IOSR Journal of Computer Engineering (IOSR-JCE)
e-ISSN: 2278-0661, p-ISSN: 2278-8727, Volume 27, Issue 6, Ser. 4 (Nov. – Dec. 2025), PP 91-98
www.iosrjournals.org

From Measurements to Defense: Strengthening IoT Security Using Cyber Threat Intelligence

Vishnu Dutt
Department of Computer Science
Hindu College of Pharmacy, Sonipat

Arvind Singh
Department of Computer Science
I K Gujral Punjab Technical University, Jalandhar

Abstract
The Internet of Things has kind of grown into one of the broadest, and also least-secured, tech ecosystems in history. With about 18.8 billion connected devices logged worldwide by the end of 2024, and IoT malware attacks climbing 400% just in the first half of 2023, the security community is staring at a problem that feels technical, organizational too, and also pushed around by policy decisions. This piece looks at how cyber threat intelligence (CTI), which is gathered from internet measurements, scanning platforms, honeypot setups, network traffic analysis, and more formalized sharing frameworks, can become a base layer for real IoT protection. It walks through the whole route from raw measurement data to useful intelligence, it checks how standards such as STIX/TAXII and frameworks like MITRE ATT&CK help make threat analysis and sharing more structured, then it weighs automated detection ideas that rely on machine learning. After that, it points at the structural obstacles: firmware fragility, sloppy patch management, and the fragmented way information gets exchanged. The overall argument is that the mismatch between what threat intelligence currently uncovers and what defenses actually apply is, by itself, the most urgent thing to fix in IoT security. Closing it, the text says, won't happen unless there's both technical innovation and some serious institutional commitment to shared situational awareness.
Keywords: *security, STIX/TAXII, MITRE ATT&CK, network measurement, cyber threat intelligence, IoT, threat sharing, anomaly detection*

I. Introduction

There's a useful thought experiment for anyone who wants to make sense of the IoT security problem. Imagine you plug in a brand-new router to your home network. Then picture that within seconds, as soon as it's online, tons of automated scanners have already spotted it and start poking at its open ports. In a few minutes, if that router is still using factory-default credentials, something (or rather, some thing that was infected months ago by some other event) is trying to log in, using a list of 60 pretty well-known username and password combinations.

And no, this isn't really hypothetical. Kaspersky's honeypot data for the first half of 2023 showed that almost 98% of network-related attacks aimed at IoT-ware went after the exposed Telnet interface (MDPI Sensors, 2024). Attackers aren't necessarily "smart" in the Hollywood sense — they're being thorough, systematic, and fast, mainly because the devices they're targeting don't change much year to year.

What's different now, though, and it's a big difference, is the quality plus sheer volume of data that defenders can access. Internet measurement platforms, honeypot networks, structured vulnerability databases, and threat-sharing protocols now produce a steady flow of real-world intelligence about who is attacking what, which tricks they're using, and how often those attempts actually work. The main challenge becomes turning that information into real defenses. That step, from measurement to protection you can actually act on , is basically what this article explores.

1.1 Scope and Structure

This article sorta focuses on the operational pipeline that ties cyber threat intelligence in with IoT defense, kind of end to end. It goes through the way measurements get gathered and put into order, and also how frameworks like STIX/TAXII and MITRE ATT&CK give that data its shape, and not just its meaning. Then it talks about how machine learning systems make use of it for automated detection, and it points out where the whole pipeline starts breaking down. Specifically it mentions firmware limitations, patching delays, information

sharing friction and the institutional coordination part that is always messy. The last sections, are more about what a more coherent, intelligence led defense architecture for IoT might actually look like.

II. The IoT Threat Landscape: Why Intelligence Is Urgently Needed

2.1 The Scale of the Problem

Numbers first, so lets just go with that. IoT Analytics reported a global count of connected IoT devices at 18.8 billion by the end of 2024, compared to 16.6 billion by the end of 2023—so basically a 15% year over year rise (IoT Analytics, 2024). And by 2030, they say it might get close to 39 billion. Each one of those devices is like a possible entry point, yet a lot of them don't really have the hardware muscle to run typical security software in the first place.

Then Zscaler's ThreatLabz looked at roughly 300,000 blocked attacks targeting IoT devices during January through June 2023, and saw a 400% year-over-year jump in IoT malware (Zscaler, 2023). The manufacturing industry took the biggest hit, pulling in 54.5% of all attacks, and averaging about 6,000 weekly attacks across the devices they monitored. The education space looked even worse, with malware attempts in schools rising by almost 1,000% over that same time window. This isnt some distant, abstract stats blip. It's more like real intrusions into real infrastructure, often through gadgets that were plugged in, and then quietly ignored.

