Journal of *Cybersecurity and Privacy*

MDPI

Article

Evolving IoT Botnet Threats and Practical Honeypot Observation: A Summary Review and Experimental Study

Rajkumar Banoth [1], Santosh Reddy Addula [2,*], Aruna Kranthi Godishala [3], Rithwik Sannapu [4], Guna Sekhar Sajja [2], Deepak Kumar [2], Vinay Kumar Kasula [2,*] and Chaitanya Tumma [2]

[1] Department of Computer Science, University of Texas at San Antonio, San Antonio, TX 78249, USA; rajkumar.banoth@utsa.edu
[2] Department of Information Technology, University of the Cumberlands, Williamsburg, KY 40769, USA; gsajja1524@ucumberlands.edu (G.S.S.); fdeepakkumar6146@ucumberlands.edu (D.K.); ctumma30994@ucumberlands.edu (C.T.)
[3] Faculty of Integrated Technologies, Universiti Brunei Darussalam, Gadong BE5101, Brunei; 21h8451@ubd.edu.bn
[4] Department of Information Systems, Saint Louis University, Saint Louis, MO 63103, USA; rithwik.sannapu@slu.edu
* Correspondence: saddula5840@ucumberlands.edu (S.R.A.); vkasula19501@ucumberlands.edu (V.K.K.)

Abstract

The rapid proliferation of Internet of Things (IoT) devices has significantly increased the attack surface for large-scale botnet operations. While previous research, including detailed analyses using Cowrie and IoTPOT frameworks, has studied IoT botnet behavior, these studies often rely on retrospective datasets, isolated protocol analyses, or hard-to-replicate setups. This paper addresses that gap with two main contributions: a structured review of ten influential IoT security studies from the USENIX Security Symposium and a confirmatory empirical experiment deploying Cowrie and IoTPOT honeypots simultaneously on a Microsoft Azure cloud-based virtual machine. Unlike earlier studies that focus on single protocols or large-scale environments, this work acts as a validation study, confirming well-known IoT botnet behaviors, including credential brute-force attacks, Mirai-style commands, and Telnet dominance, using real-time attack data collected from a reproducible, affordable cloud environment that simulates known IoT vulnerabilities (such as CVE-2016-10401, CVE-2017-17215, and CVE-2014-9222). Rather than revealing new attack methods, this study explicitly verifies the persistence of behaviors first documented almost ten years ago. The data indicates that attackers continue to exploit basic authentication flaws and reuse long-standing command sequences, confirming that core IoT vulnerabilities remain prevalent despite a decade of security research. It also highlights the ongoing gap between research progress and industry implementation. The analysis situates these findings within the broader evolution of IoT botnets, from early centralized command-and-control structures like Mirai to more resilient peer-to-peer networks that use anonymized channels and target high-wattage devices for power-grid manipulation. This study shows that small, cloud-based honeypots are valuable for continuous threat monitoring, model validation, and security assessments, providing a practical, reproducible approach for ongoing IoT security research.

Keywords: IoT botnet; honeypot; cybersecurity; Mirai; distributed denial of service; peer-to-peer botnet; command-and-control (C2); malware detection; network forensics; high-wattage attack; firmware security; Azure virtual machine

Academic Editors: Kamran Siddique, Ka Lok Man and Zahid Akhtar

Received: 16 March 2026
Revised: 18 April 2026
Accepted: 28 April 2026
Published: 2 May 2026

1. Introduction

The increasing interconnection of smart devices has expanded the functionality of modern networks but simultaneously exposed critical vulnerabilities within the Internet of Things (IoT) ecosystem. Weak default credentials, minimal firmware security, and limited user awareness have collectively enabled large-scale malware infections that weaponize IoT devices for coordinated cyberattacks. These compromised nodes form distributed botnets capable of launching denial-of-service campaigns, exfiltrating data, and conducting disruptive cyber-physical operations. The first major wave of such activity was exemplified by the Mirai botnet, which demonstrated how trivial authentication weaknesses could generate terabit-scale attacks targeting global infrastructure [1]. Subsequent research confirmed that Mirai's architectural simplicity became the foundation for numerous derivative strains and inspired new propagation models capable of evading conventional network defenses [2,3].

Over the years, IoT botnets have evolved from centralized command-and-control (C2) models toward resilient peer-to-peer infrastructures designed to sustain operations even after partial takedowns [4]. Large-scale measurements of home networks have shown that many consumer IoT devices frequently leave vulnerable services enabled by default. This creates a steady and renewable pool of targets for botnet reinfection [5]. Advanced studies, such as the "Circle of Life" analysis, detail the entire lifecycle of IoT malware, from initial exploitation to active monetization, highlighting automation in scanning and infection processes [6]. Meanwhile, high-wattage botnets, such as MaDIoT, have demonstrated the physical repercussions of cyberattacks, illustrating how compromised electrical devices can be manipulated to induce synchronized load fluctuations on power grids [3,4]. These findings underscore the progression of IoT botnets from merely causing bandwidth congestion to becoming potential instruments for cyber-physical disruption.

Detection and defense mechanisms have consequently become a central focus of research. Recent frameworks use programmable network switches and in-band telemetry to identify malicious traffic patterns in near real time [7]. Complementary approaches leverage deception technology through coordinated honeypot networks that engage attackers and record behavioral indicators for subsequent analysis [7,8]. Recent research on IoT security now includes RIS-assisted networks, focusing on secrecy-sensing optimization in full-duplex integrated systems to tackle new attack vectors in next-generation IoT infrastructures. Other research has focused on the user and organizational response layer; large-scale diagnostic services have been deployed to notify device owners of infected devices and assess the global prevalence of compromised IoT equipment [9]. In parallel, analyses of IoT market dynamics revealed that rapid consumer adoption of inexpensive connected products continues to outpace the integration of secure development practices [10]. The resulting imbalance between deployment scale and protection maturity reinforces the persistence of exploitable vulnerabilities.

Despite these advances, active threat behavior observed in open Internet environments remains an area that continues to benefit from continual validation. Academic studies often rely on controlled datasets or historical traces, leaving a gap in real-time field observations. To address this limitation, this paper combines a structured summary of major USENIX studies with an empirical honeypot experiment. Two low-interaction honeypots, namely Cowrie and IoTPOT, were deployed on a Microsoft Azure virtual machine running Ubuntu 22.04 LTS to capture contemporary attack attempts targeting common IoT protocols such as SSH and Telnet. The experiment records authentication failures, command execution patterns, and network-origin metadata, providing practical context to the theoretical and large-scale analyses found in prior work. By integrating literature insights with real-world data, this paper aims to evaluate the consistency of observed attack trends, underscore on-

going weaknesses in IoT device security, and emphasize the need for adaptive, distributed mitigation strategies.

This manuscript presents real-time data from live honeypots analyzing current IoT botnet behaviors. While scanning and attacking publicly accessible IPs are well known, our research's contribution lies in confirmatory validation: combining theoretical insights from prior literature with direct observation of real-world attack patterns to confirm the persistence of known behaviors, including credential brute-forcing, Mirai-style commands, and Telnet dominance, rather than discovering new attack mechanisms. We configured honeypots to emulate common IoT vulnerabilities documented in the CVE database, such as weak default credentials (CVE-2016-10401, CVE-2017-17215), insecure Telnet services (CVE-2014-9222), and SSH weaknesses found in consumer routers, IP cameras, and DVRs. By exposing standard IoT ports (SSH port 22, Telnet port 23) with default settings, our setup mirrors actual deployment scenarios that facilitate botnet spread. This data validates trends from earlier studies and offers insights into persistent attack behaviors. The two-week observation was intentionally brief to snapshot of current attacks. We recognize that this short period and the use of a single IPv4 address limit our ability to derive comprehensive conclusions about global attack patterns or changes over time. Nonetheless, our data confirms that fundamental vulnerabilities remain relevant and aligns with larger honeypot frameworks such as Cowrie and IoTPOT. Future work will extend the observation timeframe, deploy multiple geographically dispersed instances, and include a broader range of device types for more robust, statistically meaningful insights into IoT botnet evolution.

2. Related Work

2.1. Evolution of IoT Botnets

The study of IoT botnets began with the Mirai documentation, which exposed how trivial device weaknesses could enable Internet-scale compromise. Antonakakis et al. [1] reconstructed Mirai's full infection cycle, revealing its reliance on insecure default credentials, weak Telnet authentication, and the absence of patch management in consumer routers and cameras. Their dataset of global network telescopes showed how quickly Mirai infected heterogeneous devices, turning them into coordinated nodes in distributed denial-of-service (DDoS) attacks. This research established the methodological foundation for subsequent empirical analyses of IoT-borne malware. Kelly et al. [2] empirically extended this foundational work by subjecting consumer IoT devices to Mirai-style attack scenarios, demonstrating that targeted security interventions, such as credential replacement and firmware updates, can substantially mitigate device susceptibility to botnet infections.

