TPOT Honeypot on Azure

7 DAYS, 365K ATTACKS, AND THE ANATOMY OF THREAT INTELLIGENCE

MARCH 2025

Prepared by
PRITAM DASH

Investigation Period
26 MAR 25 - 04 APR 25

2025

TABLE OF CONTENTS

Introduction 1

Setup 2

Post-Setup Screenshots 3

Detailed Statistics & Insights 4

Port Targeting Breakdown 5

Targeting Default Logins 6

Clouds as Attack Hubs 7

Attack Origins by OS 8

Top 10 Attacker Source IPs 9

Top Exploited CVEs by Count 11

Suricata Detection Trends 13

PRITAM DASH

Introduction

Over a 7-day period, a TPOT honeypot strategically deployed on an Azure virtual machine captured and logged more than 355,000 malicious events originating from a diverse range of IP addresses worldwide. This report, TPOT Honeypot on Azure: 7 Days, 355K Attacks, and the Anatomy of Threat Intelligence, discusses those events to uncover actionable threat intelligence and patterns in real-world attacker behavior. In an era where cyber threats evolve at an unprecedented pace, understanding the nature and scale of attacks is important for building resilient defenses. This investigation provides a critical window into this dynamic landscape.

The honeypot, armed with services including SSH, HTTP, FTP, and others, acted as a decoy/honeypot, luring in automated bots, scanners, and potential threat actors. Over the course of one week, from 26th March 2025 to 2nd April 2025, it recorded an astonishing volume of malicious activity, exceeding initial expectations. The data collected reveals high-frequency brute-force attacks, port scans, exploit attempts, and early-stage malware delivery campaigns providing a snapshot of the threat landscape targeting cloud infrastructure today.

Top Stats

365K
ATTACKS

US
TOP ATTACKING COUNTRY

US
TOP ATTACKING COUNTRY

root
TOP USERNAME USED IN BRUTEFORCE

123456
TOP PASSWORD USED IN BRUTEFORCE

Setup

1.AZURE VIRTUAL MACHINE CONFIGURATION

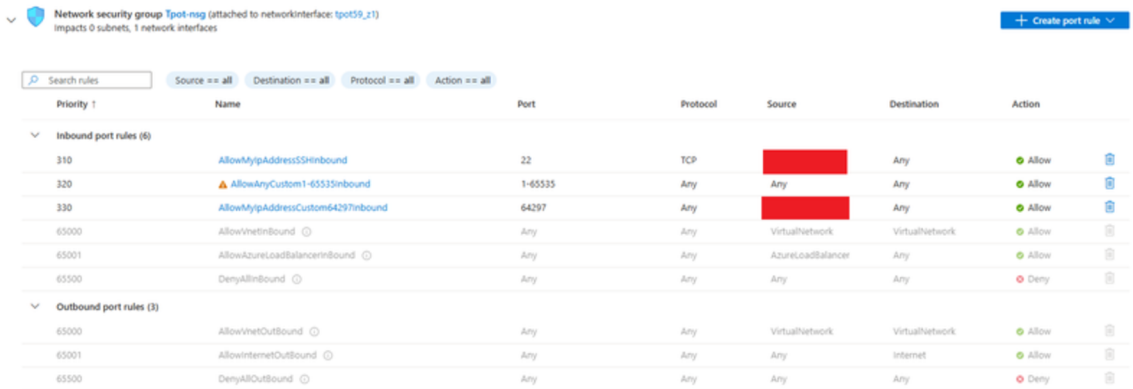
- Cloud Provider: Microsoft Azure
- VM Size: Standard D4s v3 (4 vcpus, 16 GiB memory)
- OS Image: Ubuntu 22.04 LTS
- Region: Canada Central (Zone 1)
- Storage: 128 GB Premium SSD
- Inbound Ports Opened:
  1-65535
- These ports were selected to mimic commonly targeted services (SSH, FTP, Telnet, HTTP/S, RDP, MySQL, etc.)

