Telefónica Tech

Security Status Report 2025 H2

From mobile security to vulnerability scanning, from breaking news to threat tracking, you understand the risks in today's landscape.

telefonicatech.com

Index

EXECUTIVE OVERVIEW

The purpose of this report is to synthesize the cyber security information of the last few months (from mobile security to the most relevant news and the most common vulnerabilities), adopting a point of view that covers most aspects of this discipline, in order to help the reader understand the risks of the current landscape.

Cyber security experienced two notable events in the second half of 2025.

A group of attackers calling themselves Scattered LapSus Hunters launched an unusual threat against Google in September 2025, demanding that the company dismiss two members of its Threat Intelligence Group and warning that they would leak alleged internal Google data if their demands were not met. The attackers did not show any clear evidence of having actual access to internal systems, and the threat was probably nothing more than a bluff. However, the tactic raises questions about the future of extortion. Psychological pressure on technicians (whose full names were provided) and reputational pressure on cyber security professionals instead of (only) demanding money for data. Could this virtual targeting of reputable technicians discourage ongoing investigations in the future? We will have to keep an eye on this new tactic, which opens the door to weakening the organisational response and sowing fear or mistrust towards investigation teams.

On the other hand, a few days later, Anthropic described the first cyberattack orchestrated almost entirely by AI. It is considered the first major cyberattack carried out largely by artificial intelligence with minimal human supervision.

According to the report, a state-sponsored cyber group manipulated Claude to execute offensive tasks autonomously. The model was “tricked” into doing most of the work (from system reconnaissance and exploit code generation to credential collection and data extraction) with very little human intervention. A line has been crossed: AI is no longer limited to assisting attackers with suggestions or partial automation, but can orchestrate complex campaigns on its own, taking the speed and scale of cyberattacks to a new level.

Attackers are undoubtedly taking advantage of the best of all worlds, without regulations or restrictions. Extortion not only of companies, but also of researchers. Something unheard of, which we hope will remain in the realm of anecdote. And the use of infinitely powerful tools such as AI, bent to their interests, to develop more attacks, more powerful in much less time. This potentially demonstrates that attacks can be taken to another level of intensity and depth with very few resources.

Attackers are becoming increasingly bold in their goal of undermining both the technological defences and the morale of those who maintain them. Simply because they can.

Both amateurs and professionals need to be able to keep up with relevant cyber security news: what are the most important developments? What is the current landscape? This report provides readers with a tool for understanding the state of security from different perspectives, enabling them to ascertain its current status and project possible short-term trends. The information gathered is largely based on the compilation and synthesis of internal data, cross-checked with public information from sources we consider to be of high quality. Let's go!

HIGHLIGHTS OF THE SECOND HALF OF 2025

The following are the news items that have had the greatest impact during the second half of 2025.

JULY

- The **FIA** is an international non-profit association that coordinates numerous motor racing championships, including Formula 1 and the World Rally Championship (WRC). The FIA (Fédération Internationale de l'Automobile) **stated that attackers accessed personal data after compromising several email accounts** following a phishing attack.
- ESET researchers **discovered that RomCom was exploiting a 0-day vulnerability in WinRAR discovered on 18 July 2025** and notified the team behind the popular archiving tool. ESET believes the vulnerability was used to extract dangerous executables to autorun paths when a user opens a specially crafted file.
  The vulnerability was similar to another path traversal issue in WinRAR, disclosed a month earlier and identified as CVE-2025-6218.
- French state-owned defence company Naval Group announced that it was investigating a cyberattack after **1 TB of allegedly stolen data was leaked on a data leak forum.** The company described this as an "attempt at destabilization" and an "attack on its reputation", responding by filing a complaint to protect its customers' data. On 23 July 2025, a threat actor known as "Neferpitou" published a large sample of 13 GB of data allegedly stolen from Naval Group. The data contained what appears to be a classified CMS for military vessels, technical documents, development virtual machines with simulation data, and internal communications.

AUGUST

- **Nissan Creative Box,** the creative arm of the Japanese multinational car manufacturer, **was attacked by ransomware and lost a significant amount of sensitive data in the incident.** The company is a studio specialising in satellite design that is part of Nissan's global design network.
- Google acknowledged that it **had suffered a cyberattack that resulted in the leakage of Google Ads customer data.** The company did not indicate how many customers were affected, but the attackers, "Sp1d3rHunters", claim to **have collected roughly 2.55 million records.** The attack targeted Salesforce's CRM, and the company had previously detected a phishing and vishing campaign targeting several of its employees.
- **The United States** Federal Judiciary **confirmed that it had suffered a cyberattack on its electronic case management systems, which store confidential court documents.** The organisation stated that, although most of the documents in the system are public, certain sealed files contain confidential information that is now protected by stricter access controls designed to block attackers.
- **Researchers at Nextron Systems discovered Linux malware that had evaded detection for over a year,** allowing attackers to gain persistent SSH access and bypass authentication on compromised systems. The malware exploits the PAM (Pluggable Authentication Modules) authentication infrastructure, featured obfuscation techniques, anti-debugging capabilities, encryption, and cleaned the execution environment of any traces of malicious activity.

SEPTEMBER

- **Asahi Group Holdings, Ltd (Asahi), the group responsible for Japan's best-selling beer, reported a cyberattack that disrupted several of its operations. According to the company, the incident affected its ordering and shipping activities, which were completely suspended.** Call Centre and customer service desk operations were also rendered inaccessible. Asahi holds approximately one-third of the domestic market share. It employs 30,000 people, produces 100 million hectolitres of beverages, and in 2024, the company reported annual revenues of nearly US$20 billion.
- The Google Threat Intelligence Group (GTIG) **detected malware called "Brickstorm" that has remained active for an average of 393 days on its victims.** Researchers confirmed that among the compromised organisations are the legal and technology sectors, software-as-a-service (SaaS) providers, and business process outsourcing (BPO) contractors.
- **Google released its sixth patch of the year against a 0-day vulnerability in Chrome classified as exploited.** To put this into perspective, in September 2024, 10 patches with the same characteristics were released.
- **Microsoft and Cloudflare shut down a massive phishing-as-a-service (PhaaS) operation, known as RaccoonO365,** which helped cybercriminals steal thousands of Microsoft 365 credentials. The cybercriminal group behind this service (also tracked by Microsoft as Storm-2246) **stole at least 5,000 Microsoft credentials from 94 countries since at least July 2024,** using RaccoonO365 phishing kits that included CAPTCHA pages and anti-bot techniques to appear legitimate and evade analysis.

OCTOBER

- **A foreign actor infiltrated the National Nuclear Security Administration's Kansas City National Security Campus** through vulnerabilities in Microsoft's browser-based SharePoint application, **raising questions about the need to further consolidate IT/OT security protections at the federal level.**
- **On October 6, the "Aisuru" botnet reached the highest peak ever recorded in the history of a DDoS attack: 29.69 Tbps.** The attack targeted several online gaming platforms. This same botnet, based on Mirai and which relies on infecting IoT devices, has been used to execute other large-scale attacks, such as the 22.20 Tbps recorded and mitigated by Cloudflare and the 15.72 Tbps recorded by Azure, which were directed against a single IP in Australia.
- **Spain dismantled the cybercriminal group "GXC Team" and arrested its leader, a 25-year-old Brazilian known as "GoogleXcoder".** The GXC team, which was very active in Spain, operated a crime-as-a-service (CaaS) platform that offered AI-powered phishing kits, Android malware, and voice phishing tools via Telegram and a Russian-speaking attacker forum. **The analysis of the devices seized in the initial arrests (carried out in six cities across Spain) took more than a year due to the complexity and extent of the cybercriminal group.**

NOVEMBER

- **Emergency alert systems across the United States disrupted following cyberattack OnSolve CodeRED,** a cloud-based critical events and mass notification platform. The platform recently suffered a cyberattack that forced it to shut down its environment, **as well as losing sensitive data and even a business customer. OnSolve is a service that helps organisations send urgent alerts and communications via SMS, telephone, email, push notifications and more.** It is used by state and local governments, such as the police and other emergency services.
- **A research team found thousands of credentials, authentication keys, and configuration data related to organisations in sensitive sectors within online code beautification services ("Codebeauty")** and format layouters (e.g., JSON formatters). These services incorporate the ability to store code for sharing via a pseudo-randomly generated and unindexed URL. The critical issue is that these URLs are not encrypted and do not contain any other authentication mechanism, so by guessing the URL generation structure, the content can be accessed without any problem.