Building on that, Soltan et al. [3] explored how a sufficiently large IoT botnet could manipulate energy consumption patterns. Through simulation of coordinated on-off cycles across thousands of compromised appliances, they showed that high-wattage loads could introduce instability into regional power grids. This work reframed botnets as potential cyber-physical threats rather than mere bandwidth consumers.

Subsequent research conducted by Alrawi et al. [6] in The Circle of Life expanded the view from a single incident to the broader malware ecosystem. Their longitudinal study traced the continuous reuse of Mirai-derived codebases, revealing a repeating life cycle of infection, monetization, takedown, and re-emergence under new branding. The persistence of identical propagation logic years after Mirai's disclosure demonstrated that vulnerability remediation in IoT deployments lags far behind disclosure timelines.

Complementary field observations by Kumar et al. [5] measured the prevalence of IoT devices in domestic networks and confirmed that insecure services remain widely exposed. They found that many consumer devices still broadcast management ports by

default, providing ongoing entry points for reinfection. Together, these early and mid-stage studies chart the historical trajectory of IoT botnets, from opportunistic scanning worms to resilient, economically motivated infrastructures.

2.2. Emergence of High-Wattage and Cloudless Botnets

As IoT adoption extended into industrial and energy domains, researchers identified a new class of "high-wattage" botnets capable of manipulating physical infrastructure. Shekari et al. [4] introduced the MaDIoT 2.0 framework, which models the simultaneous activation of power-intensive devices and their potential to trigger large-scale load fluctuations. Their results highlighted that a compromised fleet of high-consumption appliances could act as a distributed actuator network, blurring the distinction between cyberattacks and grid disruptions.

Parallel to these energy-centric concerns, Pan et al. [9] documented the rise of cloudless botnets in the TORCHLIGHT study. They observed adversaries increasingly hosting command-and-control (C2) servers within anonymizing systems such as the Tor network, allowing malware operations to persist without central coordination. This decentralization made conventional sinkholing and domain takedown techniques less effective. The work provided empirical evidence that attackers were deliberately migrating away from cloud-hosted infrastructure to evade platform-level defenses.

Together, these papers highlight the evolving topology of IoT botnets, from single-point C2 hierarchies to resilient, stealthy, and physically impactful networks.

2.3. Detection and Defensive Frameworks

As IoT botnets grew in scale and sophistication, researchers prioritized detection efficiency and defensive automation. Dong et al. [7] proposed HorusEye, a programmable-switch-based detection architecture that embeds anomaly analysis directly within the data plane of software-defined networks. By leveraging line-rate packet inspection, HorusEye detected abnormal IoT traffic with millisecond-level latency, outperforming traditional intrusion-detection systems that rely on centralized processing. The study demonstrated how network hardware programmability could serve as an early-warning system for botnet activity.

Complementing detection, Guan et al. [8] advanced the concept of coordinated cyber-physical deception. Their multi-honeypot framework synchronized several IoT decoys that collectively produced convincing sensor feedback and actuator responses. This approach led attackers to interact with synthetic targets while enabling researchers to harvest behavioral indicators and propagation scripts. Compared with isolated honeypots, the coordinated model captured richer contextual data about attacker decision-making and automated tool chains.

Pan et al. [9] explored remediation and user-notification efforts through a national-scale IoT Security Diagnostic Service. The system scanned networks for infected hosts, correlated findings with manufacturer data, and automatically notified device owners. Operational results showed that while notifications prompted temporary cleanup, reinfection occurred frequently when firmware updates were unavailable. This finding aligned with the "Circle of Life" observation [6] that botnet resurgence is cyclical.

Xing et al. [10] further connected security outcomes with economic incentives. By analyzing IoT device sales and privacy signal metadata, they concluded that lower-cost markets exhibit a strong negative correlation between price and security features. Their statistical evidence suggests that market pressures, not technical impossibility, drive persistent vulnerability. Hence, botnet mitigation must consider both regulatory and consumer awareness dimensions in addition to purely technical defenses.

2.4. Empirical Characterization and Lifecycle Studies

Empirical analyses have remained central to understanding IoT botnets in practice. Kumar et al. [5] performed one of the first household-level measurements, monitoring thousands of home networks to determine the density and diversity of connected devices. Their work provided evidence that a significant proportion of consumer routers, cameras, and media devices expose unnecessary services that become entry points for infection. Subsequently, Pan et al. [9] and Guan et al. [8] stated that, despite progress in awareness, these vulnerabilities persist due to outdated firmware and weak authentication protocols.

Recent diagnostic studies also highlight that while takedown campaigns can temporarily suppress botnet activity, long-term eradication remains unlikely without addressing underlying economic incentives. The consistent reappearance of repackaged Mirai variants confirms that current mitigation efforts have yet to resolve the core issue of insecure device design. These observations, drawn from longitudinal and large-scale data, set the foundation for the experimental work presented in this paper, which seeks to verify whether similar behaviors are observable in real-time traffic directed at controlled honeypots.

2.5. Transition to Contemporary Validation

The reviewed literature provides a comprehensive foundation for understanding the evolution, propagation, and defense techniques of IoT botnets. Across these ten USENIX papers, several patterns are evident: the persistence of credential-based infection vectors, the emergence of potential for cyber-physical manipulation, and the migration toward decentralized control structures. However, most existing studies depend on retrospective datasets or indirect telemetry, leaving a gap in real-time observation of current attacker behavior. To address this limitation, the present research introduces a targeted honeypot deployment that captures live interaction data from contemporary IoT threat actors. This experimental component, discussed in subsequent sections, aims to verify whether the behavioral characteristics documented in prior work remain consistent within today's Internet environment.

3. Analysis on Related Work

IoT botnet detection and mitigation have advanced considerably in recent years, showcasing both progress and ongoing challenges in tackling the growing threat from these malicious networks. Xing et al. [10] provided a comprehensive survey of botnet detection techniques, categorizing them into statistical, machine learning, and hybrid approaches. It provided an evaluation of each method's strengths and weaknesses in the context of IoT environments, where device heterogeneity and resource constraints complicated detection. Traditional detection methods often relied on analyzing network traffic patterns but struggled with the encrypted, distributed nature of modern botnets. Woodiss-Field et al. [11] examined the limitations of these methods, particularly in IoT botnets, where devices often used weak authentication protocols and could be easily compromised via default credentials. They also emphasized the challenge of detecting botnets that employ sophisticated evasion techniques, such as proxy servers and anonymized traffic. This underscored the necessity of innovative solutions that could adapt to the constantly evolving tactics of cybercriminals.

Recent research introduced more advanced detection and defense strategies, such as those proposed by Memos et al. [12], who focused on AI-powered IoT environments. They suggested that integrating AI could improve the identification of botnet activities by enabling dynamic pattern recognition and real-time anomaly detection. This aligned with the findings of Wazzan et al. [13], who recommended a combination of anomaly-based detection and behavioral analysis to address IoT botnet threats more effectively.

They stressed that while technical defenses like these were crucial, a more holistic approach involving market regulation and consumer awareness was needed to mitigate vulnerabilities in low-cost IoT devices. Furthermore, the work by Gelgi et al. [14] highlighted that despite advances in botnet detection, fundamental security flaws in IoT devices, such as inadequate firmware and outdated patches, continued to leave devices vulnerable to botnet infections. This was consistent with the findings of Asadi et al. [15], who argued that while detection methods had improved, the underlying issue of insecure device design remained a significant hurdle. In response, new strategies, such as honeypots, as discussed by Valentini et al. [16], emerged as a practical way to capture real-time attack data and analyze attacker behavior in a controlled environment. Building on these defenses for next-generation wireless systems, Illi et al. [17] demonstrate the optimization of secrecy sensing in RIS-assisted full-duplex integrated sensing and communication networks. They demonstrated how physical-layer security techniques can safeguard new IoT communication systems against eavesdropping and interference. Overall, these developments emphasize the ongoing need for adaptive, multi-layered strategies to counter IoT botnets efficiently.

In recent years, numerous studies have examined botnet detection and mitigation techniques. While numerous studies focus on theoretical models or simulation-based data, fewer address real-time empirical data collected from live honeypots. This study aims to fill that gap by capturing actual attack traffic targeting IoT devices and providing valuable insights into attack characteristics in real-world environments. In addition to academic publications, there are various technical reports and industry studies that examine IoT botnet behaviors and provide key insights into device vulnerabilities. These reports highlight issues such as outdated firmware and weak default credentials, which remain widespread across IoT devices.

The systematic review of related work, summarized in Table 1, provides a broader context for our empirical study, revealing the persistence of certain vulnerabilities (e.g., weak credentials and outdated firmware) and the need for adaptive defense mechanisms. Table 1 has been updated to include a "Validation Status in Current Study" column, which explicitly identifies which aspects of each prior study are confirmed by our honeypot results (e.g., credential brute-forcing, Telnet dominance, default account exploitation) and which aspects remain unobserved in our experiment (e.g., large-scale botnet infrastructure, peer-to-peer C2 structures, physical grid manipulation). This column directly connects the related work to the experimental evidence, strengthening the alignment between the literature review and the empirical findings.

