NOZOMI NETWORKS LABS

2025 2H REVIEW

OT/IoT Cybersecurity Trends and Insights

February 2026

About Nozomi Networks Labs

Nozomi Networks Labs is dedicated to reducing cyber risk for the world's industrial and critical infrastructure organizations. Through its cybersecurity research and collaboration with industry and institutions, it helps defend the operational systems that support everyday life.

The Labs team conducts investigations into industrial device vulnerabilities and, through a responsible disclosure process, contributes to the publication of advisories by recognized authorities.

To help the security community with current threats, they publish timely blogs, research papers and free tools.

The **Threat Intelligence** and **Asset Intelligence** services of Nozomi Networks are supplied by ongoing data generated and curated by the Labs team.

To find out more, and subscribe to updates, visit **nozominetworks.com/labs**

Table of Contents

1. Executive Summary

Top Techniques and Targets

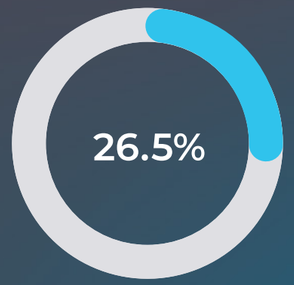

of all the raised alerts were associated with the **Adversary-in-the-Middle** technique

Transportation

Manufacturing

Government

were the three sectors that produced the highest number of alerts per organization

1 **United Kingdom**

2 **Germany**

3 **Australia**

The **UK, Germany and Australia** were the countries that produced the highest number of alerts per organization

Top Malware

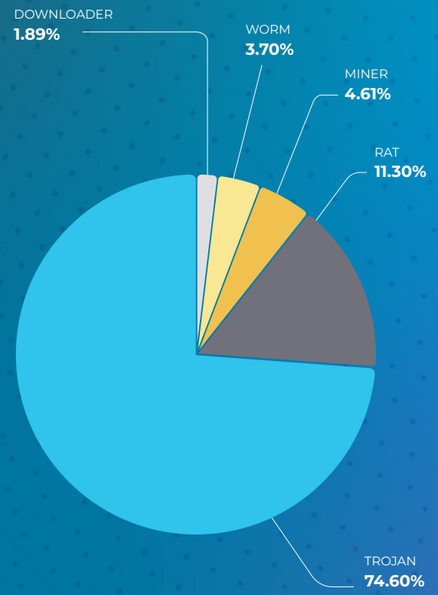

Top detected malware categories were **TROJAN, RAT, MINER, WORM and DOWNLOADER**

Top detected malware families were:

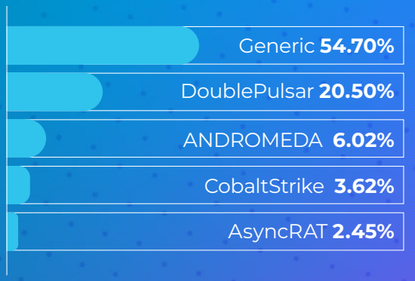

Top detected threat actors were:

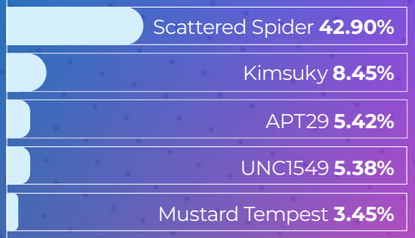

Vulnerability Statistics

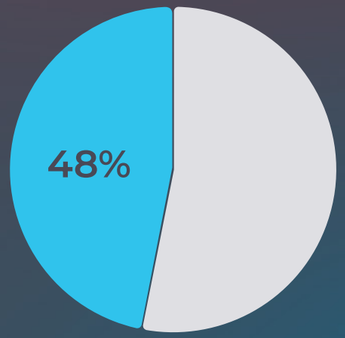

of the recent 2025 vulnerabilities present in observed environments have a CVSS score **HIGH or CRITICAL**

In the OT domain, the most commonly observed vulnerabilities discovered in 2025 were the ones affecting **Siemens, Rockwell Automation** and **Schneider Electric** devices

Among all vulnerabilities discovered in 2025 and detected in current environments, the **Use After Free** category was the most prevalent

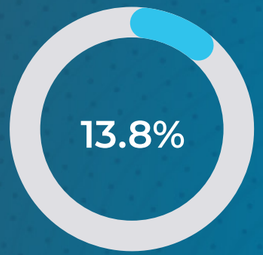

It may lead to **crashes, data corruption** or attacker-controlled **code execution**

Wireless Statistics

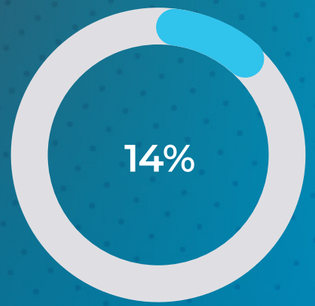

of observed networks use **open or legacy** security modes

Enterprise-grade authentication such as 802.1X is observed in only

0.3% of detected Wi-Fi networks

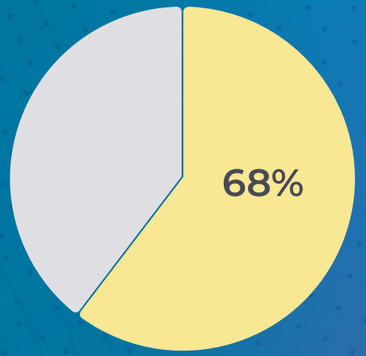

of observed wireless networks still operate **without Management Frame Protection (MFP),** which provides protection against deauth attacks

Nozomi-Discovered Zero-day Vulnerabilities

Average CVSS score: 7.0

CVSS distribution:

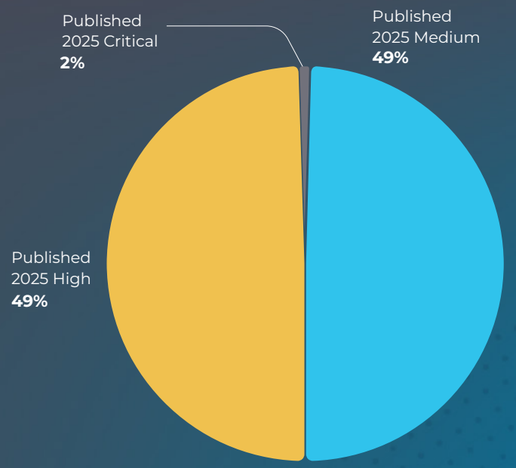

Findings were split:

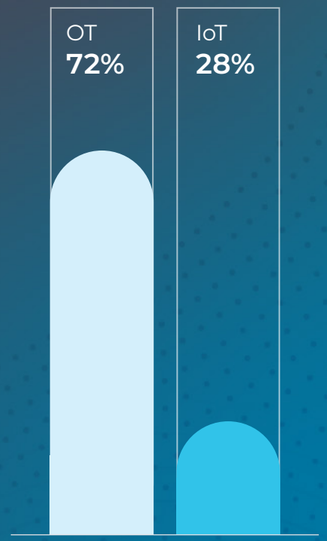

Responsible disclosure required mean time to disclosure

181 days / median 171

(-2% / -9% of previous period)

Botnet Statistics

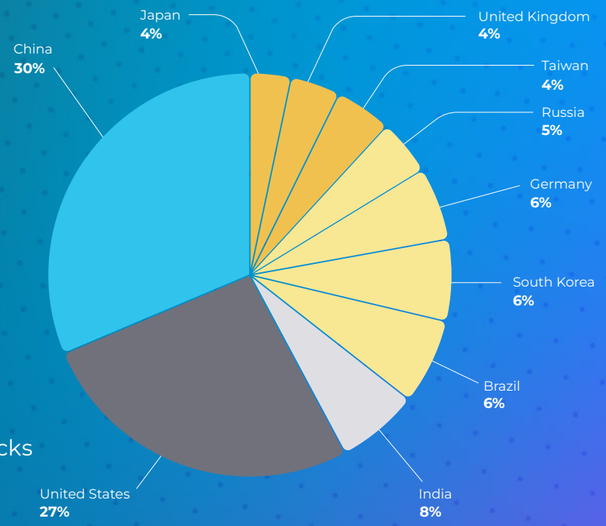

The majority of the attacks against our honeypots came from **China**

The highest botnet activity occurred on **September 2,** 2025, when attacks from **1,169** different IP addresses were recorded in one day

UPX 3.94 is still the most common packer used by attackers to protect IoT malware

2. Introduction

The second half of 2025 was an intense period in terms of the evolution of global cyber risks. Cyber adversaries, from organized crime groups to state-linked actors, continued actively targeting both public and private ecosystems, affecting various sectors including high-impact ones like transportation, government and other critical infrastructure. These developments underscore the continued transformation of cyber threats from isolated incidents into systemic risks with geopolitical and economic consequences.

One of the most disruptive incidents in late 2025 was a ransomware cyberattack on infrastructure software operated by Collins Aerospace, which led to widespread flight check-in and baggage handling disruptions at major European airports in September. The fallout highlighted how operational technology (OT) interdependencies can amplify the impact of cyberattacks on real-world services.

AI continues to be used more and more by both attackers and defenders. Here at Nozomi, apart from classic software development assistance, we actively use it in multiple dedicated research and customer-facing systems. Examples include tools to assist the discovery of new vulnerabilities, automatic security advisory parsing, spotting new strains of malware, correctly identifying customers' assets (Asset Intelligence), and also in our Vantage IQ product that, among other things, reduces the response time by facilitating more efficient communications between the user and the system. Every day, AI reshapes the cybersecurity realms and introduces new opportunities as well as new challenges. Ranging from better crafted phishing messages (which are very effective according to **recent research**) to faster malware development (a great example of which is **VoidLink**), its influence on the modern cybersecurity landscape cannot be understated. For example, in our recent **ransomware trends overview**, we mentioned 3rd parties reporting early experiments and emerging attempts to create AI-powered ransomware. Nozomi Networks Labs closely monitors the situations and **performs various steps** to be able to detect the new families of malware proactively.

Government and public sector entities were not immune either. In late 2025, the Kensington and Chelsea Council in London experienced a breach of historical resident data, prompting outreach to approximately 100,000 households due to the elevated risk of follow-on scams and identity fraud. Our telemetry also reported a rise of cyberattacks against organizations belonging to this industry, with government becoming the third most targeted sector in the second half of 2025 (see the corresponding section below).

Geopolitical tensions also reshaped the cyber threat landscape. **Poland reported an attack** against its critical infrastructure by the Russia-linked actor utilizing DynoWiper malware. Polish CERT attributed it to the Dragonfly group,

while other research groups cited Sandworm. In 2H 2025, Nozomi Networks products raised more than 20 alerts detecting Sandworm activity across the globe. The top targeted countries were the U.S., the UK and Colombia, with the Manufacturing and Transportation sectors being affected the most. The Dragonfly actor was also observed in several countries, most frequently in the US. Critical Taiwanese infrastructure reportedly endured persistent cyberattacks originating from China, with multiple daily incidents. China was also identified as the most active cybersecurity threat for the U.S. in

Executive Order 14306 released June 11, 2025, which introduced mandatory secure software development requirements and promoted initiatives such as the Cyber Trust Mark for federally procured IoT devices. In the EU, the RED Delegated Act mandated cybersecurity requirements (network, data privacy, anti-fraud) for radio IoT devices effective August 1, 2025. We will talk about wireless-specific threats in greater detail later in this report.

As for prominent new vulnerabilities, in early December 2025, a critical remote code execution (RCE) vulnerability dubbed React2Shell (CVE-2025-55182) was publicly disclosed and quickly became a widely exploited security crisis across internet-facing services worldwide. Our telemetry indicated high volumes of exploitation attempts targeting this vulnerability all over the globe. While React2Shell originated in the web application and cloud service world, its consequences impacted companies across multiple sectors, including those with OT and IoT assets due to the ongoing convergence of OT, IoT and IT domains. CISA added this vulnerability to its **Known Exploited Vulnerabilities (KEV) catalog**, indicating it is being exploited in the wild and underscoring the need to prioritize remediation, particularly for the government sector.

