Telefónica Tech

Security Status Report 2025 H1

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

telefonicatech.com

Índex

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.

This semester, two news items stand out. Although apparently unrelated, we believe they have a common denominator.

The U.S. government announced in April 2025 that it would stop funding MITRE to operate and maintain the Common Vulnerabilities and Exposures (CVE) system. This system has been managed by MITRE since 1999 with funding from the Department of Homeland Security (DHS) and the Cybersecurity and Infrastructure Security Agency (CISA). Yosry Barsoum, vice president of MITRE, warned that the contract would expire on April 16, 2025, and that, if not renewed, there would be serious consequences. The main reason pointed to be a budget cut driven by the Trump administration, although CISA and other sources did not publicly detail the exact reasons.

The potential service disruption generated alarm in the cyber security community. Fortunately, hours before the expiration, CISA stepped in and temporarily renewed MITRE's contract, avoiding an immediate disruption of CVE services. This solution was yet considered temporary and highlighted the fragility of relying on a single government funder for such a critical resource. In response, the creation of the **CVE Foundation**, an independent entity driven by members of the CVE board itself, was announced. The entity will ensure the long-term sustainability and neutrality of the system.

On the other hand, in June, what appeared to be one of the biggest cyber security incidents in history was revealed: the leak of 16 billion passwords and login credentials for services such as Google, Apple, Facebook, GitHub, Telegram... and even government portals. The investigation was led by the Cybernews team, which detected at least 30 databases of different sizes, some with more than 3.5 billion records each. It did not look like an attack on a single company, but a massive collection of thefts. It was eventually learned that the alleged "leak" of 16 billion passwords was not a new leak or a recent attack, but a compilation of credentials previously stolen over the years through infostealer-type malware, previous breaches and credential stuffing attacks. The affected sites were not recently compromised to obtain these credentials.

What happened (as other times) is that a massive database was compiled and exposed, composed of records already circulating in underground forums and marketplaces. There is no evidence that it contains unpublished data (or at least, not a relevant amount) or extracted from new breaches. The first news went unnoticed by the mainstream media. While the cyber security industry was biting its nails for 24 hours for fear of losing the CVE system, the world went on. The second news, in contrast, opened generic news in all countries, with front pages and reports. Everyone knew about the leak. In this case, the world was on edge while the industry suspected (by pure logic) that it was a compilation of old thefts of no great significance. The numbers were not so much in technical reality as in a headline.

When the continuity of CVE was compromised, the alarm was immediate: experts, companies, and security managers realized that losing this infrastructure meant flying blind in the face of

vulnerabilities, a chaos that would affect incident response, infrastructure protection and global coordination. However, the password leak, despite its enormous media coverage, did not generate the same professional concern because, in essence, it did not alter the threat landscape or introduce new risks: it was, above all, noise and recycling of already exposed data.
This phenomenon highlights a persistent gap between what really matters in the industry for digital resilience and what the headlines capture. The media tends to amplify the most spectacular or easily understood incidents, even if their real impact is limited, while structural issues - such as critical infrastructure financing and governance - are relegated or overlooked.

This not only distorts the user's perception, who often ends up believing any amplified alarm without context, but also reflects a certain immaturity in news coverage: the right experts are rarely consulted and the underlying technical consequences are rarely explained. The industry, for its part, reveals that what is really important is not always what is most visible, and that security depends more on the solidity of its foundations - such as CVE - than on the sensational nature of the incidents that make the front pages. Are we still, perhaps, too far removed from the user?

Whether you are an amateur or a professional, it is important to be able to keep up with relevant cyber security news: what is the most relevant thing going on? What is the current landscape? Through this report, the reader will have a tool to understand the state of security from different perspectives and will be able to learn about its current state and project possible short-term trends. The information gathered is based in large part on the compilation and synthesis of internal data, contrasted with public information from sources we consider to be of quality. Let's go!