Table 1. Summary of Key Related Studies on IoT Botnets.

Study	Research Focus	Methodology	Key Findings	Relevance to Current Study
Xing et al. (2021) [10]	Botnet detection techniques	Classification of detection methods (statistical, ML)	Evaluation of various detection approaches in IoT environments	Provides insights into IoT botnet detection techniques
Woodiss-Field et al. (2024) [11]	Traditional botnet detection on IoT devices	Review of traditional detection techniques	Emphasizes challenges with weak authentication in IoT botnets	Focuses on detection methods relevant to our honeypot study
Memos et al. (2026) [12]	Mitigation strategies in AI-powered IoT	AI-powered IoT environment defense architecture	Proposes an AI-based approach for botnet mitigation	Introduces innovative defense mechanisms for IoT botnets
Wazzan et al. (2021) [13]	IoT botnet detection approaches	Analysis and recommendations for future research	Highlights gaps in existing detection strategies	Offers recommendations for improving detection in real-time attacks

Table 1. *Cont.*

Study	Research Focus	Methodology	Key Findings	Relevance to Current Study
Gelgi et al. (2024) [14]	DDoS attack detection and IoT botnets	Systematic review of DDoS IoT attacks and defenses	Comprehensive review of botnet detection methods	Critical for understanding defense strategies against IoT DDoS attacks
Asadi et al. (2024) [15]	Evolution of botnet threats and defenses	Survey of evolving botnet threats and defensive strategies	Analyzes emerging IoT botnet threats and evolving defense techniques	Supports the study of emerging trends in botnet evolution
Valentini et al. (2026) [16]	Honeypot-embedded IoT hardware for real-world cyber threat collection	Low-cost IoT honeypot deployment to capture and analyze real-world attacks	Real-world attack data on IoT devices using embedded honeypots	Directly supports this study's honeypot approach for capturing real-world IoT attack telemetry

4. System Architecture and Design

The honeypot environment is deployed within a constrained Microsoft Azure subscription on a single Linux virtual machine running Ubuntu 22.04 LTS. Figure 1 shows the Microsoft Azure virtual machine running Ubuntu 22.04 LTS that hosts the Cowrie and IoTPOT honeypots used in this study (note: all personally identifiable information, public IP addresses, and sensitive infrastructure details have been redacted from screenshots to protect experimental integrity and comply with responsible disclosure practices). Network exposure is controlled by an Azure Network Security Group (NSG) attached to the VM's network interface.

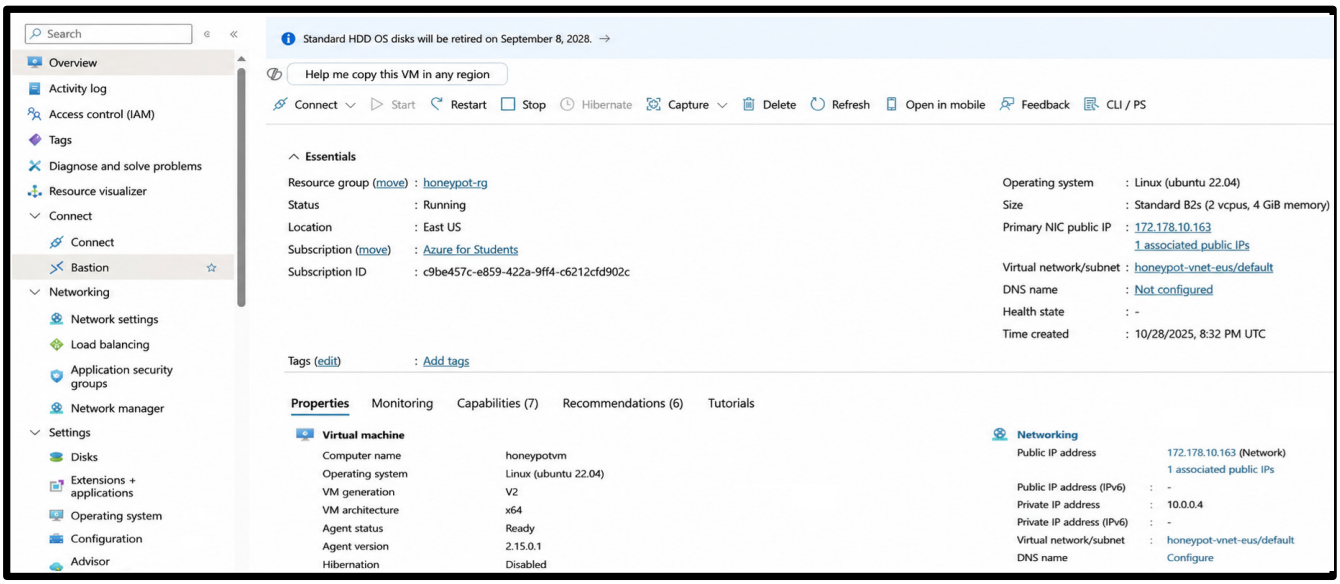

Figure 1. Azure Ubuntu Virtual Machine on which Honeypot is configured.

Figure 1 shows the Linux virtual machine used in this study, running Ubuntu 22.04 LTS. It has a public IPv4 address for direct Internet access and is the sole entry point for attacker traffic. The portal interface indicates the VM's association with an Azure Network Security Group (NSG) that controls network traffic. Sensitive details such as PII, public IPs, and infrastructure information are omitted for responsible disclosure. This VM serves as the core of the honeypot setup, hosting the Cowrie SSH honeypot on port 22 and the IoTPOT Telnet honeypot on port 23, as explained later.

22.04 LTS. This VM is assigned a public IPv4 address reachable from the global Internet and serves as the sole entry point for attacker traffic. Network exposure is controlled by an

Azure Network Security Group (NSG) attached to the VM's network interface. The NSG is configured to allow inbound TCP connections only on ports 22 and 23 (SSH and Telnet, respectively), with all other inbound ports denied. This configuration intentionally mirrors the minimal attack surface commonly observed on consumer IoT devices that expose a remote management interface while leaving other services closed. Outbound connections from the VM are restricted to the ports required for log forwarding to Azure Monitor to prevent the host from being abused for further attacks.

Within the VM, two low-interaction honeypots emulate vulnerable network services. Cowrie is bound to port 22 and presents an SSH interface that accepts arbitrary username and password combinations, records authentication attempts, and provides a simulated shell for post-login command capture. IoTPOT is bound to port 23 and emulates a generic Telnet interface like those found on embedded devices such as home routers and IP cameras. The two services share the same operating system but maintain separate log directories and configurations, which enables independent analysis of SSH and Telnet behaviors.

Telemetry collection is handled by the Azure Monitor Agent, which runs as a host process on the VM. Custom data collection rules specify the Cowrie and IoTPOT log paths and forward new entries to a Log Analytics Workspace.

Figure 2 shows the Azure Log Analytics Workspace used to collect, store, and analyze log data generated by the Cowrie and IoTPOT honeypots running on the virtual machine.

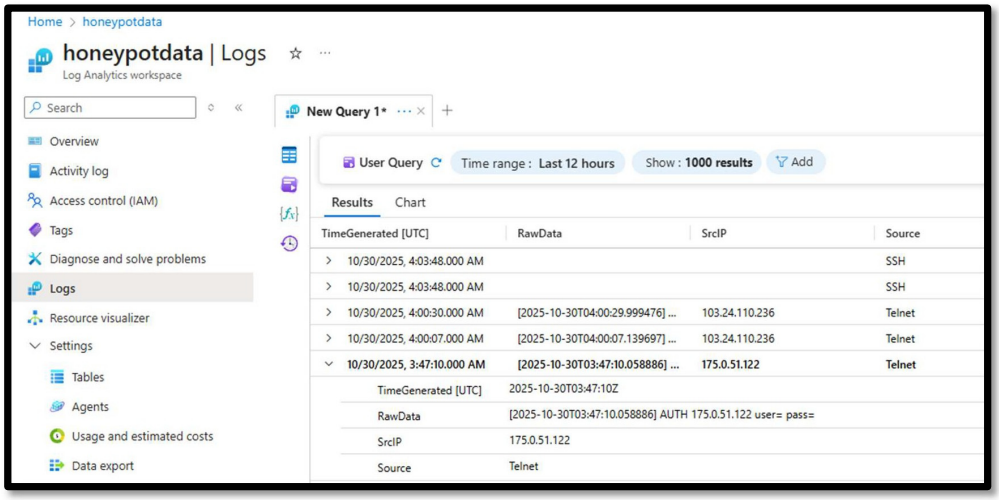

Figure 2. Log analytics workspace on Azure in which Data is being collected for the logs on the honeypot VM.

In the workspace, Cowrie events are stored in the cowrie_CL table, and IoTPOT events appear in the IoTPOT_CL table. Both tables preserve raw log content and workspace-level metadata, such as TimeGenerated, Computer, and ResourceId. This design provides a clean pipeline from attacker interaction at the network edge through to structured analytical data, eliminating the need for manual file exports from the VM.