2.TPOT FRAMEWORK INSTALLATION

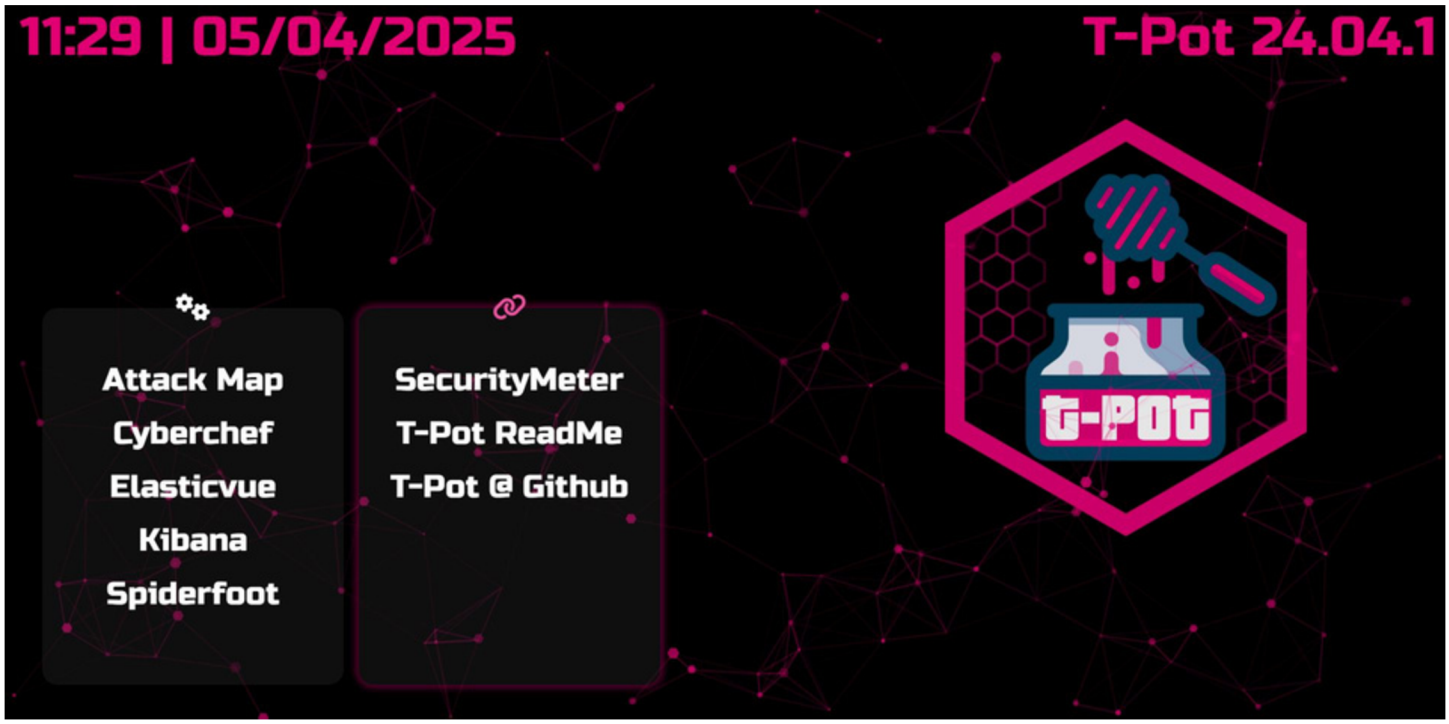
- Version: TPOT 24.04.1
- Installation Method: GitHub clone + install script
- TPOT Components Installed (Not limited to the mentioned below):
  Cowrie – SSH/FTP honeypot for credential and command capture
  Dionaea – Malware collection honeypot
  Conpot – ICS/SCADA simulation
  Suricata – IDS/IPS detection
  ELK Stack (Elasticsearch, Logstash, Kibana) – For centralized log collection, parsing, and visualization

TPOT was installed using the official script from the TPOT GitHub repository, which automates container setup, security tool integration, and data pipelines.

3. NETWORK SETUP



Post-Setup Screenshots





PRITAM DASH

Detailed Statistics & Insights

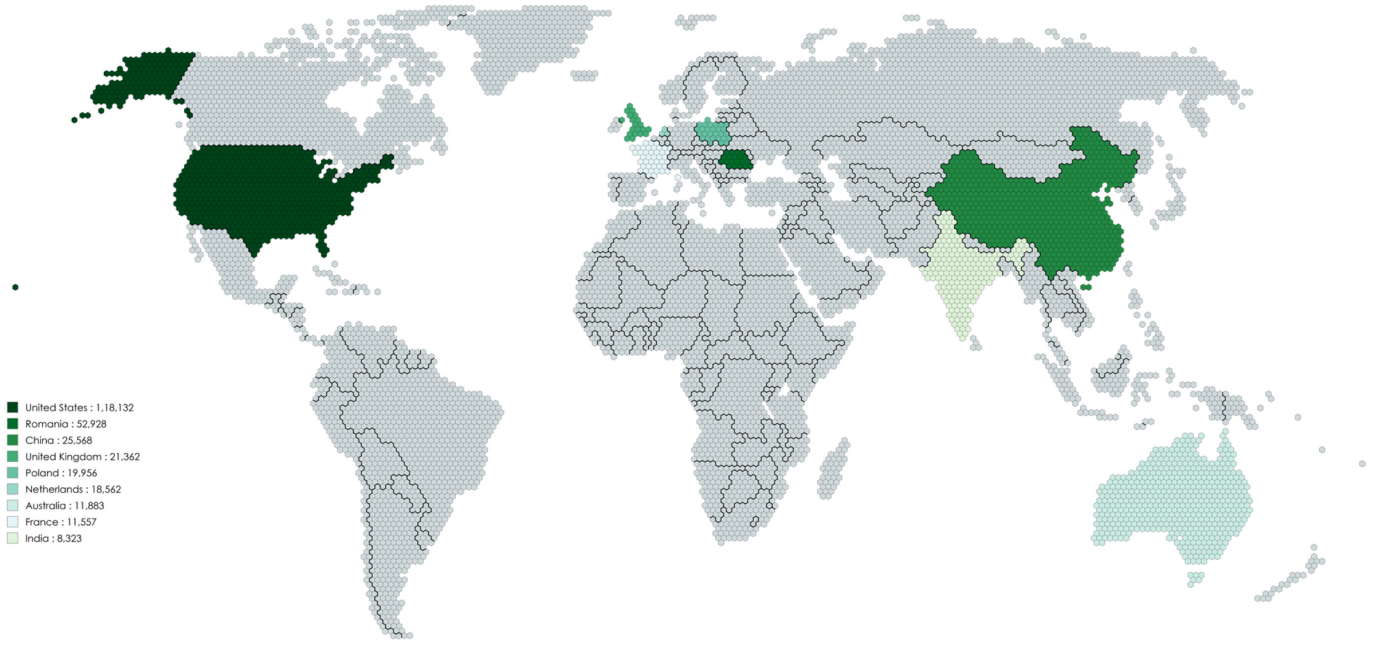
This section presents a full breakdown of attack activity observed during the 7-day honeypot deployment. The analysis includes attacker origins, methods, services targeted, exploit attempts, and suricata detections. Each insight is supported by visual data.

TOP SOURCE COUNTRIES

Analysis of the geographic origin of attacks captured during the 7-day TPOT honeypot deployment reveals distinct patterns in both volume and regional concentration.