DECEMBER

- The National Police **arrested a 19-year-old suspect in Barcelona, accused of stealing and attempting to sell 64 million records obtained through security breaches at nine companies.** The cybercriminal accessed nine different companies where he obtained millions of private personal records, which he then sold online.
- **The National Investigation Office of the South Korean National Police arrested four people for accessing IP cameras and selling compromising videos on a foreign website.** The individuals arrested were not connected to each other. One of them shared material involving underage children. Action was also taken against those responsible for the websites where the content was posted and against the users who accessed those videos. The police also contacted the victims who they were able to reach in order to warn them, provide them with all the information they had, and try to help them as much as possible.
- **MITRE published the 25 weaknesses (CWE) most frequently associated with vulnerabilities in 2025.** Cross-site Scripting ranks first, repeating its position from 2024, while SQLi and CSFR rank second and third respectively, each moving up one position as out-of-bounds writing falls to fifth place.
- **A vulnerability in MongoDB is being actively exploited, with more than 87,000 potentially susceptible instances identified worldwide.** The vulnerability is CVE-2025-14847 (CVSS score: 8.7), which allows an unauthenticated attacker to remotely leak sensitive data from the MongoDB server memory.

MOBILES

Apple iOS

From iOS 18 to iOS 26

A quantum leap? No, but it is a drastic change in the visual language of Apple's operating system.

The change in numbering, which we previewed in the last edition of this report, is purely aesthetic. Apple seems to be adopting a number that sounds more like 2026 than previous years, given how close the version numbers were to the calendar.

In addition to the change in numbering, iOS debuts a new look with an abundance of transparencies, modernised icons, etc. Beyond aesthetics, let's focus on the security improvements that iOS 26 brings.

New permission for wired devices: we can now decide whether or not to allow external devices to connect via USB when the phone is locked. The options are always asked, ask only for new devices not seen before, allow only if the device is locked, or always allow connection.

This allows us to control what external devices connect to our phone and under what conditions. The most practical use from a defensive standpoint is to eliminate or mitigate the risk of a malicious device being connected via USB in our absence.

Post-quantum key exchange in TLS 1.3: from now on, key exchange in encrypted communications will be done using quantum-resistant algorithms, which should preserve communications supposedly captured in the present from cryptographic attacks in the future, when quantum technology is expected to be available to the public.

Recovery assistant: iOS 26 includes a new feature that allows the device to be recovered if it detects faults in the system startup process, preventing the terminal from locking up if the process is aborted due to errors.

Contact blocking: A new section has been added to the privacy and security settings for managing blocked contacts.

Blocked or unwanted contacts are now treated as a shared set between certain applications and the system, such as contacts classified as telephone spam or unknown numbers. This feature aims to reduce the problem of unwanted calls.

Evolution of vulnerabilities in iOS during the second half of 2025

The first half of 2025 closed with 153 patched vulnerabilities, four of them considered high risk, with the potential to execute arbitrary code.

A similar figure to the second half of 2025, which surpasses it by a significant margin: 178 patches.

VULNERABILITIES IN IOS
Evolution of vulnerabilities per year

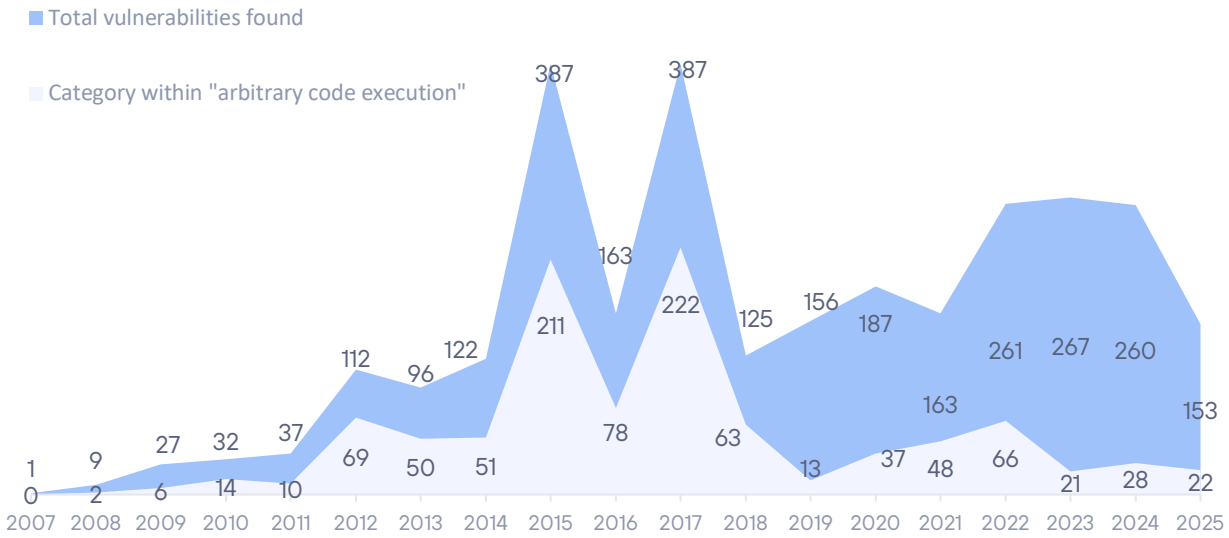

Vulnerabilities and versions released in the second half of 2025

We had to wait until the end of July for iOS 18.6 to see the light of day. The sixth iteration came with a batch of 32 security patches, many of them memory corruptions that could cause sudden application termination or information leaks, among other causes. It has no significant vulnerabilities or any that are particularly noteworthy. Fifteen days later, 18.6.1 was released without any security patches.

18.6.2 was released on 20 August. It is an emergency patch. CVE-2025-43300 is being actively exploited and is also targeting specific targets, as noted in Apple's own report. It is an error in the out-of-bounds write control of a memory buffer during image processing in the ImageIO component. A fairly quiet month.

September kicks off the season with new numbering. On 15 September, iOS 26 is released, and it is not exactly empty in the security section, with 34 patches correcting vulnerabilities. At the same time, 18.7 is released, fixing 13 vulnerabilities. On the same day, patches are released for CVE-2025-43300 in versions 15.8.5 and 16.7.12, which in theory have not been supported since March this year, although it is common for patches to be released for versions that are still “alive” if the impact is particularly serious.

Later, on 29 September, a patch is released for a vulnerability that fixes an out-of-bounds write in a buffer when processing text sources. With CVE-2025-43400, versions 18.7.1 and the first iOS 26 update, 26.0.1, appear.

And although October is the month of Halloween, ironically no security updates are released, making it an unusually quiet month, with no tricks or treats.

November brings us the first major update to iOS 26, version 26.1. It comes with no less than 61 security patches. It is followed by the new iteration for iOS 18, 18.7.2, with 38 patches. Among them, there are a multitude of privacy bugs, memory management bugs that could lead to code execution, etc.

We close December with iOS 18.7.3 with 21 patches and a new major iteration for iOS 26, 26.2 with precisely 26 patches of varying importance.

Version fragmentation during the second half of 2025

Traditionally, fragmentation has never been a problem for iOS developers. The advantage of having a homogeneous platform is undisputed and continues to yield near-identical numbers every time we review iPhone user adoption of a new version of the operating system.

The top three positions in terms of market share are occupied by iOS 18 and its different versions. It is logical, since it is still a mature, widespread version and the release of iOS 26 is in its infancy.