On top of the workspace, a set of Kusto Query Language (KQL) queries drives a monitoring dashboard. The dashboard summarizes the top attacking IP addresses, source geolocation, distributions of usernames and passwords, command execution patterns, Telnet event types, and session duration statistics.

Figure 3 shows the overall architecture of the honeypot deployment, including the Azure cloud environment, the Ubuntu virtual machine, the Cowrie and IoTPOT services, and the centralized log collection pipeline. Inbound traffic targeting SSH and Telnet ports is allowed through Azure Network Security Group rules, while all captured events are forwarded to the Log Analytics Workspace for analysis.

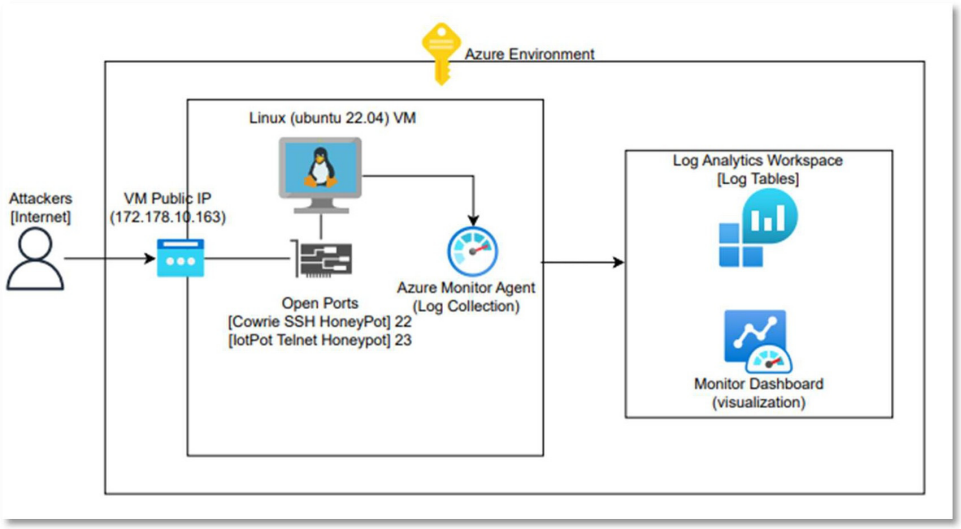

Figure 3. Architecture Diagram of HoneyPot setup.

These visualizations enable the data-driven analysis presented in the following sections, while remaining consistent with the architectural model used in prior IoT botnet measurement studies [1–5].

5. Experimental Design and Implementation

The goal of the experiment is to observe whether contemporary attack activity against exposed SSH and Telnet services exhibits the same behavioral patterns reported in earlier IoT botnet research [1–10]. The observation is designed as a validation study that focuses on qualitative similarity rather than large-scale enumeration.

The honeypot VM runs continuously for roughly two weeks. During this period, the NSG remains static, with ports 22 and 23 open to the Internet, and no additional hardening is applied beyond the isolation already provided by Azure networking. Cowrie and IoTPOT are configured with default banners and prompt strings to resemble generic Linux and embedded device shells. No active advertising, scanning, or seeding of the honeypots is performed; all interactions originate from external entities that independently discover the exposed services.

Logging is implemented at two levels. At the application level, Cowrie records each connection, authentication attempt, and executed command in structured JSON logs, while IoTPOT logs Telnet events with fields capturing the event type, credentials, and associated source addresses. At the platform level, the Azure Monitor Agent continuously ships these logs into the Log Analytics Workspace. This approach enables queries to be written directly in KQL, eliminating the need for ad hoc parsing on the VM.

For the analysis, the raw logs are transformed into higher-level features. Connection metadata is used to compute per-address attempt counts and to resolve countries using GeoIP data. Authentication events are aggregated into frequency tables of username and password guesses for both SSH and Telnet. Command input fields from Cowrie are grouped into distinct command strings, along with their corresponding hit counts. For IoTPOT, Telnet events are categorized as connection attempts, authentication events, command events, and error conditions. Long-running Cowrie sessions are identified by computing the elapsed time between the first and last event associated with a given session identifier.

The resulting feature sets provide a compact summary of attacker behavior across several dimensions: traffic origin, credential dictionaries, command usage, delivery of external payloads, and interaction style. These dimensions are then compared with the

findings of prior IoT botnet studies to determine the degree of continuity between historical and current attack patterns.

6. Results and Analysis

6.1. Cowrie SSH Honeypot

6.1.1. Source Address and Geolocation

Figure 4 shows an overview of the interaction data collected by the Cowrie SSH honeypot, including connection attempts, authentication events, and executed command patterns.

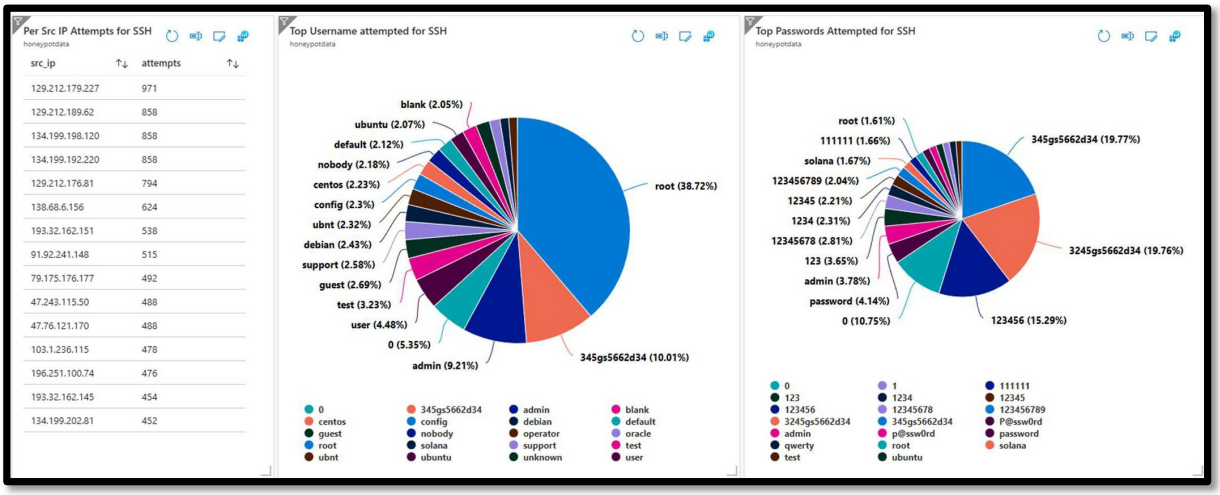

Figure 4. Data for Cowrie.

To provide a structured scientific analysis, this section is organized around three explicit analytical goals derived from the research literature: (1) comparing SSH versus Telnet attack intensity, credential reuse patterns, and command diversity; (2) measuring credential concentration by computing the proportion of attempts attributable to the top username and password tokens; and (3) evaluating behavioral consistency with prior IoT botnet studies by comparing our observations against documented attack signatures. These goals are addressed systematically in Sections 6.1–6.3. A side-by-side comparative analysis of SSH and Telnet behaviors is presented in Section 6.4. The Cowrie honeypot received repeated authentication and command attempts from a relatively small number of highly active IP addresses. The most aggressive source generated just under 1000 attempts, and several others issued between 400 and 900 each. This pattern is consistent with automated scanners or infected bots repeatedly cycling through credential lists rather than many unique human operators.

The Cowrie sources' geolocation reveals a wide geographic footprint, with a pronounced skew toward a few countries. Figure 5 shows the geographic distribution of source IP addresses observed by the Cowrie SSH honeypot during the measurement period. Addresses mapped to the United States accounted for the largest share of activity, with roughly eight to nine thousand attempts. China contributed a second significant cluster, with around 4000 to 5000 attempts. Hong Kong, the Netherlands, India, Romania, and Russia each contributed on the order of 1000 to 2000 attempts, with smaller volumes observed from South Korea, Singapore, Iran, Germany, Vietnam, Malaysia, Brazil, Taiwan, and other locations. This distribution is similar to the concentration patterns observed in large-scale IoT malware studies [5,9], which reported recurring participation by hosting providers in North America, East Asia, and Europe.

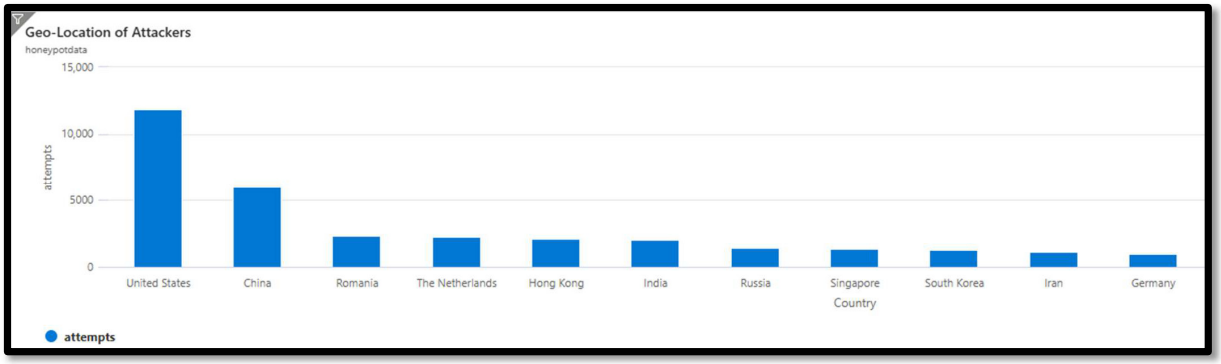

Figure 5. Geolocation data Cowrie.