The United States emerged as the leading source, responsible for 40% of all attacks. While this may seem counterintuitive given that the U.S. is frequently a target of cyber threats. It highlights the prevalence of compromised systems, open proxies, and botnet activity operating from within the country. These systems are often unknowingly leveraged by attackers, making the U.S. a major hub for attack origination.

Romania accounted for 17% of attacks, a surprising figure given its relatively smaller global cyber footprint. This high volume likely points to the presence of compromised systems, open proxies, or data centers being abused for malicious activity, rather than direct involvement in state-sponsored operations. Romania has occasionally been linked to cybercrime groups, but the data here more plausibly reflects the use of local infrastructure as part of larger, globally distributed botnets.

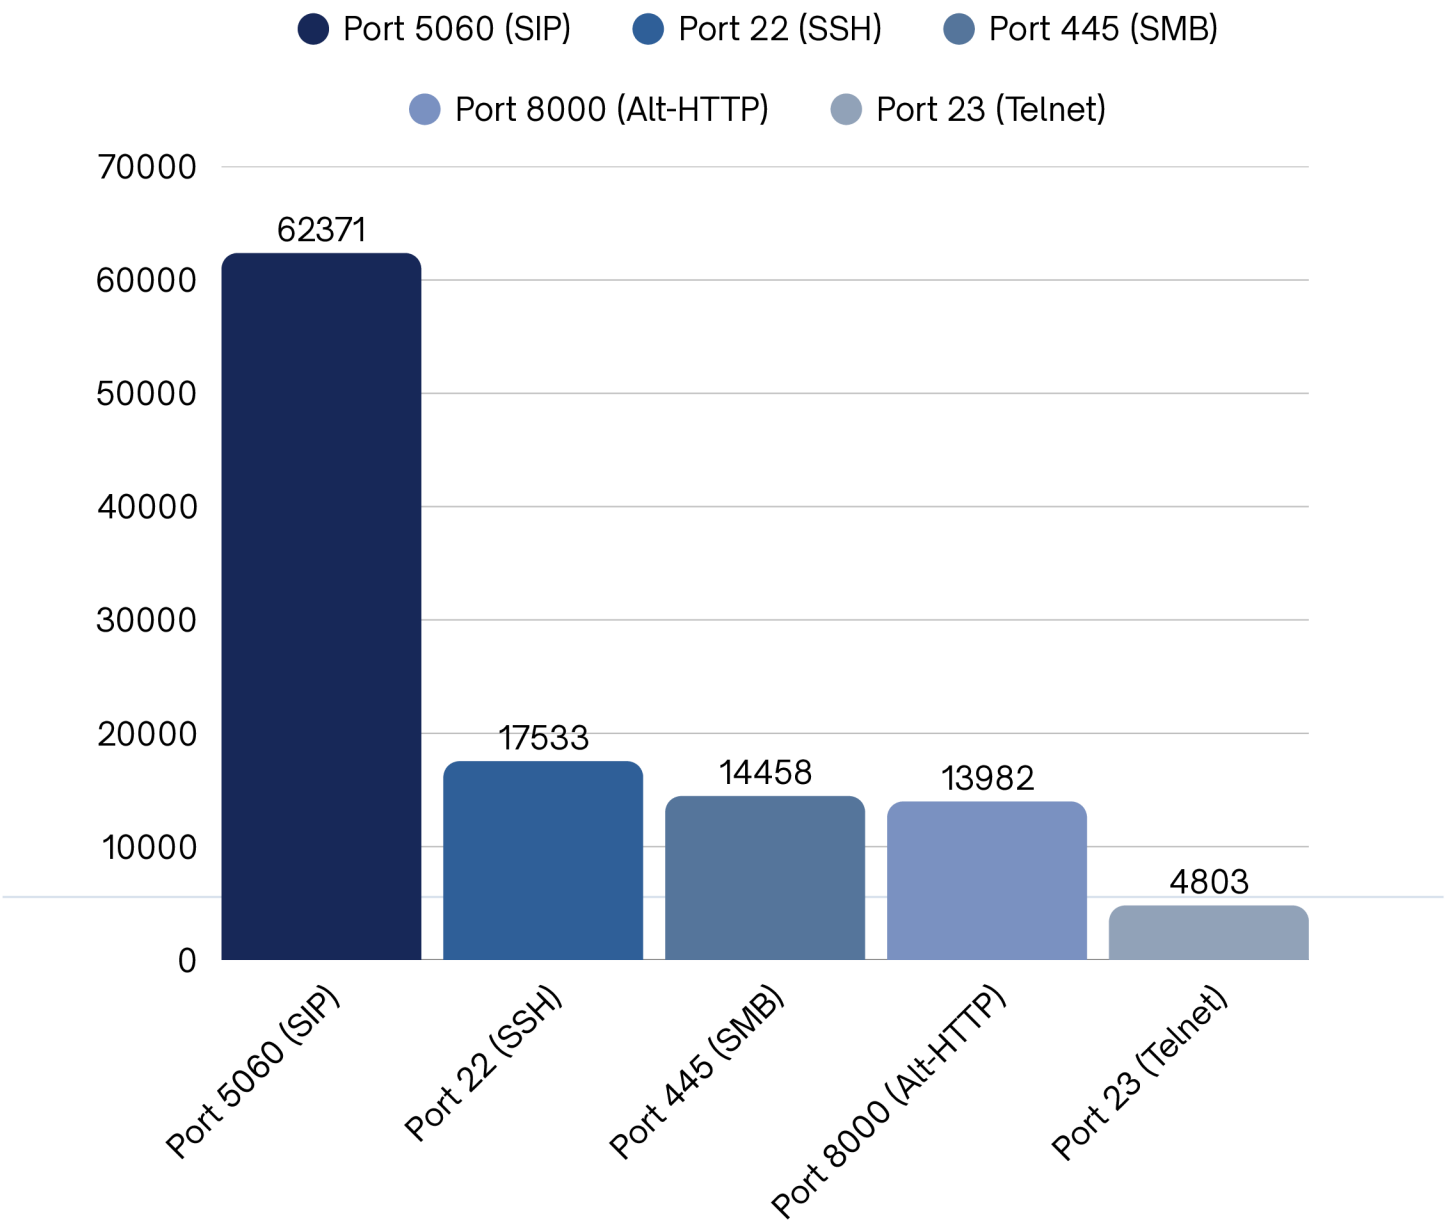