Collectively, these events underscore a strategic inflection point in which attackers leveraged both traditional techniques, such as data theft and ransomware extortion, and advanced techniques involving vulnerability exploitation and the use of AI. This period reinforces the importance of proactive security governance with a focus on visibility, real-time threat detection, and comprehensive resilience planning as organizations should prepare for an even more complex risk landscape heading into 2026.

3. Telemetry and Threat Trends

The anonymized telemetry collected from a large and diverse set of participating customers worldwide comprises a core component of our threat intelligence research. The data comes from production environments across multiple industries and regions and is collected exclusively from deployments of our visibility and cybersecurity platform. This telemetry is leveraged to enhance detection capabilities, refine signatures, identify emerging threats and, for this report, produce aggregated statistics that provide a broad and empirically grounded view of observed malicious activity.

The dataset reflects years of iterative refinement aimed at reducing noise and increasing analytical accuracy. All telemetry undergoes extensive validation, filtering, and de-duplication processes to ensure that false positives are excluded from the analysis. Artificially generated traffic, including laboratory simulations and traffic replays, is explicitly omitted.

While this dataset is not intended to be exhaustive of all global malicious activity, it provides a consistent, high-fidelity perspective on threats observed in real-world environments. The following sections summarize the key findings derived from this dataset during the second half of 2025 (2H 2025).

3.1 Top Techniques Used by Attackers

We continuously analyze the MITRE ATT&CK® techniques that are most commonly used by attackers trying to breach the organizations we protect

Here we aggregate the activated alerts raised in our customers' environment to understand the current trends in the OT and IoT cybersecurity landscape.

Global Trends

Here are the techniques that were used the most between June and December 2025, compared to the first half of 2025:

Top 10 MITRE ATT&CK Techniques

ID	Technique name	Tactics	%	% 1H 2025
T1557	Adversary-in-the-Middle	Credential Access; Collection	**26.5%**	**16.00%**
T1110	Brute Force	Credential Access	**10.2%**	**7.36%**
T1498	Network Denial of Service	Impact	**9.54%**	**17.60%**
T0814	Denial of Service	Inhibit Response Function	**9.39%**	**17.40%**
T1565	Data Manipulation	Impact	**8.36%**	**4.11%**
T0846	Remote System Discovery	Discovery	**7.13%**	**11.40%**
T0841	Network Service Scanning	Discovery	**7.13%**	**11.40%**
T0812	Default Credentials	Lateral Movement	**5.81%**	**5.27%**
T0859	Valid Accounts	Persistence; Lateral Movement	**5.81%**	**5.27%**
T1071	Application Layer Protocol	Command And Control	**3.62%**	**1.33%**

The infamous Adversary-in-the-Middle (AiTM) technique (also known as Man-in-the-Middle or MiTM) took first place, associated with over one-quarter of all alerts. AiTM is generally used to sniff sensitive information, including credentials, which can later be used in other stages of the attack. Compared to 1H 2025, it moved up two positions, signifying an increased interest from attackers. For example, a recent **Evilginx-based AiTM phishing campaign** targeted multiple U.S. universities.

The second most frequent technique was Brute Force, commonly used to guess valid credentials by trying many different possible combinations. Like the AiTM technique, it went up a few positions compared to the beginning of last year. The increase in brute-force activity is consistent with the high activity of IoT botnets, which routinely attempt automated logins at scale using password-guessing to propagate and grow in scale. As more compromised IoT devices are added to these botnets, attackers gain additional distributed infrastructure to probe exposed services continuously, driving up brute-force volumes.

Network Denial of Service (DoS) and Denial of Service were each used in slightly less than 10% of attacks, down from the ~17% that we observed between January and June 2025. These techniques pose an increased risk for OT environments, which typically are extremely sensitive to downtime and unavailability of services. DoS cyberattacks carried out by Russia on Danish election-related infrastructure are an example.

Another popular technique was Data Manipulation, associated with violating the integrity of the data whether it is stored, transferred or runtime. It rose from #9 to #5 in our chart. An increase in data manipulation attacks is especially dangerous because it can silently alter critical values and records, leading to wrong decisions, unsafe operations, compliance breaches and costly downstream impacts before anyone realizes the data can't be trusted.

The dominance of AiTM and credential-centric techniques over several reporting periods likely represents a structural shift rather than a tactical fluctuation. Attackers are increasingly optimizing durable access over immediate disruption, favoring techniques that scale across converged IT/OT/IoT environments and remain effective even in partially monitored networks.

Industry Insights

To minimize any bias towards the number of Nozomi customers in different industries, we identified the most targeted industries by the number of alerts per customer in each industry, not just by the total number of alerts raised in each industry. We also excluded industries in which the number of participating customers was not yet statistically significant enough to draw any solid conclusions.

With that in mind, here are the top three at-risk sectors in the second half of 2025 according to our telemetry:

Transportation

Manufacturing

Government

Transportation was the most targeted industry in both the first and second halves of 2025. This is alarming, given how much our daily lives and operations depend on it. The KNP incident in July 2025, involving a 158-year-old company in the UK with 700 employees that went bankrupt because of a cyber incident, is a grim reminder of the potential consequences.

The Manufacturing sector again came in second, reminding us how attractive it is for attackers. A successful cyberattack against manufacturers may cause tremendous financial losses, including the downtime associated with the investigation and remediation steps.

Finally, the number of attacks against the Government sector spiked, replacing Business Services at #3. Given how geopolitical tensions have increased across many different regions, this isn't surprising. Not only state-sponsored APTs but also financially motivated actors from different countries may attempt to target rival government infrastructure as a form of hacktivism. Cyberattacks on government may be especially harmful because they can disrupt essential public services, expose sensitive citizen data and undermine trust in institutions that communities rely on.

The diversity of attacker behavior across sectors reveals that adversaries don't pursue a one-size-fits-all strategy. Government environments are being aggressively probed, mapped and stressed, with discovery and DoS techniques dominating the activities. These patterns indicate reconnaissance at scale and intent to disrupt, signaling preparation for future campaigns or geopolitical pressure rather than immediate monetization.

In contrast, the transportation and manufacturing sectors show a far more balanced threat profile. Attackers are prioritizing credential theft there, laying out the groundwork for long-term persistence. These sectors appear to be viewed as high-value operational targets where stealth, durability and timing will matter more than having an immediate impact.

Looking at each of them in greater detail, here are the techniques most widely used by attackers when targeting them.

Transportation

ID	Technique name	Tactics	%
T1557	Adversary-in-the-Middle	Credential Access; Collection	**20.40%**
T1110	Brute Force	Credential Access	**19.70%**
T1565	Data Manipulation	Impact	**13.70%**
T0846	Remote System Discovery	Discovery	**8.38%**
T0841	Network Service Scanning	Discovery	**8.38%**

Just as in the global overview, the AiTM and Brute Force techniques lead the chart here, together comprising about 40% of the alerts seen by our transportation customers. AiTM was also the most common technique used against this sector during 1H 2025. Instead of running DoS attacks, attackers preferred to explore the environment by sticking to the Remote System Discovery and Network Service Scanning techniques.

Manufacturing

ID	Technique name	Tactics	%
T1557	Adversary-in-the-Middle	Credential Access; Collection	**30.30%**
T1565	Data Manipulation	Impact	**8.76%**
T1110	Brute Force	Credential Access	**8.58%**
T1498	Network Denial of Service	Impact	**8.14%**
T0814	Denial of Service	Inhibit Response Function	**7.95%**

AiTM was also the top technique in manufacturing, up from #3 during 1H 2025. Its jump underscores a clear shift in attacker focus toward credential interception and access-driven tactics. The distribution of other techniques is also quite similar to what we observe in the global overview.

Government

ID	Technique name	Tactics	%
T0846	Remote System Discovery	Discovery	**24.50%**
T0841	Network Service Scanning	Discovery	**24.50%**
T1498	Network Denial of Service	Impact	**21.80%**
T0814	Denial of Service	Inhibit Response Function	**21.80%**
T1557	Adversary-in-the-Middle	Credential Access; Collection	**3.55%**

The Government sector looks very different. Here, Discovery tactics and DoS attacks are the most popular, while AiTM comprised only 3.5% of all observed attacks.

One explanation may be that currently about half of attackers are still exploring the environments they eventually intend to attack, while the other half is already attempting to bring the systems they discovered down.

Regional Drilldown

To minimize any bias towards the number of customers that Nozomi Networks has in each country, we again focused on the areas with the highest number of alerts per customer, not just the total number of alerts in each country. We also excluded countries where we still don't have a statistically significant number of customers contributing anonymized telemtry.

With that in mind, here are the top three countries:

1  **United Kingdom**

2  **Germany**

3  **Australia**

The most surprising observation is the UK jumping to #1, with its OT and IoT environments now producing the highest number of alerts per organization

according to our telemetry. This is both unexpected and potentially alarming, as a spike in alerts is often associated with either weakened cybersecurity posture or an increased interest of attackers. This is the moment for companies to turn cybersecurity from a checklist into a strategic advantage. By actively reviewing defenses, addressing weaknesses, and staying ahead of emerging threats, organizations can protect operations and build long-term resilience.

Germany stayed in second place for the full year, indicating that attackers continue to see the country's highly connected industrial organizations as high-value, high-impact targets. That stability is also consistent with the shift observed in 2H 2024, when Germany moved up one position from third place and then stayed there, suggesting the change wasn't a one-off spike but a sustained focus. This is consistent with the broader telemetry-backed themes we have observed there, including persistent discovery behaviors and high volumes of DoS-related activity.

Finally, Australia moved from #4 to #3 during this period. It has consistently ranked among the most targeted countries in our data, and its continued upward movement indicates sustained attacker focus there — underscoring the need for appropriate measures to keep environments secure and operational.

Now, let's see a drilldown of the most commonly used attack techniques in each of these countries.

United Kingdom

ID	Technique name	Tactics	%
T1557	Adversary-in-the-Middle	Credential Access; Collection	**88.20%**
T0812	Default Credentials	Lateral Movement	**3.10%**
T0859	Valid Accounts	Persistence; Lateral Movement	**3.10%**
T1110	Brute Force	Credential Access	**1.62%**
T1565	Data Manipulation	Impact	**0.78%**

Here we can clearly see that the number of AiTM-related alerts was abnormally high in the region during this period, well above the global average of 26.5%, and coming from multiple organizations located there. In this case, the alerts were related to cleartext passwords observed in unencrypted traffic at a very large scale. Given this issue, U.K. organizations should take extra care to ensure that, in addition to using strong credentials, they are handled securely, since passwords of any length are weak once leaked.

Germany

ID	Technique name	Tactics	%
T1565	Data Manipulation	Impact	**38.80%**
T0812	Default Credentials	Lateral Movement	**10.30%**
T0859	Valid Accounts	Persistence; Lateral Movement	**10.30%**
T0846	Remote System Discovery	Discovery	**7.93%**
T0841	Network Service Scanning	Discovery	**7.93%**

Compared to the UK, attackers stuck to other techniques in Germany, with Data Manipulation leading the charts. Attackers may prioritize data manipulation because silently altering data can mislead decisions and operations, causing real-world impact while remaining harder to detect than outright disruption. It was followed by Default Credentials and Valid Accounts techniques, both taking slightly more than 10%. Here, attackers may target devices with preconfigured, well-known usernames and passwords to access more systems performing by moving laterally. Once inside, they may attempt to pivot through these compromised assets to escalate privileges, expand access, and ultimately reach high-value systems and sensitive data.