6.1.2. Credential Probing Behavior

The SSH authentication logs strongly confirm that attackers continue to rely on trivial or default credentials. Figure 6 shows representative raw log entries captured by the Cowrie SSH honeypot, illustrating source IP addresses and attempted username–password combinations used during authentication. This figure demonstrates the systematic nature of botnet reconnaissance, where each log entry represents an automated brute-force attempt from geographically distributed IP addresses. The data reveals that attackers employ credential dictionaries derived from known IoT device defaults, with common patterns including "root/root," "admin/admin," and numeric sequences. The timestamp progression in the logs shows the relentless, round-the-clock nature of these scanning campaigns, confirming that IoT devices with publicly accessible IP addresses are continuously probed for authentication regardless of geographic location or time zone. This empirical evidence validates the theoretical models of botnet propagation presented in the literature review and underscores the continued importance of default-credential persistence as the primary infection vector for IoT botnets.

```
root@honeypotvm:~# sudo jq -r 'select(.eventid=="cowrie.login.failed" or .eventid=="cowrie.login.succ
) \t \(.src_ip) \t \(.username) \t \(.password)"' /opt/cowrie-src/var/log/cowrie/cowrie.json | tail -
2025-10-29T04:28:19.112039Z     68.183.9.147    root    qwerty
2025-10-29T04:28:53.764387Z     68.183.9.147    root    letmein
2025-10-29T04:29:26.666180Z     68.183.9.147    root    Password1
2025-10-29T04:29:59.505172Z     68.183.9.147    root    123123
2025-10-29T04:30:34.384671Z     68.183.9.147    root    111111
2025-10-29T04:31:09.085412Z     68.183.9.147    root    default
2025-10-29T04:31:42.384911Z     68.183.9.147    root    system
2025-10-29T04:32:17.375593Z     68.183.9.147    root    root
2025-10-29T04:32:51.850173Z     68.183.9.147    admin   123456
2025-10-29T04:33:25.419602Z     68.183.9.147    admin   password
2025-10-29T04:33:58.736271Z     68.183.9.147    admin   admin
2025-10-29T04:34:32.192274Z     68.183.9.147    admin   admin123
2025-10-29T04:35:06.198927Z     68.183.9.147    admin   12345
2025-10-29T04:35:41.122157Z     68.183.9.147    admin   123456789
2025-10-29T04:36:14.336936Z     68.183.9.147    admin   passw0rd
2025-10-29T04:36:16.150155Z     66.154.119.223          admin   admin
2025-10-29T04:36:16.701379Z     213.209.143.51          admin   admin
2025-10-29T04:36:47.490218Z     68.183.9.147    admin   12345678
2025-10-29T04:37:19.557799Z     68.183.9.147    admin   Administrator
2025-10-29T04:37:52.688638Z     68.183.9.147    admin   1234
2025-10-29T04:56:36.393201Z     139.19.117.131          root    null
2025-10-29T04:56:36.495296Z     139.19.117.131          root    null
2025-10-29T05:37:04.797156Z     196.251.114.15          admin   1234
2025-10-29T05:37:07.176162Z     196.251.114.15          debug   admin
2025-10-29T05:37:08.715315Z     196.251.114.15          admin   12345
```

Figure 6. Raw logs for Cowrie SSH showing source IP and username–password combinations.

The most common username guess was "root," accounting for roughly 5900 attempts and about one-third of all username trials. The next most frequent entries were an opaque string of the form 345gs5662d34 and the generic administrative account 'admin', each representing approximately 10 percent of the total. Additional popular usernames included 0, user, test, support, ubnt, debian, guest, centos, and ubuntu, along with configuration-related labels such as config, nobody, and default. These names align closely with lists extracted from Mirai and its descendants [1,5].

Password guesses exhibited a similar pattern. Two closely related strings, 345gs5662d34 and 3245gs5662d34, dominated the distribution with over 1600 attempts each, followed by common passwords such as 123456, 0, password, admin, 123, and several incremental numeric sequences. The overlap between username and password dictionaries suggests that many attackers simply reuse the same token on both sides of the credential pair. This behavior aligns with previous observations that IoT malware authors optimize for breadth rather than the strength of their guesses by testing short, recurring patterns [1,3].

Figure 6 illustrates the actual raw log entries documented by the Cowrie SSH honeypot during the observation period. Each log entry corresponds to a single authentication attempt and includes four essential fields: (1) a UTC timestamp indicating the precise moment of the connection attempt, (2) the source IP address of the attacking host, (3) the username employed by the attacker, and (4) the password submitted alongside that username. These entries demonstrate the automated, methodical nature of IoT botnet credential attacks. The progression of timestamps suggests that probing activity is continuous throughout the day and night, with no pause, irrespective of time zone differences. The source IP addresses originate from various countries and autonomous systems, aligning with the characteristics of a distributed botnet conducting scans from geographically dispersed infected nodes. The username–password combinations directly mirror well-known default credentials for IoT devices (e.g., root/root, admin/admin), which are hardcoded in Mirai and its derivatives. In summary, Figure 6 provides tangible empirical evidence of brute-force credential spraying, which constitutes the primary initial infection vector for IoT botnets, as documented in previous scholarly work [1,5].

6.1.3. Commands and Payload Delivery

Once the honeypot accepted a simulated login, attackers issued a small set of recurring command sequences. Figure 7 shows the most frequently executed commands observed in sessions captured by the Cowrie SSH honeypot during the experiment.


Command Execution Attempts
honeypotdata

input	hits
uname -s -v -n -r -m	8.417 K
/bin/./uname -s -v -n -r -m	4.151 K
cd ~; chattr -ia .ssh; lockr -ia .ssh	2.395 K
cd ~ && rm -rf .ssh && mkdir .ssh && echo \	2.379 K
cat /proc/uptime 2 > /dev/null | cut -d. -f1	1.38 K
uname -s -v -n -m 2 > /dev/null	1.38 K
export PATH=/usr/local/sbin:/usr/local/bin:/usr/sbin:/usr/bin:/sbin:/bin:$PATH; uname=$(uname -s -v -n -m 2>/dev/null); arch=$(uname -m 2>/dev/null); uptime=$(cat /proc/uptime 2>/dev/null | cut -d. -f1); cpus=$((nproc || grep -c \	1.38 K
uname -m 2 > /dev/null	1.38 K
uname -a	278
whoami	221

Figure 7. Top Command Execution Cowrie.

The most frequent command was uname -s -v -n -r -m, which appeared over seven thousand times and represented more than one third of all recorded command inputs. Variants that prefixed the command with relative paths such as ./bin/./uname and that redirected output to/dev/null were also common. Additional commands queried/proc/uptime and/proc/cpuinfo, counted processor cores using grep and wc-l, and inspected the operating system with uname -a and whoami. These actions match the environment profiling phase described in earlier botnet analyses [1,3,6].

A second notable cluster of commands attempted to manipulate SSH configuration directories with statements such as ‘cd ~’; ‘chattr -ia .ssh’; and variants that removed and recreated the .ssh directory. These operations are typically associated with establishing or repairing persistent access by modifying key stores. Complementing these sequences,

the honeypot recorded multiple attempts to download external binaries using HTTP. Representative URLs included paths ending in/linux and/bins/x86/hosted on IP addresses located in cloud and hosting networks. While binaries were never executed, their presence reinforces the persistence of a download–execute model for IoT malware distribution.

In addition to the observed command sequences, the honeypot also captured multiple attempts by attackers to download external binaries via HTTP. The representative URLs, such as paths ending in/linux and/bins/x86/, were hosted on IP addresses in cloud and hosting networks, suggesting the use of external resources for payload delivery. Although these binaries were never executed in this experiment, their presence aligns with the well-established download-and-execute infection model commonly seen in IoT botnet activity, particularly in Mirai and its derivatives. To confirm whether these binaries are indeed Mirai derivatives, a more in-depth static analysis should be performed. Hashing these downloaded files and comparing them against known malware databases, such as VirusTotal, would provide insights into their origin and possible links to Mirai or other botnet families. Additionally, reverse-engineering these binaries could help understand their exact functionality, whether they contain the same exploitative features or propagation mechanisms as the Mirai botnet, or introduce new tactics. By conducting such analyses, we would be able to verify if the observed binaries are legacy artifacts or part of newer botnet variants, offering a clearer picture of the evolving IoT malware landscape.