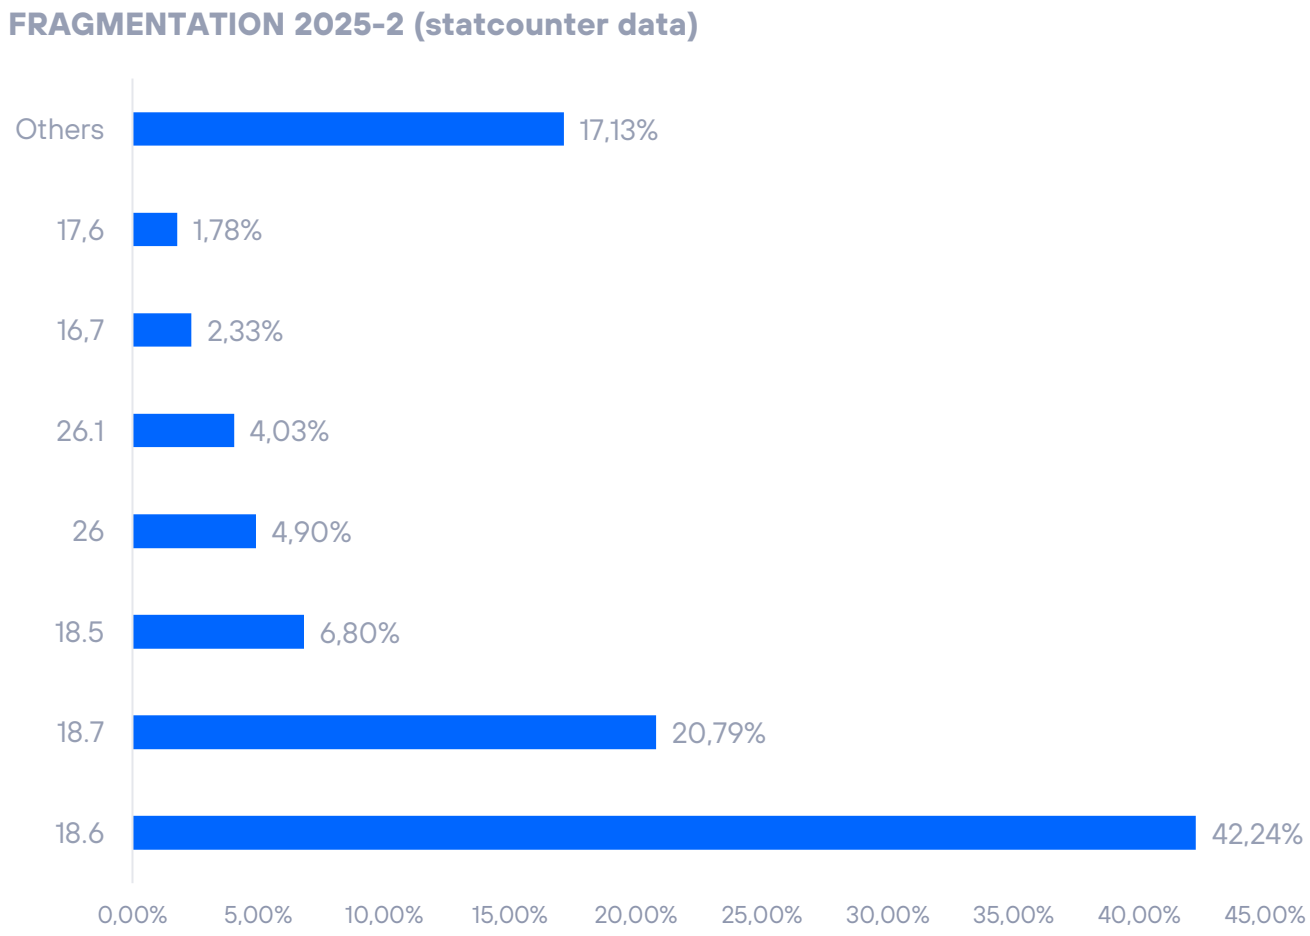

We have a 17.13% share that does not specify a version, which means it is a bag of versions that we cannot identify and may or may not be supported.

iOS 26 accounts for close to 10%. Given that the data is from November, we will have to wait until the next semester to see the real impact at the end of the year.

The latest version supported by Apple is iOS 18, released in September 2024. The support for 16 and 15 expired last semester.

Android

Android 16 consolidates its position

We will not have a new version of Android until next year, which is expected to be version 17, but we will not see it until almost halfway through 2026 if Google's mobile operating system continues to be released at the same pace.

Vulnerabilities

Android releases a set of patches every month, usually during the first week. In this second half of 2025, six bulletins have been published with the following distribution of vulnerabilities for each month:

Month	CVEs	Critical or CER
July	0	0
August	6	1
September	110	3
October	0	0
November	2	1
December	106	7

However, some CVEs may not have associated impact information as of the date of publication of this report, so the number of CVEs may subsequently be higher than indicated.

There has been a total of 224 patches this semester (compared to 253 last semester), 12 of which are considered critical (compared to 22 last semester). The total for 2025 has been 482 patches, 34 of which have been critical.

It should be noted that many of these flaws affect software or firmware from certain manufacturers in particular, which means that the same vulnerability does not necessarily affect the entire Android device fleet, but only those with the affected components.

ANDROID VULNERABILITIES:

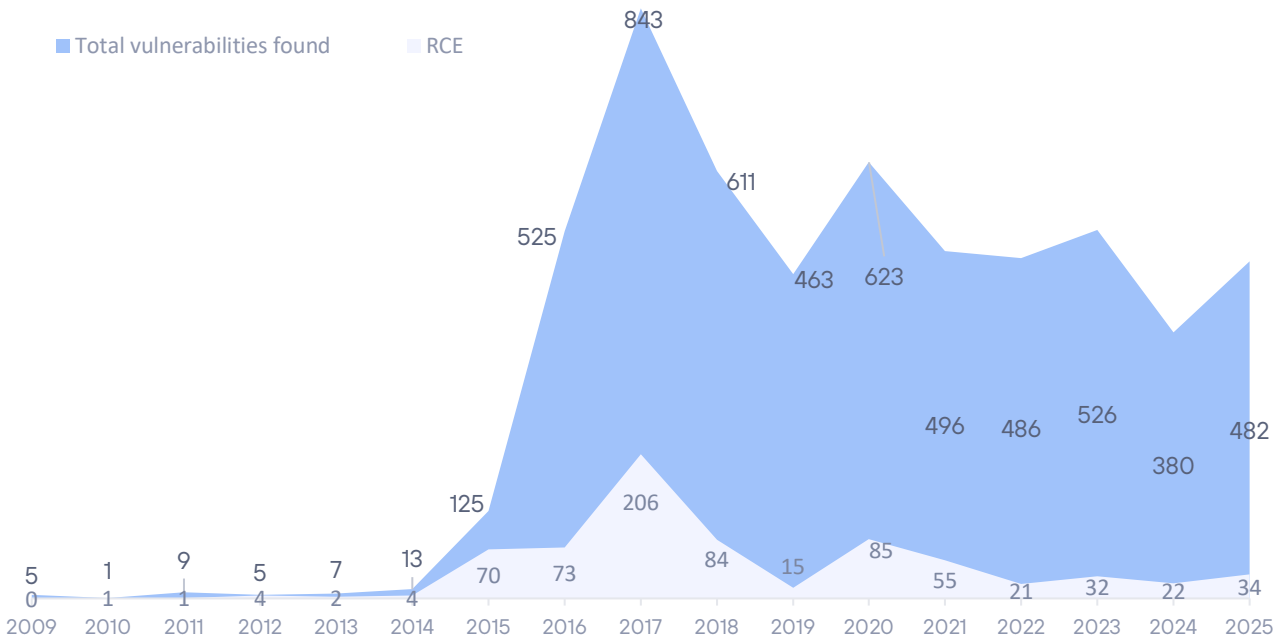

Fragmentation in Android systems

The latest publication by Statcounter at the time of writing this report indicates that the most widely used version of Android is 15, with a share of 28.27%. Despite the release of Android 16 in June 2025, as usual, new versions of Android are slow to catch on and accumulate a user base.

The ranking continues to show a worrying share of unsupported Android versions. This means that almost 35% of Android systems do not receive updates.

Support for Android 12 ended last March, so only versions higher than this receive updates, fulfilling the 3-year support period since the operating system was released.