Australia

ID	Technique name	Tactics	%
T0812	Default Credentials	Lateral Movement	**18.00%**
T0859	Valid Accounts	Persistence; Lateral Movement	**18.00%**
T0846	Remote System Discovery	Discovery	**14.20%**
T0841	Network Service Scanning	Discovery	**14.20%**
T1498	Network Denial of Service	Impact	**13.20%**

The situation was even more extreme in Australia, where these two techniques contributed to ~36% of all the raised alerts. For the companies located there, it is definitely worth checking their IoT devices, making sure they are not an easy target for IoT botnets (see the dedicated section below).

Sudden regional increases in alert density should not be interpreted solely as dependent on an attacker's origin or on targeting preferences. In many cases, they also reflect variations in exposure, legacy configurations and protocol usage that create optimal conditions for certain techniques to surface. Keeping a close eye on alerts that manifest more frequently in specific regions is as valuable as attributing them to whoever is behind them.

Now, let's see what we can learn about the malware we detected in our customers' environments during this period.

3.2 Top Malware Observed in the Wild

Malware is a serious threat to organizations regardless of their location or industry. By aggregating and analyzing the telemetry we collected (associated with both **Nozomi's standard threat intelligence** (TI) detections and those coming from our **TI Expansion Pack Powered by Mandiant**), we can highlight the malware families and categories most frequently encountered. Our goal is to provide a concise, evidence-driven view of what is actively impacting organizations now, helping security teams prioritize threat hunting, security hardening and incident response efforts where they will have the greatest effect..

Most Prevalent Malware Categories

As usual, to determine the top detected malware categories during 2H 2025, we carefully curated the underlying telemetry, making sure to exclude confirmed false positives and test data.

Top 5 Malware Categories

Malware category	%
TROJAN	**74.60%**
RAT	**11.30%**
MINER	**4.61%**
WORM	**3.70%**
DOWNLOADER	**1.89%**

The Trojan category accounts for almost three-quarters of all the detected malware. This isn't surprising because Trojan is the most universal detection name used by various security vendors (including Nozomi), signifying any unwanted high-severity code found on a victim system. It is commonly used when the main functionality of malware is costly to figure out, and sometimes if malware has several different categories of features.

The prevalence of Remote Access Trojans (RAT) can be attributed to its versatility. Once successfully deployed (for example, following the initial access phase), attackers have great flexibility to move onto various next stages of the attack, using this implant as an orchestrator. We have observed increased activity of such RAT families as CobaltStrike, AsyncRAT, XWorm and Remcos in our telemetry and have made sure the corresponding detections always remain up to date.

The Miner and Worm categories took third and fourth place, which can be linked to the specifics of IoT environments. It is increasingly common for attackers to pair these two functionalities together to achieve persistence across a large number of IoT assets, then misuse their computational power to mine cryptocurrencies. Even though IoT devices may not be very powerful on their own, together they can generate healthy profits, incentivizing more attackers to leverage them.

Finally, even though the prevalence of other malware categories was lower, they shouldn't be ignored. For example, ransomware accounted for only 0.40% of all detections, but the impact of it, if undetected, is not something any organization would like to experience. This is one of the main reasons why the Nozomi **Arc endpoint sensor** is now able to not just alert on, but also carefully block the detected threats.

Most Active Malware Families

Now, let's look at the malware families detected the most during the 2H 2025 period.

Top 5 Malware Families

Malware category	%
Generic	**54.70%**
DoublePulsar	**20.50%**
ANDROMEDA	**6.02%**
CobaltStrike	**3.62%**
AsyncRAT	**2.45%**

Generic accounted for more than half of all detections, for the same reasons as Trojan above. It's a placeholder name used in various detections where the malware family is unknown or costly to deduce, or when the same indicator of compromise is associated with multiple different malware families. Defaulting to Generic (or Agent) is a standard practice among the cybersecurity community. It's important to understand that these detections are as valid as any other and should be taken just as seriously.

DoublePulsar is still detected across different environments, reminding us how costly it may be for affected organizations to completely remediate a threat used at scale. ANDROMEDA has historically functioned as a distribution/loader component, and it is still active as loader-style threats tend to persist longer because they're reused as staging tools for many different payloads.

CobaltStrike and AsyncRAT, both tracked as RATs, appear at #4 and #5. As we discussed in the previous section, their popularity can be explained, among other reasons, by the versatility that RAT tools provide to operators who prioritize interactive control.

Top Observed Threat Actors

Threat actors are individuals or groups, ranging from financially motivated cybercriminal crews to state-sponsored teams, that conduct malicious activity with distinct goals and tooling. Tracking them over time matters because adversaries don't stand still: they reuse proven infrastructure and malware, evolve techniques in response to defenses, shift targeting as opportunities change, and often rebrand or fragment across campaigns, blurring the attribution and making detections more difficult to maintain.

This section summarizes the threat actors most frequently observed in our telemetry during the reporting period. We identify who is active and what malware they currently favor so defenders can prioritize their detection, hunting and risk-reduction efforts. All the findings come from our telemetry, based on detections of malware linked to threat actors who exclusively use it.

Here are the results for 2H 2025:

Top 5 Threat Actors

Threat actor	Main country of origin	% Associated actor-linked alerts
Scattered Spider	US and UK	**42.90%**
Kimsuky	North Korea	**8.45%**
APT29	Russia	**5.42%**
UNC1549	Iran	**5.38%**
Mustard Tempest	No strong attribution	**3.45%**

According to our data, Scattered Spider was the most active actor in 2H 2025, associated with 42.9% of all actor-linked alerts. This aligns with broader reporting that Scattered Spider remained highly active throughout the year, often leveraging social engineering to gain initial access. Most of the detections were triggered in Sweden, followed by the U.S. and UK taking second and third place respectively. Because Scattered Spider focuses on people and processes, as defenders harden endpoints, they can now frequently detect this threat actor's intrusion at the access stage.

North Korea-linked Kimsuky (or APT43) took second place in terms of the frequency of detected payloads. While they originally targeted mostly South

Korean government organizations and individual experts, they have expanded operations to also target international government-adjacent organizations and remain a **constant threat** to various entities across the globe.

Notorious APT29, aka Cozy Bear, was the third most-detected actor during the reported period. It is a well-resourced, mature espionage actor widely attributed to Russia's Foreign Intelligence Service (SVR), known for targeting government bodies, diplomats, policy and research organizations, and major technology and cloud service providers. APT29 remains highly active, evidenced by the **watering hole campaign** disrupted by Amazon's threat intelligence team in August. We predict that this actor will be among the most active at least while the current geopolitical situation persists.

The Iranian threat actor UNC1549 (overlapping with CURIUM/Tortoise Shell/ Crimson Sandstorm) took fourth place. Iran has been the focus of everyone's attention since 2025, mostly because of the increased tensions with the U.S. and Israel. The actor is believed to represent the interests of its government, infamous for targeting IT service providers in the Middle East, as well as other sectors across the globe. Our findings confirm that the actor seems to remain an active threat that should be monitored. According to our monthly telemetry, December was its most active month with the highest number of alerts.

Finally, Mustard Tempest took fifth place. It is linked to the infamous SocGholish malware observed in use against multiple sectors globally for initial access, primarily through drive-by-downloads. As long as it remains popular, we expect to continue to see this actor's activity in our telemetry.

To conclude, 2H 2025 telemetry shows a threat landscape dominated by both opportunistic initial-access groups and highly persistent, state-linked actors, with activity spanning social engineering, espionage and drive-by infection chains. Continued investment in early-stage detection, user/process hardening and rigorous monitoring for associated TTPs will remain critical to disrupting these campaigns before attackers can establish a deeper foothold.

Latest Ransomware Trends

Last October we published a blog post covering **ransomware trends between March and September.** In this section, we will complement that data with the telemetry insights for the remainder of 2025, when activity spiked even higher.

Number of Ransomware-Related Alerts

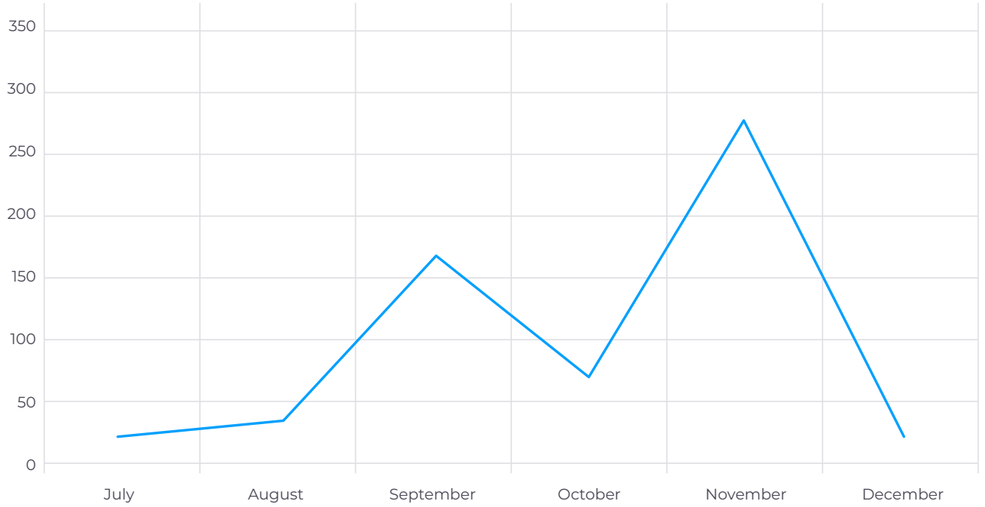

The figure above clearly shows us that the number of ransomware-related alerts increased significantly in September and especially November, during which the Nozomi platform raised close to 300 new ones. To explain this increase, let's see which malware families they were related to.

Top Ransomware Families

Compared to the six months covered in our August blog post, Elpaco ransomware overtook BlackSuit in terms of most alerts during this reported period. The decline of BlackSuit was expected due to a **takedown operation by the U.S. government** to hamper its operations in August. We anticipate a further decrease in its detections in the following months, as the remaining actors behind it will likely switch to other projects.

Elpaco, an evolution of the Mimic family, has become increasingly common, contributing around 12% of all the ransomware-related alerts during 2H 2025. It uses a legitimate library called **Everything** to quickly iterate through the victims' files in order to encrypt them and demand a ransom. Interestingly, it acts as a local search engine and comes with a GUI allowing the operator to customize its behavior.

Surprisingly, in third place we got multiple alerts for the Ryuk ransomware, which has not been active recently. However, its code has been re-used in more modern ransomware families like Conti, and the same detection rule can detect both of them, thus obscuring the exact attribution of the malware family spotted.

We will continue to monitor ransomware trends to ensure our detections remain accurate and can detect both existing and future strains, thanks in no small part to our mature, native AI engine.

Top 3 Targeted by Ransomware Countries

Here are the top three countries where the highest number of ransomware alerts were raised:

Country	%
US	**40.32**
UK	**15.48**
CA	**13.55**

U.S.-based companies alone accounted for more than 40% of all ransomware attacks in 2H 2025. We believe many ransomware groups view the U.S. as a prime target for a mix of geopolitical and economic reasons. For instance, some Russian-based operators oppose U.S. support for Ukraine in the ongoing war, adding a hacktivist dimension to their activity. The high concentration of large enterprises and critical infrastructure, often with greater ability to pay, can also make U.S. organizations especially attractive.

The UK was second with 15.48% of reported incidents, followed by Canada with 13.55%. It's unlikely to be a coincidence that the most-targeted countries are predominantly English-speaking. Since English is widely taught worldwide,

many ransomware operators may rely on it as their primary second language, making it easier to communicate with victims directly and reducing the need for automated translation tools.

Overall, the 2H 2025 data underscores both a rising alert volume and a shifting mix of dominant ransomware families, with activity concentrated in a small set of highly exposed, high-value geographies. As families adapt, whether by migrating to new tooling after disruption efforts or by refining established strains like ElPaco, defenders must maintain strong visibility and continuously evolving their detection prowess to staying ahead of both current campaigns and the next wave of variants.

