

GCA Internet Integrity Papers: IoT Policy and Attack Report

OCTOBER 2021



Foreword

This **IoT Policy and Attack Report**, the first of the **GCA Internet Integrity Papers** series, is a crucial milestone in a long journey that started in 2018, when the Global Cyber Alliance (GCA) set the AIDE (Automated IoT Defense Ecosystem) platform in motion. This report is also the first—but not intended to be the last—large-scale documentation of work that has been crafted by GCA’s collective technical team for several years already.

It is also a fundamental step for GCA’s newly formed **Internet Integrity Program**, which addresses cybersecurity issues that can best be seen and treated in the heart of the Internet, and brings together key players to identify top priorities for addressing cybersecurity issues.

This report provides a review of what real-life IoT attack data looks like, from the vantage point of a multi-hundred node honeyfarm, as well as from specific deceptions set up using GCA’s own honeypot technology (ProxyPot).

In addition to the key takeaways highlighted in the executive summary, the reader should observe the impact that this type of combined broad-scale survey of Internet activity, coupled with specific targeted deceptions, can yield important insights for policy makers, device manufacturers, and network operators the world over. We look forward to partnering with more organizations to scope out the reality of IoT attacks, and find ways to prevent and mitigate them.

GCA is releasing this report at the beginning of Cyber Security Awareness Month. Be cyber security aware and enjoy the reading!

October 4, 2021

Leslie Daigle

Chief Technical Officer,
Director, Internet Integrity Program
Global Cyber Alliance



Executive Summary

The **IoT Policy and Attack Report** shows the results of a research project conducted jointly by the **Global Cyber Alliance (GCA)** and **Microsoft** in mid 2021.

Based on data about real IoT attacks coming from GCA's **AIDE** platform and **ProxyPot** infrastructure, the project aimed at providing factual evidence on the validity of the most widespread policies, recommendations, and standards on IoT security around the globe.

The result of the research strongly suggests that policymakers are correct in emphasizing **secured access** when turning standards into policy. **Strong passwords**, **vulnerability disclosure policies**, and **efficient patching** should be guiding principles of any baseline strategy for IoT protection.

- The ETSI principles of “no default passwords,” “implement a vulnerability disclosure policy,” and “keep software updated” operate as a golden rule for IoT security
- Secured access and strong passwords go hand in hand
- Telnet is as a perfect gateway for IoT attacks (use SSH instead); Mirai is, by far, the largest source of Telnet-based attacks
- AIDE and ProxyPot work as reliable testing tools to validate IoT policies,

Table of Contents

- 7** | Introduction
- 10** | Overall Recommendations to Policymakers
- 15** | Methodology
- 18** | Findings
- 35** | Conclusions
- 36** | Appendix: Geo IP Locations

Policymakers around the world are acknowledging the profound implications of IoT security for privacy, safety, critical infrastructure, and trustworthy digital transformation. Different approaches to IoT security policy range from voluntary programs for industry to mandatory security requirements. The range of IoT device types, growing number of devices, and the volume of interactions between devices, the physical world, and the internet make developing effective and relevant cybersecurity policy a complex task. To tackle the IoT threat landscape, the global community of IoT manufacturers and cybersecurity experts have developed sets of best practices for IoT device cybersecurity through standards. These standards have demonstrated effectiveness against common attacks, and industry and policymakers alike can leverage them for immediate improvements to the global state of IoT device security for consumer, enterprise, and government users.

The increased availability of standards for minimum IoT security baselines is a promising start to improve global cybersecurity health.

Prominent examples of international standards include the **European Telecommunications Standards Institute (ETSI)** standard for consumer IoT security released in June 2020, ETSI EN 303 645. The ETSI standard is now a public set of resources that governments and companies around the world can use to enhance the security of IoT devices, through governance and technical recommendations. It was created through a rigorous and collaborative multistakeholder process involving experts from industry, government, and academia.

Similarly, the **National Institute of Standards and Technology (NIST)** released NISTIR 8259A in May 2020, which details a baseline set of device capabilities necessary for common cybersecurity controls. The **International Organization for Standardization (ISO)** and the **International Electrotechnical Commission (IEC)** are also developing a minimum-security baseline.

These standards are intended to provide guidance for manufacturers, developers, and users to identify and adopt best practices for device security.

Policy can help manufacturers take up international standards in a consistent way to improve security across a range of consumer products and promote an advanced state of security in critical applications.

Policy initiatives based on standards include the proposed consumer IoT device labelling program in the US to be based on NISTIR 8259A. Other examples include the proposed mandatory security requirements in the UK or the voluntary device labelling schemes in Singapore and Finland, all based on the ETSI standard. Policymakers can also help to ensure that mandatory requirements are harmonized and mutually recognized across regions to avoid fragmentation that would work against IoT innovation, interoperability, and security.

Standards and laws for minimum IoT device security baselines share many commonly recommended or required controls. The first three provisions of the ETSI standard for consumer IoT security are the most highly recommended by ETSI and make appearances in several national-level policies: **“no default passwords, implement a vulnerability disclosure policy, and keep software updated.”** In the US, State laws in California and Oregon both prohibit default passwords for IoT devices. Similarly, recommendations or requirements for updated software and use of secure communications protocols like HTTPS also frequently appear in standards and policies around the globe.

To demonstrate the effectiveness of commonly recommended controls, the Global Cyber Alliance (GCA) conducted research to test how well they prevent attacks.

Using dedicated honeypot infrastructure, it is possible to consider the effects of real Internet attacks on variably configured devices— such as those conforming, or not conforming, to measurable controls in standards and policy. This provides some insight into the effectiveness of the control in real world environments.

This analysis can be used to help policymakers and industry leaders understand the benefit of best practice approaches to IoT device security, and as proof points for manufacturers to adopt standards and comply with policy.

The result of the research strongly suggests that policymakers are correct in emphasizing secured access when turning standards into policy. The following items are strongly suggested as recommendations for IoT devices:

- **Implement SSH over Telnet.** Telnet is less secure than SSH due to Telnet communication being unencrypted. Attackers can sniff traffic and obtain usernames and passwords (regardless of the strength of the password).
- **Do not allow continuous use of default passwords.** Require immediate password changes during initial device setup.
- **Require the use of strong passwords along with SSH usage.**

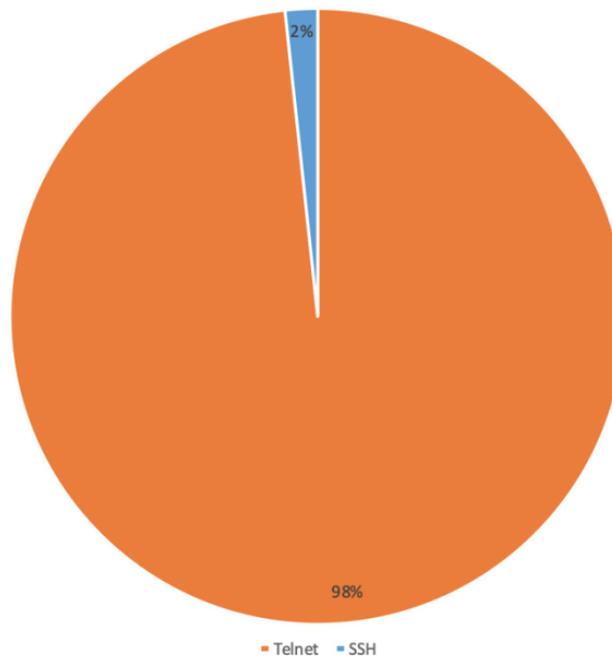
This report will go into further details around why the above recommendations need to be considered.

Overall Recommendations to Policymakers

Policymakers should emphasize secured access policies to mitigate Telnet-based Mirai attacks

GCA data revealed that, for the past 2.5 years, IoT attacks that have resulted in successful logins have been carried out over the Telnet protocol rather than the SSH protocol. This is clearly shown by the graph below:

AIDE Traffic by Protocol - Attacks with Logins



Our analysis has shown that Mirai is by far the largest source of Telnet-based attacks in the GCA platform. The length of our data collection and the geographical dispersion of our AIDE sensors has allowed us to see just how prevalent Mirai attacks are globally: they have appeared consistently throughout our 2.5 years of data collection, and across all geographical locations in which our hundreds of sensors are located. Many different variants of Mirai have been present.

