41] Nataliia Neshenko, Elias Bou-Harb, Jorge Crichigno, Georges Kaddoum, and Nasir Ghani. "Demystifying IoT security: an exhaustive survey on IoT vulnerabilities and a first empirical look on internet-scale IoT exploitations". In: *IEEE Communications Surveys & Tutorials* 21.3 (2019), pp. 2702–2733.

[42] Bruce Schneier. "The internet of things is wildly insecure-and often unpatchable". In: *Schneier on Security* 6 (2014).

[43] Sarthak Grover and Nick Feamster. "The internet of unpatched things". In: *Proc. FTC PrivacyCon* (2016).

[44] Hikvision. *DS-2TD1117-2/PA | HeatPro Series | Hikvision*. 2021. URL: https://www.hikvision.com/en/products/Thermal-Products/Security-thermal-cameras/heatpro-series/ds-2td1117-2-pa/ (visited on 03/24/2021).

[45] Avtech. *AVTECH - Leader in Push Video HDCCTV, IP Camera, CCTV camera, DVR, IVS Network camera, EagleEyes mobile surveillance, NVR, NAS and CMS total solution.* 2014. URL: https://www.avtech.com.tw/EOL.aspx (visited on 03/24/2021).

[46] Pardis Emami-Naeini, Yuvraj Agarwal, Lorrie Faith Cranor, and Hanan Hibshi. "Ask the Experts: What Should Be on an IoT Privacy and Security Label?" In: *2020 IEEE Symposium on Security and Privacy (SP)*. IEEE. 2020, pp. 447–464.

[47] JM Blythe and SD Johnson. "The Consumer Security Index for IoT: A protocol for developing an index to improve consumer decision making and to incentivize greater security provision in IoT devices". In: *IET* (2018).

[48] P. Morgner, C. Mai, N. Koschate-Fischer, F. Freiling, and Z. Benenson. "Security Update Labels: Establishing Economic Incentives for Security Patching of IoT Consumer Products". In: *2020 IEEE Symposium on Security and Privacy (SP)*. 2020, pp. 429–446.

[49] Pardis Emami-Naeini, Henry Dixon, Yuvraj Agarwal, and Lorrie Faith Cranor. "Exploring How Privacy and Security Factor into IoT Device Purchase Behavior". In: *Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems*. CHI '19. Glasgow, Scotland Uk: Association for Computing Machinery, 2019, 1–12.

[50] The Mozilla Foundation. *Mozilla - *privacy not included*.* 2020. URL: https://foundation.mozilla.org/en/privacynotincluded/ (visited on 01/11/2021).

[51] Eireann Leverett, Richard Clayton, and Ross Anderson. "Standardisation and Certification of the 'Internet of Things'". In: *Proceedings of WEIS*. 2017, pp. 1–24.

[52] European Commission. *The EU cybersecurity certification framework | Shaping Europe's digital future*. 2020. URL: https://ec.europa.eu/digital-single-market/en/eu-cybersecurity-certification-framework (visited on 03/21/2021).

[53] Sandhya Aneja, Nagender Aneja, and Md Shohidul Islam. "IoT Device Fingerprint using Deep Learning". In: *2018 IEEE International Conference on Internet of Things and Intelligence System (IOTAIS)*. IEEE, 2018, pp. 174–179.

[54] Arunan Sivanathan, Hassan Habibi Gharakheili, and Vijay Sivaraman. "Can We Classify an IoT Device using TCP Port Scan?" In: *2018 IEEE 9th International Conference on Information and Automation for Sustainability, ICIAfS 2018*. Institute of Electrical and Electronics Engineers Inc., 2018.

[55] Emeline Marechal and Benoît Donnet. "Network Fingerprinting: Routers under Attack". In: *IEEE International Workshop on Traffic Measurements for Cybersecurity (WTMC)*. 2020.

[56] Xuan Feng, Qiang Li, Haining Wang, and Limin Sun. "Acquisitional Rule-based Engine for Discovering Internet-of-Things Devices". In: *27th USENIX Security Symposium (USENIX Security 18)*. Baltimore, MD: USENIX Association, Aug. 2018, pp. 327–341.

[57] Talha Javed, Muhammad Haseeb, Muhammad Abdullah, and Mobin Javed. "Using Application Layer Banner Data to Automatically Identify IoT Devices". In: *SIGCOMM Comput. Commun. Rev.* 50.3 (July 2020), 23–29.