The ranking is as follows:

FRAGMENTATION 2025-2

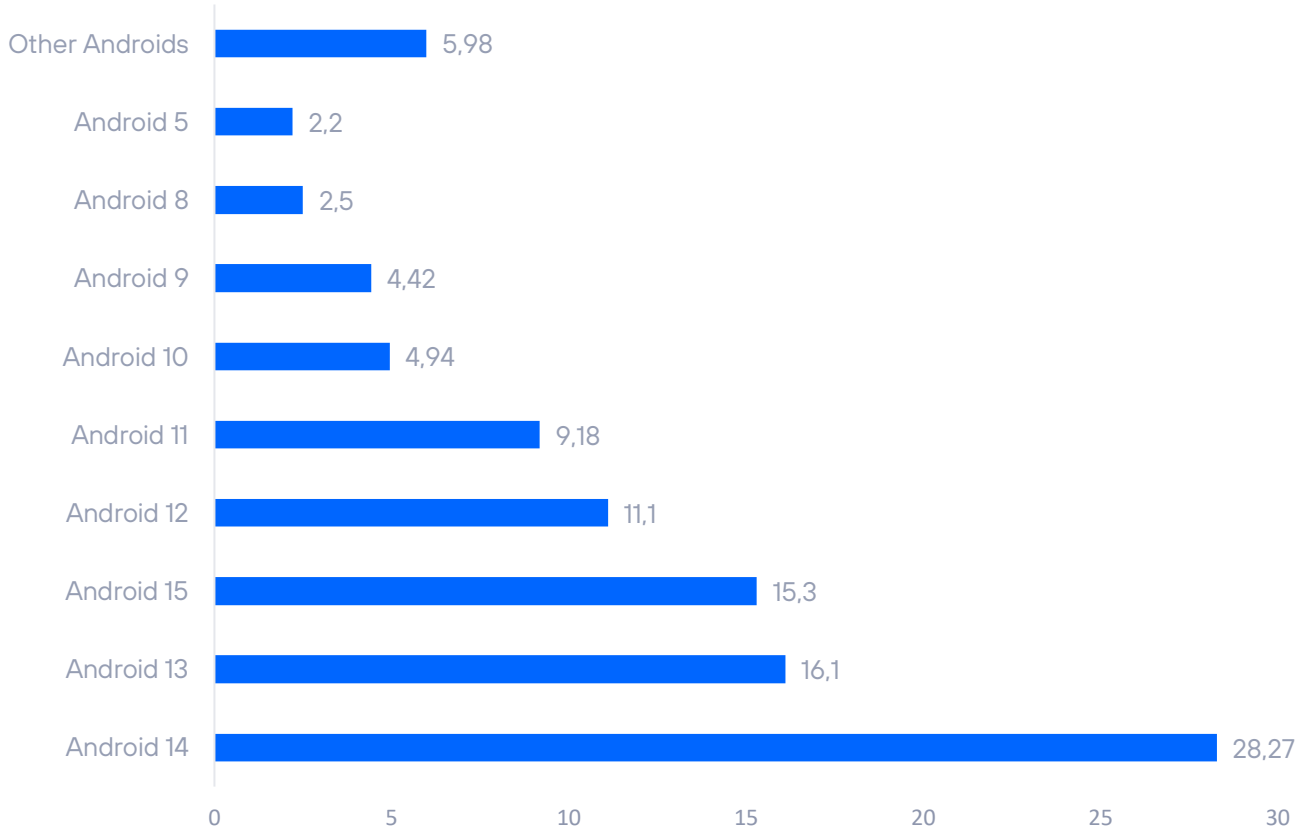

We might reasonably ask ourselves, why are such old versions of Android still on the market? We must bear in mind that many Android devices with a long lifespan are still in use in countries with less developed economies. These are inexpensive devices with modest features, but they still serve a basic function for people in those regions.

SIGNIFICANT VULNERABILITIES

In this section we will discuss what we consider to be some of the most notable vulnerabilities in the second half of 2025.

CVE ID	TARGET	DESCRIPTION	SCORING
CVE-2025-55182	React Server Components (React2Shell)	It allows full control of the software via HTTP, which can be easily exploited remotely.	10
CVE-2025-8284	Packet Power Monitoring	The web interface for monitoring and controlling packet power does not implement authentication mechanisms. This vulnerability could allow unauthorised users to access and manipulate the monitoring and control functions.	9.3
CVE-2025-9696	PVS6 de SunPower	A vulnerability that could allow attackers to gain full access to the device, enabling them to replace the firmware, modify settings, disable the device, create SSH tunnels, and manipulate connected devices.	9.4
CVE-2025-10659	Telenium Online Web Application de Megasys Enterprises	A vulnerability that could allow an unauthenticated attacker to inject arbitrary operating system commands through a specially crafted HTTP request, resulting in remote code execution on the server in the context of the affected system account.	9.3
CVE-2025-12108	License Plate Recognition (LPR) Camera de Survision	There is no default password protection, which allows immediate access to the configuration wizard without having to log in or verify credentials.	9.3
CVE-2025-13510	Productos iHUB e iHUB Lite de Iskra	There is a lack of authentication in existing iHUB and iHUB Lite products in the web administration interface. This allows unauthenticated users to access and modify critical device settings.	9.3
CVE-2025-10035	License Servlet of Fortra's GoAnywhere MFT	A vulnerability in the servlet could allow an attacker with a forged licence response signature to deserialise an arbitrary object, leading to command injection. If the log files contain “SignedObject.getObject”, it is likely that the instance has been affected by this vulnerability.	9.8

CVE-2025-6218	RARLAB WinRAR	This vulnerability allows remote attackers to execute arbitrary code. User interaction is required to exploit this vulnerability, as the target must visit a malicious page or open a malicious file. The flaw is found in the handling of file paths within compressed files. A manipulated file path can cause the process to access unwanted directories. An attacker can exploit this vulnerability to execute code.	7.8
CVE-2025-8088	WinRAR	A path traversal vulnerability affecting the Windows version of WinRAR allows attackers to execute arbitrary code by creating malicious files.	8.4

Vulnerabilities in figures

The distribution of CVEs published by risk level (scoring based on CVSSv3), in terms of number of vulnerabilities discovered, was as follows.

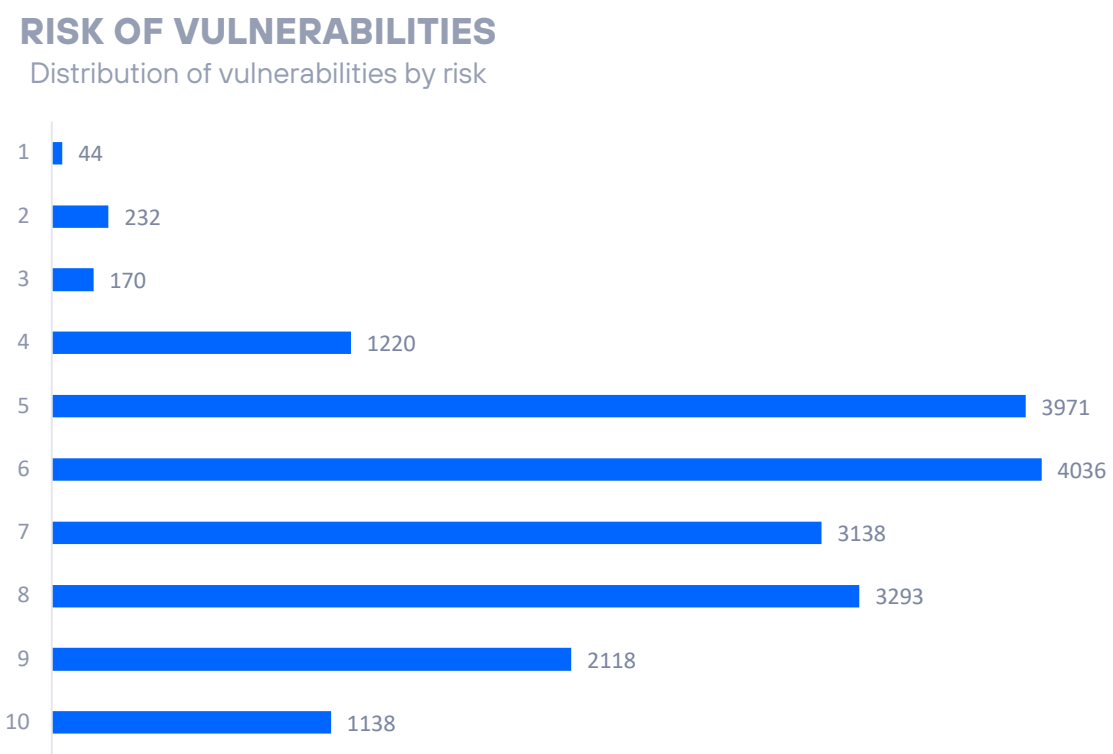

Top 25 companies with the most accumulated CVEs

In the second half of 2025, Microsoft led by some distance in terms of the number of known vulnerabilities, followed by Linux and Apple. It is now commonplace to see them in the top three positions.

The rest of the manufacturers are made up of large organisations, medium-sized software developers and even small companies.