4. Vulnerability Landscape and Wireless Threats

As industrial environments become more interconnected, the vulnerability landscape continues to expand, driven by legacy systems, unpatched devices and increasing exposure to external networks. At the same time, the growing use of wireless technologies introduces new attack surfaces that are often difficult to monitor and control. This section examines the most significant vulnerabilities discovered in 2025 affecting OT and IoT environments at the end of the reported period, as well as emerging wireless threats that can bypass traditional security controls and put critical operations at risk. Once again, all the insights here come from our customers' anonymized telemetry.

4.1 Telemetry-based Vulnerability Statistics

Telemetry provides a real-world view of vulnerability exposure in OT and IoT environments, showing not just which CVEs exist but where they appear, how widespread they are and which assets are most at risk. By analyzing observed data, organizations can move from static lists to **prioritizing remediation actions** based on multiple risk factors. This section summarizes key trends that took place in the second half of last year, highlighting the distribution of vulnerabilities discovered in 2025 and their unique characteristics.

Top 2025 CVEs Found in Modern Environments

First, let's look at the most common recently discovered vulnerabilities found in customer telemetry across industries and regions. To measure their prevalence, we looked at which vulnerabilities appeared in the most organizations, broken down by domain.

IT Domain

CVE ID	CVSS	CWE	Affected product(s)
CVE-2025-13042	8.8	CWE-787: Out-of-bounds Write	Google Chrome
CVE-2025-13223	8.8	CWE-843: Access of Resource Using Incompatible Type ('Type Confusion')	Google Chrome
CVE-2025-13224	8.8	CWE-843: Access of Resource Using Incompatible Type ('Type Confusion')	Google Chrome
CVE-2025-13226	8.8	CWE-843: Access of Resource Using Incompatible Type ('Type Confusion')	Google Chrome
CVE-2025-13227	8.8	CWE-843: Access of Resource Using Incompatible Type ('Type Confusion')	Google Chrome

The tremendous popularity of the Google Chrome browser is reflected in it accounting for the top five IT vulnerabilities across all observed organizations. Moreover, all five top entries have a High CVSS score of 8.8. Browsers are versatile tools that may allow users to access sensitive data. They may also be involved in the attack process (for example, when used by a victim to click on a malicious link). We strongly encourage all companies to embrace automatic updates and use solid vulnerability management solutions to remediate them as soon as possible.

OT Domain

CVE ID	CVSS	CWE	Affected product(s)
CVE-2025-24811	8.7	CWE-404: <Improper Resource Shutdown or Release>	Siemens SIMATIC S7-1200 CPU family
CVE-2025-24812	7.1	CWE-1286: <Improper Validation of Syntactic Correctness of Input>	Siemens SIMATIC S7-1200 CPU family
CVE-2025-7353	9.3	CWE-1188: <Initialization of a Resource with an Insecure Default>	Rockwell Automation ControlLogix Ethernet Modules
CVE-2025-6625	8.7	CWE-20: <Improper Input Validation>	Schneider Electric Modicon M340 and Communication Modules
CVE-2025-0631	8.7	CWE-319: <Cleartext Transmission of Sensitive Information>	Rockwell Automation PowerFlex 755

In the OT realm, the popularity of the vulnerable assets is directly linked to the distribution of the associated vulnerabilities. Of course, the relationship is not linear, as different organizations have different policies related to vulnerability management. In some of them, the assets are updated promptly and regularly, while in others they may stay outdated forever. In addition, some assets are easier to update for various reasons, like the straightforwardness of the process or the low impact of this asset going down briefly to perform the update. Finally, some vulnerabilities may be discovered after the product has reached its EOL, and some of them will never be addressed by the vendor. Here, we see that CVE-2025-24811 with a CVSS score of 8.7 (High) ranked #1. Siemens SIMATIC S7-1200 PLCs are both versatile and popular; they are used across industries and run various processes with various levels of criticality. Just below it is CVE-2025-24812 with a lower CVSS score of 7.1 (High).

CVE-2025-7353 took third place. Given its CVSS score of 9.3 (Critical), its impact shouldn't be underestimated. This vulnerability affects multiple different Rockwell Automation ControlLogix assets.

Finally, CVE-2025-6625 and CVE-2025-0631 took fourth and fifth place. Both have High CVSS scores and affect popular Schneider Electric Modicon M340 and Rockwell Automation PowerFlex 755 solutions, respectively. As above, when circumstances and processes allow, patching these vulnerabilities is essential for maintaining the long-term health of the affected environment.

IoT Domain

CVE ID	CVSS	CWE	Affected product(s)
CVE-2025-20175	7.7	CWE-805: <Buffer Access with Incorrect Length Value>	Cisco IOS Software and Cisco IOS XE Software
CVE-2025-20363	9.0	CWE-122: <Heap-based Buffer Overflow>	Cisco Secure Firewall Adaptive Security Appliance (ASA) Software, Cisco Secure Firewall Threat Defense (FTD) Software, Cisco IOS Software, Cisco IOS XE Software, and Cisco IOS XR Software
CVE-2025-20352	7.7	CWE-121: <Stack-based Buffer Overflow>	Cisco IOS Software and Cisco IOS XE Software
CVE-2025-20169	7.7	CWE-805: <Buffer Access with Incorrect Length Value>	Cisco IOS Software and Cisco IOS XE Software
CVE-2025-20173	7.7	CWE-248: <Uncaught Exception>	Cisco IOS Software and Cisco IOS XE Software

Like Google Chrome in the IT realm, the popularity of Cisco products is reflected in them accounting for all top five IoT vulnerabilities in 2H 2025. Here, four out of the five most widespread vulnerabilities affected Cisco IOS and IOS XE assets and had a CVSS score of 7.7 (High). Most notably, the infamous CVE-2025-20363 affecting other Cisco products came in at #2. With a CVSS score of 9.0 (Critical), this vulnerability **continues to be targeted**, according to Cisco.

Top 2025 Vulnerability Categories Affecting OT/IoT Environments

Here is the distribution of categories for vulnerabilities discovered in 2025 and found in customer environments at the end of last year. In this and the following sections, we evaluated their prevalence by the total number of instances when the associated vulnerability was identified in customers' environments.

CWE	%
CWE-416: Use After Free	13.80%
CWE-125: Out-of-bounds Read	8.05%
CWE-843: Access of Resource Using Incompatible Type ('Type Confusion')	7.50%
CWE-122: Heap-based Buffer Overflow	5.91%
CWE-451: User Interface (UI) Misrepresentation of Critical Information	5.29%
CWE-701: Weaknesses Introduced During Design	4.05%
CWE-20: Improper Input Validation	3.68%
CWE-79: Improper Neutralization of Input During Web Page Generation ('Cross-site Scripting')	3.54%
CWE-200: Exposure of Sensitive Information to an Unauthorized Actor	3.22%
CWE-362: Concurrent Execution using Shared Resource with Improper Synchronization ('Race Condition')	3.06%

Again, as in June 2025, the leading category was Use After Free. This vulnerability occurs when a program continues to use a pointer or reference to memory after it has been freed, which can lead to crashes, data corruption or attacker-controlled code execution. Thanks to AI assistance, it has become much easier for researchers to spot such mistakes in the code, which may explain why the number of vulnerabilities of this type remains consistently high in the last few years.

The Out-of-bounds Read category also took second place again. This occurs when a program reads data past the end (or before the start) of a buffer, potentially leaking sensitive information or causing unexpected behavior. The infamous Heap-based Buffer Overflow dropped from third to fourth place, giving way to CWE-843, commonly known as Type Confusion. It occurs when a program treats a piece of data as the wrong type. The rise of this category may be tied to increasing system complexity, in which some parts make assumptions about others, and improvements in modern fuzzers (including ML/AI assistance).

Owing to this shift, the User Interface (UI) Misrepresentation of Critical Information category also slipped from #4 to #5. We hope this information will help vendors create better, safer products that are immune to these types of vulnerabilities.

CVSS Score Variability Across Different Businesses

CVSS scores offer a standardized way to estimate the potential severity of vulnerabilities, helping teams compare issues and communicate risk consistently. In OT and IoT environments, however, a high CVSS score doesn't always translate to the highest operational risk. Asset criticality, network exposure, exploitability and impact on safety or downtime often matter just as much. Still, the score provides valuable context when prioritizing remediation actions and is a key variable in multi-factor risk scoring.

In December, the distribution of vulnerabilities discovered in 2025 and detected across our monitored OT/IoT environments by CVSS scores looked like this:

Distribution of 2025 Vulnerabilities Found in Real Environments by CVSS Score

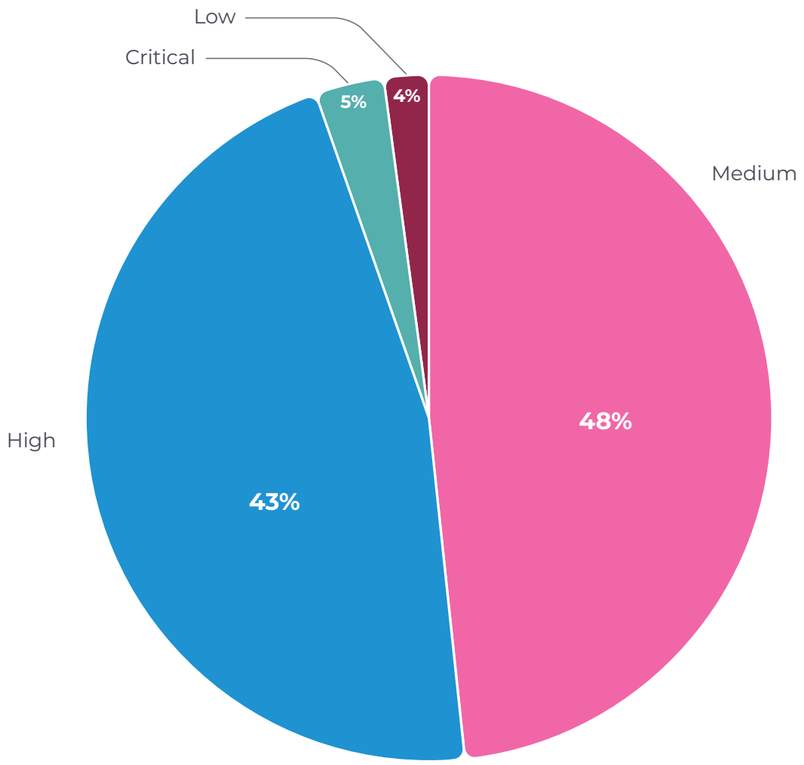

Firstly, the low percentage of vulnerabilities with a Low CVSS score is expected, because many researchers don't even report them to vendors when they're spotted. More importantly, the percentage of vulnerabilities with a High CVSS score is substantial (43%). These vulnerabilities generally have a meaningful impact and feasible exploitation conditions, presenting real risk that requires remediation.

The stakes are even higher for the 5% of vulnerabilities with a Critical CVSS score. While low in number, each of them poses significantly more risk where they are found due to their potential impact if exploited successfully. Exploitation is also significantly more probable because these vulnerabilities often allow unauthenticated remote code execution (RCE). Understanding a vulnerability's CVSS score and its potential impact in exploited in the context of your environment is a powerful approach to efficient vulnerability management.

Top Exploitability Metrics

In addition to CVSS scores, many IT and cybersecurity professionals responsible for patching vulnerabilities also pay attention to dedicated exploitability-related metrics. EPSS (Exploit Prediction Scoring System) is a data-driven score that estimates the probability a specific vulnerability (CVE) will be exploited in the wild within a defined time window. In addition to severity, it uses signals from public and threat-intel data. In vulnerability management, EVSS scores are used to prioritize patching and mitigations by focusing efforts on vulnerabilities that are most likely to be exploited, often helping teams reduce CVE backlogs more effectively than CVSS scores alone.

Here is the distribution of vulnerabilities published in 2025 in monitored environments based on their exploitability at the end of 2H 2025.