[58] Michael Dodson, Daniel Thomas, and Alastair R Beresford. "When will my PLC support Mirai? The security economics of large-scale attacks against Internet-connected ICS devices". In: (2020). ECRIME 2020 – SYMPOSIUM ON ELECTRONIC CRIME RESEARCH.

[59] Cencys. *Home*. 2020. URL: https://censys.io/ (visited on 03/20/2021).

[60] Xuan Feng, Qiang Li, Haining Wang, and Limin Sun. "Characterizing industrial control system devices on the Internet". In: *2016 IEEE 24th International Conference on Network Protocols (ICNP)*. 2016, pp. 1–10.

[61] Kaspersky. *New Mirai botnet is targeting enterprise IoT | Kaspersky official blog*. 2019. URL: https://www.kaspersky.com/blog/mirai-enterprise/26032/ (visited on 02/18/2021).

[62] Harm Griffioen and Christian Doerr. "Examining Mirai's Battle over the Internet of Things". In: *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*. 2020, pp. 743–756.

[63] Matheus Xavier Ferreira, S Matthew Weinberg, Danny Yuxing Huang, Nick Feamster, and Tithi Chattopadhyay. "Selling a Single Item with Negative Externalities". In: *The World Wide Web Conference*. 2019, pp. 196–206.

[64] Erin Kenneally. "Economics and Incentives Driving IoT Privacy and Security, Pt. 1". In: *IEEE Internet of Things Magazine* 2.1 (2019), pp. 6–7.

A Appendices

	Manufacturer	Device	FW/SW	Password changing procedure	Advice to protect the device
1	ABUS	DVR	Yes	Yes	No
2	Advanced Multimedia Internet Technology (AMIT)	WIP-300 Router	No	Yes	No
3	ASUS	RT-AC5300	Yes	Yes	No
		RT-N10U	Yes	No	No
		RT-AC58U	Yes	Yes	Yes
		RT-N10 + B1	Yes	No	No
		RT-AC54U	Yes	Yes	Yes
		RT-AC87U	Yes	Yes	Yes
		RT-N14U	Yes	Yes	No
		RT-N13U.B1	Yes	No	No
		RT-G32	Yes	No	No
		RT-N10	Yes	No	No
		DSL-N10	Yes	No	No
		WIRELESS-AC1200	Yes	Yes	Yes
4	AVM GmbH	FritzBox Router	Yes	Yes	Yes
5	AirTies	Air4920-2 SetTopBox	No	No	No
		Air7120 SetTopBox	No	No	No
6	Amlogic	SetTopBox S905L	No	No	No
7	Asustor	NAS	Yes	No	No
8	Avtech	IP Camera, DVR	Yes	Yes	No
		IP Camera	Yes	Yes	Yes
9	Bab Technologie	Unknown	NA	NA	NA
10	Beijer Electronics	QTERM Panel	No	No	No
11	Broadcom	BCM Router	No	No	No
12	Ceru Co. Ltd	vu+ Solo2	No	No	No
13	Cisco	Docsis Gateway	No	No	No
14	D-Link	Router	Yes	Yes	Yes
15	Devolo	Microlink Dlan Wireless	Yes	No	No
16	Digicom	RAW300L-A05 Router	Yes	Yes	Yes
17	Domoticz	Home Automation	Yes	Yes	No
		Domoticz Machinon	Yes	Yes	No
18	DrayTek	Vigor 2860 Router	Yes	Yes	Yes
		Vigor 2925 Router	Yes	Yes	Yes
		Vigor 2760 Router	Yes	Yes	Yes
		Vigor 2960 Router	Yes	Yes	Yes
		Vigor 2926 Router	Yes	Yes	Yes
		Vigor 2133F Router	Yes	Yes	Yes
		Vigor 2862 Router	Yes	Yes	Yes
19	Dream Multimedia	Dreambox DVB Satellite	No	No	No
20	Fibaro	Home Centre	Yes	No	No
21	Flying Voice Technology	FWR9601 VoIP Router	Yes	No	No
22	Foscam	Foscam	Yes	Yes	Yes
23	Freebox	SetTopBox	Yes	No	No
24	GNSS	Receiver Net-G5 GNSS	Yes	No	No
25	Grandstream	UCM6202 IP PBX	Yes	Yes	Yes
26	Hichan Technology	Router WiDisk	No	No	No
27	HikVision	IP Camera	Yes	Yes	Yes
		DVR	Yes	Yes	Yes
28	Hisilicon	Hi3798MV300 SetTopBox	No	No	No
29	Huawei	Router	Yes	Yes	No
		SetTopBox	No	No	No
		Home Gateway	No	Yes	Yes
		HG659	No	Yes	Yes
30	Inim Electronics	Smartlan Fire Control System	No	No	No
31	Innbox	VDSL2 modem	No	No	No
32	Interlogix	TruVision NVR	Yes	Yes	Yes
33	Level One	WBR-6005 Router	No	Yes	Yes
34	Lifetrons	FG1060N Wifi Router	No	No	No
35	Linksys	Router	Yes	Yes	Yes
		Linksys LRT214	Yes	Yes	Yes