What's particularly interesting is that the top countries aren't necessarily those with the most advanced cyber warfare capabilities but rather those with high levels of internet penetration, lax cybersecurity hygiene, and widely distributed unpatched devices, which are ideal for abuse by threat actors.

04/04/2025

Port Targeting Breakdown

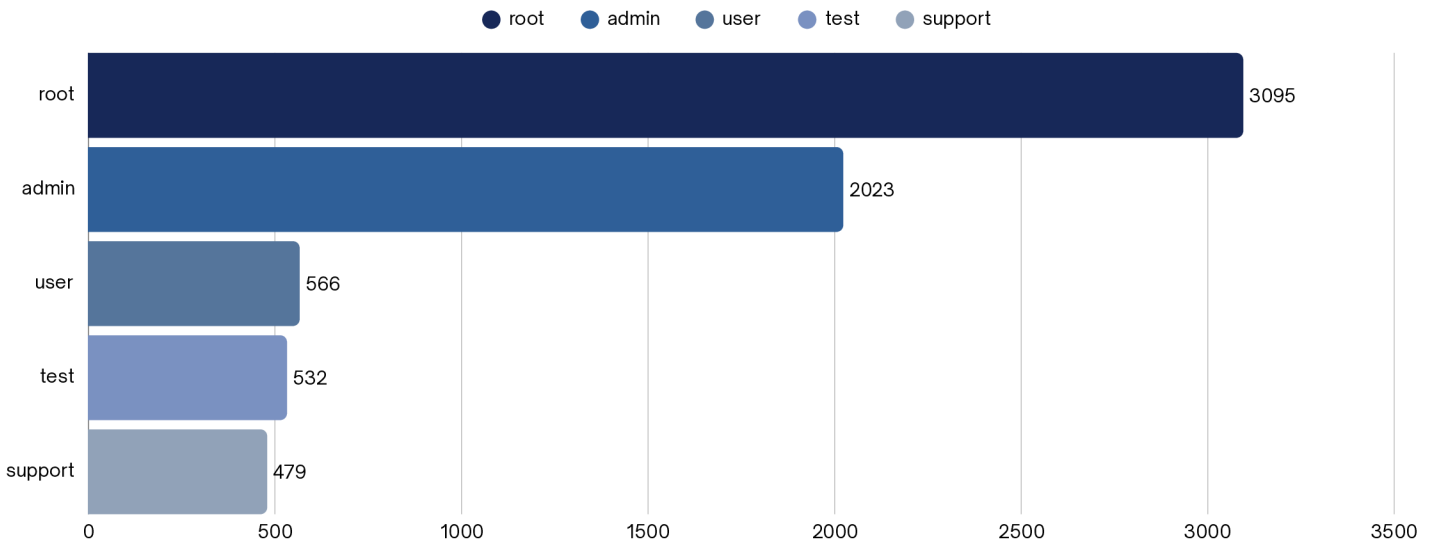
Port 5060 (SIP) saw the highest attack volume, including a spike of over 62,000 attacks in 12 hours highlighting VoIP systems as a major target for call interception and telephony exploits. SSH on port 22 faced sustained brute-force attempts, with peaks over 17,000, indicating automated credential attacks on remote servers. Port 445 (SMB), despite being a legacy protocol, remains actively targeted particularly for ransomware delivery with spikes like 14,400+ attacks in one window. Port 8000, often used for alternate HTTP services, recorded a sudden surge of 13,900+ attacks, suggesting a focused campaign on web dashboards or proxies. Lastly, port 23 (Telnet), while deprecated, continues to be exploited via IoT botnets, showing the lingering threat of unsecured embedded devices.