Distribution of 2025 Vulnerabilities Found in Real Environments by EPSS Score

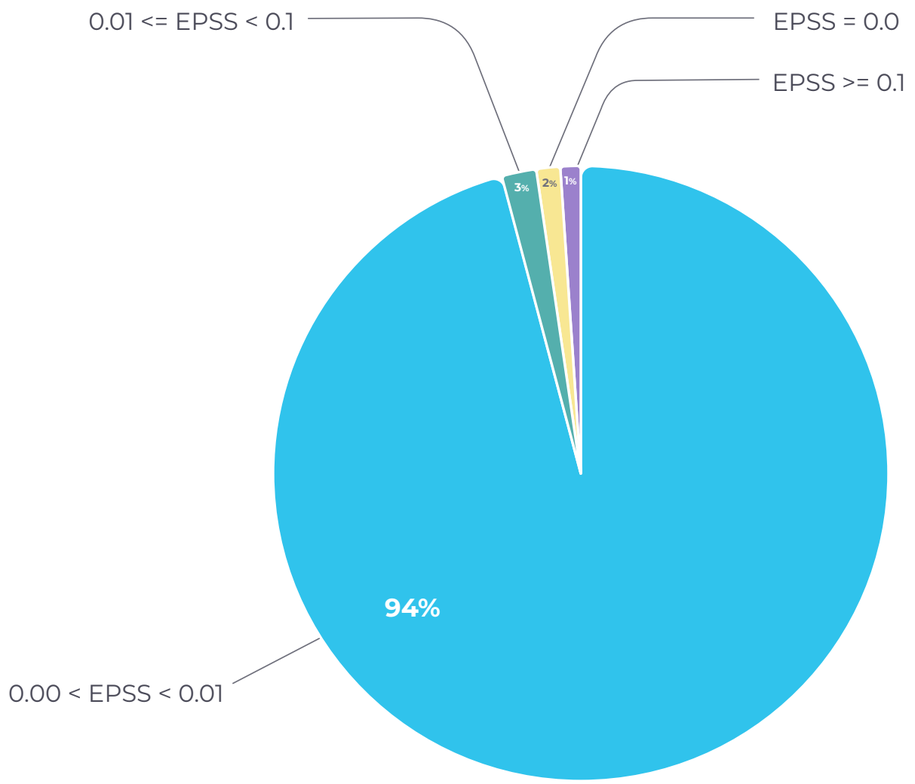

For 2H 2025, 94% of the associated vulnerabilities have an EPSS score of less than 1%, on par with the 92% we observed in June 2025. The 2% decrease suggests that the attack surface has likely decreased. The remaining 6% with an EPSS score above 1% still pose a risk to affected organization and may require attention, based on other risk factors and compensating controls. The same applies to Known Exploited Vulnerabilities (KEVs).

Percentage of 2025 Vulnerabilities Found in Real Environments Marked as KEV

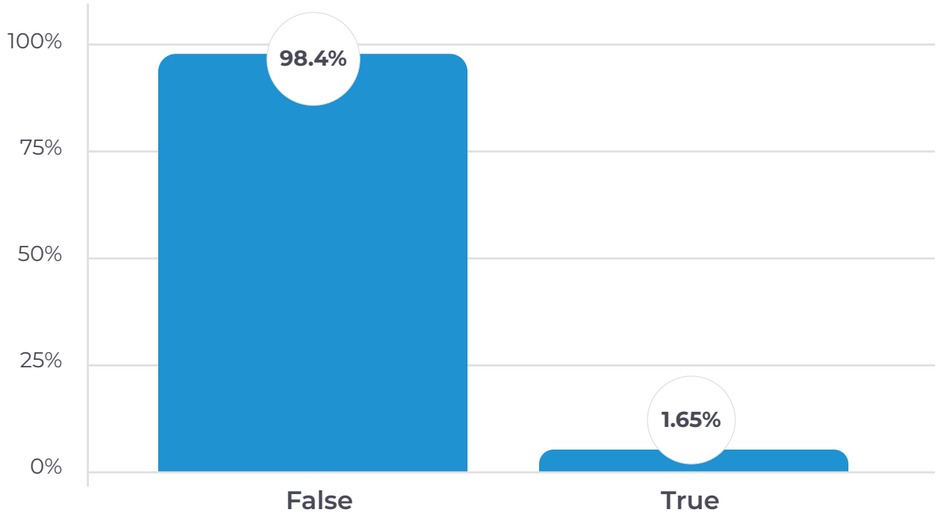

At first glance, the low prevalence of KEVs among organizations operating OT/IoT environments may seem surprising. This may be related to how designations are often driven by activity observed in U.S. government contexts. Given the differences in asset mix and deployment patterns in commercial vs. government deployments, there would naturally be fewer KEVs in customer environments. Despite their low number, by definition any KEV found in any environment should be considered for remediation.

4.2 Wireless Exposure in Industrial Environments

Wireless communications are increasingly present in industrial and critical infrastructure environments, often without formal design or attention to security. The radio spectrum is now a persistent and relevant attack surface, not merely an extension of wired risk.

Wireless systems span personal devices, industrial equipment and temporary tools, creating a fluid environment that challenges static inventories and perimeter-based security models.

Telemetry collected over time shows that the scale of wireless activity continues to grow. Across reporting periods, environments described as “fully wired” continue to show persistent wireless activity. Although specific devices change, the pattern remains constant: Bluetooth-equipped laptops, mobile smart tools, and embedded radio modules commonly operate near critical systems regardless of intended architecture. This activity is not inherently malicious, but its persistence exposes a structural gap. Devices often communicate without centralized visibility, occasionally bypassing established monitoring and governance and connecting air-gapped networks to the internet. Security teams are forced into a reactive cycle of adding tools to chase an ever-present condition, treating wireless exposure as an anomaly rather than a stable reality.

The result is a durable layer of implicit connectivity that remains largely unaddressed by security models built on static assumptions. Without continuous wireless visibility, organizations operate with an incomplete understanding of their environment, leaving unprotected spaces where attackers can act undetected.

Ambient Wireless Activity and Boundary Erosion

Long-term telemetry shows a persistent presence of Bluetooth devices across IT, OT and IoT ecosystems. Although vendors change, the continued coexistence of personal and industrial wireless platforms reflects a lasting convergence between human-operated and machine-operated systems.

Top 10 Bluetooth device vendors at the end of 2H 2025

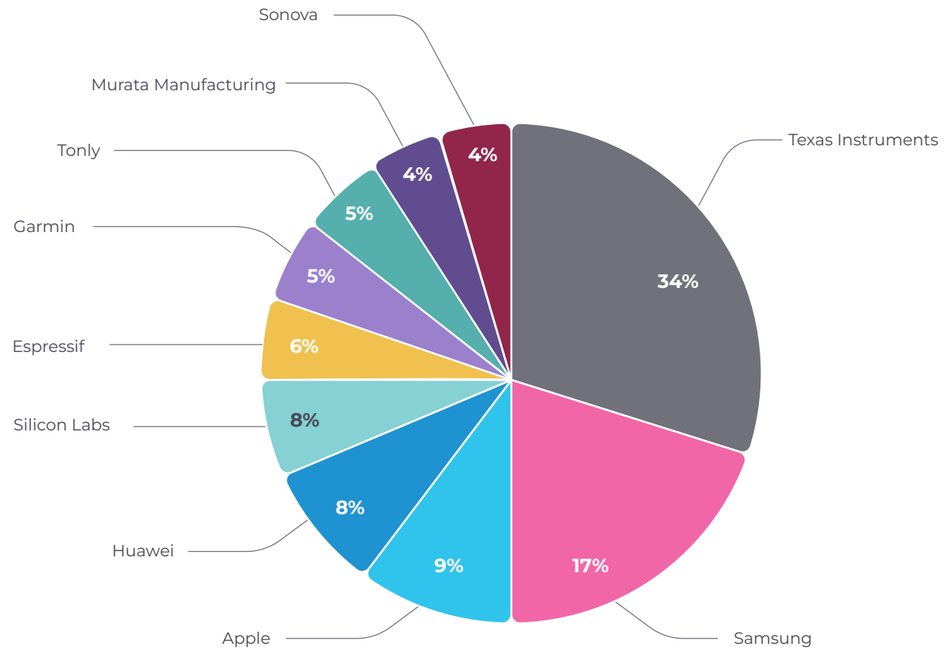

This convergence weakens traditional trust boundaries. Convenience-oriented and low-power industrial protocols such as Zigbee and WirelessHART found in IEEE 802.15.4 networks operate within industrial environments, often without centralized oversight or ownership. In this context, vulnerabilities such as CVE-2024-3043, where unauthenticated IEEE 802.15.4 coordinator realignment frames can disrupt network operation, and CVE-2024-6350, which enables denial-of-service via malformed 802.15.4 traffic in widely deployed Zigbee stacks, exemplify the risks introduced by these low-power industrial networks if left unmonitored.

Distribution of 802.15.4 based protocols at the end of 2H 2025

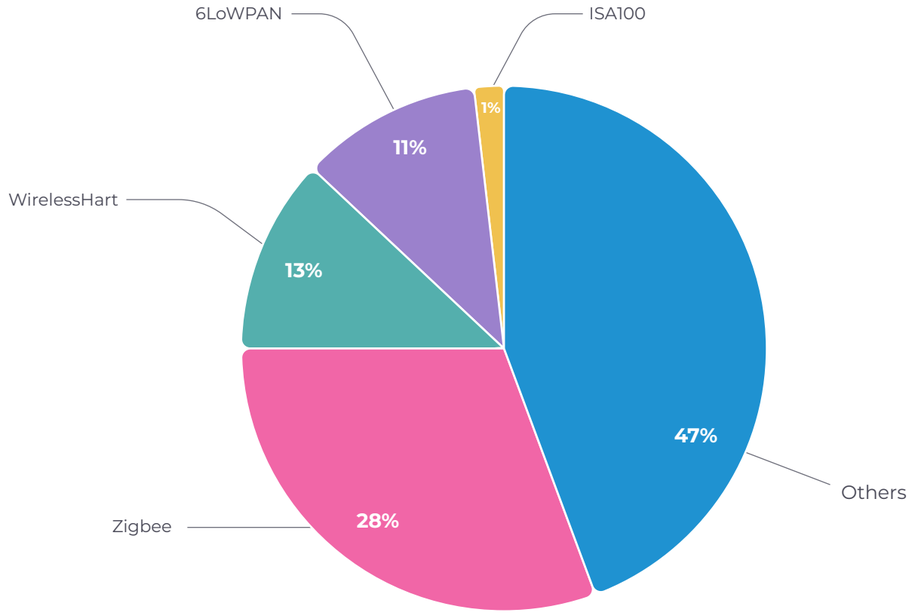

As a result, they create a persistent and largely unmanaged attack surface that exists alongside governed infrastructure.

Wireless Security Posture: Improvements Without Assurance

At the protocol level, WPA3 adoption increased by roughly eight percentage points compared to the previous period, largely due to OS defaults and hardware refresh cycles, indicating a gradual shift toward stronger cryptographic protections.