Also, Verizon's 2024 Data Breach Investigations Report noted that one in three breaches now includes an IoT device (Verizon, 2024). And the 2024 IoT Security Landscape Report from Bitdefender and NETGEAR stated that home network devices deal with an average of 10 attacks every 24 hours. In the meantime, Bitdefender's systems blocked around 1,736 IoT threats each minute (Bitdefender/NETGEAR, 2024).

2.2 Why Conventional Security Falls Short

IoT devices are quite different from computers and smartphones in ways that really matter a lot for security. They usually have limited processing power, fixed firmware, very little or no user interface at all, and they keep running for years after the security support stops. The 2024 IoT Security Landscape Report pointed out buffer overflow (28.25%) and denial of service (27.20%) as the two most common vulnerability types. Both are pretty well known categories that people have been exploiting in IT settings for decades, and yet they keep showing up in devices that do not include basic protections (Bitdefender/NETGEAR, 2024).

More recently, research in MDPI Sensors (2024) noted that 34 out of the 39 most frequently exploited IoT vulnerabilities were on average over three years old. So, attackers are not really doing a "brand new" playbook, they are mostly putting energy into scale and automation. This imbalance is exactly why defense that is intelligence-driven is so important because defenders can't just wait for incidents to find out what is already known.

III. Cyber Threat Intelligence: From Raw Data to Actionable Insight

3.1 What Cyber Threat Intelligence Actually Is

CTI is kind of a broad term, but in practice it really does mean something fairly specific. Basically, it points to evidence based knowledge, about current or possibly upcoming threats, that then guides security decisions — like, who is attacking, how they operate, what kind of infrastructure they use, and what their objectives are. In the most plain sense, CTI contains indicators of compromise, the IoCs for short: like IP addresses, domain names, file hashes, and protocol signatures tied to nasty activity. Then at a more in-depth level, it also covers tactics, techniques, and procedures (TTPs): those behavioral patterns that sort of define how adversaries carry out their work.

The difference between IoCs and TTPs matters a lot especially for IoT. IoCs tend to be fleeting, attackers rotate IP addresses and domains pretty quickly. TTPs stick around, because they mirror how attackers think, and what circumstances they try to take advantage of. For example, an adversary using Mirai variants to hit Telnet on port 23 with default credentials, that is showing a TTP. So even if the actual IPs flip daily, the deeper behavior pattern can still be spotted, expected, and prevented.

Then there is the CTI lifecycle — direction, collection, processing, analysis, dissemination, and feedback — which gives you that organized method for turning raw measurements into defensive action (SANS Institute, 2023). Using this same lifecycle for IoT usually means some careful adjustments, mainly around collection (because IoT traffic doesn't look the same as enterprise IT traffic) and dissemination (IoT operators are often small manufacturers, and they may have limited security capacity, or just not enough personnel).

3.2 Structured Threat Information: STIX and TAXII

Sharing threat intelligence at scale needs kind of common standards, and the main ones in the cybersecurity community are STIX (Structured Threat Information Expression) and TAXII (Trusted Automated

Exchange of Intelligence Information). They were created by MITRE, and now they are treated as OASIS standards, so STIX 2.1 and TAXII 2.1 are the ones that got approved and released on June 10, 2021 (OASIS, 2021).

STIX acts like a modular, machine readable language for describing threat intelligence, it can reach observable events, malware characteristics, plus threat actor profiles and even attack campaigns. TAXII, on the other hand, is the transport mechanism for moving that information around via HTTPS, it also supports both a hub-and-spoke setup and a peer to peer distribution style. In practice, they end up being the technical spine for most today's threat intelligence sharing systems.

Now, an NDSS 2024 investigation that followed shared STIX material reported that by April 10, 2023, people had shared 10,392,889 STIX objects across major platforms, starting from the first object that appeared back in October 2014. That comes out to an average of about 3,371 objects each day (NDSS, 2024). In terms of formats, 69.74% of those objects were in STIX 1 form, while 30.26% used STIX 2, which is more expressive. So yes, this suggests real uptake, but it also highlights the growth headache, the whole ecosystem still has to migrate more toward STIX 2 so it can portray the full nuance of today's IoT threats.

As illustrated in Figure 1, the CTI pipeline for IoT defense follows a structured flow from raw measurement sources through intelligence processing to operational defensive controls.