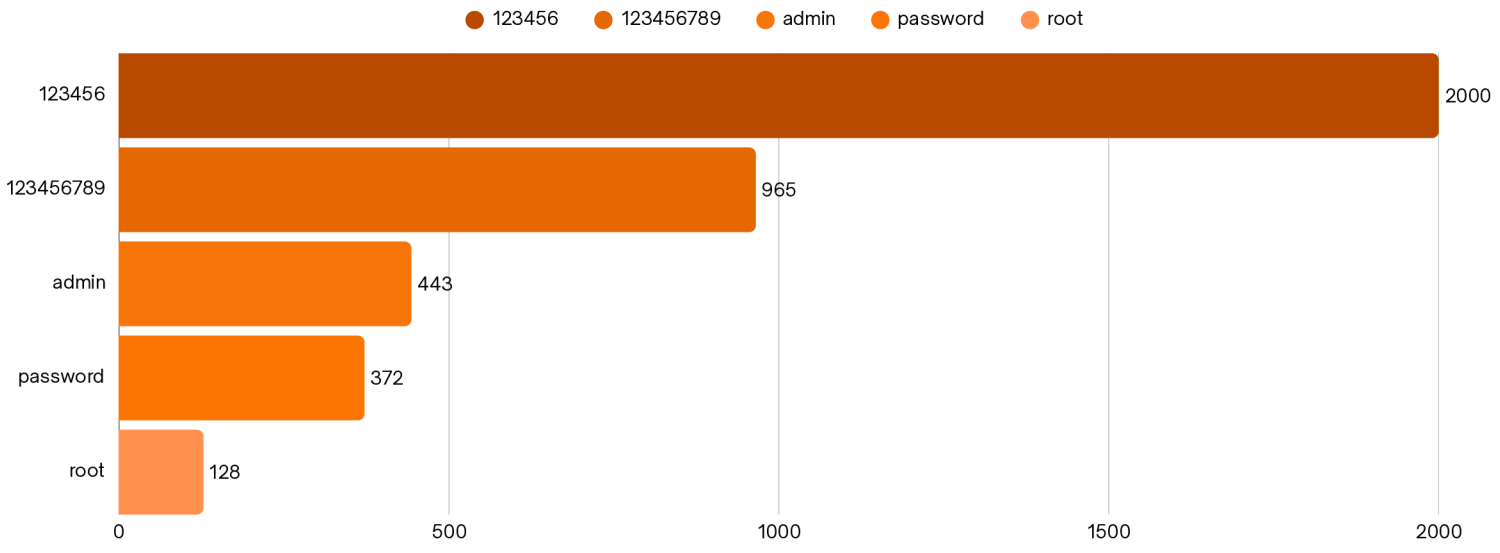
PRITAM DASH

Targeting Default Logins

Brute-force attacks observed through Cowrie honeypots reveal that attackers are still heavily focused on systems using weak or default credentials. The username “root” was by far the most frequently targeted, followed by “admin”, “test”, “guest”, and service-related accounts like “mysql”, “postgres”, and “support”. On the password side, combinations like “123456”, “password”, “admin123”, and “123456789” dominated the attack attempts, illustrating attackers’ reliance on widely known, easy-to-guess credentials. This trend underscores how default or weak configurations remain a critical vulnerability especially on exposed SSH services making it crucial for organizations to enforce strong authentication policies and disable default accounts wherever possible.



Most Targeted Usernames



Most Targeted Passwords

04/04/2025

Clouds as Attack Hubs

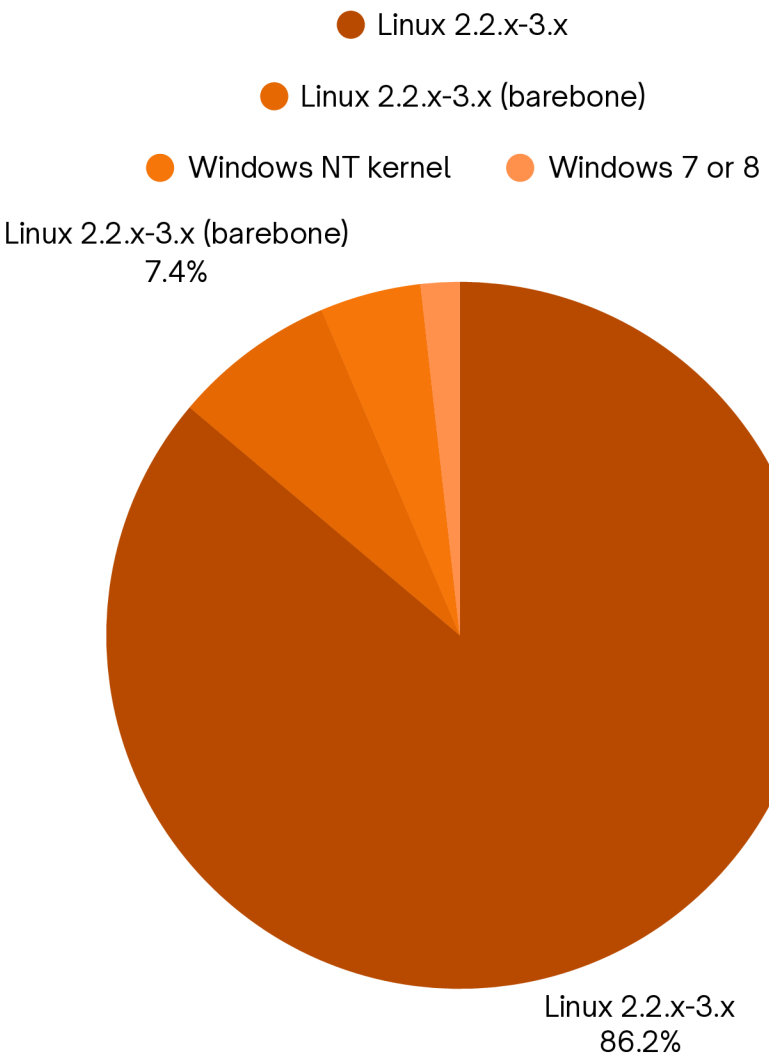
A notable pattern emerged in the analysis as many attacks originated from cloud-based infrastructure. Platforms such as DigitalOcean, Google Cloud, and OVH were among the top autonomous Systems linked to malicious traffic. This highlights a growing trend where threat actors leverage the flexibility and scalability of cloud services to run their attack operations. With quick setup and minimal oversight, cloud environments are increasingly being misused to host attack tools or initiate scanning and exploitation campaigns.

AS	ASN	Count
396982	GOOGLE-CLOUD-PLATFORM	72,141
47890	Unmanaged Ltd	52,593
14061	DIGITALOCEAN-ASN	30,198
210558	1337 Services GmbH	18,390
4134	Chinanet	15,110
63949	Akamai Connected Cloud	11,899
207566	Chang Way Technologies Co. Limited	10,803
401120	CHEAPY-HOST	10,516
398324	CENSYS-ARIN-01	9,204
16276	OVH SAS	9,097

Attack Origins by OS

Analysis of attacker infrastructure revealed that Linux systems, particularly older versions like Linux 2.2.x to 3.x, dominate as the source of attacks, accounting for over 860,000 attempts. A notable portion came from minimal “barebone” setups and systems with stripped timestamps, likely to evade detection.