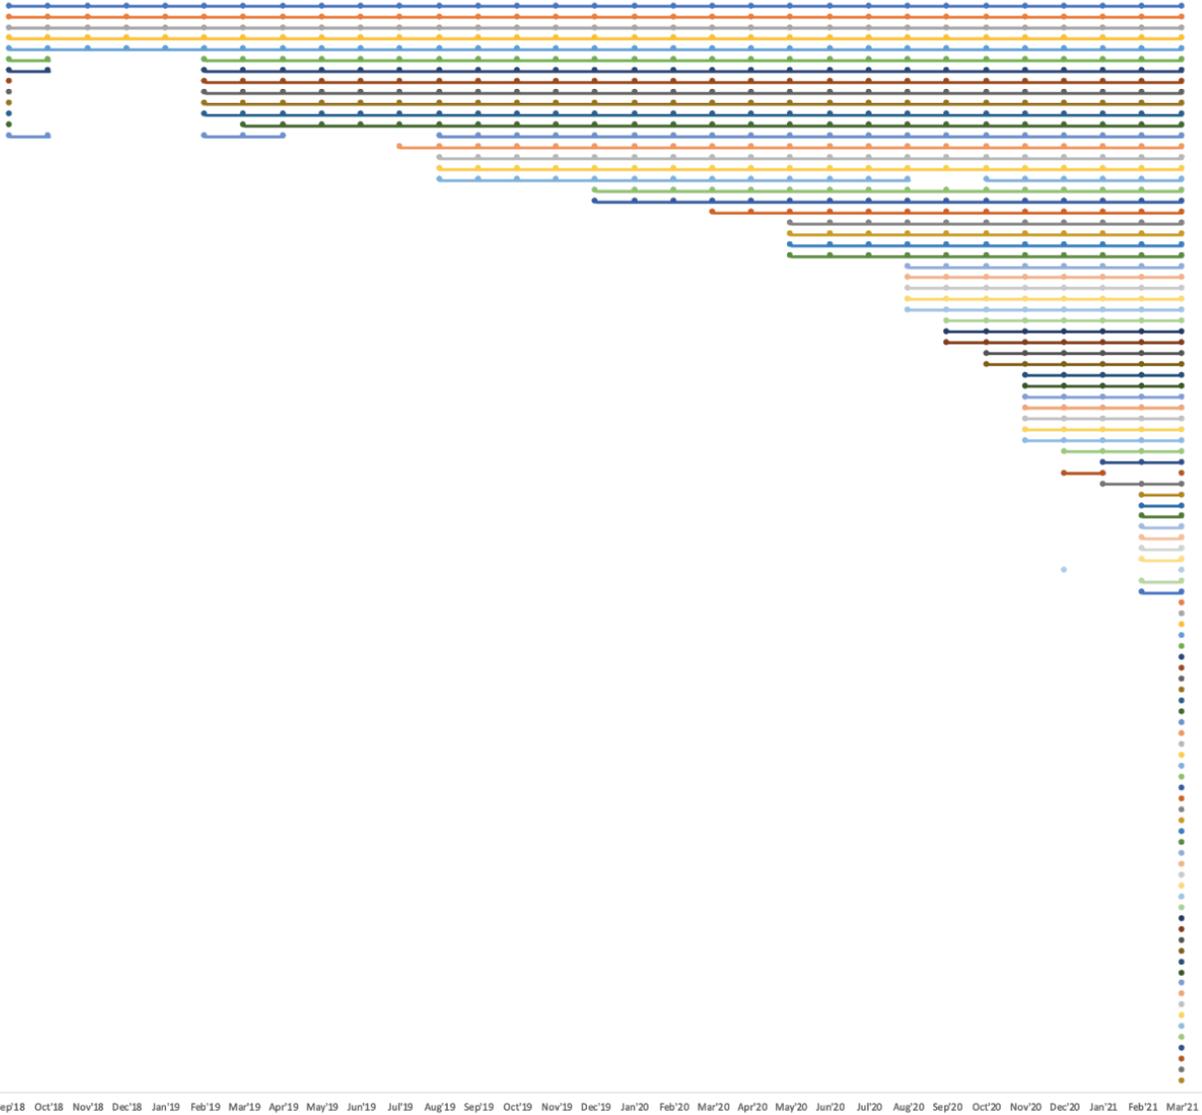
Since Mirai works by scanning for open Telnet ports and then brute-forcing the login credentials using common default credentials, policies which mandate secured access could potentially have a large effect upon decreasing the frequency of these Telnet-based attacks.

GCA's analysis of IP longevity suggests the importance of implementing policies that ensure secure access in order to prevent Mirai reinfection. This is because our data identifies a large number of IP addresses which have consistently been issuing high numbers of Mirai attacks every single month since AIDE data collection began. Since Mirai infection does not survive reboots, these IPs could represent either machines that have been running without reboots for a long time, or machines that have been reinfected.

Since reinfection is enabled by keeping weak or default passwords, maybe the sustained frequency of attacks from IPs as these could be disrupted if policies mandated secure access.

To understand whether policymakers should prioritize certain types of devices when crafting secured access policies, more research could be done on the types of devices within AIDE that are most often targeted by Mirai malware for brute-force credential attacks.

AIDE Data - Monthly Trends: Top IPs in March, 2021



As a secondary measure to secured access, policymakers could emphasize a shift away from Telnet

Because the majority of Mirai attacks within AIDE occurred via Telnet, it may be tempting to draw the conclusion that IoT policymakers should emphasize policies that do not allow Telnet to be enabled on systems— such as banning the use of Telnet, and not allowing devices to be sold which utilize Telnet.

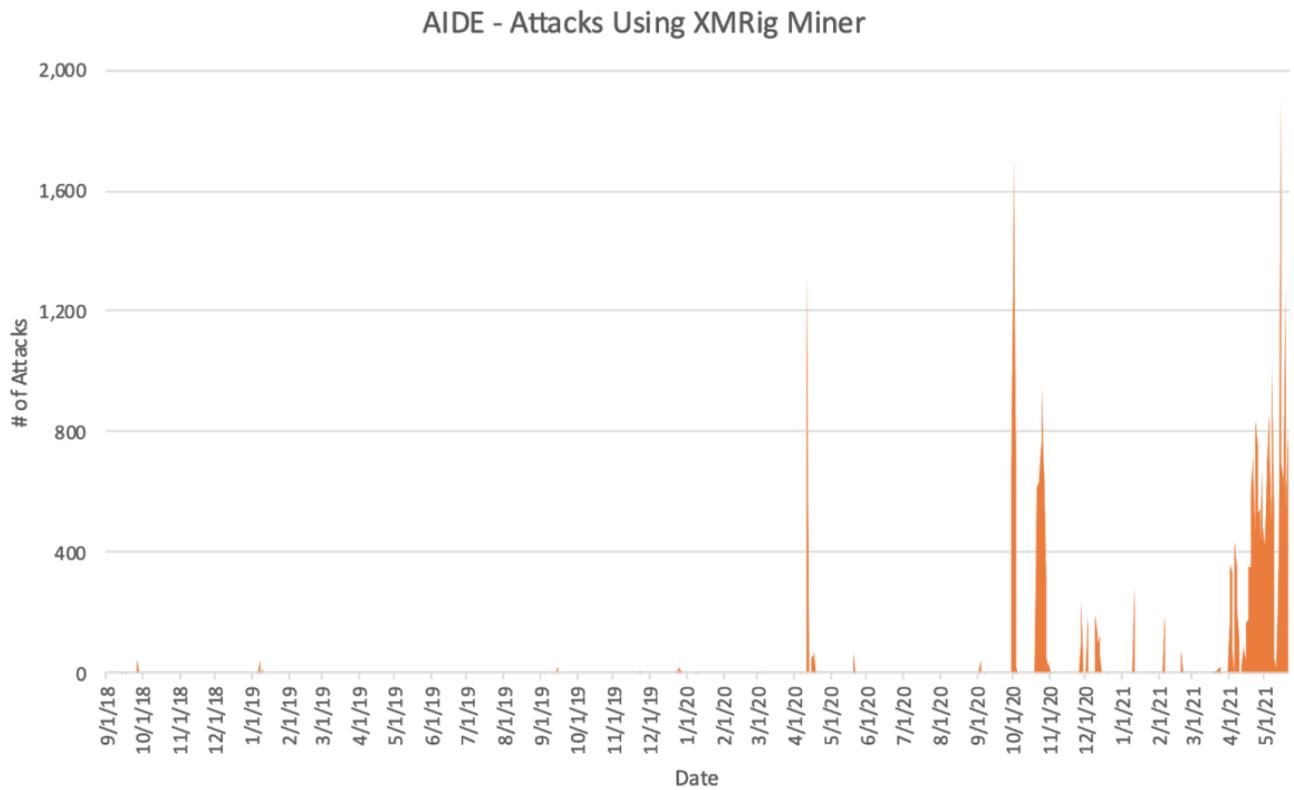
However, such recommendations should be secondary to the implementation of secured access, since devices with weak passwords are vulnerable to intrusion over both Telnet and SSH. Ensuring secured access will help secure devices no matter what protocol they use.

The AIDE open source honeypot is primarily configured to analyze attacks from Telnet, which is why we are picking up large amounts of information about Telnet-based Mirai attacks. Policymakers should work to better understand the comparisons between SSH and Telnet attacks before deciding to ban Telnet altogether.

SSH may be a more reliable alternative to Telnet, yet SSH attacks may be on the rise

SSH may be a more reliable alternative to Telnet because, unlike Telnet traffic, SSH traffic is encrypted. Because Telnet is an unencrypted protocol, data can be sniffed in transit, potentially allowing even strong passwords to be compromised.