Figure 1: End-to-End Cyber Threat Intelligence Pipeline for IoT Defense

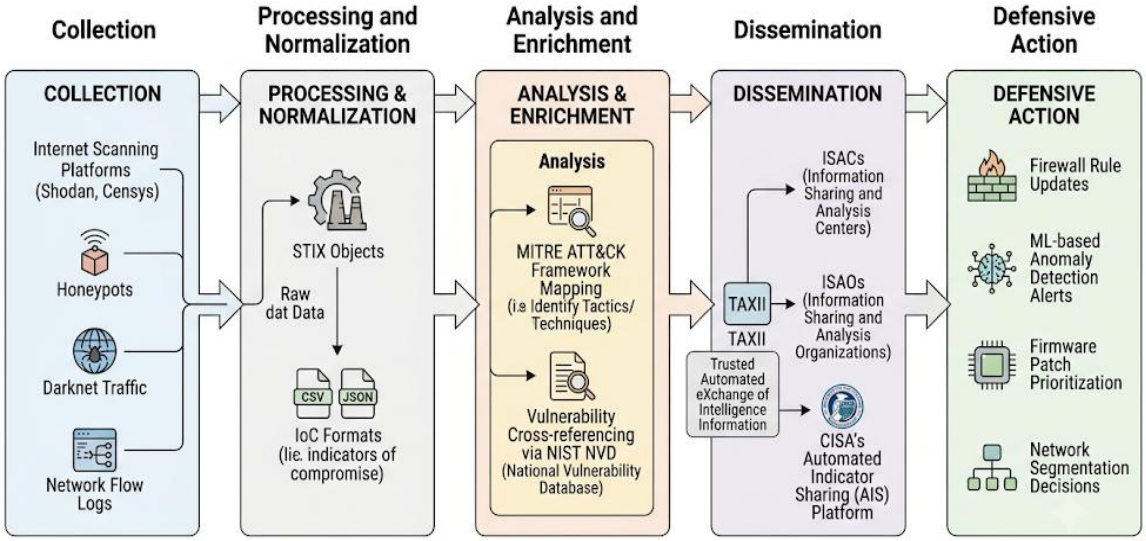

This flowchart shows the CTI pipeline like it goes left to right, basically as a process with five stages. Stage one , (Collection) shows the data inputs: internet scanning platforms (Shodan, Censys) , honeypots, darknet traffic, and network flow logs. Stage two (Processing and Normalization) shows the raw stuff being tidied up into STIX objects and IoC formats. Stage three (Analysis and Enrichment) is where they apply the MITRE ATT&CK framework mapping and then do vulnerability cross-referencing using the NIST NVD. Stage four (Dissemination) covers the sharing side, via TAXII channels to ISACs, ISAOs, and CISA's Automated Indicator Sharing (AIS) platform. Stage five (Defensive Action) finally shows the outputs: firewall rule updates, ML-based anomaly detection alerts, firmware patch prioritization, and even network segmentation decisions. The main idea here is that each stage adds value — and at the same time introduces possible weak spots — when you go from raw threat data into real device-level protection. Data flows referenced from CISA (2023), NDSS (2024), and Zscaler ThreatLabz (2023).

IV. Frameworks for Making Sense of IoT Threats

4.1 MITRE ATT&CK and Its Application to IoT

If STIX/TAXII gives the plumbing for sharing threat intelligence, MITRE ATT&CK sort of provides the grammar for trying to understand it. ATT&CK is this globally recognized knowledge base about adversary tactics and techniques, built from real-world observations, not just theory. The framework groups attacker behavior into tactics (the actual objective) and techniques (the method). It does this across the Enterprise, Mobile, and ICS (Industrial Control Systems) matrices. By 2023, the framework lists more than 300 individual techniques across 14 tactics in the Enterprise matrix alone (ATT&CK/MITRE, 2023).

Now for IoT and operational technology setups, the ICS matrix is especially relevant. It includes techniques that have been seen in real attacks against industrial control systems, and a lot of those follow

architecture and connectivity patterns that also show up in wider IoT deployments. IriusRisk's review of the ICS matrix pointed out that MITRE introduced the "Hardcoded Credentials" technique after it turned up as one of the top vulnerabilities mentioned in CISA ICS advisories during 2022 (IriusRisk, 2022). That sort of ongoing adjustment, tied to what's actually observed, is part of why ATT&CK feels genuinely useful as an intelligence framework, and not just a static checklist.