VULNERABILITIES BY MANUFACTURER

Top 25 manufacturers by cumulative CVE

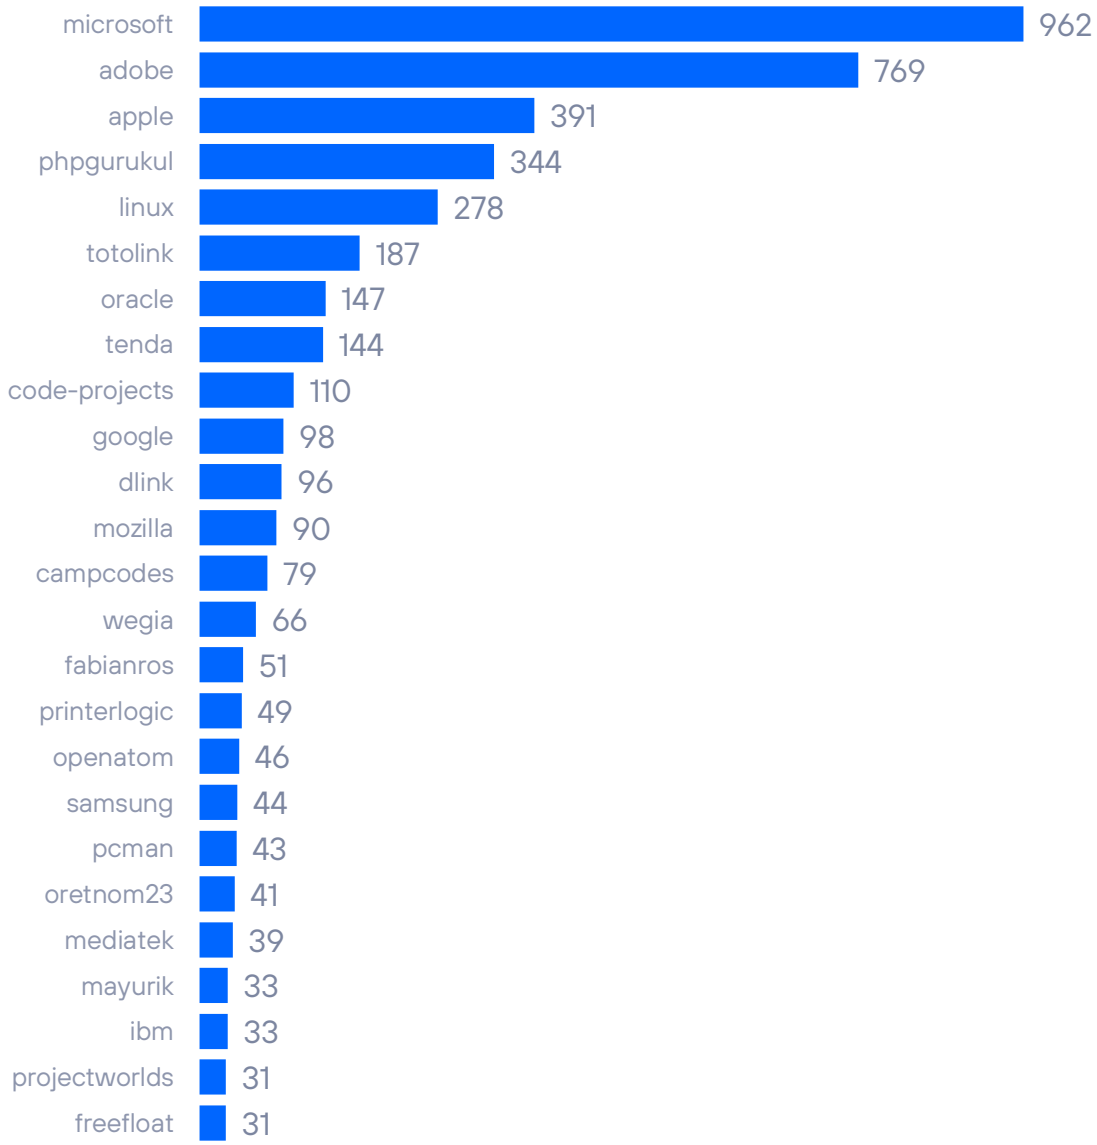

APT OPERATIONS, ORGANIZED GROUPS, AND ASSOCIATED MALWARE

We review the activity of the various groups attributed with the authorship of APT operations or noteworthy campaigns.

We must point out that the attribution of this type of operation, as well as the composition, origin and ideology of the organized groups, is complex and cannot necessarily be completely reliable. This is due to the capacity for anonymity and deception inherent in this type of operation, in which the actors may use means to manipulate information in such a way as to conceal their true origin and intentions. It is even possible that in certain cases they may act with the modus operandi of other groups to divert attention or harm the latter.

Remarkable APT activity detected during the second half of 2025

APT36 – Transparent Tribe: A leopard hard to spot

12 years operating on behalf of Pakistan's interests, specifically targeting India. Recently, in line with this plan, a cyber espionage campaign targeting Indian government institutions was detected.

The campaign targets Linux Boss (Bharat Operating System Solutions) systems, an Indian distribution based on Debian. Its entry vector is a phishing email (PDF information) with a link to a ZIP file (or the ZIP file itself) located on a service such as Google Drive or a dedicated staging server. The ZIP file executes the DeskRAT RAT.

This actor also launched an additional phishing campaign during this semester, albeit a different one, in which they requested information from their victims under the pretext of a meeting. Among other information, they requested the Kavach code. Kavach is a 2FA application developed by the National Informatics Centre (NIC) to improve the security of government email services in India by generating time-based OTPs. Together with the user's password and email address, the attackers then had everything they needed to authenticate themselves on Indian government systems.

More information *https://blog.sekoia.io/transparenttribe-targets-indian-military-organisations-with-deskrat/*

Lazarus: Get up, walk... and succeed

Lazarus returns to the spotlight for the second consecutive semester.

In this case, they were discovered targeting European defence companies involved in the development of drones: a metallurgical company, an aeronautical component manufacturer and a defence contractor. ESET researchers have grouped this activity into a campaign called Dream Job, in which several groups from various countries are participating.

They compromised the security of their targets in the following way: social engineering (job offers) in PDF format that forced the opening of a 'Trojanised' reader. From there, the RAT ScoringMathTea, first detected in campaigns in Portugal in 2022, is installed and takes control of the systems to collect information about them.

More information: https://www.eset.com/us/about/newsroom/research/north-korean-lazarus-group-targets-drone-sector-europe/

Primitive Bear: keeps evolving

Another actor that appears twice in our 2025 report. In this case, it has been detected participating in espionage activities against Ukraine. Its efforts have focused on extracting high-level intelligence from the Ukrainian government.

In order to survive, the bear keeps evolving. In this campaign, it employed a whitelist domain camouflage technique to construct malicious URLs. This method uses legitimate syntax (username:password@host). From there, they use Microsoft's Tunnelling Service for Developers infrastructure to obtain a valid TLS certificate (issued by MS) and hide their activity among legitimate traffic.

More information: https://cn-sec.com/archives/4411359.html

Volt/Salt Typhoon: Two typhoons threaten Australia

Australian intelligence chief Mike Burgess said at least two Chinese state-sponsored groups are positioning themselves for future sabotage and espionage operations against Australian networks, raising serious national security concerns about pre-emptive cyber intrusion and potential disruption of critical infrastructure.

According to Burgess, the Volt Typhoon group is targeting electricity, water and transport networks, while Salt Typhoon is focusing on Australian telecommunications networks.

These groups are not new and, in fact, have already been recognised by media outlets such as the Wall Street Journal as part of retaliatory actions against the United States for its support of Taiwan.

More information: https://securityaffairs.com/184540/intelligence/australias-spy-chief-warns-of-china-linked-threats-to-critical-infrastructure.html

OT THREAT ANALYSIS

The following information comes from the OT threat capture and analysis system, Aristeo. **It** incorporates a network of **decoys, made of real industrial hardware,** which appear to be industrial systems in real production. **They behave as such** but are extracting all the information about the threats accessing the system. Aristeo uses the information from all the devices deployed in the different decoy-nodes to apply relationships and intelligence to go beyond the data, being able to proactively detect campaigns, targeted or sectorized attacks, 0-day vulnerabilities, etc.