6.1.4. Session Duration and Interaction Style

Most Cowrie sessions were short and automated; however, a small number of connections remained active for extended periods. The longest sessions persisted for between seven and nine hours and originated from addresses in Singapore, Russia, India, Malaysia, Hong Kong, Bangladesh, Réunion, Bosnia and Herzegovina, China, Sweden, and Romania. The durations suggest either poorly coded automation scripts that failed to close sessions or semi-interactive tooling that retains an open channel while background tasks execute. Similar long-lived sessions have been reported in honeypot studies that captured human-in-the-loop activity [7].

6.2. IoTPOT Telnet Honeypot

The IoTPOT honeypot was deployed to emulate vulnerable Telnet services, which are commonly targeted by IoT malware. During the observation period, it received a high volume of unsolicited Telnet connection attempts from diverse source addresses.

Figure 8 shows representative raw log entries captured by the IoTPOT Telnet honeypot, highlighting the source IP addresses associated with unsolicited connection attempts.

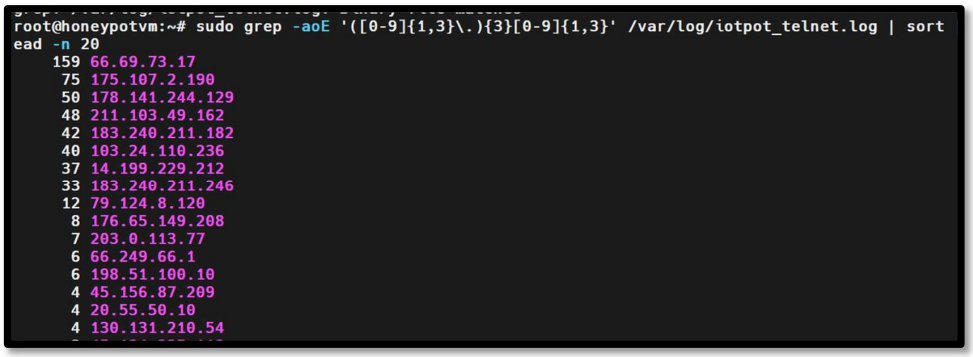

Figure 8. Raw logs for IoTPOT showing source IP of attacks.

These raw logs confirm widespread scanning behavior and provide the basis for subsequent geographical and temporal analysis.

Figure 9 shows an overview of the interaction data collected by the IoTPOT Telnet honeypot, summarizing connection attempts, authentication activity, and command events.

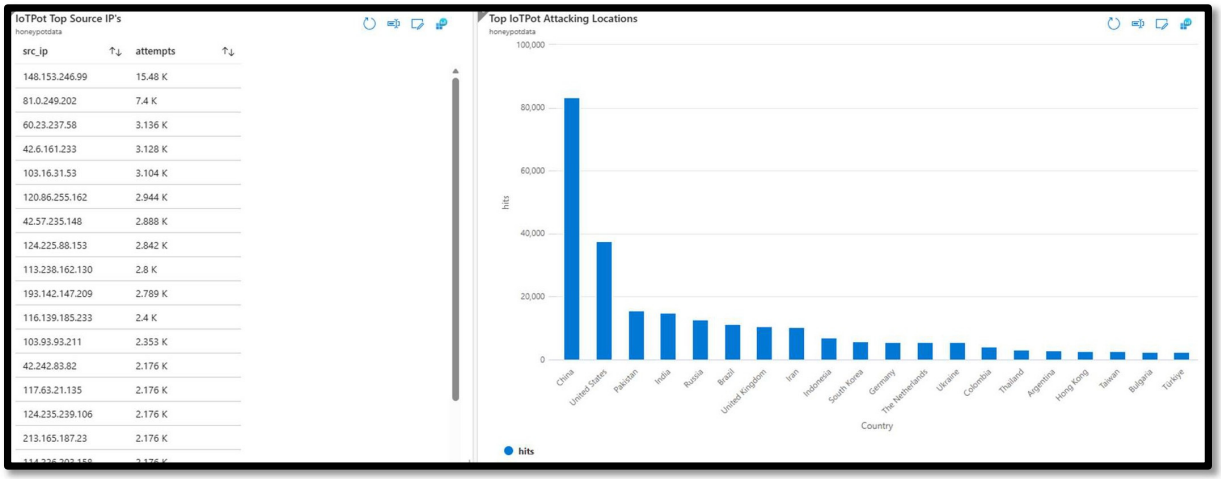

Figure 9. IoTPOT Data.

6.2.1. Source Addresses and Locations

The IoTPOT service, which emulates a Telnet interface, attracted substantially more traffic than the SSH honeypot. The most active individual source IP generated more than 15,000 events, with several additional addresses producing between 2000 and 8000 events each. When grouped by country, China clearly dominated the IoTPOT dataset with over seventy thousand hits. The United States ranked second-largest, with more than 30,000 hits, followed by Pakistan, Russia, India, the United Kingdom, Brazil, and Iran, each with roughly 10,000 to 12,000 hits. Smaller but non-trivial activity volumes were observed from Indonesia, Germany, Ukraine, South Korea, the Netherlands, Colombia, Argentina, Hong Kong, Bulgaria, Taiwan, Turkey, and Malta. The particularly high concentration from China and a long tail of other countries is consistent with earlier measurements of Telnet-focused IoT scanning infrastructure [5,8,9].

6.2.2. Telnet Credentials

Username guesses against IoTPOT were even more skewed than those against Cowrie. The username 'root' accounted for more than 23,000 attempts and represented more than half of all Telnet username trials. The generic admin account formed the second-largest group, with over 11,000 attempts. At the same time, other values, such as guest, GET, supervisor, user, support, default, and administrator, appeared at much lower but still noticeable frequencies. This confirms that attackers treat Telnet endpoints as embedded devices that are expected to ship with privileged default accounts.

The Telnet password distribution displayed a mixture of trivial sequences and slightly more complex strings. The most frequent password guess was admin with roughly 2400 attempts, followed by 1234, 12345, password, 123456, and pass. Additional common passwords included 888888, aquario, 54321, 1111, user, 666666, 5up, 7ujMko0admin, and Win1doWS, along with a small number of non-ASCII tokens. Many of these values match lists previously reported in studies of Telnet-based IoT malware [1,3], reinforcing the notion that default device credentials remain widespread and exploitable.

6.2.3. Event Type Distribution

IoTPOT classifies Telnet activity into several event types. Figure 10 shows the distribution of event types recorded by the IoTPOT Telnet honeypot, including connection attempts, authentication events, command executions, and error conditions.

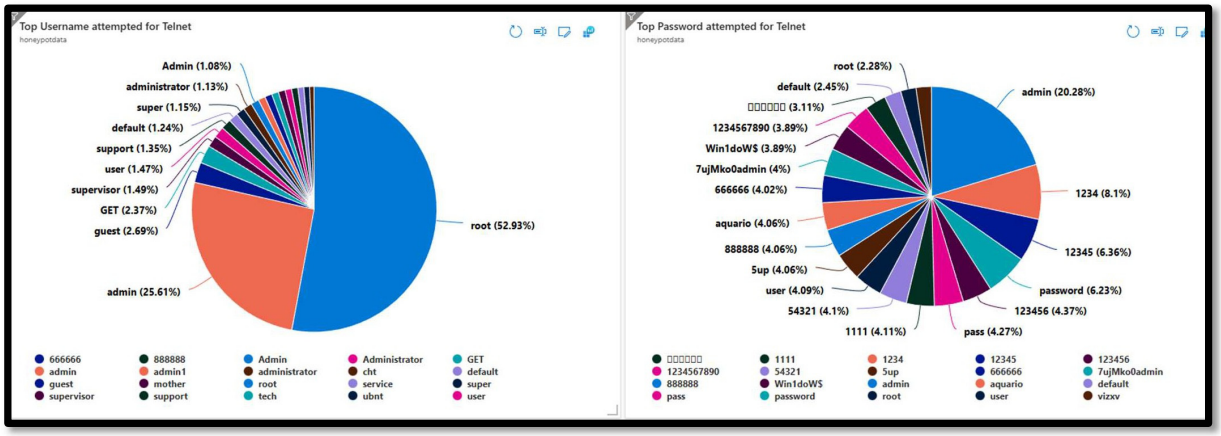

Figure 10. Distribution of event types recorded by the IoTPOT Telnet honeypot, including connection attempts, authentication events, command executions, and error conditions.

Over the observation period, the honeypot recorded approximately 88,000 connection events, 82,000 authentication events, around 28,000 command events, and 34,000 error events. Connection and authentication accounted for more than 70 percent of all IoTPOT events, indicating that most interactions did not progress past login. The smaller fraction of command events suggests that only a subset of attacking infrastructure was successfully authenticated, or that many tools terminate quickly after failed logins. This behavior aligns with high-volume credential spraying rather than deep, interactive exploitation.