Wi-Fi Encryption Protocols Distribution at the end of 2H 2025

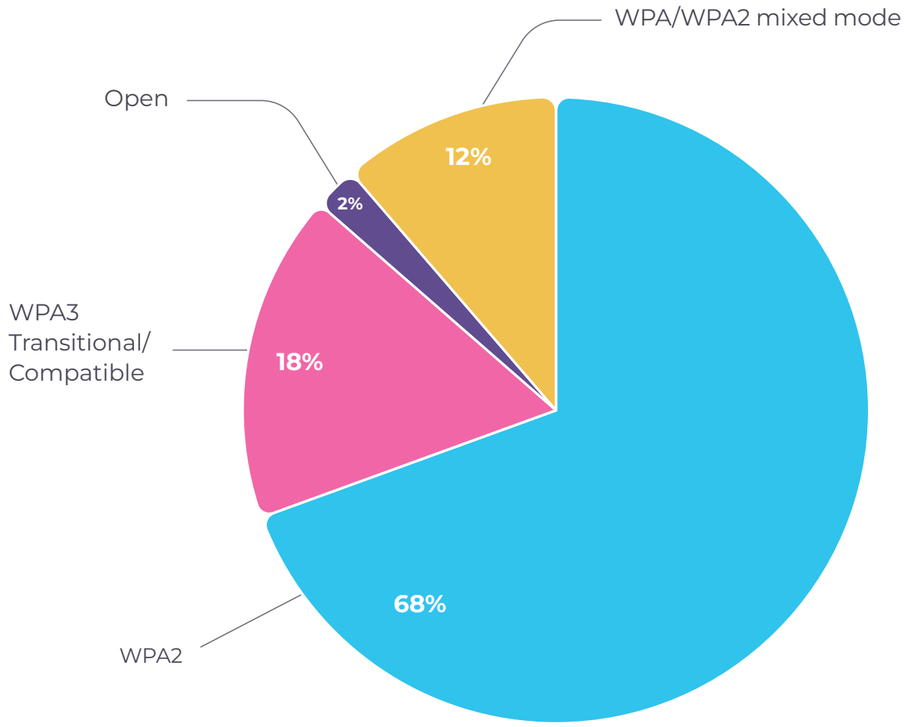

About 14% of observed networks still use open or legacy security modes such as WPA or WPA/WPA2 mixed configurations, reflecting accumulated technical debt. WPA2 remains the dominant standard, while roughly 18% of networks have adopted WPA3 (including transitional mode).

In addition, roughly 30% of observed networks use WPA/WPA2 mixed mode, WPA3 transitional mode or similar configurations that allow multiple security standards within a single SSID. While these modes support gradual migration and backward compatibility, they inherently limit the adoption of stronger security mechanisms, because legacy clients prevent full use of modern encryption and key-management algorithms.

As a result, multicast and broadcast traffic in such networks is encrypted with the weakest common cipher to ensure interoperability. For instance, in WPA/WPA2 mixed mode, legacy WPA devices require the WPA cipher for all traffic, and although WPA2 clients may use stronger ciphers for unicast traffic, multicast and broadcast frames still rely on the weaker WPA cipher for compatibility.

When looking beyond protocol labels and focusing on effective protections, the picture becomes more nuanced. Despite improvements in encryption standards, effective control-plane protection has not advanced at the same pace.

802.11w MFP Configuration Status Distribution at the end of 2H 2025

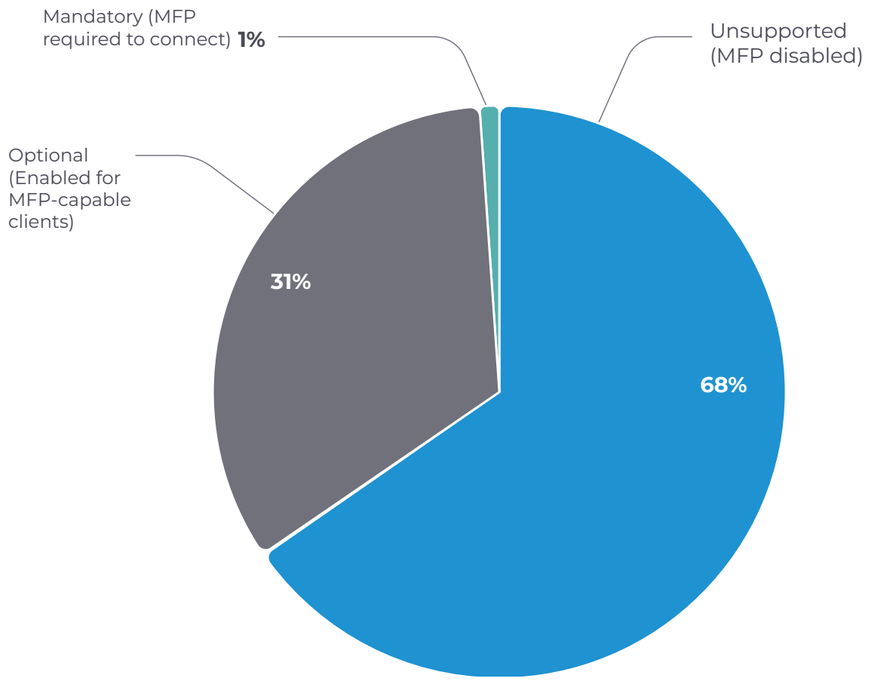

Approximately 68% of observed wireless networks still operate without Management Frame Protection (MFP). While this is a notable reduction compared to the previous period (83%), it nevertheless leaves the majority of control-plane traffic unprotected even when modern encryption is enabled. About 31% of networks now advertise optional MFP support, a 14% increase compared to the previous reporting period (17%). However, only a negligible fraction of networks mandate MFP, limiting the real-world impact of this improvement.

Limited adoption of MFP leaves networks exposed to downgrade and control-plane attacks that exploit unauthenticated management frames to disrupt connectivity or impersonate access points, without needing to break encryption.

Authentication Practices as a Structural Weakness

Telemetry shows that nearly all of observed wireless networks rely exclusively on Pre-Shared Key (PSK)–based authentication, making it by far the de facto standard in operational environments.

Wi-Fi Authentication Method Distribution at the end of 2H 2025

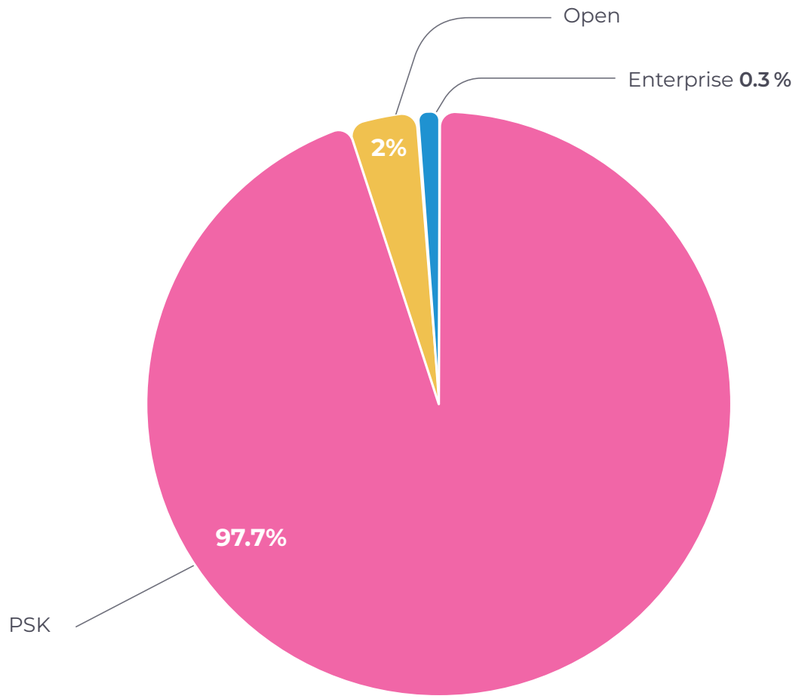

As a result, shared credentials remain valid entry points, downgrade paths are preserved and attackers can target the weakest supported method rather than the strongest one. Enterprise-grade authentication such as 802.1X remains rare, found in only 0.3% of deployments. When present, it is often limited to specific segments or it coexists with PSK-based networks, reducing its effectiveness in lowering risk.

In operational environments, shared credentials remove accountability and enable long-term reuse, making it hard to distinguish legitimate access from misuse once exposed. This favors attackers seeking persistence, because a single compromise can enable silent re-entry and lateral movement, turning wireless connectivity into a durable access mechanism.

Mobility and Asset Location: A New Dimension of Visibility

Unlike wired devices, which are (typically) fixed to known infrastructure, wireless assets can appear, disappear and reappear in different physical locations without changes to configuration, credentials or identifiers. This mobility fundamentally changes how asset tracking and risk assessment must be performed. A device observed near a sensitive system one day may operate elsewhere the next, crossing physical and logical boundaries that are difficult to traverse in wired environments. As a result, traditional asset inventories based on static placement quickly become inaccurate when applied to wireless systems.

From a security standpoint, this creates operational and investigative challenges. Knowing where a device was at a specific moment is critical for incident response, forensic analysis and policy enforcement, yet without spatial context, visibility remains partial. Advances in wireless telemetry address this gap by correlating signal characteristics across sensors to infer zone-level location. While not equivalent to precise geolocation, this approach enables consistent zone-level and area-level positioning within industrial environments.

Triangulation UI on Vantage powered by Guardian Air

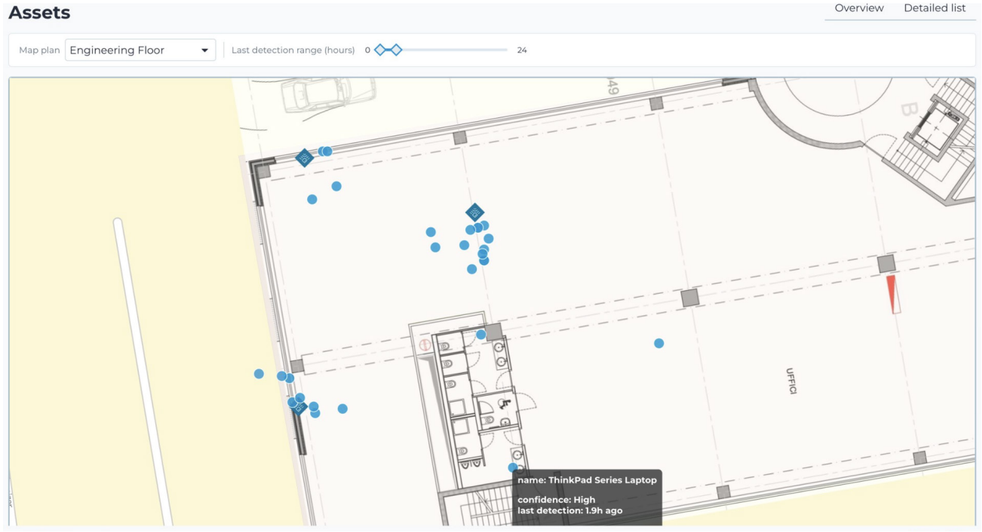

Continuous location-aware wireless monitoring turns mobility into context, revealing movement patterns, detecting proximity, and tracking assets across operational spaces. This spatial insight deepens visibility by showing not only how devices communicate but where they operate.

4.3 Nozomi-Discovered Zero-Day Vulnerabilities

Vulnerability research findings provide a practical view of security conditions in the OT and IoT ecosystem, revealing not only which issues exist, but which weakness patterns recur across technologies, device classes and vendors. This section summarizes vulnerabilities that Nozomi Networks published over the past six months through its 0-day research program, based on in-depth analysis of industrial, IoT and edge devices deployed in operational environments.

Beyond individual findings, the data highlights trends in severity, CWE categories, discovery methods, disclosure timelines and vendor engagement. Together, these observations offer insight into how the OT/IoT attack surface is evolving and where systemic engineering weaknesses persist, supporting more targeted research and more effective coordinated disclosure.

Vulnerability Severity Distribution

Severity is expressed using CVSS and categorized as Critical, High or Medium. Low-severity issues are excluded because we do not track or report them to vendors. We consider average CVSS and severity distribution to assess risk and hardening across devices and product families, identify recurring weaknesses and guide research priorities.

The average CVSS during this period was 7.0, with Critical + High totaling 51%, indicating a substantial share of findings that require urgent attention. This supports the need to prioritize rapid mitigations for exploit-prone device families and allocate triage resources to ensure timely validation and disclosure of high-severity issues.