In IoT threat environments, security teams tend to apply ATT&CK in four main ways: threat modeling (to rehearse attack scenarios before an incident), gap analysis (to spot detection blind spots across different device categories), adversary emulation (to validate defenses against documented TTPs), and incident response, which speeds things up by giving responders a shared vocabulary for describing attacker behavior. Those activities become much more practical once the raw CTI is structured, enriched, and organized through a standardized framework, instead of being left as scattered notes.

4.2 Threat Sharing Infrastructure: ISACs, ISAOs, and AIS

Threat intelligence only really helps the organizations that actually have it. Like a decent chunk of the IoT ecosystem, and we're talking small manufacturers, municipal governments, healthcare providers, plus the smaller enterprises too, just don't have the internal know-how to collect, process, and then act on raw CTI. This is why sharing infrastructure ends up mattering more than people expect, even if it feels a bit abstract at first.

CISA's Automated Indicator Sharing (AIS) program, it was started in 2016 under the Cybersecurity Act of 2015, gives a way for real-time exchange of machine-readable IoCs and defensive measures across both government and private sector players, and at no cost (CISA, 2023). Then AIS 2.0 came out in March 2022, and it basically refreshed the platform's interface, also the way submissions are handled. The whole program runs on STIX and TAXII as its technical spine, and it links participants through ISACs (Information Sharing and Analysis Centers) and ISAOs (Information Sharing and Analysis Organizations).

Still, OIG data shows something kinda bleak. The number of AIS participants slid from 304 in 2020 to 135 in 2022, and IoC sharing via the program dropped 93% in that same stretch (SANS, 2023). That's troubling, and I'd say not a small change. It doesn't automatically mean threat sharing is finished—ISACs and sector-specific platforms keep moving and they operate pretty robustly— but it does suggest that a centralized government-style sharing channel is finding it harder to keep people engaged. For IoT security, specifically, where threats cut across every sector, and also across device categories, fragmented sharing infrastructure ends up creating fragmented situational awareness.

V. Detection Technologies: Putting Intelligence to Work

5.1 Anomaly Detection and Machine Learning in IoT Networks

The most direct use of CTI in IoT defense is kind of simple, you take enriched threat intel and feed it into detection systems that can spot hostile actions in real time. The older style, signature-based intrusion detection, sort of works ok for threats that are already cataloged. Basically it matches network traffic against known attack "templates" or patterns. But it breaks down, sometimes badly, when the attackers roll out novel variants, zero-day exploits, or those botnet campaigns that change slowly over time. They'll tweak their signatures just enough to slip past the rules and still look almost familiar.

Then machine learning-based anomaly detection takes a more roundabout strategy: it studies what "normal" IoT network traffic looks like, and later it raises an alert when something deviates from that learned baseline. A 2021 anomaly detection study in IEEE NetSoft looked at a system that dynamically profiles all connected IoT devices, and then keeps watching for behavioral oddities. When it was tested on the Cyber-Trust testbed, using real normal along with malicious traffic, the authors report overall detection accuracy around 98.35%, with a false-positive rate of about 0.98% (Rose et al., IEEE NetSoft, 2021). It sounds quite strong, although in actual deployments there's more background noise, so the "clean" results from controlled test environments often don't transfer perfectly.

Deep learning approaches have pushed performance yet again. Research published in Springer's Journal on Information Security evaluated several models using the CIC IoT-DIAD 2024 dataset, which is a broad collection of flow-based IoT network traffic covering both benign behavior and attack scenarios. A 1D Convolutional Neural Network (CNN) reportedly reached multi-class attack classification accuracy of 99.12%, while the LSTM model followed closely at 98.98% (Springer, 2025). These numbers are impressive, but training on a specific dataset can make generalization tricky, and that means real field deployment will have to handle things like distribution shift, calibration, and ongoing updates, instead of assuming the lab setup will stay the same.

Machine learning is kinda especially good for CTI- enhanced IoT defense because it can fold in threat intelligence right into the training signal. So, when fresh IoC patterns, or TTP signatures get discovered via intelligence sharing, and then they're used to create labeled training data, the models tend to adjust faster than

humans would by hand, writing those signature rules. That loop, yknow, between collecting CTI and improving detection, is probably one of the most promising pockets of active development right now.

5.2 The Firmware Problem

Even the best detection system, running on a network gateway, cannot fully compensate for what happens inside a compromised device. Firmware—the low level software that sort of steers how an IoT device operates— is often the real point of exploitation. The 2024 MDPI Sensors firmware review noted that nearly 98% of network related attacks against IoT devices in the first half of 2023 went after the unsecured Telnet interface, and they did so by directly triggering firmware-level authentication failures (MDPI Sensors, 2024).