Windows systems were also prominent, with more than 70,000 attempts traced to NT-based kernels and versions like Windows 7 or 8, suggesting their use in botnets or via compromised endpoints. Attacks from Linux 3.x and 2.4.x were less frequent, and Mac OS X remained nearly absent confirming it's not a common attacker platform.



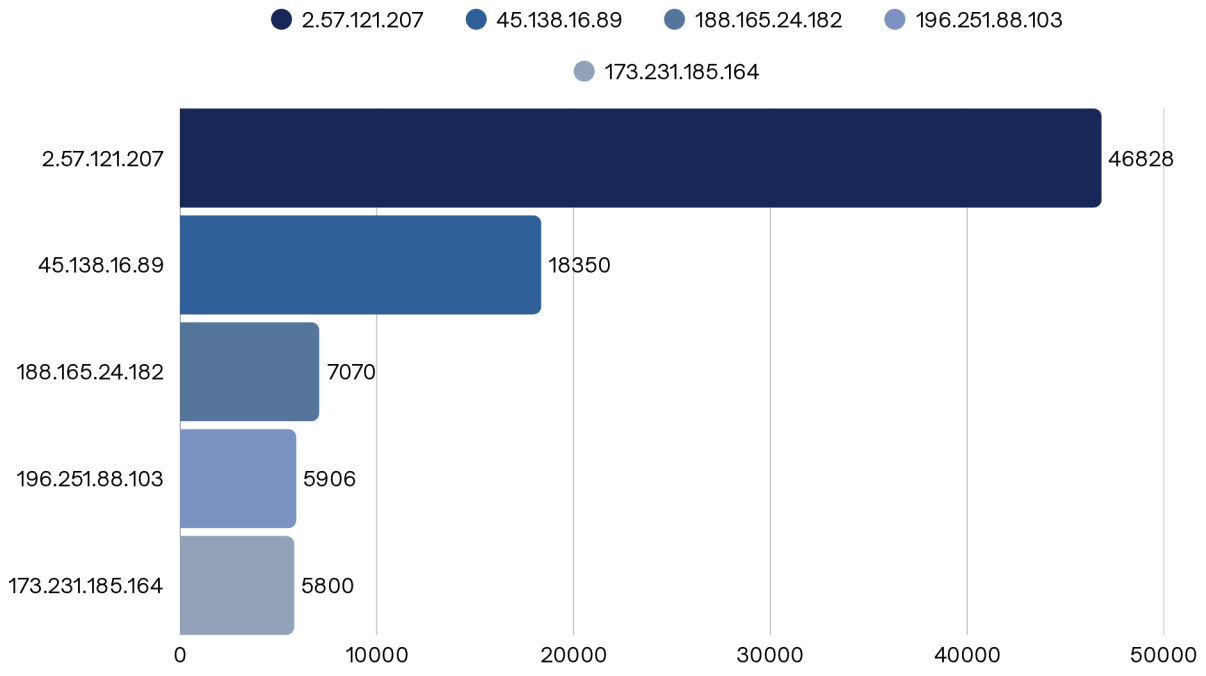
Top 10 Attacker Source IPs

Source IP	Count
2.57.121.207	46,828
45.138.16.89	18,350
188.165.24.182	7,070
196.251.88.103	5,906
173.231.185.164	5,800
138.124.60.132	5,297
182.40.102.60	4,041
92.255.57.129	3,678
185.7.214.200	3,386
105.100.97.19	3,146

MOST ACTIVE SOURCE ADDRESSES DRIVING MALICIOUS TRAFFIC

A close look at source IP data reveals that a small group of IPs generated a disproportionately large volume of malicious traffic. The top offender, 2.57.121.207, was responsible for over 46,800 attempts alone—far outpacing the next most active IP, 45.138.16.89, with 18,350 attempts. These IPs were followed by others such as 188.165.24.182, 196.251.88.103, and 173.231.185.164, each contributing thousands of hits.

These IPs are often tied to data centers or unmanaged hosting services, suggesting attackers are leveraging VPS infrastructure or compromised cloud instances to launch sustained campaigns. This not only provides them flexibility but also makes takedown efforts more difficult.



Top Threats from Attacker Source IPs

Source IP	Count	Potential Threat
2.57.121.207	46,828	Brute-force (Various)
45.138.16.89	18,350	Malware (AsyncRAT)
185.165.24.182	7,070	VoIP Brute-force
196.251.88.103	5,906	Port Scanning
138.124.60.132	5,297	Potential Malware

The above table summarizes potential security threats detected from various source IPs. The data highlights significant activity, with the IP 2.57.121.207 exhibiting the highest count (46,828) linked to widespread brute-force attempts. Other notable threats include malware (AsyncRAT) from 45.138.16.89 (18,350 counts), VoIP brute-force attacks, and multiple instances of port scanning and network sweeping. These activities suggest a range of malicious behaviors targeting the network.



Top Exploited CVEs by Count

Attackers are relentlessly exploiting vulnerabilities—some over 20 years old—proving that outdated and unpatched systems still offer fertile ground for compromise.

Legacy issues, such as CVE-2002-0013 and CVE-1999-0517, affecting SNMP, were among the most frequently observed in attacks, with detections numbering in the hundreds. This persistent interest in historical vulnerabilities highlights the importance of maintaining effective patch management and decommissioning unsupported software.

However, the threat landscape isn't limited to the old. CVE-2023-46604, a recent Apache ActiveMQ flaw, and CVE-2021-3449 in OpenSSL also ranked high in exploit attempts, highlighting how quickly adversaries pivot to weaponize newly disclosed bugs.