Still, GCA research reveals that attacks on SSH-enabled devices may be on the rise. When GCA investigated the small percentage of meaningful SSH sessions, we discovered that some of them were attacks connected to cryptocurrency mining, which have been increasing since April 2020, as per the chart below. These cryptocurrency mining-related attacks have the potential to take up large amounts of bandwidth on various systems, resulting in out-of-pocket costs for businesses and consumers who pay for bandwidth, and/or reduced data speeds for consumers who work from home.



GCA's research on the prevalence of Telnet-based Mirai attacks suggests that policymakers may wish to promote SSH as a more secure protocol for IoT devices than Telnet. Yet the prevalence of these cryptocurrency attacks suggests that attackers are also shifting techniques. Thus, further research into SSH attacks is needed.

Methodology

ProxyPot Infrastructure and GCA Honeyfarm

GCA has developed home-grown honeypot technology, referred to as ProxyPot, designed to help alleviate the resource intensive task of creating, customizing, and managing environments to collect data on attacks to systems and devices on the internet.

The general intention in building a honeypot is to create an enticingly-accessible target for would-be attackers in order to learn as much as possible about the existence, source and prevalence of attacks. With its ease of configuration, ProxyPot is particularly well-suited for experimentation and hypotheses testing, such as conducting A/B tests with controlled and test examples.

A notable feature of ProxyPot is its ability to act as a pass-through (“proxy”) for traffic to and from devices. It supports three modes of emulation that allow a wide range of devices to be monitored and, therefore, make it ideal as an IoT proxypot: static, network, and virtual.

The static method is a simplistic emulation of the device accomplished by the use of manufactured static web pages representing the major functionality of the device. The network method allows for actual, physical devices to be monitored. The virtual method involves running the devices under a virtual operating system over a hypervisor.

ProxyPot captures web traffic in two distinct ways: first, the data exchanged between the attacking client and the ProxyPot (and the IoT device it is proxying for) is captured. This includes host and peer IPs, ports, timestamps, as well as any associated application traffic. Additionally, packet capture (PCAP) collects all application-level traffic between the attacker and the device.

The ProxyPot infrastructure was utilized to assess the effectiveness in reducing the risk of unauthorized attacks on IoT devices of several controls recommended by applicable standards. Leveraging the unique capabilities of ProxyPot outlined above, the research was conducted as a series of A/B tests:

- The “secured access” control (“no default passwords”) was tested with a honeynet comprising 70 honeypots emulating open source firewalls, network-attached storage (NAS) solutions, and operating systems commonly found in IoT devices. Half of the honeypots were deployed with default passwords and the other half were hardened with strong passwords.
- Additional research into the “patchability” control was conducted with a separate honeynet of 16 honeypots that compared patched/most recent vs. unpatched/older versions of IoT software.
- The “secured access” honeynet ran longer than the “patchability” honeynet, which was started after and ran concurrently with the “secured access” honeynet.

In addition to the highly configurable, quickly deployable ProxyPot infrastructure, GCA also operates a large honeyfarm of several hundred emulated IoT devices distributed worldwide.

The GCA honeyfarm was designed to provide at scale, long-term real attack data for the purpose of understanding trends and changes in IoT attack methodologies, and as an early alert system on new threats. It captures network traffic targeting console-based attacks (i.e., those over Telnet and SSH). This includes host and peer IPs, ports, timestamps, login credentials, console/shell commands, and hashes of files downloaded to the honeypots

The GCA honeyfarm has been running uninterrupted since mid 2018 and has amassed over 6TB of real attack data. The ProxyPot infrastructure and the GCA honeyfarm complement each other and together provide a best-in-class honeypot environment:

- ProxyPot captures web traffic (HTTP/HTTPS); the GCA honeyfarm captures Telnet and SSH traffic.
- ProxyPot is optimized for quick build-up and tear-down; the GCA honeyfarm is optimized for long-term operation.

- ProxyPot is designed to support different honeypot configurations operating simultaneously; the GCA honeyfarm is made of a large pool of identical honeypots.
- ProxyPot typically runs within a geography and network; the GCA honeyfarm was implemented to ensure a global footprint in both geography and networks.

Data from the GCA honeyfarm was utilized in this IoT security policy research project, particularly for the “data in transit is protected” control (“use secure communications protocols”).

The data collected by both the ProxyPot infrastructure and the GCA honeyfarm is fed back into GCA’s Automated IoT Defense Ecosystem (AIDE) framework, which consumes these feeds into an extract, transform, load (ETL) pipeline and is then fed via Kafka to Elasticsearch for indexing.

Findings

Secured access

The “secured access” control (“no default passwords”) was tested with a ProxyPot-controlled honeynet comprised of 70 honeypots emulating open source firewalls, network-attached storage (NAS) solutions, and operating systems commonly found in IoT devices: FreeNAS, OpenMediaVault, OpenWrt, pfSense, XigmaNAS, M0n0Wall, and SmallWall.

For each of the 7 emulations, 10 honeypots were deployed, 5 with default passwords and 5 hardened with strong passwords. On the unhardened devices, default passwords were chosen over even weaker credentials, such as root/root, admin/admin, root/<empty>..., to provide a realistic test environment. Notice that this test also included an element of a “patchability” control (“keep software updated”) by utilizing a mix of supported systems (FreeNAS, OpenMediaVault, OpenWrt, pfSense, and XigmaNAS) and unsupported systems (M0n0Wall and SmallWall).

As it is common in honeypot environments, the ProxyPot devices were discovered and started getting attack traffic within a few hours of being deployed. Traffic remained robust for the 2-month period of data collection, with a sustained base level of about 300 sessions per hour and multiple peaks of thousands of sessions per hour [Fig. 1].

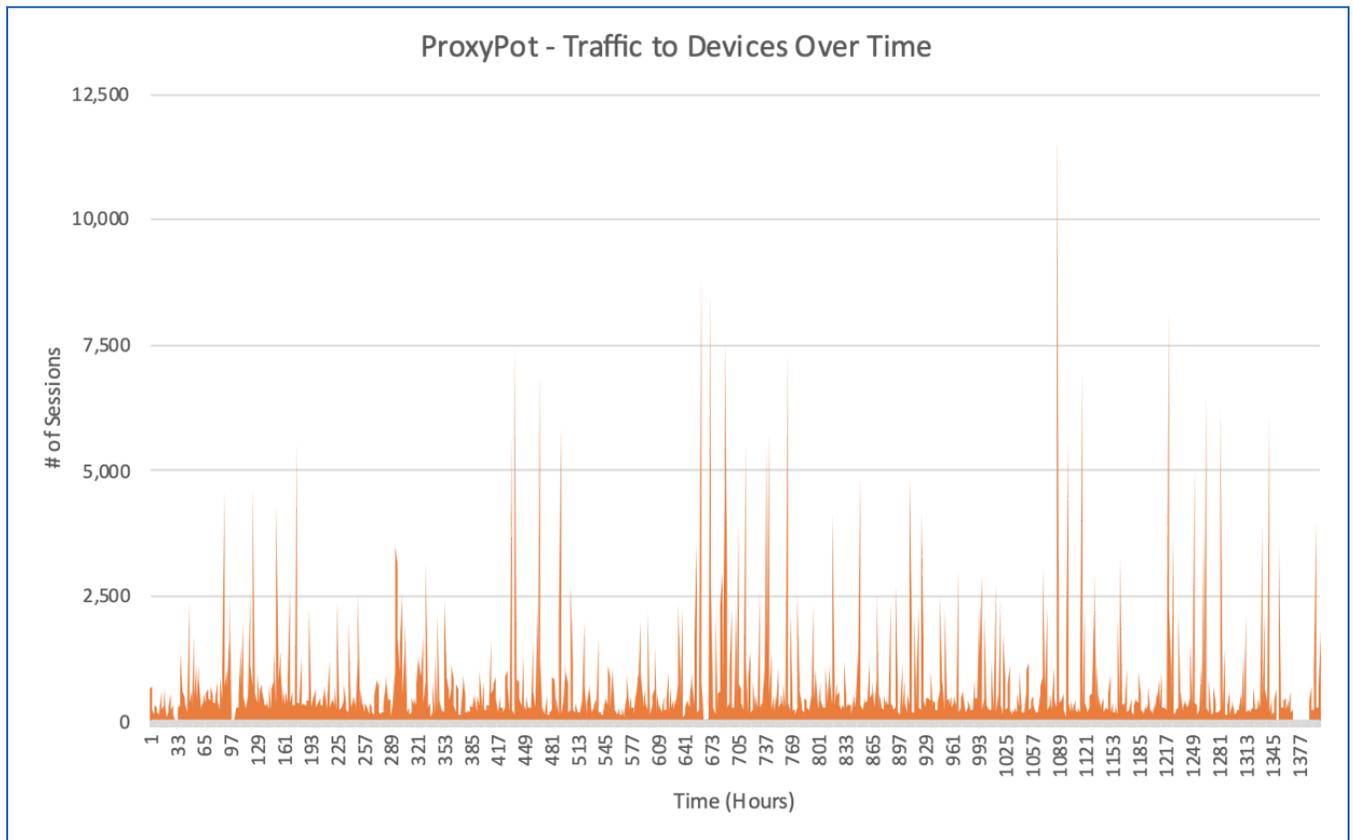


Fig. 1

Between April 5 and June 3, 2021, the system recorded 786,086 sessions, which resulted in 1,113,729 HTTP requests and 1,083,277 responses. A small number (6,432) of those sessions were legitimate scans by search bots. The remaining 779,654 sessions were classified as “attacks.”