A lot of IoT devices keep running firmware for years after the manufacturer support window is basically over. Compared with operating systems on normal computers, IoT firmware usually lacks automated update mechanisms, it also misses cryptographic integrity checks, and in most deployments there is no rollback safety at all. Research in Springer's Discover Internet of Things (2023) documented that delayed patching plus a lack of security testing by manufacturers are the main reasons IoT devices stay active targets for large scale attacks, and in many cases reverse engineering is needed to uncover vulnerabilities hidden inside closed source firmware images (Springer, 2023).

So the CTI practical take away is this: intelligence about firmware vulnerabilities only becomes useful when there is some kind of delivery path for remediation. When 60% of IoT devices still carry unpatched known vulnerabilities older than two years, and 75% of IoT devices have no automated update mechanisms (Wifitalents compilation, 2024), then those vulnerability details don't really turn into patching behavior. In other words, the measurement-to-defense pipeline stalls here—not because the intelligence is missing, but because the update infrastructure is missing, or it's too weak to act.

Table 1: Key IoT Attack Vectors, Intelligence Sources, and Associated Defense Mechanisms

Attack Vector	Prevalence / Data Point	Primary CTI Source	Defense Mechanism	Source
Default/weak credentials (Telnet, SSH)	98% of IoT network attacks via unsecured Telnet (H1 2023)	Honeypot captures, Shodan scans	Credential policy enforcement; port blocking	MDPI Sensors, 2024
Unpatched firmware vulnerabilities	60% of IoT devices carry CVEs older than 2 years	NIST NVD, CVE databases	OTA firmware update mechanisms; vulnerability triage	Wifitalents, 2024
Botnet malware (Mirai, Gafgyt)	66% of IoT attack payloads in H1 2023	STIX/TAXII IoC feeds, darknet traffic	ML anomaly detection; network segmentation	Zscaler ThreatLabz, 2023
Lateral movement via compromised IoT	1 in 3 breaches involves an IoT device (2024)	Network flow analysis, ATT&CK ICS mapping	Zero trust architecture; VLAN isolation	Verizon DBIR, 2024
Shadow IoT devices in enterprise networks	Proliferation across education and manufacturing in 2023	Device fingerprinting, traffic profiling	Continuous asset discovery; network monitoring	Zscaler ThreatLabz, 2023

VI. Structural Barriers: Where the Pipeline Breaks

6.1 The Shadow IoT Problem

Intelligence-led defense kind of assumes you already know what devices are sitting on your network. In a lot of enterprise or institutional setups, that idea doesn't really survive for long. Zscaler's ThreatLabz 2023 report, for example, calls out the growth of shadow IoT devices—connected gizmos that show up on enterprise networks without any formal IT knowledge or approval—as a big risk driver for those first attacker entry points (Zscaler, 2023). Like when a school staff member brings in a personal smart speaker and simply connects it to the institutional network, or when a hospital worker plugs in a consumer-style connected camera, the thing becomes part of the network attack surface, without ever going through an inventory record.

CTI becomes extremely hard to use, when the security team can't even name the devices that exist. That's why continuous network discovery and device fingerprinting—basically comparing traffic patterns to device type profiles—tends to be one of the most workable routes to surface shadow IoT. In fact, Zscaler's own ThreatLabz reporting approach used device fingerprinting to sort out 850+ distinct device types across 3 trillion IoT transactions over a three-month observation window. If organizations use that same method on the defensive side, they can get visibility into what's actually present on their networks, before attackers arrive and get there first.

6.2 Intelligence Quality and Sharing Fragmentation

Not all CTI is the same, like, not even close. The NDSS 2024 study on STIX sharing, found that 37.89% of the shared STIX objects had substantial redundancy even when it was just from a single provider. And it also says that security providers, together, only make about 2,063 unique threat intelligence objects per day across the main sharing platforms, which the researchers basically called "inadequate for increasing cyber

threats" (NDSS, 2024). So yeah, duplicated and low-quality intelligence brings in more noise, it pulls analyst focus away without actually improving detection results.

The AIS participation decline mentioned earlier shows a kind of related structural issue: threat sharing infrastructure needs real community upkeep, not just "having a system." An NDSS 2024 analysis of AIS reported that IOC sharing via the platform fell 93% from 2020 to 2022. Even when the technical pieces are there, people still tend to drop off if they feel the contribution payoff isn't there. For smaller IoT device manufacturers or operators, who often do not have dedicated security teams, the cost and benefit of contributing to these threat sharing platforms, usually never pencils out.