Each node-signature has its own characteristics and reproduces a different process. Therefore, protocols, devices, productive sectors... change in each of them. In addition, the nodes are alive, which means that they can undergo alterations in their configuration at the discretion of the team of researchers working with them, or of the customer who has temporary or permanent use of them. This variability may generate slight discrepancies in the data shown in this section when compared between semesters.

Information analysis

We have some exciting news for Aristeo this semester. In addition to its natural capabilities for extracting information from industrial devices and networks (through its network of decoys), it now has the ability to extract information from 5G decoys. Aware that private 5G networks are increasingly used in industrial environments, we thought it was time to take Aristeo one step further and introduce this concept into its network of decoys.

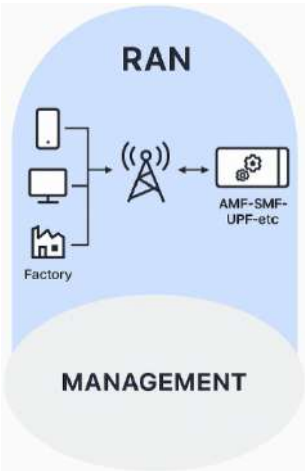

The decoy could receive interaction from attackers from the management environment (via remote access) and from the radio spectrum. However, to use the latter attack vector, the attacker would need to be within range of the antenna. In environments where the customer does not seek to monitor their surroundings (a typical operating environment), the decoy emulates traffic from devices on the 5G network.

Naturally, the system is designed to avoid superfluous elements in the decoy network and give it an appearance of legitimacy. It is integrated as follows:

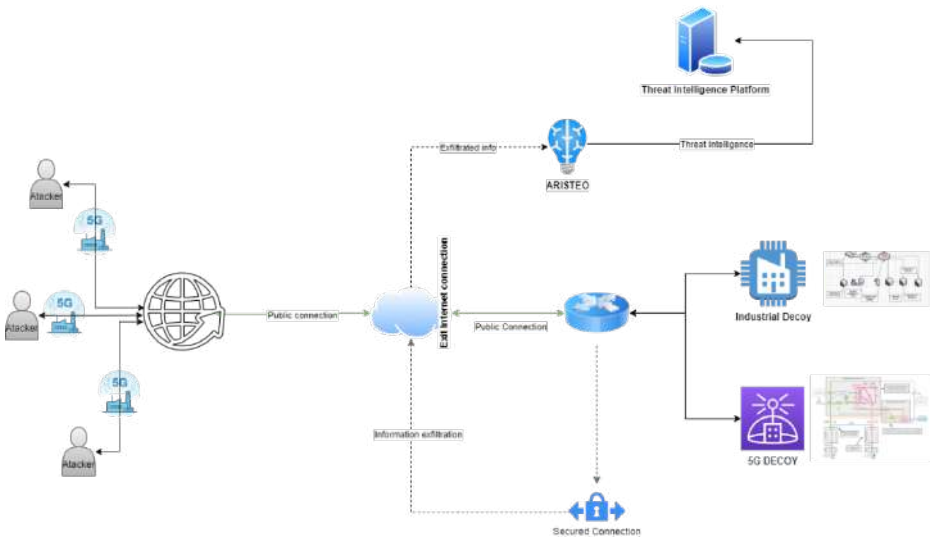

This shows that the 5G decoy is located within the infrastructure of an industrial decoy. This makes it seem like a private environment for deploying a 5G bubble in that industrial environment. An attacker could additionally attempt to gain access by moving laterally through the compromise of one of the industrial decoy devices. As the system is operational, the attacker can register devices, alter configurations, collect traffic between devices and the antenna (whether physical or virtual, as indicated above) or attempt to manipulate the values of industrial devices within the same network.

Let's move on to the general statistics of the information recorded. In the second half of 2025, 43 million cyber security events were detected. At this point, and for the last time (as we have already mentioned several times), it is worth remembering that these are complex events and that, thanks to Aristeo 2.0, these events are now associated with each other, which means that the more than 214 million "simple" events we have had this semester have been converted into that figure of 43. This figure represents a sharp drop in activity compared to last semester and the same period (second semester) in 2024.

Part of the reason for this decline is likely due to a targeted attack by a professional actor on one of our decoys permanently connected to these statistics. We provided more details about this targeted attack in the previous half-yearly report and in a post on the Telefónica Tech blog, but its activity continued until July 2025. This led us to pause the activity of this decoy while we implemented what we had learned from this actor after an in-depth forensic analysis.

Focusing now on the statistical analyses specific to this report, the distribution by country would be as follows

Interactions - 2025H2

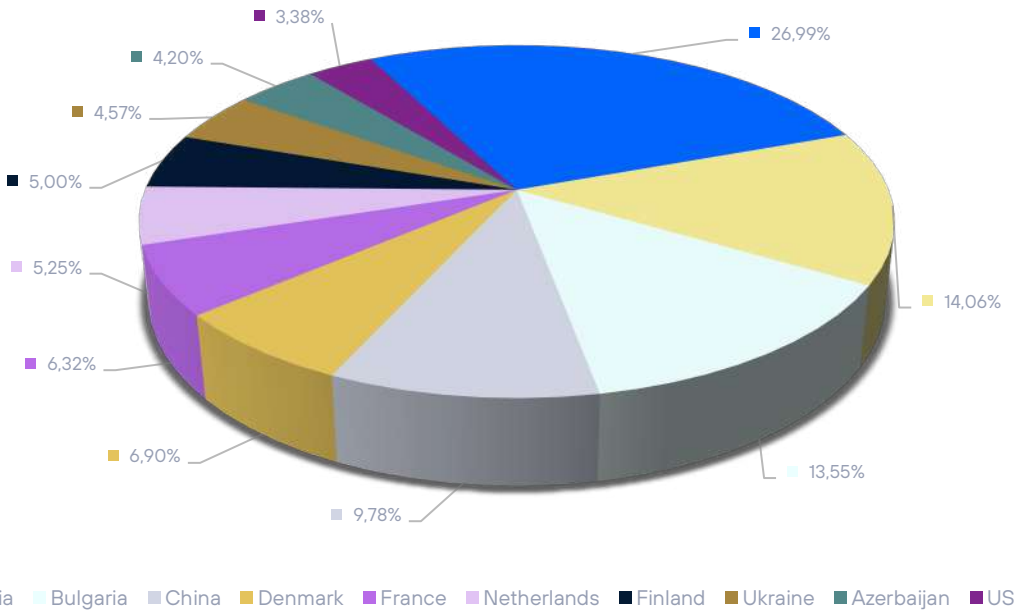

This semester, Russia leads the top 10, while Bulgaria, last semester's winner, drops to second place. The distribution has changed compared to last semester, when the top-ranked country accounted for 36% of events and the others became less representative. This semester, the situation is different.

Now let's look at the ten IP addresses with the most interaction with the Aristeo system. This semester, the countries with the most visits to Aristeo are the ones that rank highest on the list. In this case, we could speak of a certain stability. However, we could also say that the TOP 10 (we see the IPs without obfuscation) is made up of blocks of IP addresses and their activity is related to more automated and coordinated actions of approximation and discovery. It seems that the cycle of searching and reviewing the security of devices floating around the internet is coming back.