The distribution of traffic by emulation and type (default password vs hardened password) indicates that some devices are preferred over others by attackers. In particular, FreeNAS and pfSense honeypots were attacked 2.1x more than SmallWall and M0n0Wall devices, which are the two discontinued systems [Fig. 2: connections = requests + responses].

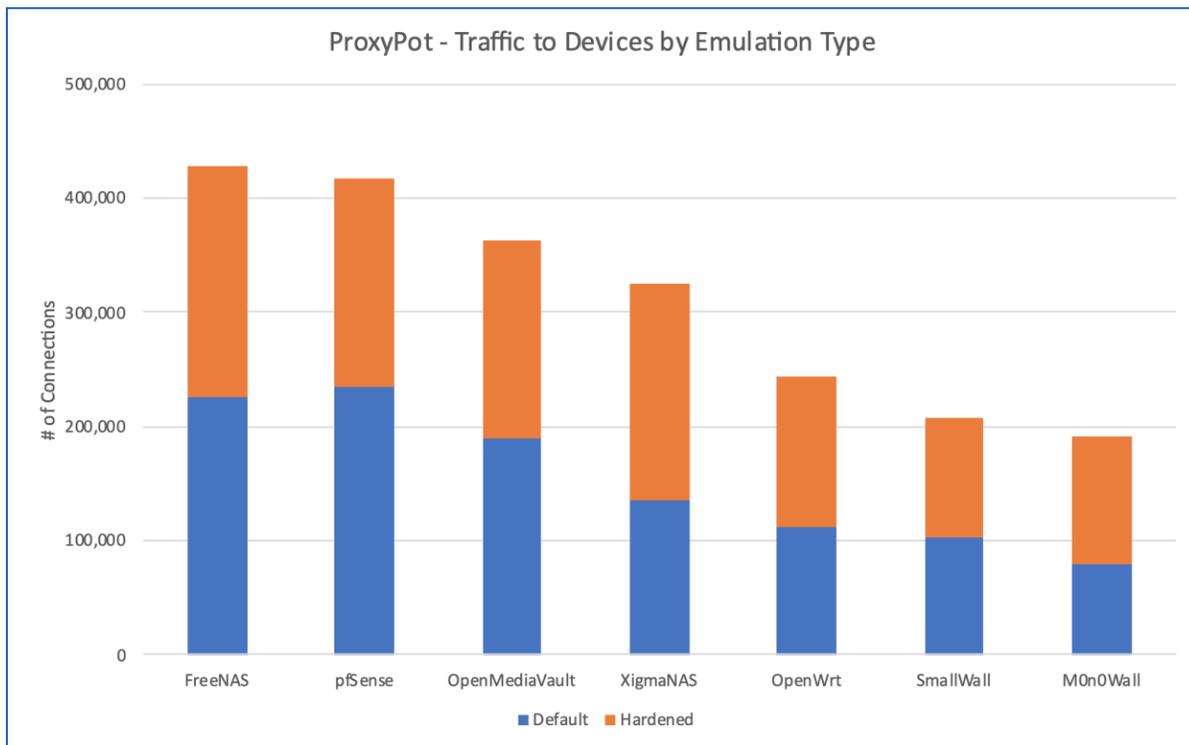


Fig. 2

It is believed that this finding indicates that attackers favor active systems, perhaps as a way to ensure a wider, more relevant population of attack targets.

Of the 779,654 attacks, a large number of the sessions (511,128, 66%) were intended as probes to try to learn something about the target devices (e.g., scans to detect open ports, presence of a particular software stack...) or to keep alive the communications between the attacking system and ProxyPot.

That leaves 268,526 (34%) “meaningful,” active attacks against the devices, which can be broken down into the following categories:

● ThinkPHP exploit attempts:	122,497	(46% of active attacks)
● SQL exploit attempts:	60,052	(22%)
● Login attempts:	47,218	(18%)
● WordPress exploit attempts:	20,428	(8%)
● Misc. botnets:	14,860	(6%)
● Mozi botnet:	2,803	(1%)
● Apache Axis2 exploit attempts:	590	
● GeoVision camera exploit attempts:	78	

The breakdown of attacks shows that attackers routinely target vulnerabilities in common ingredients of the software stack of web servers, such as PHP, SQL, WordPress, and Axis2. Although some of those ingredients are not common in embedded/IoT systems, attempts at using them by attackers is explained by the fact that all the ProxyPot devices included a web interface.

The attempts to exploit vulnerabilities in the GeoVision IP camera shows that the ProxyPot devices were misidentified by one attacker. Another notable finding is that the Mozi botnet, active since 2019 and peaking in the first half of 2020, when it accounted for 90% of IoT attacks, is still active.

Particularly relevant to this research is the analysis of the 47,218 login attempts. The first notable finding is that only 7,578 (16%) of those were attempts to log into the device. The majority (39,640 or 84%) were attempts to log into, again, two components of web servers: PHP and the Boa embedded web server. The default username/password combinations on the unhardened devices are as follows:

Device	Username	Password
FreeNAS	admin	freenas
pfSense	admin	pfsense
OpenWRT	root	<blank>
OpenMediaVault	admin	openmediavault
SmallWall	admin	small
MOn0Wall	admin	mono
XigmaNAS	admin	xigmanas

Although brute-forcing those credentials is an easy task—a point painfully made by the Mirai botnet (more on this in the **Findings, Data in transit is protected** section)—, it is hypothesized that attackers prefer vectors that target ubiquitous components of the software stack over device-specific techniques to maximize their return on investment.

The analysis of the 7,578 attempts to penetrate the devices supports the recommendation that changing default passwords is an effective way to reduce the risk of a successful attack. Although login attacks were launched against default and hardened devices in comparable numbers, they were successful in breaking into the devices 79 times, all of them on devices with default passwords. There were no successful attempts at breaking into devices with strong passwords [Fig. 3].

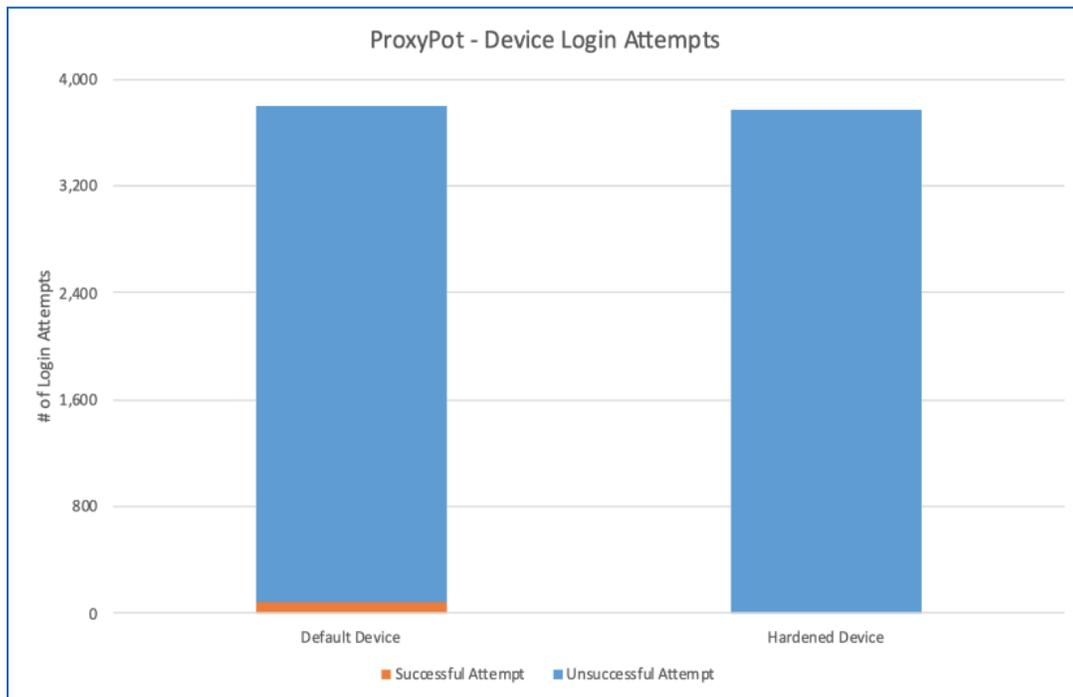


Fig. 3

Analysis of the successful login attempts on the unhardened devices shows that the OpenMediaVault, OpenWrt, pfSense, and SmallWall devices were cracked, whereas FreeNAS and MOnOWall not [Fig. 4]:

Device	Attempts	Cracked
FreeNAS	810	0
M0n0Wall	429	0
OpenMediaVault	529	20
OpenWrt	513	20
pfSense	636	28
SmallWall	595	8
XigmaNAS	214	3