6.3 The Telemetry Gap in IoT Devices

Consumer and enterprise IT systems can run endpoint detection agents, record huge amounts of telemetry, and send questionable binaries to a central analysis place. A smart thermostat, an industrial sensor, or a connected camera just can't. IoT devices usually produce only a small amount of network telemetry — mostly traffic patterns rather than detailed packet content, and timing hints rather than process logs — so the threat intelligence you can pull from that weak signal is sort of limited too.

Because of that telemetry gap, CTI for IoT ends up being very network-focused. It tells you what devices are talking to, not what is going on "inside" them. This kind of view works nicely for spotting botnet command-and-control activity and odd outbound connections. It does less for catching firmware tampering, local credential stealing, or staged attacks that don't trigger any weird network behavior while an attacker is doing early reconnaissance. Work on hybrid IoT/OT honeypots has been especially useful here, documenting the gap, and showing that attacker actions inside already compromised devices often leave no network visible trace (ACSAC, 2022).

VII. Discussion

7.1 Building a Coherent CTI-Driven Defense Architecture

The most effective IoT security architectures kind a treat CTI not as an add-on, but as the whole organizing thing of the defense posture. Like you start with continuous asset discovery and network baselining — because you can't protect what you can't see. Then, you push enriched threat intelligence straight into automated detection systems, so the time between when an IOC gets published and when a firewall rule finally shows up drops from days to minutes. And network segmentation, especially VLAN-based separation of IoT devices from the core enterprise network, directly tackles the lateral movement issue described in Verizon's DBIR— basically containing the blast radius when a device gets compromised.

Then there's the zero trust architecture idea Zscaler's CISO Deepen Desai advocates, "never trust, always verify, and assume breach" — and that sort of shifts CTI into a continuous verification posture, not a perimeter defense posture (Zscaler, 2023). In the IoT world, this means no device is given automatic trust just because it lives on some network segment, and all device communication is checked continuously against expected behavioral profiles. This doesn't magically erase the telemetry gap inside many IoT devices themselves, but it compensates for it by making the network the sensor, instead.

Also, when you map the MITRE ATT&CK ICS matrix to an organization's specific IoT deployment setup it can point out which attack techniques are actually plausible against which device categories. That helps you invest defensively in a prioritized way, rather than running generic security checklists. And in manufacturing— which is Zscaler's most targeted sector in 2023 — teams have extra reasons to go for this more structured method, since OT attacks in industrial environments can lead to physical consequences, that go way beyond simple data loss.

7.2 Prioritizing What Gets Acted On

Given the volume of available IoT threat intelligence, prioritization is increasingly the core, competency. Like, not every vulnerability can be patched immediately, not even. And no, not every IoC really warrants a block rule. The NIST risk based triage approach—putting enrichment resources on vulnerabilities that show up in CISA's Known Exploited Vulnerabilities (KEV) catalog—kind of reflects this idea. For IoT defenders, the equivalent move is to aim patching and configuration remediation efforts at those device categories and vulnerability classes that keep appearing most often in the current threat intelligence, instead of trying to fix everything known at once, or, all issues simultaneously.

As shown in Figure 2, the relationship between IoT attack vector prevalence and current defense adoption reveals a persistent gap: the most exploited vectors are not always the most actively defended.