The pattern suggests heavy use of automated scanners sweeping the internet for known CVEs, from remote code execution to denial-of-service exploitation. For analysts, this means visibility and proactive hardening are essential.

SOC teams should combine timely patching with threat-hunting strategies:

- Monitor for CVE-specific indicators of compromise
- Deploy honeypots to capture early-stage probes
- Use SIEM rules and network segmentation to isolate vulnerable assets
- Restrict outdated protocols and limit external exposure

CVE ID	Count
CVE-2006-2369	703
CVE-2002-0013 CVE-2002-0012	692
CVE-2002-0013 CVE-2002-0012 CVE-1999-0517	436
CVE-2001-0414	190
CVE-2023-46604 CVE-2023-46604 CVE-2023-46604	128
CVE-2019-11500 CVE-2019-11500	116
CVE-2002-1149	93
CVE-2021-3449 CVE-2021-3449	59
CVE-1999-0183	25
CVE-2006-3602 CVE-2006-4458 CVE-2006-4542	13

PRITAM DASH

CVE Exploitation Summary

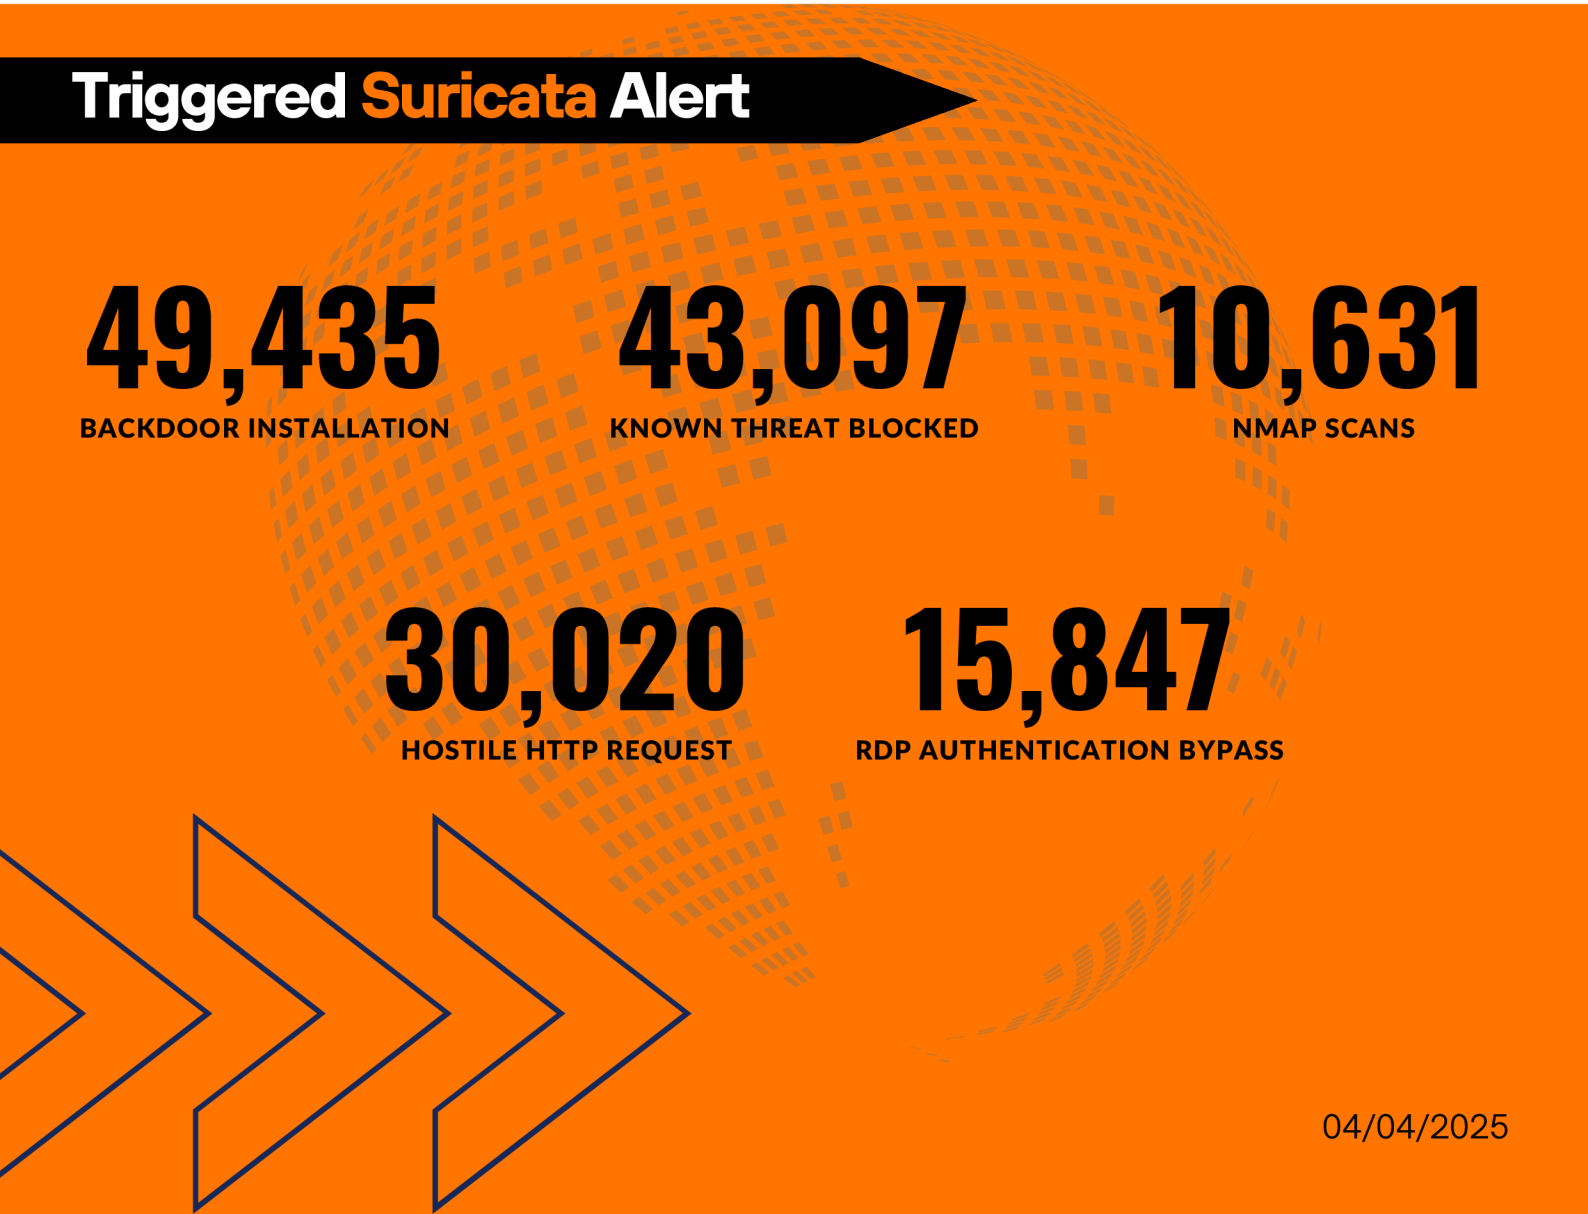
CVE ID(s)	Vulnerability	Exploit	Severity	Count
CVE-2006-2369	RealVNC Remote Authentication Bypass	Allows remote attackers to bypass authentication in RealVNC due to a protocol flaw.	High	703
CVE-2002-0013 CVE-2002-0012	SNMPv1 Vulnerabilities	Exploits in SNMPv1 implementations allowing unauthorized access or DoS attacks.	Critical	692
CVE-2002-0013 CVE-2002-0012 CVE-1999-0517	SNMP and Default Community String	Combines SNMP flaws with default community string "public" for network device access.	Critical	436
CVE-2001-0414	BIND Buffer Overflow	Buffer overflow in BIND DNS software enabling remote code execution.	High	190
CVE-2023-46604	Apache ActiveMQ RCE	Remote code execution in Apache ActiveMQ via deserialization flaws (listed thrice).	Critical	128