TOP-10 IP attackers

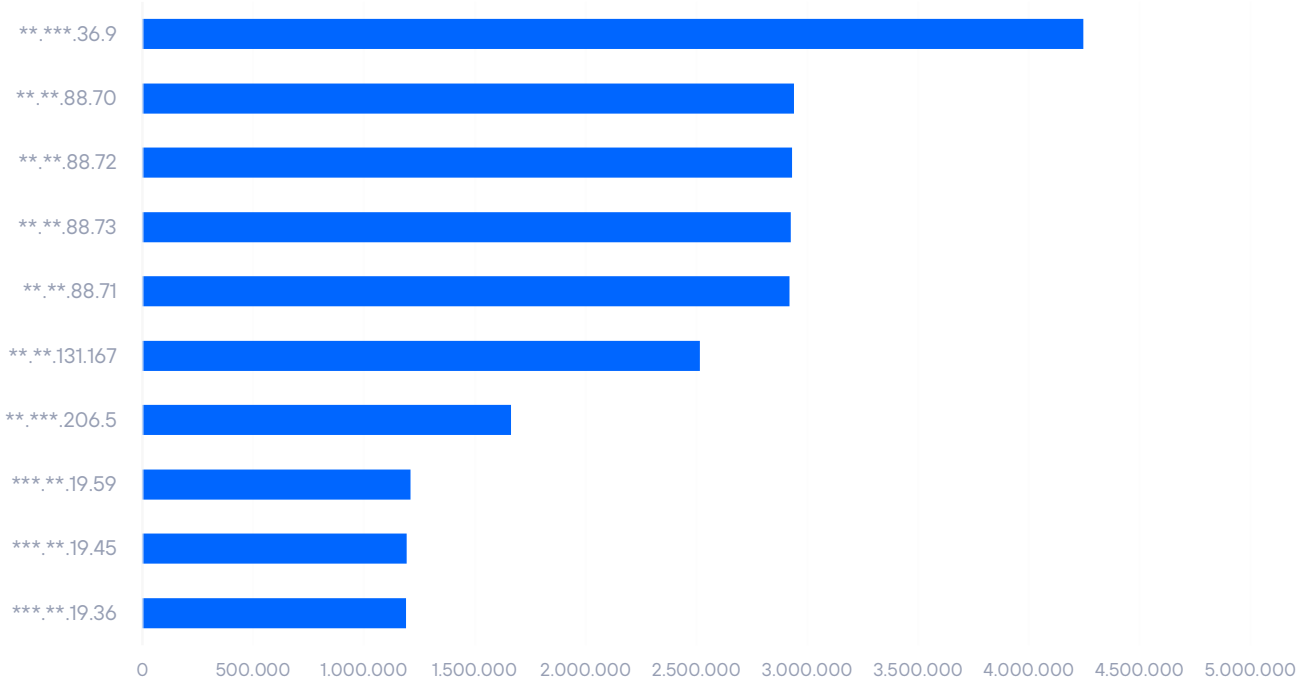

We can see below how the top 10 countries are distributed. As mentioned above, interactions on Aristeo have been more evenly distributed than in the previous semester. This is due to increased activity from Russia and China, two countries that had a very low-profile last semester.

Top 10 countries

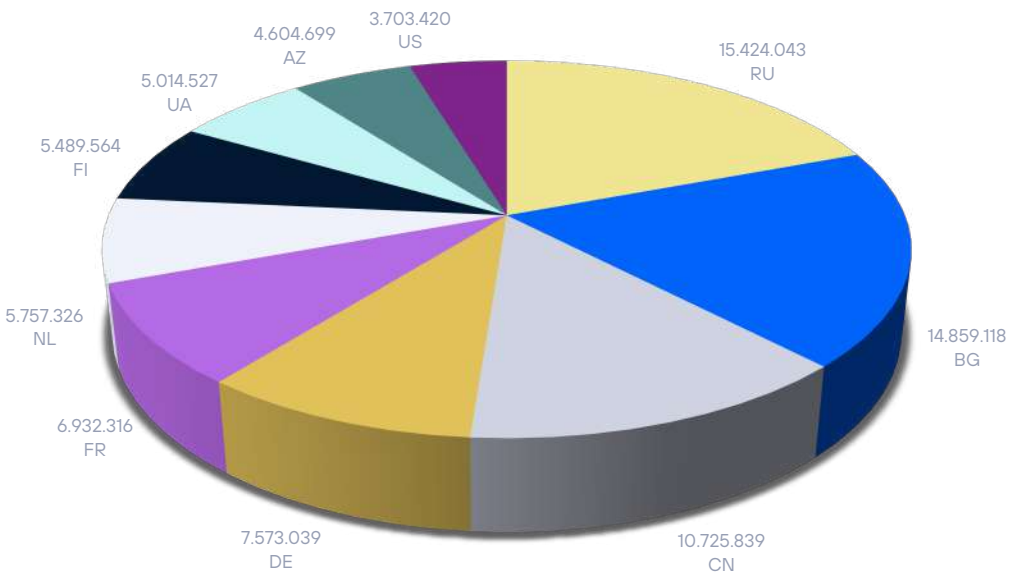

Thanks to Aristeo 2.0, this semester we are once again analysing the TTPs (Tactics, Techniques and Procedures) most exploited by attackers, but with a different graph from last semester.

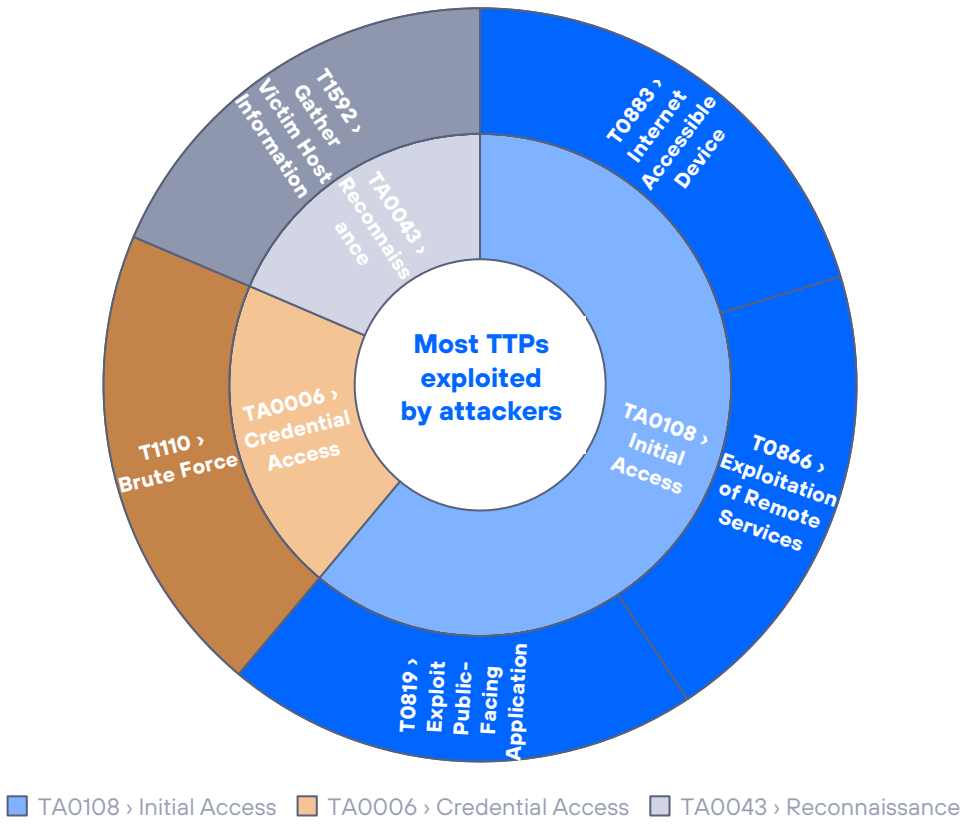

It can be seen that most of the activity focuses on initial actions, such as brute force access attempts or attacks targeting accessible services and devices. Other actions outside the TOP 5, such as exfiltration and information gathering, are less common because Aristeo's decoys are not an open bar. As a good deception environment, the decoys are properly configured, and only high-level attackers can access and continue to demonstrate their TTPs. Regarding references to physical media, Aristeo environments are not usually physically accessible to anyone, but they can still be used as learning environments for employees or as bait if the customer wants to detect potential insiders.

STUDY OF THREATS BY INDICATOR

In collaboration with **Maltiverse**, we have conducted a ranking study of the indicators of compromise detected on their platform. That is, to indicate interesting attributes of maliciousness detected in IP addresses, domain names and URLs over the last six months.

A total of 334,335 IP addresses, 228,954 domains and 486,004 URLs have been studied regarding the various IOCs involved.

What type of maliciousness do the URLs studied involve?

As we know, URLs allow us to access resources, they describe a protocol, a machine on the Internet (either directly through an IP or indirectly from a domain) and within that machine a resource is specified through a path.

In the end, in the context of malware, every IP and domain will be part of a URL to request a resource. Whether it is a URL that directs us to a phishing site and has a domain very similar to the original one, or it may be that the URL serves as a download point for malware.