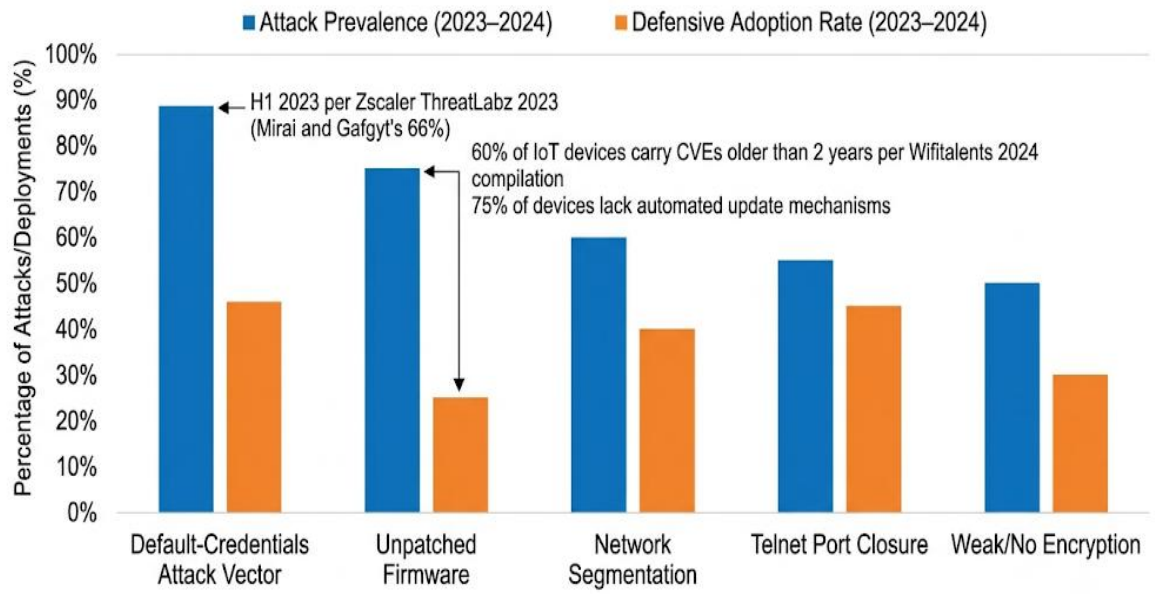

Figure 2: IoT Attack Vector Prevalence versus Current Defensive Adoption Rate (2023–2024)

This paired bar chart compares, for five big IoT attack vectors, the percent of documented attacks that lean on each vector (left bar, blue) versus the estimated percent of IoT deployments where that same vector is actively addressed (right bar, orange). The default credentials vector looks like it has the most action, with the highest attack prevalence (it is referenced in almost all major botnet campaigns, like Mirai and Gafgyt, and together they made up 66% of IoT attack payloads in H1 2023 per Zscaler ThreatLabz 2023) but the mitigation side is only kind of moderate, not really matching that level. Unpatched firmware also sits high, it shows strong prevalence (60% of IoT devices still have CVEs that are older than 2 years, per Wifitalents 2024 compilation), yet mitigation adoption is pretty low, since 75% of devices do not have automated update mechanisms in place. Meanwhile, network segmentation and Telnet port closure are getting better, though adoption is still incomplete. The main message kinda lands on this, the biggest mismatch between threat prevalence and defensive effort lines up exactly with the least costly and most practically doable mitigations, which suggests the real limiter is organizational, not technical.

VIII. Conclusion

The intelligence is there, basically. That's the kind of thing you want to say early on before you jump to any conclusion about CTI and IoT security. You know, scanning platforms, honeypots, botnet trackers, structured sharing protocols, and machine learning detection systems are all together producing more and better threat intelligence about IoT attack patterns than at any time in the technology history. The 400% surge in IoT malware that Zscaler documented in 2023, the 66% share of attack payloads tied to Mirai and Gafgyt variants, the near universal focus on default Telnet credentials — it's all already known, mapped out, and shared.

What still feels genuinely hard is using that intelligence in a systematic way, and at scale, across a device ecosystem that was never built with security intelligence pipelines in mind. Firmware update infrastructure is kinda inadequate. "Shadow IoT" devices slip around asset inventory systems. Participation in threat sharing on centralized platforms is trending downward. And the telemetry IoT devices produce is just too thin, like, not enough for deep behavioral analysis inside the device itself.

Meanwhile, the regulatory frameworks coming out in the EU and the UK — requiring lifecycle security support, vulnerability disclosure, and incident reporting — they lay down the institutional scaffolding that CTI-led defense needs to work at scale. When manufacturers are legally required to ship security updates and report exploited vulnerabilities, the whole intelligence ecosystem gets a critical extra input: manufacturer originated, device specific threat data that no outside scanner can realistically reproduce.

Ultimately, the journey from measurement to defense in IoT security is kind of a team sport, you know. Researchers mapping exposure patterns, vendors sharing threat indicators, regulators mandating baseline security, manufacturers building update mechanisms, and operators rolling out network segmentation plus zero trust controls all of these things have to work together, not just in theory. If any one of them acts alone the effect stays limited, even if the effort is solid. The real question is whether the coordination infrastructure — technical,

regulatory, and institutional — can mature fast enough to stay ahead of an attacker community that is already coordinated really well.