6.2.4. Side-by-Side Comparative Analysis: SSH vs. Telnet

To systematically compare SSH and Telnet behavior, Table 2 provides a side-by-side overview of key metrics observed during the measurement period across both honeypots. Telnet (IoTPOT) experienced significantly higher traffic than SSH (Cowrie), with the most active Telnet source producing over 15,000 events, compared to about 1000 for the top SSH source, a ratio of approximately 15:1. This suggests that Telnet remains a more favored attack vector for automated IoT scanning. Credential usage was notably more concentrated in Telnet sessions: over 50% of Telnet username attempts used "root," whereas SSH had about 33%, indicating that malware targeting Telnet relies on a smaller set of credentials. The chances of successful authentication were similarly low for both protocols, as neither honeypot rejected credentials; however, the higher volume of Telnet authentication attempts (82,000 vs. fewer for SSH) shows that Telnet botnets use more aggressive credential spraying. Command usage was much more limited in Telnet sessions, mostly ending at login, whereas SSH sessions displayed a wider variety of post-authentication commands, including environment profiling, configuration adjustments, and binary downloads. In conclusion, both protocols confirm that brute-force credential attacks are common, but Telnet experiences higher attack intensity and less sophistication, whereas SSH shows deeper post-authentication activity.

While the study successfully observed and captured real-time attack patterns, it is important to acknowledge the limited observation period of two weeks. This short duration restricts the ability to draw large-scale, statistically meaningful conclusions, and we view this as a preliminary investigation into the current state of IoT botnets. The study provides a snapshot of attack behaviors rather than a comprehensive, long-term analysis. We will address this limitation in future studies by extending the observation period and incorporating data from additional devices and IP addresses.

Table 2. Side-by-Side Comparison of SSH (Cowrie) and Telnet (IoTPOT) Honeypot Metrics.

Metric	SSH-Cowrie	Telnet-IoTPOT
Peak source activity (top IP)	~1000 events	>15,000 events
Relative traffic volume	1× (baseline)	~15× higher
Top username ("root" share)	~33% (root)	>50% (root)
Total authentication attempts	Lower volume	~82,000 events
Credential spraying intensity	Moderate	High (aggressive)
Post-authentication command diversity	High (env profiling, config changes, binary downloads)	Low (mostly ends at login)
Attack sophistication	Higher (deeper post-auth activity)	Lower (high volume, less depth)
Primary attack vector	Brute-force credential attacks	Brute-force credential attacks

A critical aspect of IoT botnet persistence is the role of software vulnerabilities, particularly outdated and non-updatable firmware. These vulnerabilities leave devices exposed to exploitation, enabling attackers to easily hijack IoT devices for botnet purposes. Despite advances in detection techniques, this issue remains a significant challenge for IoT security. Many devices continue to ship with weak default credentials, and the lack of patch management exacerbates the problem. In the context of our study, the persistence of certain attack behaviors, such as credential brute-forcing, highlights the ongoing risks associated with these vulnerabilities. A more thorough discussion of firmware-related issues and security patching would add value to the manuscript by addressing the systemic nature of IoT security flaws.

6.3. Discussion of Protocol Deviation and Fuzzing Attempts

In addition to the typical attack behavior observed, attackers often attempt to exploit vulnerabilities within the Telnet protocol by sending malformed negotiation options or fuzzing inputs to crash or destabilize the service. Telnet, a widely used protocol for IoT devices, is prone to protocol deviations, where attackers intentionally craft invalid or unexpected Telnet negotiation sequences to exploit implementation weaknesses. These fuzzing attempts aim to trigger unexpected behaviors, such as buffer overflows or service crashes, that could allow attackers to gain control of the target device or network.

While our current study primarily focused on credential brute-forcing and command-execution patterns, future work could expand the scope to include detailed monitoring of malformed Telnet packets. This could involve analyzing the Telnet negotiation phase, where attackers might send out-of-specification negotiation options to cause disruptions, and checking for any signs of service crashes or abnormal terminations of the Telnet service. These types of attacks, often used in botnet exploitation, can be difficult to detect but are crucial for understanding the full range of techniques attackers use to compromise IoT devices. By incorporating protocol fuzzing analysis into future honeypot deployments, we can gain a more comprehensive understanding of the methods attackers use to exploit IoT vulnerabilities and improve mitigation strategies against these increasingly sophisticated attack vectors.

6.4. Discussion of Source Attribution and GeoIP Limitations

While the geographic analysis of attack sources provided valuable insights into the global distribution of IoT botnet activity, it is important to acknowledge the limitations of relying solely on GeoIP data for source attribution. This approach does not account for the extensive use of VPNs, proxy servers, and other anonymizing technologies employed

by botnet operators to obfuscate their true geographic locations. As a result, the GeoIP-based attribution presented in this study may not accurately reflect the actual origins of the attacks, as the observed IP addresses could be routed through various intermediaries, distorting the true locations of the attackers.

Given the growing use of anonymizing services in cybercrime, we acknowledge that the geographic distribution identified in our analysis is likely biased. Although GeoIP data provides a general overview of attack source locations, it should be interpreted with caution. The observed high volumes of attacks from countries such as China and the United States could reflect the presence of proxy networks or compromised machines rather than the true locations of the botnet operators.

6.5. Threats to Validity

While GeoIP data provided valuable insights into the geographic distribution of attack sources, it is important to recognize the limitations of relying on this method for source attribution. IP-spoofing and the use of proxy servers or VPNs by botnet operators can significantly distort the geographic analysis, as attackers often anonymize their location to evade detection and attribution. This means that the geographic data captured in this study may not accurately reflect the true origins of the attacks, as the observed IP addresses may be routed through various intermediaries, potentially masking the attackers' real locations.

7. Conclusions

This study is explicitly framed as a confirmatory or validation study rather than one introducing new attack insights. Its primary contribution is the demonstration that well-established IoT botnet behaviors, credential brute-forcing, Mirai-style command sequences, and Telnet protocol dominance persist unchanged in current Internet environments, confirming findings from foundational research conducted nearly a decade ago. This study bridged a decade of IoT botnet research with contemporary attacker behavior by combining a comprehensive review of influential security studies with the deployment of a live, cloud-based honeypot. The empirical data collected through the Cowrie and IoTPOT honeypots demonstrated that core characteristics of IoT botnet activity remained largely unchanged, despite years of academic research and defensive innovation. Weak authentication, large-scale automated scanning, predictable infection routines, and globally distributed attack sources continue to dominate the threat landscape.

Both the SSH and Telnet honeypots predominantly observed authentication attempts targeting default, privileged accounts such as root and admin, with these usernames accounting for the majority of login attempts. This pattern highlighted that attackers continued to exploit poor credential hygiene in deployed IoT devices. Password guessing also followed similar trends, with attackers relying on short numeric sequences, factory-default strings, and reused tokens, behavior consistent with long-established botnet tactics. The persistence of these attack strategies suggests that attackers prioritize speed and volume over sophistication, exploiting ongoing vulnerabilities in IoT devices rather than developing novel brute-force methods.

Post-authentication actions further underscore the consistency of these attack behaviors, with the most frequent commands focused on system reconnaissance, such as operating system identification, CPU checks, and environment profiling, as well as attempts to manipulate SSH configuration files and download external binaries. These command sequences reflect a well-established download-and-execute infection model, highlighting minimal evolution in malware operational logic over the years.

Geographical analysis of attack sources revealed a globally distributed attack surface, with a high volume of activity originating from China and the United States, followed

by significant traffic from regions across Asia, Europe, and the Middle East. Notably, the Telnet-focused IoTPOT received significantly more traffic than the SSH-based Cowrie honeypot, reaffirming that Telnet remained a preferred attack vector due to its widespread exposure in legacy and embedded devices.

The experimental findings supported three key observations: first, credential brute-forcing remained the primary infection vector for IoT botnets; second, contemporary malware continued to reuse long-established command sequences with little technical innovation; and third, global automated scanning remained the dominant reconnaissance method. These findings suggested that the persistence of IoT botnets was not driven by increasingly sophisticated attackers, but rather by the enduring weaknesses in device configurations, patch management, and supply chain security practices.

Furthermore, this research demonstrated that small-scale, cloud-hosted honeypots constitute an effective and economical means of capturing representative Internet of Things (IoT) attack telemetry. The selection of cloud hosting was both imperative and strategically advantageous for several reasons: (1) cloud platforms provide authentic public IPv4 addresses that are indistinguishable from legitimate IoT devices to external scanners, thereby ensuring that attack traffic accurately reflects actual botnet behavior rather than artificial laboratory conditions; (2) cloud infrastructure affords elastic scalability, enabling researchers to swiftly deploy, configure, and replicate honeypot instances across multiple geographic regions without the need for physical hardware investments; (3) cloud service providers offer high-bandwidth, always-on network connectivity essential for capturing continuous attack streams and would be prohibitively costly with dedicated on-premises infrastructure; (4) cloud-based deployment mitigates the risks associated with exposing physical organizational networks to malicious traffic, providing network isolation while maintaining realistic Internet exposure; and (5) cloud platforms facilitate reproducible experimental conditions through infrastructure-as-code, allowing other researchers to validate findings and extend the methodology. These environments provide substantial value for validating theoretical models, monitoring evolving threat landscapes, and conducting ongoing security assessments without necessitating extensive infrastructure. The cloud-hosted method is particularly advantageous for academic research teams and security practitioners who require longitudinal IoT threat studies with constrained budgets while upholding methodological rigor and data authenticity.