Distribution of the Vulnerabilities Disclosed by Nozomi Networks in 2H 2025 by CVSS Score

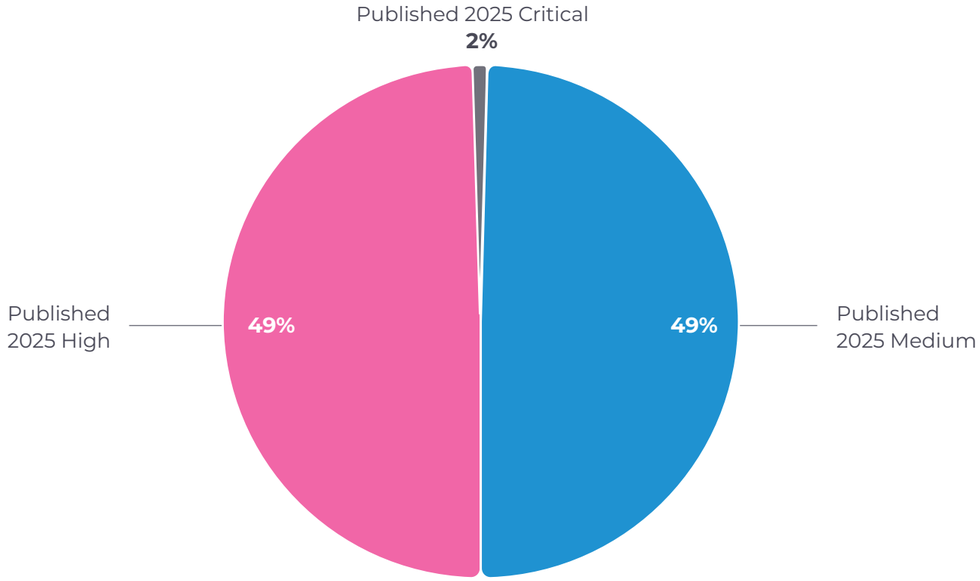

CNA Activity and Vendor Maturity

Nozomi Networks Labs is a CVE Numbering Authority (CNA); that is an authorized organization that identifies, assigns CVE IDs to, and publishes information about security vulnerabilities within our scope. CNA activity tracks who publishes the vulnerability advisory such as Nozomi Networks or the affected vendor. In OT and IoT, many vendors are still not registered as CNAs; in those cases, Nozomi

Networks ensures CVE assignment and disclosure aligned with industry best practices. Tracking this metric over time provides a view of vendor maturity: sustained increases in vendor CNA adoption typically reflect greater awareness, the establishment of PSIRTs and stronger internal vulnerability-handling processes.

During 2H 2025, we published as CNA for 15.3% of advisories, while vendors published 84.6% as CNA, a 41.5% increase in vendor-as-CNA share over the previous period. This upward trend suggests improved vendor security processes and greater ownership of vulnerability management, at least for the vendors involved in this reporting period's disclosures.

Domain Balance: OT vs. IoT

Domain balance is the split of findings between OT and IoT within an increasingly converged OT, IoT and IT landscape. OT comprises long-lived industrial devices with strict availability and limited patching windows, while IoT is inherently more connected and exposed, expanding the attack surface. Based on customer telemetry and direct engagement, our researchers have observed a growing IoT presence in operational environments, driving Nozomi Networks to expand research beyond traditional OT to include IoT, IT and edge devices. We track this metric as an indicator of sustained research focus across converging domains.

The domain balance split for this period was OT 72% / IoT 28%, reflecting an increasing alignment with customer exposure. OT remains dominant in customer environments and requires safety-aware disclosure and mitigation suitable for long-lived assets. At the same time, the growing IoT share confirms that poorly secured IoT deployments can create new entry points for attackers, especially at scale.

Weakness Trends

We use CWE to classify the engineering weaknesses that enable exploitation and understand why vulnerabilities exist at a design or implementation level. Tracking recurring CWE patterns helps identify systemic issues and guide proactive research across devices sharing similar technologies, architectures, purposes or vendors. The most frequent weakness categories were the following:

- **CWE-23: Relative Path Traversal**
- **CWE-732: Incorrect Permission Assignment for Critical Resource**
- **CWE-59: Improper Link Resolution Before File Access**

Together, these three categories accounted for 25.6% of findings. Notably, this reflects a shift from previous quarters, where authentication and authorization issues were more common, toward file handling and filesystem access control weaknesses. We believe this change may be influenced by our increased focus on IoT devices, which often exhibit different and more complex security issues than OT counterparts. Overall, the results support targeted improvements in secure design, input validation and access control mechanisms rather than one-off fixes.

Discovery Methods

We classify vulnerabilities by discovery technique: automated fuzzing or expert-driven manual analysis. As automation and AI-assisted tools become more common, we track this dimension to understand their contribution to discovery. While automation can accelerate parts of the workflow, effective fuzzing still depends on researcher expertise, and many vulnerability classes—especially logic flaws, authentication flows and architectural weaknesses—still require human judgment and context.

Fuzzing accounted for just 10% of vulnerabilities uncovered by Nozomi Networks, while 90% resulted from expert-driven, AI-enhanced research. This distribution shows that automation can efficiently surface certain issue classes, but most impactful vulnerabilities still require hands-on investigation and contextual reasoning. Overall, automation complements manual research by enabling broader exploration while relying on expert analysis to uncover complex vulnerabilities.

Time-to-Disclosure and Coordination

Time-to-disclosure is the time between initial vendor reporting and public advisory release. We track mean and median values to capture typical timelines and identify delays while ensuring responsible disclosure balances transparency with user protection. Nozomi Networks follows a rigorous disclosure process with no exceptions on CVE assignment: all validated vulnerabilities are published.

Mean time-to-disclosure was 181 days (median 171), reflecting the coordination required among researchers, vendors and CNA processes. This is an improvement versus the first half of 2025 (206/189) and may indicate increasing vendor maturity and responsiveness for the vendors involved in this reporting period.

Together, these metrics illustrate the scale, impact and complexity of vulnerabilities addressed through Nozomi Networks' 0-day research, highlighting recurring weakness patterns, discovery dynamics and disclosure characteristics across OT and IoT environments. They underscore the continued importance of expert-driven analysis and coordinated disclosure to effectively manage risk in increasingly complex operational ecosystems.

Latest IoT Botnet Activity

IoT botnets continue to evolve, adding more propagation methods to their arsenals and targeting more and more devices. A good example is the recent RondoDox botnet constantly updating its collection of exploits. In this section, we will continue tracking this endless process, sharing with you the most interesting and actionable insights. This time, all the data comes from our chain of globally distributed honeypots, completely unrelated to our customers' environments. These honeypots enable us to monitor trends and provide us with a daily batch of fresh Indicators of Compromise (IoCs) that we use to improve our TI offering.

5.1 Where Do Attacks Originate?

When an attack against a vulnerable device is initiated, it comes from a particular IP address associated with a country. This allows us to precisely trace the growth of self-propagating botnets by identifying the location of each vulnerable device when it was compromised by IoT malware and became part of the botnet. While request proxying is possible in theory, in practice, botnets don't deploy it at scale given that its high cost yields few advantages, making this attribution method more reliable.

Here is what the data looks like in 2H 2025:

Distribution of Attacks Based on IP Address Origin

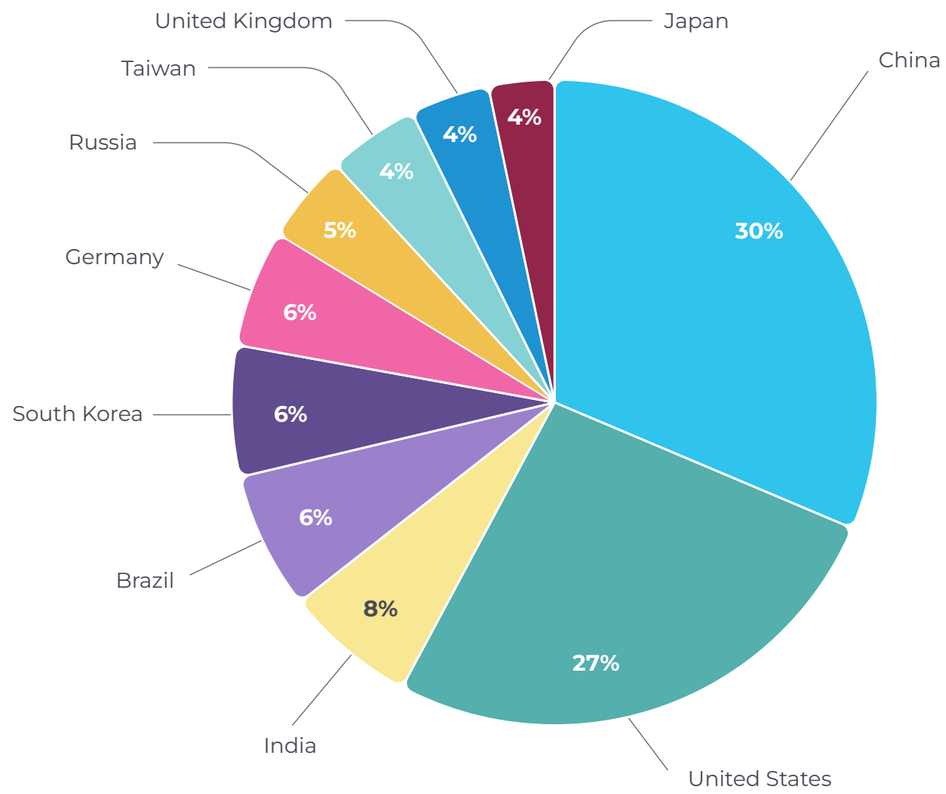

China again took first place in terms of the number of devices that initiated attacks against our honeypots, overtaking the U.S, which briefly took first place in the previous 6-month period. The level of automatization and the number of smart devices keep increasing in this country, which might have contributed to the number of devices compromised by botnets there. India and South Korea remained at third and fifth place respectively, signifying a relatively stable concerning situation in these regions. This persistence suggests entrenched infrastructure and remediation gaps, underscoring the need for sustained, coordinated mitigation efforts across ISPs, enterprises and national CERTs in both regions.

We saw other notable changes in Brazil and Japan. Brazil jumped from ninth to fourth place, likely because of an increase in exploitable devices, decrease in cybersecurity posture or both. After dropping out of the top 10 chart during 1H 2025, Japan reappeared with 4% of IP address origins, tying with Taiwan and the UK. These trends suggest a potential escalation in botnet activity. Organizations in the affected regions should confirm they have full visibility into their IoT assets and that their assets are not being misused to launch attacks against others.

5.2 Attack Intensity Trends

Apart from providing us with actionable IoCs, our constant botnet monitoring enables us to promptly spot activity spikes that may reveal new or significant changes in an existing botnet. When we plot the number of unique IP addresses from which attacks against our honeypots originated in 2H 2025, we see the following picture:

Daily Volume and Activity from Unique Attacker IP Addresses

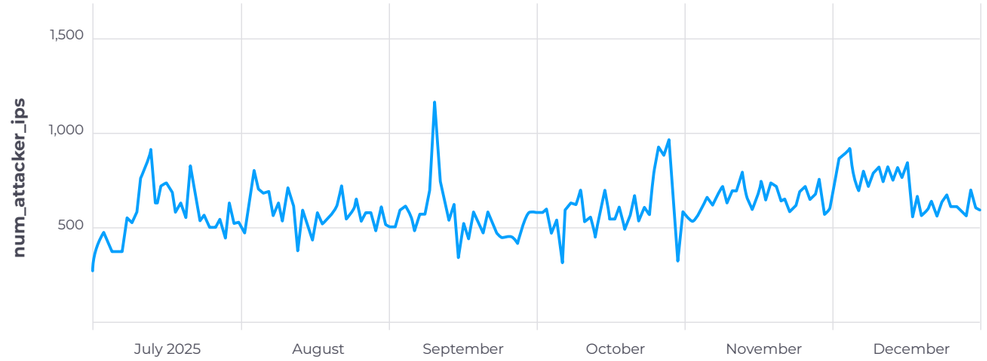

The biggest spike happened on September 2, 2025, when, in one day, 1,169 attack sources were recorded. It was related to an upgrade of the Mirai clone, and we immediately made sure that our customers were provided with solid protections against it. This example highlights how important it is to invest in visibility and embrace OT/IoT-focused threat intelligence as this investment pays off later with big dividends.