PRITAM DASH

Suricata Detection Trends

The most frequent signature, tied to the DoublePulsar backdoor (ID 2024766, 49,435 hits), is a post-exploitation tool associated with the EternalBlue exploit, highlighting persistent attempts to leverage legacy SMB vulnerabilities in unpatched systems. Other prominent signatures include TCP stream anomalies (ID 2210020, SURICATA STREAM ESTABLISHED packet out of window, 43,357 hits) and traffic from blocklisted IP sources (ID 2402000, ET DROP Dshield Block Listed Source, 43,097 hits), both indicative of scanning, probing, or suspicious outbound connections from compromised or malicious hosts. Additional alerts, such as Terminal Server traffic on non-standard ports (ID 2023753, 37,417 hits), HTTP requests on unusual ports (ID 2006408, 30,020 hits), and RDP authentication bypass attempts (ID 2034857, 15,847 hits), underscore active efforts by attackers to exploit weak or misconfigured services. Classic reconnaissance techniques, including NMAP -sS scans (ID 2009582, 10,631 hits) and external IP lookups via TLS SNI (ID 2047703, 8,508 hits), further suggest adversaries are mapping networks for vulnerabilities.

SOC teams should prioritize tuning their SIEM systems to escalate these signatures, enabling rapid triage and investigation to mitigate potential breaches stemming from these pervasive threats.



04/04/2025

Conclusion

The analysis presents a clear picture: attackers are leveraging both old and new tactics with alarming persistence. From brute-force attempts using default credentials like root/123456, to exploitation of decade-old vulnerabilities such as CVE-1999-0517, the threat landscape remains riddled with risks many of which root from misconfigurations, unpatched systems, and exposed services.

Cloud providers like Google Cloud, DigitalOcean, and OVH have emerged as major sources of attack traffic, reflecting how threat actors continue to abuse scalable cloud infrastructure for launching campaigns. Meanwhile, OS fingerprinting revealed that Linux 2.2.x–3.x systems dominate as the attacker's platform of choice, likely due to ease of automation and stealth. The presence of stripped timestamps and lightweight builds further indicates efforts to evade detection.

The data also highlights a surge in alerts related to backdoor installations, brute-force RDP attempts, and malicious HTTP traffic, emphasizing attackers' focus on remote access and web-based attack vectors. Suricata signatures like DoublePulsar backdoor, Nmap scans, and authentication bypasses accounted for tens of thousands of hits.

The persistent exploitation of both legacy and modern CVEs demonstrates that attackers are employing automated scanners to probe networks broadly and relentlessly. Organizations that delay patching, overlook exposure points, or neglect logging and alerting are likely to remain vulnerable.

Final Thoughts

Security researchers must be as advanced as their adversaries. This means:

- Patching vulnerabilities without delay, especially known CVEs still being actively exploited.
- Monitoring for brute-force and scanning activity, particularly around RDP, SSH, and HTTP.
- Segmenting networks and limiting remote access, especially for critical systems.
- Deploying honeypots and threat intelligence integrations to gain early visibility into attack behavior.

The bottom line? Threat actors don't need new exploits when old ones still work. It's up to us to close the gaps before they do.

NAME
Pritam Dash

EMAIL
pritamdash1997@gmail.com

LINKEDIN
/dashpritam

WEB
https://brutefoss.com