Ultimately, the persistence of these attack patterns pointed to a systemic issue rather than a technical one. The IoT security challenge was no longer simply a matter of insufficient research or lack of defensive techniques, but rather one of inconsistent enforcement across manufacturing, deployment, and maintenance processes. Critically, the stagnation in attacker behavior can be causally attributed to three reinforcing structural factors: first, the economic incentive structure of the IoT market, where manufacturers face little financial penalty for shipping insecure devices, removes the impetus to adopt secure-by-design principles; second, the absence of mandatory, enforceable patching obligations means that millions of already, deployed devices remain permanently vulnerable, creating a persistent and renewable infection pool that sustains botnet operations without requiring attackers to develop new techniques; and third, the decentralized nature of IoT ownership, spanning consumers, small businesses, and ISPs, creates accountability gaps that no single actor has the authority or incentive to resolve. These causal factors collectively explain why attack strategies documented at Mirai's emergence in 2016 remain effective and dominant in 2025: the underlying vulnerability conditions have not been resolved, only better described. Until secure-by-design principles, mandatory credential management, and reliable patching mechanisms are widely adopted, the attack patterns first observed nearly a decade ago will continue to define the IoT security landscape. The evidence presented in this study

underscores the continued relevance of lessons learned from past research and highlights the urgent need for action from manufacturers, regulators, and network operators alike.

Author Contributions: Conceptualization, R.B. and A.K.G.; methodology, S.R.A. and V.K.K.; software, R.S. and S.R.A.; validation, V.K.K., A.K.G. and G.S.S.; formal analysis, D.K. and C.T.; investigation, G.S.S. and D.K.; resources, R.S., G.S.S. and D.K.; data curation, D.K. and S.R.A.; writing—original draft preparation, R.B., A.K.G. and S.R.A.; writing—review and editing, S.R.A. and V.K.K.; visualization, R.S. and C.T.; supervision, R.B. and A.K.G.; project administration, S.R.A. and V.K.K.; funding acquisition, D.K., G.S.S. and S.R.A. All authors have read and agreed to the published version of the manuscript.

Funding: This research received no external funding.

Institutional Review Board Statement: Not applicable.

Data Availability Statement: The original contributions presented in this study are included in the article. Further inquiries can be directed to the corresponding authors.

Acknowledgments: The authors acknowledge the support and resources provided by the University of the Cumberlands, the University of Texas at San Antonio, Universiti Brunei Darussalam, and Saint Louis University, whose conducive academic and research environments were instrumental in the successful completion of this study. The authors also extend their sincere appreciation to the editorial team of this journal for the opportunity to publish this work and for their constructive feedback throughout the review process. Grammarly was used for grammar, readability, and stylistic refinement of the manuscript; all technical content, conceptual development, modeling decisions, and scientific conclusions remain solely the responsibility of the authors.

Conflicts of Interest: The authors declare that they have no conflicts of interest in this work.

References

1. Antonakakis, M.; April, T.; Bailey, M.; Bernhard, M.; Bursztein, E.; Cochran, J.; Durumeric, Z.; Halderman, J.A.; Invernizzi, L.; Kallitsis, M.; et al. Understanding the Mirai Botnet. In Proceedings of the 26th USENIX Security Symposium (USENIX Security '17), Vancouver, BC, Canada, 16–18 August 2017; pp. 1093–1110.
2. Kelly, C.; Pitropakis, N.; McKeown, S.; Lambrinoudakis, C. Testing and hardening IoT devices against the Mirai botnet. In *Proceedings of the 2020 International Conference on Cyber Security and Protection of Digital Services (Cyber Security), Dublin, Ireland, 15–19 June 2020*; IEEE: Piscataway, NJ, USA, 2020; pp. 1–8. [CrossRef]
3. Soltan, S.; Mittal, P.; Poor, H.V. BlackIoT: IoT Botnet of High Wattage Devices Can Disrupt the Power Grid. In Proceedings of the 27th USENIX Security Symposium (USENIX Security '18), Baltimore, MD, USA, 15–17 August 2018; pp. 15–32.
4. Shekari, T.; Cardenas, A.A.; Beyah, R. MaDIoT 2.0: Modern High-Wattage IoT botnet attacks and defenses. In Proceedings of the 31st USENIX Security Symposium (USENIX Security 22), Boston, MA, USA, 10–12 August 2022; pp. 3539–3556.
5. Kumar, D.; Gupta, S.; Li, B.; Mason, J.; Wetherall, D.; Anderson, T. All Things Considered: An Analysis of IoT Devices on Home Networks. In Proceedings of the 28th USENIX Security Symposium (USENIX Security '19), Santa Clara, CA, USA, 14–16 August 2019; pp. 1169–1185.
6. Alrawi, O.; Lever, C.; Valakuzhy, K.; Snow, K.; Monrose, F.; Antonakaki, M. The Circle of Life: A Large-Scale Study of IoT Malware Evolution. In Proceedings of the 30th USENIX Security Symposium (USENIX Security '21), Virtual Conference, Online, 11–13 August 2021; pp. 3505–3522.
7. Dong, Y.; Li, Q.; Wu, K.; Li, R.; Zhao, D.; Tyson, G.; Peng, J.; Jiang, Y.; Xia, S.; Xu, M. HorusEye: A realtime IoT malicious traffic detection framework using programmable switches. In Proceedings of the 32nd USENIX Security Symposium (USENIX Security '23), Anaheim, CA, USA, 9–11 August 2023; pp. 571–588.
8. Guan, C.; Cao, G. Cyber-Physical Deception Through Coordinated IoT Honeypots. In Proceedings of the 34th USENIX Security Symposium (USENIX Security 25), Seattle, WA, USA, 13–15 August 2025; pp. 529–545.
9. Pan, Y.; Ling, Z.; Zhang, Y.; Wang, H.; Liu, G.; Luo, J.; Fu, X. TORCHLIGHT: Shedding LIGHT on Real-World Attacks on Cloudless IoT Devices Concealed within the Tor Network. In Proceedings of the 34th USENIX Security Symposium (USENIX Security 25), Seattle, WA, USA, 13–15 August 2025; pp. 1053–1072.
10. Xing, Y.; Shu, H.; Zhao, H.; Li, D.; Guo, L. Survey on botnet detection techniques: Classification, methods, and evaluation. *Math. Probl. Eng.* **2021**, *2021*, 6640499. [CrossRef]

11. Woodiss-Field, A.; Johnstone, M.N.; Haskell-Dowland, P. Examination of Traditional Botnet Detection on IoT-Based Bots. *Sensors* **2024**, *24*, 1027. [CrossRef] [PubMed]
12. Memos, V.A.; Stergiou, C.L.; Bermperis, A.I.; Plageras, A.P.; Psannis, K.E. A Novel Architecture for Mitigating Botnet Threats in AI-Powered IoT Environments. *Sensors* **2026**, *26*, 572. [CrossRef] [PubMed]
13. Wazzan, M.; Algazzawi, D.; Bamasaq, O.; Albeshri, A.; Cheng, L. Internet of Things Botnet Detection Approaches: Analysis and Recommendations for Future Research. *Appl. Sci.* **2021**, *11*, 5713. [CrossRef]
14. Gelgi, M.; Guan, Y.; Arunachala, S.; Rao, M.S.S.; Dragoni, N. Systematic Literature Review of IoT Botnet DDoS Attacks and Evaluation of Detection Techniques. *Sensors* **2024**, *24*, 3571. [CrossRef] [PubMed]
15. Asadi, M.; Jamali, M.A.J.; Heidari, A.; Navimipour, N.J. Botnets Unveiled: A Comprehensive Survey on Evolving Threats and Defense Strategies. *Trans. Emerg. Telecommun. Technol.* **2024**, *35*, e5056. [CrossRef]
16. Valentini, E.P.; Franco, T.A.; Gottsfritz, E.N.; Estrella, J.C.; Júnior, L.A.P.; Meneguette, R.I. Honeypot Embedded in Low-Cost IoT Hardware for Collecting and Analyzing Real-World Cyber Threats. In *Proceedings of the 18th IEEE/ACM International Conference on Utility and Cloud Computing (UCC '25), Nantes, France, 1–4 December 2025*; Article No. 64; ACM Digital Library: New York, NY, USA, 2026; pp. 1–6. [CrossRef]
17. Illi, E.; Bazzi, A.; Qaraqe, M.; Ghrayeb, A. On the Secrecy-Sensing Optimization of RIS-Assisted Full-Duplex Integrated Sensing and Communication Network. *IEEE Trans. Wirel. Commun.* **2025**, *25*, 9530–9547. [CrossRef]