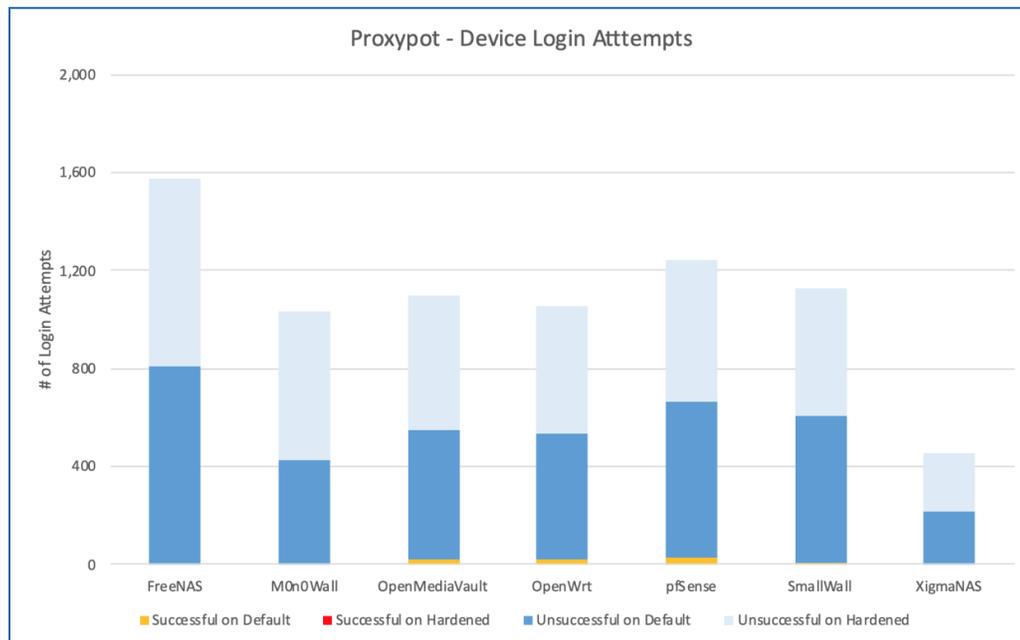


Fig. 4

The low number of login attempts on XigmaNAS compared to the other devices is somewhat of an anomaly that should be investigated.

Analysis of credentials

GCA conducted a comprehensive analysis of the credentials attempted on the large GCA honeyfarm for the almost 3 years of uninterrupted operation. In total, 2.6B username/password pairs were analyzed.

Unlike the ProxyPot honeypots, all 1,200 honeypots in the GCA honeyfarm were configured in a less-secure way to learn as much as possible about the nature of IoT attacks. As a result, the honeypots accept almost all passwords with username “root,” as well as a few other usernames. Although a very high number of credentials were used to break into the devices, the top-100 pairs accounted for over 90% of all logins. The results, which separate access over Telnet and SSH, are presented below:

	Telnet	SSH
Successful logins	870,566,082	1,735,606,873
Unique credentials	4,382,627	224,493
% attempts with top-100 credentials	93.9%	99.2%

Note that over SSH, a relatively small number of credentials was responsible for a very large number of logins. This was caused by two username/password combinations that were particularly effective: an empty password on the “root” account was responsible for 1,035,710,976 logins, and the “admin” password on the “root” accounted for an additional 569,775,879 logins [Fig. 5 and 6].

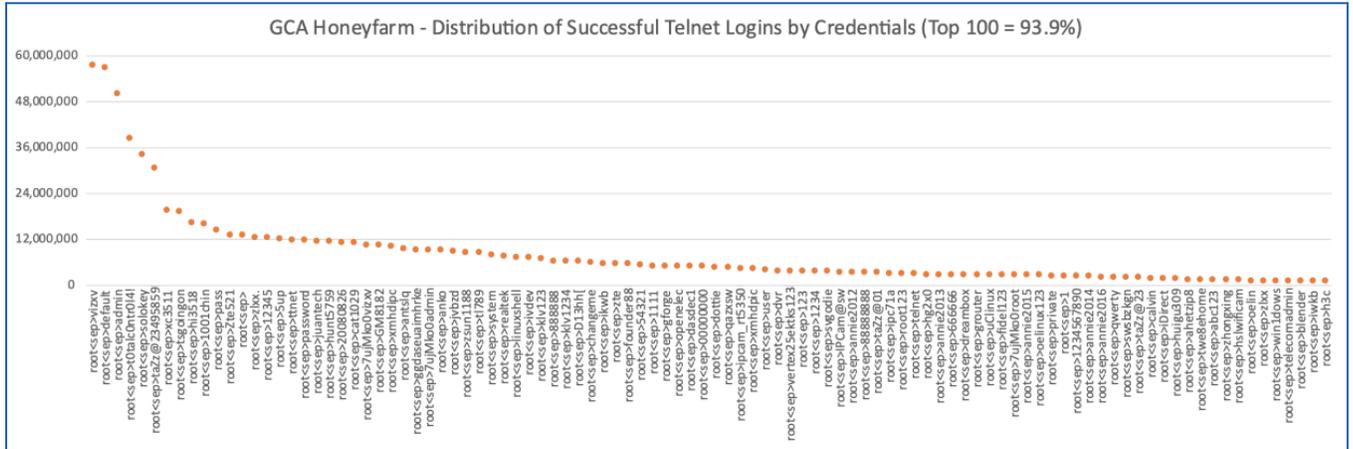


Fig. 5

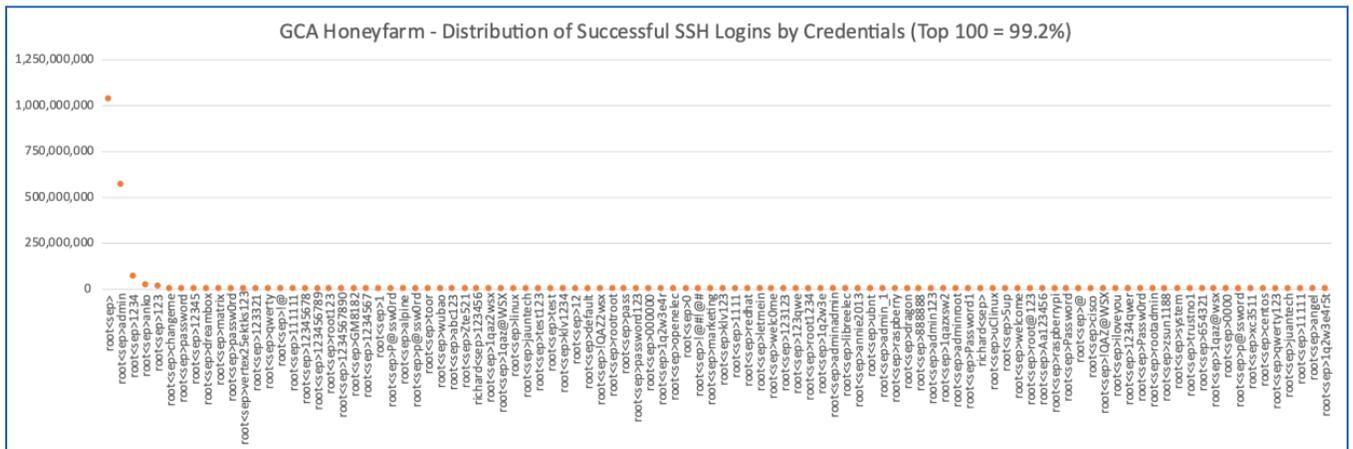


Fig. 6

The top-20 credentials by number of successful logins are given below in the format **username<sep>password**:

	Telnet	SSH
#1	root<sep>vizxv	root<sep>
#2	root<sep>default	root<sep>admin
#3	root<sep>admin	root<sep>1234
#4	root<sep>t0talc0ntr0l4!	root<sep>anko
#5	root<sep>solokey	root<sep>123
#6	root<sep>taZz@23495859	root<sep>changeme
#7	root<sep>xc3511	root<sep>password
#8	root<sep>tsgoingon	root<sep>12345
#9	root<sep>hi3518	root<sep>dreambox
#10	root<sep>1001chin	root<sep>matrix
#11	root<sep>pass	root<sep>passw0rd
#12	root<sep>Zte521	root<sep>vertex25ektk123
#13	root<sep>	root<sep>123321
#14	root<sep>zlxx.	root<sep>qwerty
#15	root<sep>12345	root<sep>!@
#16	root<sep>5up	root<sep>111111
#17	root<sep>ttnet	root<sep>12345678
#18	root<sep>password	root<sep>123456789
#19	root<sep>juantech	root<sep>root123
#20	root<sep>hunt5759	root<sep>1234567890

For at least the 18 credentials underlined above there are formal common vulnerabilities and exposures (CVEs) of the “default password for ‘root’ account” kind. This clearly shows the importance of changing default passwords, and how prevalent attacks on IoT devices using default passwords are.

Data in transit is protected

Traditional honeypot technology is not a particularly effective way to measure the cyber risk associated with the usage of non-secured communications protocols (Telnet, HTTP) and/or the impact of using secured/encrypted protocols (SSH, HTTPS). Techniques like network traffic interception and TLS fingerprinting—technologies beyond the capabilities of the GCA ProxyPot and honeypot infrastructure—are better approaches to measure the vulnerability of data in transit.