References

[1]. Ahmed, M., Panda, S., Xenakis, C., &Panaousis, E. (2022). MITRE ATT&CK-driven cyber risk assessment. *Proceedings of the 17th International Conference on Availability, Reliability and Security*, 1–10. https://doi.org/10.1145/3538969.3539004

[2]. Annual Computer Security Applications Conference (ACSAC). (2022). *Interaction matters: A comprehensive analysis and a dataset of hybrid IoT/OT honeypots*. ACM. https://dl.acm.org/doi/fullHtml/10.1145/3564625.3564645

[3]. Bitdefender, & NETGEAR. (2024). *The 2024 IoT security landscape report*. NETGEAR. https://www.netgear.com/hub/network/2024-iot-threat-report/

[4]. Chorfa, W., Ben Youssef, N., & Jemai, A. (2023). Threat modeling with MITRE ATT&CK framework mapping for SD-IoT security assessment and mitigations. *2023 IEEE Symposium on Computers and Communications (ISCC)*, 1323–1326. https://doi.org/10.1109/ISCC58397.2023.10218246

[5]. Cybersecurity and Infrastructure Security Agency (CISA). (2023). *Automated indicator sharing (AIS)*. U.S. Department of Homeland Security. https://www.cisa.gov/topics/cyber-threats-and-advisories/information-sharing/automated-indicator-sharing-ais

[6]. El-Kosairy, A., & Azer, M. A. (2023). A survey on IoT and embedded device firmware security: Architecture, extraction techniques, and vulnerability analysis frameworks. *Discover Internet of Things*, 3(1), 1–38. https://doi.org/10.1007/s43926-023-00045-2

[7]. European Union Council. (2024, October 10). *Cyber Resilience Act: Council adopts new law on cybersecurity requirements for digital products*. Council of the European Union. https://www.council.europa.eu/en/news/detail/en/PRESS_1024_24

[8]. Greynoise Intelligence. (2024, December). *Checking it twice: Profiling benign internet scanners — 2024 edition*. Greynoise Blog. https://www.greynoise.io/blog/checking-it-twice-profiling-benign-internet-scanners----2024-edition

[9]. Intelligence and National Security Alliance (INSA). (2024). *The importance of information sharing across sectors in defending U.S. cyberinfrastructure*. INSA. https://industrialcyber.co/threat-landscape/insa-paper-highlights-importance-of-information-sharing-across-sectors-in-defending-us-cyberinfrastructure/

[10]. IoT Analytics. (2024). *State of IoT summer 2024: Number of connected IoT devices growing 13% to 18.8 billion globally*. IoT Analytics Research. https://iot-analytics.com/number-connected-iot-devices/

[11]. IriusRisk. (2022). *MITRE ATT&CK for industrial control systems in IriusRisk*. IriusRisk Blog. https://www.iriusrisk.com/resources-blog/mitre-attck-for-industrial-control-systems-in-iriusrisk

[12]. MITRE Corporation. (2023). *MITRE ATT&CK: Enterprise matrix*. MITRE. https://attack.mitre.org/matrices/enterprise/

[13]. Network and Distributed System Security Symposium (NDSS). (2024). *Sharing cyber threat intelligence: Does it really help?* NDSS 2024 Proceedings. https://www.ndss-symposium.org/wp-content/uploads/2024-228-paper.pdf

[14]. OASIS Open. (2021, June). *STIX 2.1 and TAXII 2.1 approved as OASIS standards*. OASIS. https://www.oasis-open.org/2021/06/10/oasis-approves-stix-v2-1-and-taxii-v2-1-as-standards/

[15]. Rose, J., Swann, M., Bendiab, G., Shiaeles, S., & Kolokotronis, N. (2021). Intrusion detection using network traffic profiling and machine learning for IoT. *2021 IEEE 7th International Conference on Network Softwarization (NetSoft)*. https://doi.org/10.1109/NetSoft51509.2021.9492685

[16]. SANS Institute. (2023). *Reflections on the U.S. government's OIG report on CISA's Automated Indicator Sharing program*. SANS Blog. https://www.sans.org/blog/reflections-on-the-us-governments-oig-report-on-cisas-automated-indicator-sharing-program/

[17]. Tuptuk, N., Hazell, P., Watson, J., & Hailes, S. (2024). A review of IoT firmware vulnerabilities and auditing techniques. *Sensors*, 24(2), 708. https://doi.org/10.3390/s24020708

[18]. Verizon. (2024). *2024 data breach investigations report*. Verizon Business. https://www.verizon.com/business/resources/reports/2024-dbir-data-breach-investigations-report.pdf

[19]. Zscaler ThreatLabz. (2023, October). *ThreatLabz 2023 enterprise IoT and OT threat report*. Zscaler, Inc. https://www.zscaler.com/press/zscaler-threatlabz-finds-400-increase-iot-and-ot-malware-attacks-year-over-year-underscoring