HIGHLIGHTS OF THE FIRST HALF OF 2025

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

JANUARY

- A flaw that allows hibernate images to be retrieved in clear text is identified, potentially exposing passwords and sensitive data if an attacker has physical access to the BitLocker-encrypted hard drive (CVE-2025-21210). This won't be the only problem with this technology this year.
- **Two serious vulnerabilities are discovered in SAP:** CVE-2025-0070, Improper Authentication in SAP NetWeaver AS for ABAP and ABAP Platform, which allows an authenticated attacker to gain illegitimate access to the system by exploiting improper authentication checks. In addition to CVE-2025-0066, an information disclosure in SAP NetWeaver AS for ABAP and ABAP Platform (Internet Communication Framework).
- A so-called **"Belsen Group"** leaks configuration files, IP addresses, and VPN credentials of more than 15,000 FortiGate devices from around the world, both public and private sector, for free on the dark web. The 1.6 GB dump includes sensitive information such as clear text passwords, private keys and firewall rules, and has been collected after exploiting vulnerability CVE-2022-40684 in 2022.
- **Moxa,** a provider of industrial communications and networks, issues an urgent **advisory on January 3 about two high-impact breaches.** There are no mitigation measures available for some of the affected devices at the time. One of these vulnerabilities allows control of affected devices and systems to be taken over remotely.
- Xlab publishes research on a Mirai-based botnet that is exploiting 0-day vulnerabilities in industrial and home routers. The botnet handles about 15,000 bots daily and most of them are located in the USA, China, Russia, Iran, and Turkey. The most common activity of this botnet is DDoS attacks.
- Although it occurred in October, it becomes known in January. **Cloudflare mitigates the largest DDoS attack ever recorded,** which peaks at 5.6 terabits per second. The attack is launched by a Mirai-based botnet with 13,000 compromised devices, it targets an Internet service provider in East Asia and lasts only 80 seconds.

FEBRUARY

- **A 0-day vulnerability in 7-Zip (CVE-2025-0411) was detected in September 2024** that allowed Russian attackers to bypass "Mark of the Web" (MoTW) protection in Windows using double-nested compressed files. This flaw is exploited during February in phishing campaigns targeting Ukrainian government agencies and companies, allowing the execution of malware such as SmokeLoader without displaying security warnings to the user.
- **At least 28 apps on Google Play and Apple's App Store included malware,** known as SparkCat or SparkKitty, designed to steal cryptocurrency wallet recovery phrases using optical character recognition (OCR) techniques. This malware, hidden in legitimate apps and downloaded more than 242,000 times, accesses the device's image gallery and searches for screenshots with keywords associated with cryptocurrencies, allowing attackers to empty victims' wallets. This is the first time such a stealer has managed to infiltrate the App Store.

- The FBI blames North Korea, through the APT Group "Lazarus", for the **theft of 1.5 billion dollars from the cryptoasset exchange Bybit.** In this operation, which the FBI called "TraderTraitor" and is listed as the largest cryptoasset theft in history, the criminals manages to breach one of the cold wallets, considered secure, accessing the funds and splitting them into multiple wallets to make them harder to trace.
- **Lee Enterprises,** one of the largest newspaper groups in the United States, **says a cyberattack affecting its systems** caused a disruption in February and impacted its operations. These operations include worker access to the system, making it impossible to deliver articles, print, and deliver dozens of newspapers.
  This company, which publishes 77 daily newspapers with a daily circulation of 1.2 million, 350 weeklies and digital editions with more than 44 million unique visitors, was already attacked in 2020 by Iranian cybercriminals during the presidential election campaign.
- **According to chainAnalysis, in 2024, ransomware payments fell by 35%** compared to the previous year, totaling $813.55 million compared to $1.25 billion in 2023, despite record attack volumes. Only 30% of victims who negotiated with cybercriminals eventually paid up.
- **Two critical vulnerabilities are discovered in OpenSSH:** one, identified as **CVE-2025-26465,** allows (MitM) type attacks against OpenSSH customers when the VerifyHostKeyDNS option ware enabled. This flaw has existed since 2014 and particularly affected FreeBSD systems where the option was enabled by default. The second one, CVE-2025-26466, introduced in 2023, allows denial-of-service (DoS) attacks before authentication.

MARCH

- The **GreyNoise research group detects a backdoor campaign affecting thousands of ASUS** routers. According to the researchers, this movement would be focused on the establishment of a future botnet.
- Proofpoint researchers identify a highly targeted email campaign aimed at fewer than five Proofpoint customers in the United Arab Emirates related to aviation and satellite communications, as well as critical transportation infrastructure. This campaign exposes a backdoor called Sosano and the use of polyglots to obfuscate the payload content. These files are specially designed so that different applications interpret them as different file types.
- **Truffle Security researchers discover nearly 12,000 API keys** and valid passwords (including AWS, MailChimp and WalkScore credentials) within the Common Crawl dataset, widely used to train artificial intelligence models from companies such as OpenAI, Google and Meta, many of them hardcoded in HTML and JavaScript.
- **YouTube warns of a sophisticated phishing campaign** in which cybercriminals use an AI-generated video simulating the platform's CEO, Neal Mohan, announcing fake changes to the monetization policy. The attackers send this deepfake as a private video to content creators, along with emails that include links to a fake YouTube Partner Program verification page. Attempting to "confirm" the new terms, victims hand over their credentials.
- Researchers identify a **massive malware campaign called "Steam", in which 331 malicious apps on Google Play,** disguised as utilities such as health trackers, battery optimizers or QR scanners, achieves more than 60 million downloads. These initially legitimate apps to bypass

Google's controls, download malicious code after installation to display intrusive ads, steal credentials and card data through phishing, and make them difficult to remove by hiding their icon and activity.