It is important to determine what is at the end of the URL and categorize it properly to know what type of threat we are dealing with. This is precisely what we have asked to the Maltiverse database and we have found these results in the top 10:

Phishing	212.556	43,74%
Malware Download	95.334	19,62%
Whatsapp Phishing	20.885	4,30%
Naver Phishing	19.251	3,96%
Trojan.generic	18.427	3,79%
Trojan.phishing.pdf	17.246	3,55%
Booking Phishing	6.590	1,36%

Malicious.PDF	6.001	1,23%
Facebook phishing	4.675	0,96%
Rakuten phishing	4.641	0,95%

There are no surprises regarding the two categories with the highest number of indicators: phishing and malware download. Because if there is a classic in cyber security regarding what awaits us at the end of a URL, it is precisely these two major categories

However, these are categories that group or assimilate a large part of what we find in the long queue. The rest of the categorizations are more explicit and even indicate to which malware family they belong to.

It is worth highlighting the presence of new types of threats that we were already aware of, phishing attacks specialising in specific brands, such as Facebook, Rakuten, Booking, etc. These are by far the most common. The strategy is clear, to impersonate these well-known brands that inspire trust in order to lower the guard of potential victims.

The rest are divided, as we can see, into generic categories, which include all kinds of vectors and threats: malware inserted into documents or various types of trojans.

Which domains are most commonly used by URLs marked as malicious?

This edition we have consulted with Maltiverse to find out which domains appear most frequently in the URLs studied.

It is interesting to note which services, mostly legitimate, are the most employed by malware writers and their associated campaigns.

In the end, a URL will have a hosting or redirection and needs an executable web space or application that at some point it will use for its purposes. It is the domain that will “tell us” where it has been hosted and what service it has made (illegitimate) use of.

ru.com	21.067	4,33%
google.com	9.881	2,03%
pages.dev	7.618	1,57%
vercel.app	6.847	1,41%
za.com	6.836	1,41%

sa.com	5.902	1,21%
checkin-arrivals.com	5.823	1,20%
weebly.com	4.873	1,00%
github.io	4.494	0,92%
duckdns.org	2.563	0,53%

As usual, the top spots belong to online services that allow free web content hosting. The rise of 'ru.com' is surprising, as it has never been seen before with these figures, as is "google.com", with the rest being "classics" in this ranking.

It is a common pattern: why take a risk on private hosting or compromised servers when you are offered free and anonymous hosting?

There are also domains associated with these malicious URLs that use dynamic domain resolvers: duckdns.org. In other words, they are actually naked IPs that, through a free DNS service, can be resolved to a particular subdomain, and even if they need to migrate the malicious infrastructure, they move the IP address and continue to resolve to the new location.

As we can see, in both types of service, the trend is always the same: free and anonymous. These are two characteristics that are sought after and eagerly exploited by cybercriminals.

Which countries are the IP addresses detected with malicious activity?

Before answering the question, it should be clarified that just because a country appears in this ranking does not mean that there is any malicious intent with respect to that country. Many countries stand out from the rest because they have more services and hosting companies, which translates directly into a greater fraudulent use. A server can be hosted in one country and the criminal organization using it can come from another nationality.

United States	67.172	20,09%
China	31.376	9,38%
India	24.325	7,28%
Brazil	14.227	4,26%

The Netherlands	12.312	3,68%
Vietnam	12.115	3,62%
Germany	11.478	3,43%
Russia	10.730	3,21%
Singapore	10.199	3,05%
Canada	9.879	2,95%

There are no major variations in this aspect in recent years. These are countries with large technological infrastructures and, therefore, as mentioned above, they have a proportionally greater potential to be used by cybercrime.

What kind of maliciousness do IP addresses engage in?

Suspicious host	157.456	47,10%
Malicious host	156.157	46,71%
HTTP Spammer	97.928	29,29%
Mail Spammer	86.382	25,84%
SSH Attacker	46.492	13,91%
Bruteforce	40.668	12,16%
DDoS Attacker	37.612	11,25%
HTTP Attacker	36.094	10,80%
Port Scanner	35.594	10,65%
Hacking	34.796	10,41%

Crowning the top 10 ranking is a general category: "Suspicious host". It is a categorization that practically overlaps half of the dataset since it is awarded whenever there are indications of suspicious activity although the operation observed from that IP address is not yet known in detail.

When a label is added later on, with the detail of why: spam, indiscriminate scans, etc., the suspicious host label is not removed as it is a further refinement. Another type of generalist labeling is found in "Malicious host". Identical meaning, although it adds a little more certainty in the preliminary diagnosis.

If we aggregate the tags by specific IP address activity, we see that SPAM, both HTTP and Mail, top the ranking with almost 80% of the tags. As a reminder, tags overlap, so the same IP can contain several of them. For example, a general "suspicious" and "HTTP Spammer", and even the same IP can be used for port scanning because it has been a detected activity at some point in time.

SSH Attacker is a unique category. It almost certainly belongs to groups of infected hosts coordinated by a Mirai-type botnet. Mass scanning for easy access via SSH (Secure Shell) has been a constant for decades on the Internet (as was Rlogin or telnet in its early days). Almost 13,91% of IP addresses have been observed performing attacks on SSH (mostly dictionary attacks on the login).

Similarly, "Bruteforce" refers to the continuous attempt to perform brute-force authentication (actually, again: common username and password dictionaries). This category accounts for 12.16%, slightly less than the previous figure.

We can find another subcategory (10.65%), indiscriminate scans, which include Port scanner. IP addresses that have been detected performing massive scans on complete ranges or multiple ports on certain hosts. In other words, horizontal scans searching for certain ports or vertical (in-depth) scans on a group of hosts.

We find the ‘hacking’ category with 10.41% closing the ranking. These are nodes that have been observed performing attacks in general, either trying to find SQL vulnerabilities or launching exploits. These are often vulnerability scanners used indiscriminately and, of course, without authorisation.

What are the top-level domains (TLDs) with the most malicious domains?

As we know, a domain resolves to an IP address. Domains are extremely important in the world of cybercrime, as they allow criminals to use them and change the IP address if the server that is active at that moment ceases its malicious activity.

A domain is composed of several levels. If we look closely, we can see that they are strings separated by dots. If we take these groups from right to left, they form a hierarchy. The one on the far right is the highest-level domain.

This allows us to group domains categorised as malicious by their highest-level domain. The top 10 results are as follows:

com	71.573	31,26%
top	25.381	11,09%
shop	19.612	8,57%
ru	14.465	6,32%
app	8.268	3,61%
dev	7.534	3,29%
click	7.419	3,24%
xyz	6.721	2,94%
org	5.453	2,38%
cn	5.108	2,23%

'com' once again tops our TLD ranking this semester, dethroning "xyz", undoubtedly driven by the aforementioned 'ru.com'. It does so with force, almost tripling the figures of its closest competitor.

'shop' is rising strongly. The TLD specialising in e-commerce is a common target for phishing scams that impersonate online retailers, gaining the trust of victims, and not many domains are registered under this TLD, which is also quite inexpensive.

The '.app' TLD is particularly interesting, as Google paid ICANN more than $25 million in February 2015 to take control of it. Furthermore, HTTPS traffic is mandatory for this TLD.

The rest of the domains are made up of the usual suspects that always make it into the ranking.

What malicious categorization do the studied domains possess?

Domains are closely linked to URLs (of which they form part) and also, of course, to the IP addresses to which a domain resolves.

Lastly, let's see how the top 10 have been categorised over the last six months.

Phishing	149.014	65,08%
ClearFake	141.800	6,19%
Generic Malware	9.394	4,10%
Malware Download	7.590	3,32%
MetaStealer	5.039	2,20%
Necurs	4.015	1,75%
WhatsApp Phishing	3.866	1,69%
Command and Control	3.560	1,55%
Orchard	3.090	1,35%
Phishing Allegro	2.488	1,09%

As we have already mentioned, there is a very close relationship between domains and URLs. This can be seen in the top 10 categories: phishing and malware in general. The rest belong to malware families that have had an impact.

USEFUL LINKS

Do not just stay in the top layer of cyber security analysis, the semi-annual reports are both cumulative and summarized. Telefónica Tech's cyber security blog has much more information and news which may be interesting for you. Here are our most relevant articles.

CYBER SECURITY

Cómo hacer frente a los fraudes telefónicos: así se intentan proteger España y otros países de Europa contra este tipo de estafas – (Only Spanish)

Cyber Intelligence in OT: staying ahead of the attack

From paper to practice: how to build an effective OT cybersecurity roadmap

ARTIFICIAL INTELLIGENCE

AI sandbox: secure environments for evaluating and protecting Artificial Intelligence models

Quantum Machine Learning: the next revolution in AI?

Can you trust that AI? Verifiable credentials are your guarantee