The previous caveat notwithstanding, the large corpus of real IoT attack data from the GCA honeypot was looked at for the purpose of gaining insights into the impact of protocol security on IoT device vulnerability. The GCA honeypot exposes honeypots with both the primary Telnet and SSH ports open, so it provides valuable information on attacker preferences over communications protocols.

As precursors to attacks, scans are performed by attackers on the honeypots to identify open ports through which to launch attacks [Fig. 7]. Scans on the GCA honeypot indicate to would-be attackers that both Telnet and SSH ports are open. What they do with this information is analyzed next. The spike in scans probing on port 22 (SSH) is a strong anomaly that should be further investigated.

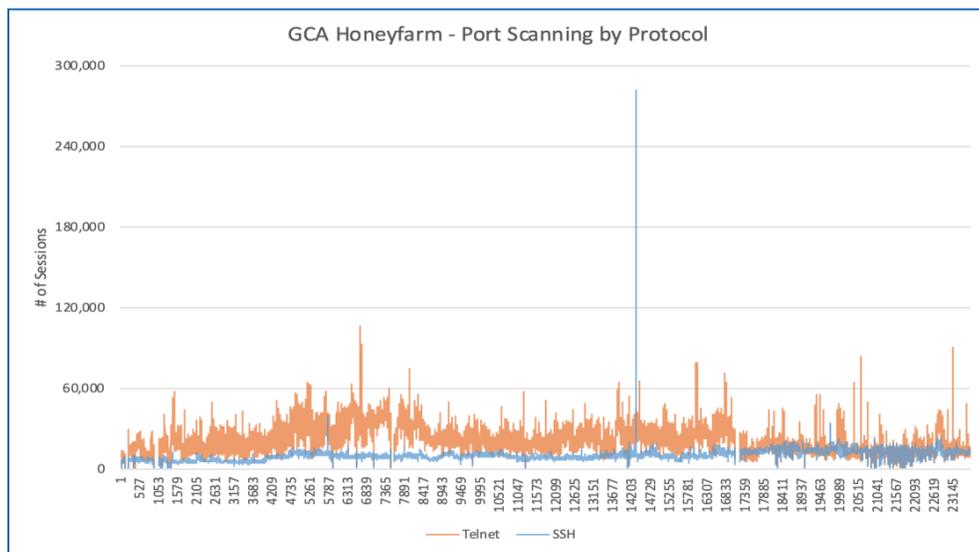


Fig. 7

Figure 8, which plots the volume of traffic over Telnet and SSH, shows that Telnet was the preferred protocol to launch attacks against the IoT honeypots for the first two years of honeyfarm operation. Those results are a reflection of the global IoT malware landscape, dominated since 2016 by the infamous, Telnet-based Mirai botnet and its many variants. Mirai attacks are by far the largest source of Telnet attacks on the GCA honeyfarm.

[Side note: The DDoS Mirai malware has a direct connection to the study of the secured access control, in that it infects target devices by conducting a brute-force access attempt through open Telnet ports using a small dictionary of 61 default username/password combinations from common IoT devices. Although the malware does not survive a system reboot, devices will get reinfected as long as the default password is not changed. The fact that after 5 years there still is a significant amount of Mirai activity is an indication of how widespread default passwords and open Telnet ports are.]

Mirai attacks have decreased during the last year, although flare-ups are still being observed. Again, this is reflected in the Telnet traffic observed on the GCA honeyfarm. Traffic over SSH ports has increased steadily since the start of the GCA operations.

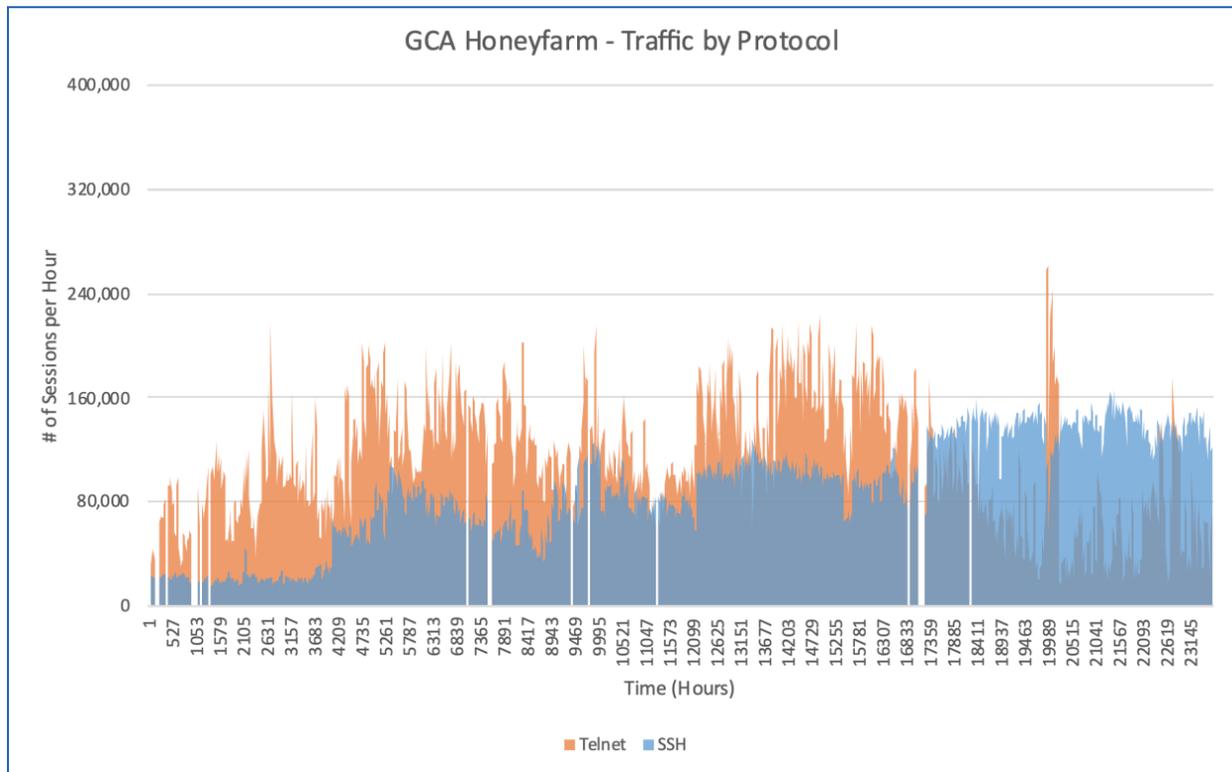


Fig. 8

Unlike Telnet traffic, which can be attributed almost entirely to the Mirai malware by analyzing the sequence of commands executed on the target devices and the hashes of the malware files downloaded, attribution of SSH traffic needs further research. Only a very small percentage of the SSH sessions leave attack fingerprints on the honeypots, such as commands or downloaded files [Fig. 9]. When present, those markers were analyzed and discovered activities like crypto currency mining and attempts to exploit known vulnerabilities in the software stack of Linux servers.