5.3 Top Credentials Used by Botnets

As we saw in the first section, the MITRE ATT&CK techniques T1110 (Brute Force), T0812 (Default Credentials) and T0859 (Valid Accounts) are still among the most common ones used by attackers to target various environments across the globe. IoT botnets generally implement them via a list of hardcoded credentials that

they use to try to log in to victim devices and move laterally. Here are the most common pairs of usernames and passwords botnets used to try to get access to our honeypots:

- **admin:admin**
- **root:root**
- **root:123456**
- **root:admin**
- **root:password**
- **user:user**
- **root:12345**
- **admin:password**
- **root:1234**
- **admin:1234**

Not surprisingly, credentials giving high-privilege access to devices continue to be the most favored by malware. With such credentials, the malware can easily access more system directories and utilize commands that require elevated privileges. As researchers, educators, governance bodies and regulators consistently point out, these are the actual default credentials that manufacturers (and later their customers) commonly leave on their devices to facilitate maintenance. We remind everyone about the importance of changing them and making sure that your security solutions will raise an alert if they detect them in network traffic.

5.4 Top Executed Post-Compromise Commands

Once IoT malware gets access to the compromised device, it will generally continue the attack by launching various commands, aiming to achieve persistence, evade defenses or continue moving laterally. Here are the most common commands that we recorded being used by them in 2H 2025:

- **sh**
- **system**
- **shell**
- **enable**
- **linuxshell**
- **"/bin/busybox BOTNET"**
- **"/bin/busybox"**
- **"ping ;sh"**
- **ncorrect**
- **"/bin/busybox echo -ne ""\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00\x00"" >> .d"**

While the first seven are pretty generic and are just needed to set up the environment, the last three are quite interesting. For example, the ping command is often used by botnets to check for the presence of the working internet and thereby detect simulated researchers' environments. The "ncorrect" string is commonly observed in the encrypted configuration of various Mirai clones.

Finally, the last command uses the echo tool to deliver the next-stage payload (in this case, under the filename ".d") to the compromised system. The whole binary payload is split into chunks and written on a disk one block after another. Here, we can see the block with zeros being the most common one as zeroes comprise a big part of executables and therefore will be written down on a disk more often than other bytes delivered this way.

5.5 Top Payload Types and Anti-Debugging Techniques

Tracking the payloads that malware deploys in compromised environments allows us to ensure our detections always work, regardless of the architecture they are compiled for. Here is the distribution of payload types that we observed in 2H 2025. This time, we excluded multi-architecture shell scripts that work on various architectures and focused only on executable binaries.

Top Binary Payload Types by Targeted Architecture

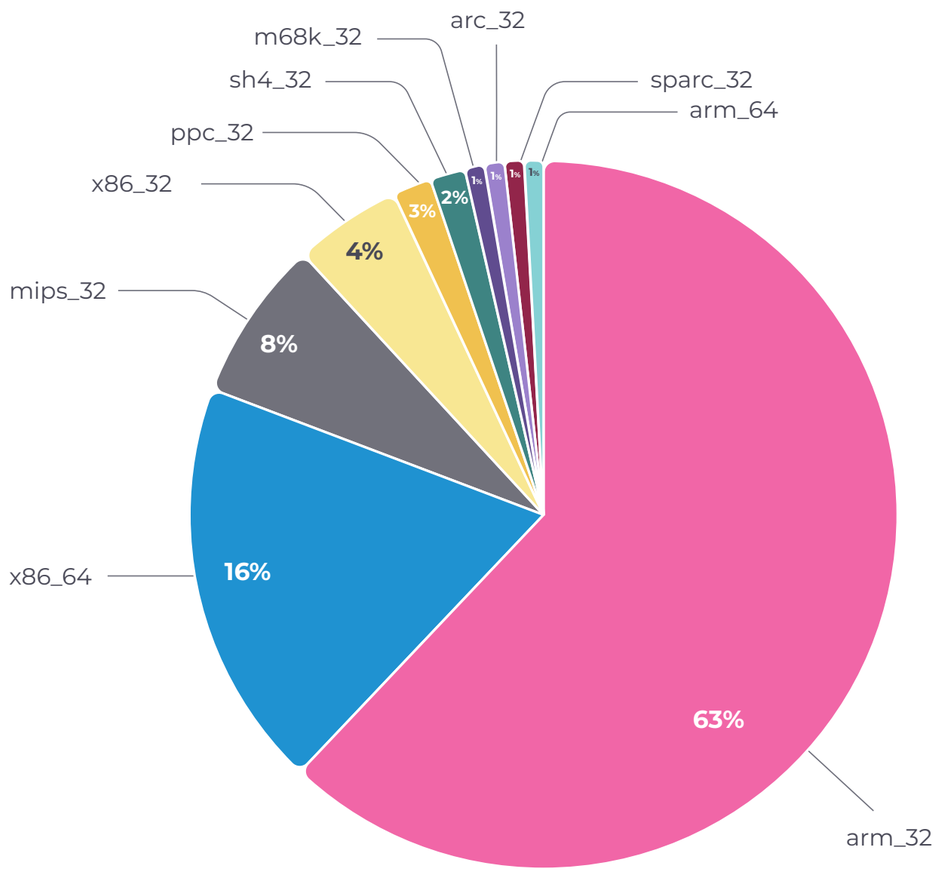

The 32-bit ARM remains the most popular architecture among attackers, followed by a 64-bit version of x86 and a 32-bit MIPS. Although other architectures may be targeted less frequently, devices that are based on them aren't inherently more secure. Similar to the infamous "security through obscurity" approach, that strategy will eventually fail. Instead, solid visibility and the use of cybersecurity solutions fully aware of these architectures and their unique features provide a much more solid long-term defense.

Turning to the packers that malware uses to protect itself against researchers and cybersecurity solutions, we observe the following distribution:

Top Packers Used to Protect Malicious Payloads

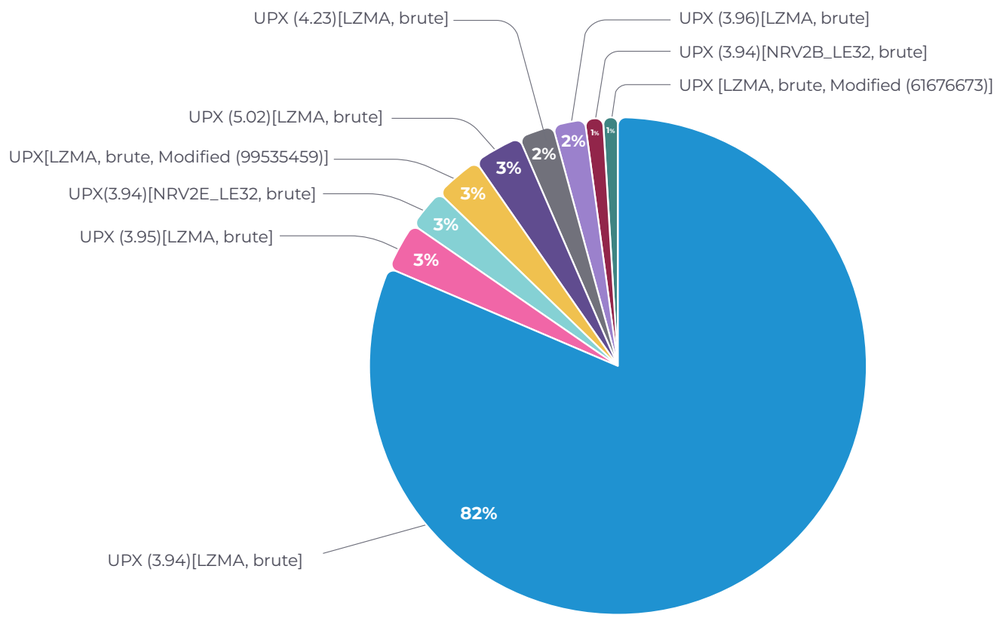

Despite the availability of alternatives like kiteshield (at least for x86), the old versions of UPX continue to be the most common packers used by IoT malware creators. This is likely because:

- **The old versions of UPX can be easily installed as packages on the relatively old versions of OSs that malware developers use**
- **Attackers want to simplify their operations by using one packer for multiple payloads targeting different architectures.**

Regardless of these reasons, cybersecurity solutions must be able to readily unpack malware payloads (including augmented variations with patched UPX structures) before applying pattern-based signatures like YARA. To facilitate this, we developed a public open-source tool that can fix such edited UPX fields, and in this way make malware samples easily unpackable by automation. We continuously maintain the tool and invite the community to use it, contribute improvements and share feedback to make automated unpacking more robust for everyone.

6. Recommendations

Establish complete asset and network visibility across OT and IoT

Organizations must achieve continuous, real-time visibility into all OT and IoT assets, communications and dependencies. Comprehensive visibility is the foundation of effective risk management, enabling security teams to identify unknown devices, understand normal operational behavior and detect anomalies before they escalate into incidents.

This directly addresses the visibility gaps observed in credential exposure, wireless activity and botnet propagation highlighted throughout this report.

Leverage AI-driven security systems to detect anomalies and threats

Given the scale and complexity of modern industrial environments, AI and machine-learning–based security platforms are essential. These systems can automatically baseline normal behavior, identify subtle deviations indicative of cyber or operational risk, surface the most critical issues and provide relevant context and guidance, thereby dramatically improving detection accuracy while increasing SOC efficiency. It's therefore essential to adopt a multi-layered defense strategy and deploy solutions capable of identifying abnormal behavior that deviates from your established baselines.

Prioritize risk-based vulnerability management for OT environments

Traditional vulnerability management approaches are insufficient for industrial networks that may contain hundreds or even thousands of OT and IoT devices from a variety of vendors, many of which are insecure by design—lacking authentication, encryption and other security standards. Organizations should adopt risk-based vulnerability management that correlates asset criticality, exploitability and operational impact, allowing teams to focus remediation efforts on the vulnerabilities that pose the greatest real-world risk without disrupting operations.

Correlating exploitability metrics and asset criticality is essential to address the high concentration of High and Critical vulnerabilities still present in operational environments.

Strengthen malware prevention and detection in industrial networks

Malware targeting OT environments continues to increase in both frequency and sophistication. Security programs must include advanced malware detection and blocking capabilities that inspect industrial protocols, monitor lateral movement and identify malicious payloads, including previously unseen ones. The goal is to successfully recognize and stop malware at the earliest stage possible to minimize the impact.

Detect and monitor wireless threats in OT and IoT environments

Wireless technologies such as Wi-Fi, Bluetooth, cellular and proprietary RF protocols introduce significant blind spots if left unmonitored. Organizations should deploy dedicated wireless threat detection to identify rogue access points, unauthorized devices, misconfigurations, weaknesses like the lack of MFP and wireless-based attack paths that could bypass traditional network defenses.

Wireless exposure repeatedly emerged as a silent enabler across multiple attack stages, making it a foundational control rather than a niche capability.

Enable intelligence sharing to improve collective cyber resilience

Threat intelligence sharing—across industries, regions and vendors—is critical to staying ahead of attackers. By contributing to and consuming actionable OT-specific threat intelligence, including telemetry sharing, organizations can detect emerging threats faster, improve defensive posture, and strengthen overall ecosystem resilience against large-scale or coordinated attacks.

Cybersecurity for OT, IoT and Critical Infrastructure

Nozomi Networks protects the world's critical infrastructure from cyber threats. Our platform uniquely combines network and endpoint visibility, threat detection, and AI-powered analysis for faster, more effective incident response. Customers rely on us to minimize risk and complexity while maximizing operational resilience.

NOZOMI NETWORKS

nozominetworks.com