36	MAGINON	Camera, camcorders, other electronics	Yes	Yes	Yes
		IPC-250HDC	Yes	Yes	Yes
		Security Camera	Yes	Yes	No
37	Merit Lilin	NVR	No	Yes	No
38	MikroTik	Router	Yes	Yes	Yes
		Router v6.12	Yes	Yes	Yes
		Router v6.43.12	Yes	Yes	Yes
39	Netcomm	VDSL2 N300 WiFi Router	Yes	Yes	No
40	Netis	Router	Yes	No	No
41	Opendreambox	SetTopBox	No	No	No
42	Phicomm	Router	No	Yes	Yes
43	QNAP	QNAP QTS	Yes	Yes	Yes
		Network Attached Storage	Yes	Yes	Yes
		QNAP QTS 4.3.3.1098	Yes	Yes	Yes
		QNAP QTS 4.4.2.12.62	Yes	Yes	Yes
		QNAP QTS 4.3.4.1129	Yes	Yes	Yes
		QNAP QTS 4.2.6	Yes	Yes	Yes
		QNAP QTS 4.2	Yes	Yes	Yes
44	Reolink	NVR	Yes	Yes	No
45	Ricoh	Aficio MP 301 Printer	No	No	No
46	Samsung	DVR	Yes	Yes	No
47	Sansco	NVR Security Camera	No	Yes	No
48	Siera	Siera Panther DVR	No	No	No
49	Sompy	Alarm System	No	No	No
50	Sony	Ipela SNC-CH160	Yes	Yes	Yes
51	STMicroelectronics	Unknown	NA	NA	NA
52	Strong	Extender 1600	Yes	No	No
53	Synology	Disk Station	Yes	Yes	Yes
		Disk Station DS916	Yes	Yes	Yes
54	TOTOLink	Router	Yes	Yes	No
55	TP-Link	Router	Yes	No	Yes
56	Tecom	AH2322 ADSL Router	No	No	No
57	Ubiquiti	Aircube AC	Yes	No	No
58	Uniview	Unv IP Camera	No	No	No
59	Upvel	UR 313N4G Router	Yes	Yes	No
		UR-321BN Router	Yes	Yes	No
60	VACRON	NVR	Yes	No	No
61	Vimar	Elvox Video Door entry	Yes	No	No
62	X10 Wireless Technology Inc	IP Camera AirSight Xx34A	No	No	No
63	XPO Tech	ZEM560 Fingerprint	No	No	No
64	Xiong Mai	White labeling DVR, White labeling NVR	Yes	No	No
		DVR	Yes	No	No
		NAS	No	No	No
65	ZKTeco	ZEM560 Fingerprint	No	No	No
		ZMM220	No	No	No
66	ZNDS	Smart TV Box	No	No	No
67	ZTE	Router	No	No	No
		F620V2 Router	No	No	No
68	Zhejiang Dahua Technology Co., Ltd.	IP Camera (IR PTZ Dome Camera)	Yes	Yes	Yes
		IP Camera	Yes	Yes	Yes
69	Zhone Technologies	ZNID-GPON-2426A-NA Router	No	Yes	No
70	Zyxel	ADSL gateway	No	No	Yes
		WAP5705 Media Streaming Box	No	Yes	Yes
		NSA325 v2	Yes	Yes	Yes