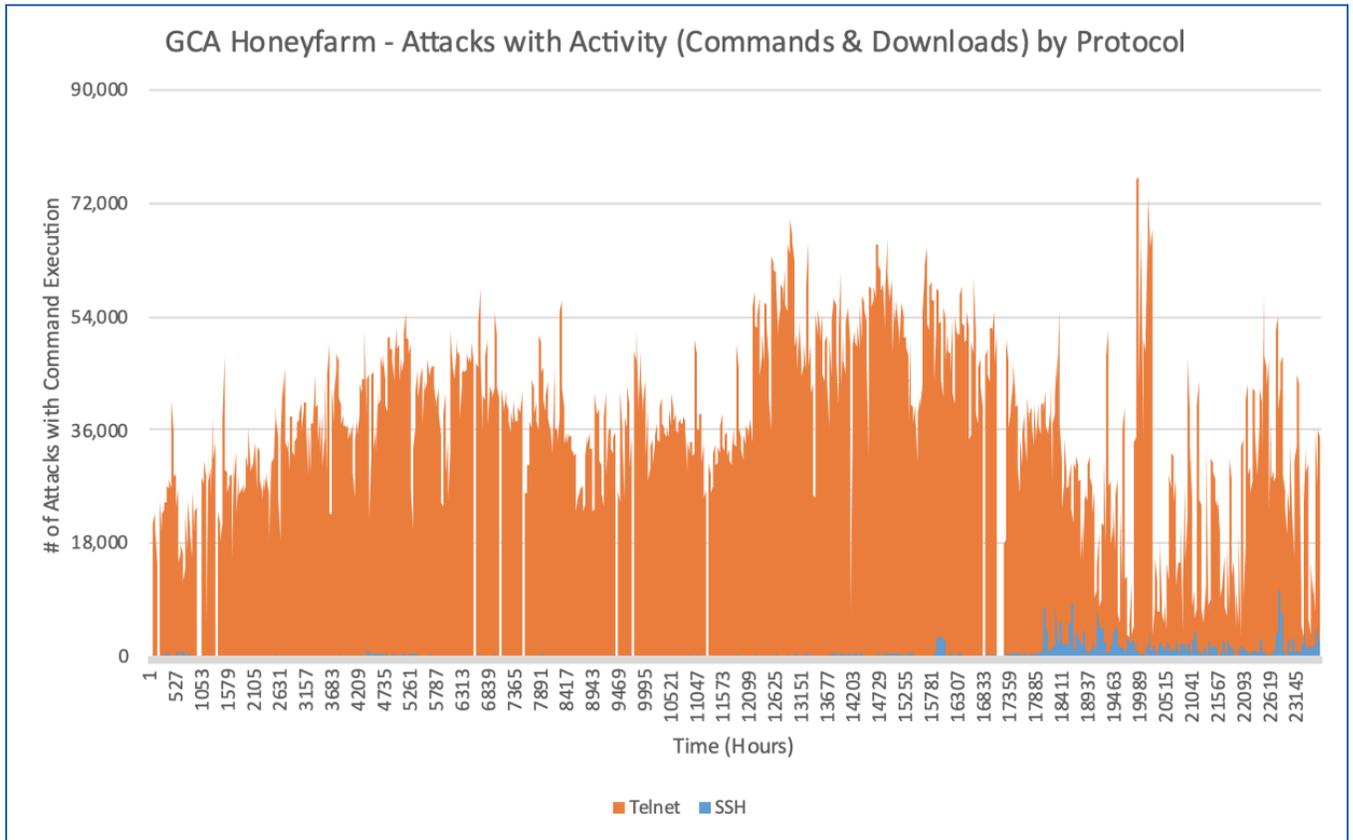


Fig. 9

Patchability

The “patchability” control (“keep software updated”) was tested with a ProxyPot-based honeynet comprised of 16 honeypots that compared patched/most recent vs. unpatched/older versions of IoT software: FreeNAS 11.3 and 8.0, OpenMediaVault 5.2 and 1.9, OpenWrt 19.07 and 12.09, XigmaNAS 12.2.0 and 9.30.

Two identical honeypots were deployed for each of the 8 different emulations, for a total of 16 honeypots, 8 of them patched and the other 8 unpatched. The SmallWall and M0n0Wall emulations used in the evaluation of the “secured access” control were dropped from the evaluation of the “patchability” control because those systems were discontinued and are no longer supported/patched. pfSense was also removed because its developer routinely removes older versions, a practice that was felt would impact the results of the experiment.

This honeynet ran for a shorter period of time than the “secured access” honeynet. Between May 29 and June 4, 2021, the honeynet recorded 23,737 sessions, which resulted in 33,865 HTTP requests and 33,355 responses. A very small number (31) of those sessions was confirmed to be legitimate scans by search bots. The remaining 33,834 sessions were classified as “attacks.”

The distribution of traffic by emulation and type (patched vs unpatched) indicates that there is a slight preference towards targeting the most recent configurations, which are the patched devices. This finding is consistent with a similar observation from the “secured access” exercise, where attackers were found to favor actively supported systems over discontinued systems. Again, it is believed that this pattern represents an attempt by attackers to secure a wider, more relevant set of targets.

The discrepancy in the volume of attacks to the four OpenWrt devices, something not observed in the “secured access” analysis, requires further investigation. It was confirmed that the emulations were correctly configured and that there were no unexpected downtimes. An investigation is currently being conducted with the aid of IoT search engines to determine if there is something in the configuration of

the OpenWrt emulations that may allow attackers to identify it as a honeypot and, therefore, a target not worth pursuing [Fig. 10: connections = requests + responses].

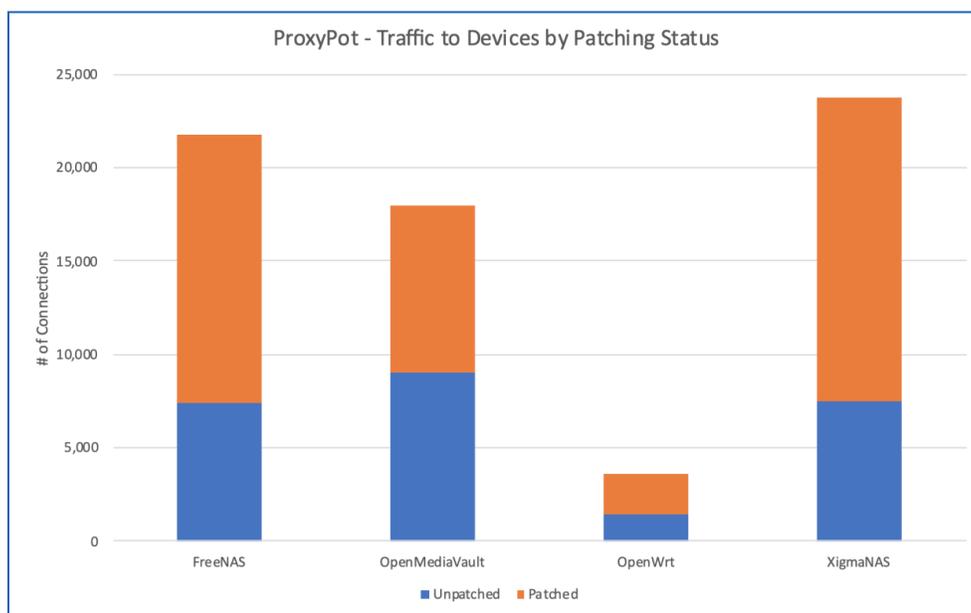


Fig. 10

Of the 33,834 attacks, 20,057 (59%) were exchanges of messages intended to keep alive the communications between the attacking system and the target device, detect open ports, identify specific items in the software stack... The remaining 13,777 (41%) “meaningful,” active attacks against the devices can be categorized as follows:

- **WordPress exploit attempts:** **5,661** (41% of active attacks)
- **Login attempts:** **3,638** (26%)
- **SQL exploit attempts:** **3,538** (26%)
- **ThinkPHP exploit attempts:** **402** (3%)
- **Misc. botnets:** **286** (3%)
- **Mozi botnet:** **243** (3%)
- **Apache Axis2 exploit attempts:** **9**

A comparison of the types of active attacks in the “patchability” control to those in the “secured access” control shows similarities and a few differences. WordPress and SQL exploits remain common attack vectors. Less botnet activity and fewer attempts at exploiting ThinkPHP vulnerabilities, which have a strong geographical component (i.e., mostly Asia), were observed.

Focusing on the 3,538 login attempts, the vast majority (3,591 or 99%) were attempts to log into PHP and the Boa embedded web server. Only 47 (1%) of the login attempts were on the actual devices. The analysis of those 47 attempts to penetrate the devices shows that more attempts were made on newer (patched) devices, but the only successful break-ins happened on unpatched devices [Fig. 11 and 12].

There are some indications that the attacks on unpatched devices were manual and the attacks on patched devices were scripted. That would indicate an interesting pattern of attack behavior. Given the small data sample, further research would be needed to confirm those observations.

Another thing that this brings to light is the scope of “patching.” Typically, keeping a device up to date is thought of in terms of the native (operating system) software of the device. Clearly, it is important to be able to update the application software running on devices, as shown by the prevalence of attacks on PHP and the Boa embedded web server.

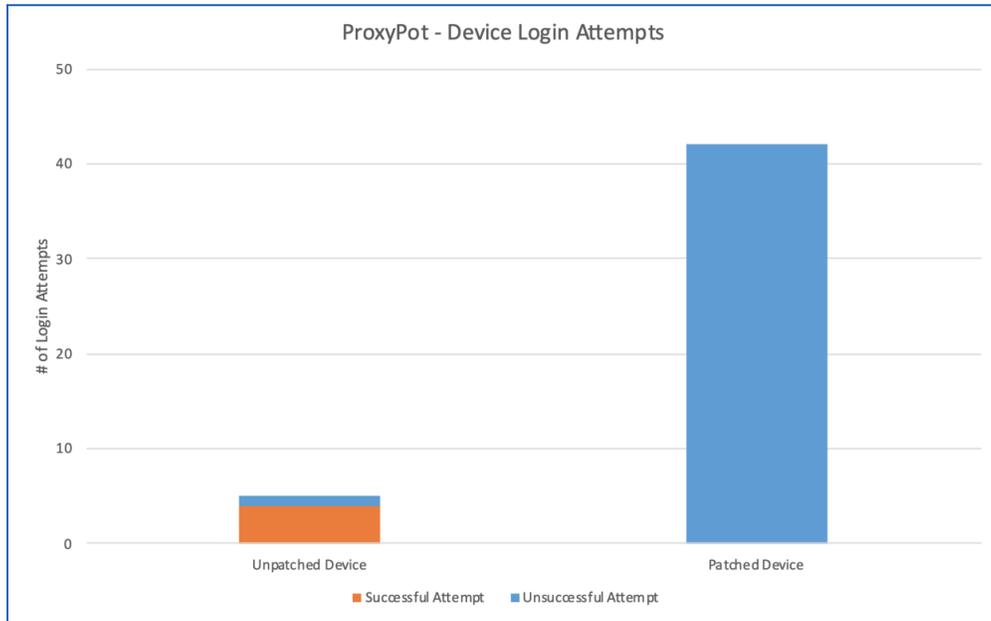


Fig. 11

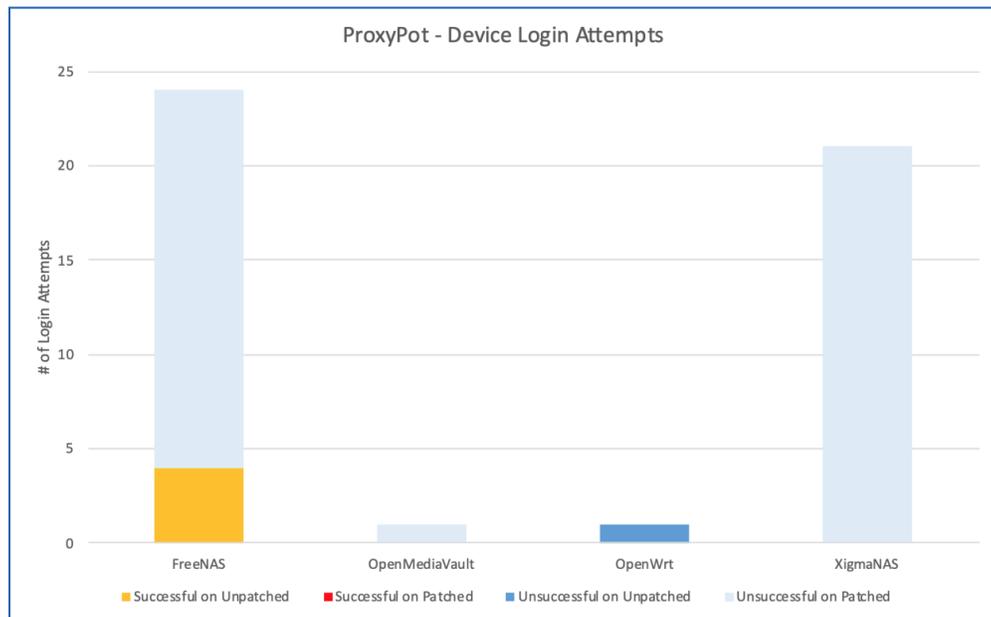


Fig. 12

Conclusions

GCA's analysis of real attack data shows that default passwords factory-set by device manufacturers and never changed by users, and weak passwords set by users together represent the biggest risk and is the most exploited vulnerability on IoT devices.

Adoption of coordinated international standards for IoT device security with mandatory requirements to prohibit default passwords and require strong authentication will have an immediate benefit in reducing device compromise and creating a safer, more trustworthy internet. Policy and regulatory frameworks can help drive adoption and harmonize implementation of the requirements in the standards.

In addition to vulnerabilities inherent in default and weak passwords, GCA's research also showed widespread prevalence of attempts to exploit security vulnerabilities in the software stacks of IoT devices. For example, attacks to try to exploit vulnerabilities in the Chinese-made ThinkPHP framework were extremely common. Although the GCA honeypots did not include devices vulnerable to these attacks, the observed behavior indicates that this type of attacks is a common, widely used vector whose risk of exploitation should be mitigated.

Inclusion in the standards of requirements for manufacturers to disclose vulnerabilities in their devices using common frameworks should be strongly considered. Examples of such common frameworks include The Common Vulnerabilities and Exposures (CVE) system maintained by the US National Security FFRDC and operated by MITRE, and the open industry standard Common Vulnerability Scoring System (CVSS). Additional requirements for manufacturers to disclose the support status of their products should also be considered.

Appendix: Geo IP Location

WARNING: The IP information provided in this section is only indicative. Many internet users concerned about privacy utilize anonymizing technologies, such as VPN, proxies and Tor, to enable anonymous communications on the Internet. For obvious reasons, cybercriminals and malicious actors also try to hide or obfuscate their identities and locations. Lookups on IPs assigned by anonymizers incorrectly interpret the IP of the intermediary servers as that of the original client. Even in the absence of anonymization, geolocation and autonomous system (AS) information is based on data obtained directly from the regional internet registries (RIR) and there are several cases where the information provided by the RIR might be misaligned with network reality— for example, IP addresses assigned to organizations with international presence are often geolocated incorrectly. Additionally, merger and acquisition of organizations is a key source of IP geolocation inaccuracies, and legacy address space (addresses allocated before the development of the RIR system for management) is often effectively untraceable .

MaxMind was used to perform lookups on peer (i.e., attacking) IP addresses to get an idea of the possible origins (geographical location and network) of the attacks.

Figure A1 shows the countries hosting the highest numbers of IPs launching attacks against the ProxyPot devices.

Figure A2 ranks the countries by the number of attacks launched from them. Countries near the top of the first list but not near the top of the second list have more attacking IPs, but these tend to be less prolific. Countries near the top of the second list but not near the top of the first list have fewer but very active attacking IPs.

Figure A3 shows the location of the attacking IPs on a map.

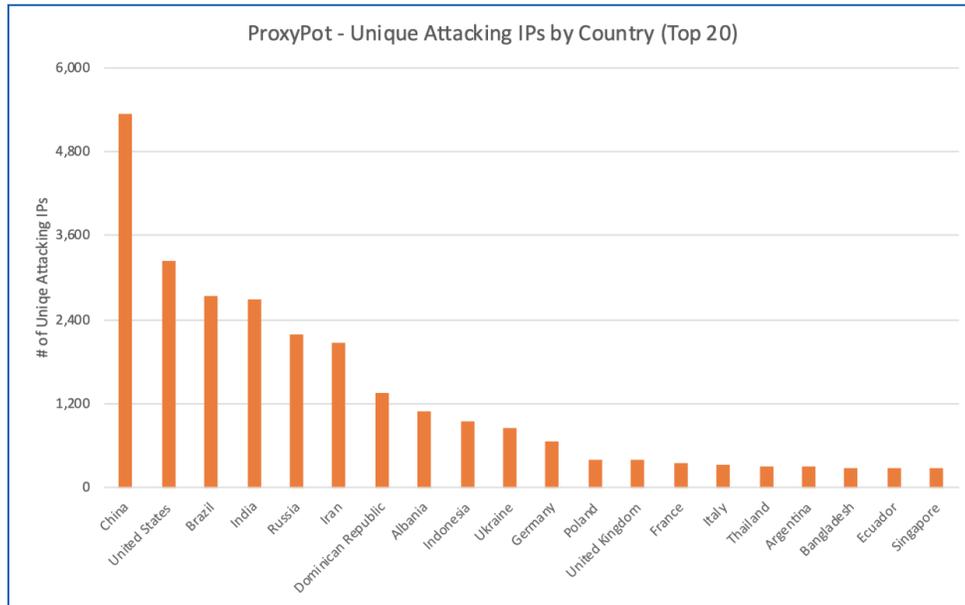


Fig. A1

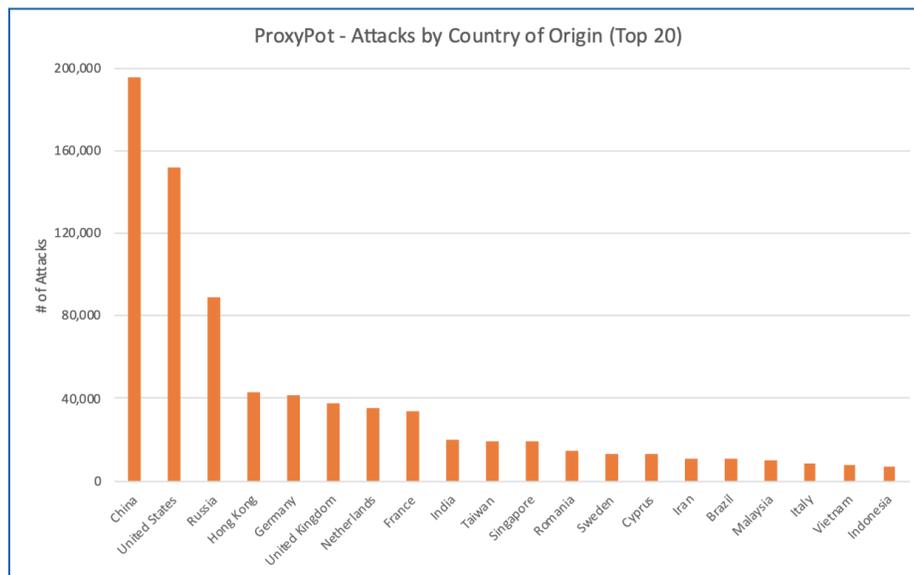


Fig. A2



Fig. A3

New York

731 Lexington Avenue
New York, NY 10022
UNITED STATES

London

City of London Police
3rd Floor, Guildhall Yard East
London, EC2V 5AE
UNITED KINGDOM

Brussels

Scotland House,
City Office in Brussels
(c/o Global Cyber Alliance)
Rond Point Schuman 6
1040 Brussels
BELGIUM

www.globalcyberalliance.org

Copyright 2021 Global Cyber Alliance