- Researchers warn about the massive **exploitation of the critical vulnerability CVE-2024-4577 in PHP for Windows,** which allows remote code execution on servers using PHP in CGI mode. Although the flaw was patched in June 2024, attackers quickly exploited the exploit, especially in Japan, and then globally, with more than 1,000 unique IPs attempting to exploit the vulnerability in January alone.

APRIL

- The FBI, through the Internet Crime and Complaint Center (IC3), states that the **U.S. lost more than $16.6 billion in 2024.** As a reference, in 2023 the reported losses were $12.5 billion. It should be cautioned that, although the figure is very striking, it only reflects activity detected by the IC3 or reported by victims, so it represents only a portion of the final amount.
- Renal dialysis company **DaVita discloses that it suffered a ransomware attack that encrypted parts of its network and impacted some of its operations.** Days later, the cybercriminal group Interlock claims credit for the attack and announced that it had 20 terabytes of the company's confidential information. According to the same group, negotiations with the company were unsuccessful, so they put the information up for sale. DaVita has annual revenues in excess of $12.8 billion.
- **GitHub detects more than 39 million leaked secrets in repositories,** including API keys and credentials, exposing users and organizations to serious security risks. Despite measures such as push protection enabled by default on public repositories.
- **A critical remote code execution (RCE) vulnerability has been discovered in Apache Parquet,** a data storage format widely used in big data and analytics environments, especially on platforms such as Hadoop, AWS, Google Cloud and Azure. It is identified as CVE-2025-30065 and has a maximum CVSS score of 10.0.
- A massive campaign is discovered in April in which at least **ten malicious extensions for Visual Studio Code, published on Microsoft's official marketplace, infected thousands of Windows users with the XMRig cryptominer** for Monero.These extensions, which pretended to be legitimate development tools such as “Prettier”, “Discord Rich Presence” or “Solidity Compiler”, accumulates hundreds of thousands of downloads, many of them artificially inflated to appear popular.
- **Google Chrome 136 fixes a vulnerability that had been allowing websites to identify users'** browsing history by analyzing the color of links visited with the CSS :visited selector for more than 20 years. This flaw allows malicious sites to track, profile and even launch phishing attacks by detecting which links a user has visited on other sites.

MAY

- **A joint police operation** carried out by Dutch and US authorities **dismantles a botnet consisting of some 7,000 Internet of Things (IoT) devices and others in an End of Life (EoL) state.** The

network was rented to provide anonymity to the attackers. Attackers paid between $9.95 and $110 per month, generating revenues of more than $46 million.

- **Arla Foods, the Danish food industry giant, confirms that it was the target of a cyberattack that disrupts its production operations** at a dairy plant in Germany. This attack forced the company to inform its customers of delays in their orders.
- **A supply chain attack targeting Linux servers is detected,** where three malicious Go modules published on GitHub contained obfuscated code that downloaded and executed a Bash script capable of completely overwriting the main disk (/dev/sda) with zeros. This action causes total and irreversible data loss and renders the systems unusable. The modules, which imitated legitimate projects, are quickly removed.
- **The LockBit criminal organization suffers an internal breach when its administration and affiliate panels on the dark web were attacked and replaced with a warning message** along with a link to download a copy of its MySQL database. The dump exposes key operational details: more than 59,000 bitcoin addresses, attack configurations, affiliate credentials and, especially, more than 4,400 records of negotiations between victims and extortionists since December 2024.
- **Researchers uncover a massive campaign that introduced more than 100 malicious extensions to the Chrome Web Store,** masquerading as legitimate tools such as VPNs, AI assistants and utilities from well-known brands such as Fortinet and YouTube. These extensions, promote through fake websites and malvertising, offer some of the promised functionality, but in reality, they steal cookies, credentials and browsing data, can execute remote scripts, modify traffic and turn the browser into a proxy for the attackers.

JUNE

- **WestJet, Canada's second largest airline, confirms a cyberattack that disrupted access to some internal systems.** The attack also prevents users from logging into the website and mobile app, services that were restored by now.
- A new attack named '**SmartAttack' uses smartwatches as covert ultrasonic signal receivers to extract data** from physically isolated systems through the "Airgap" technique, used in many industrial environments. This technique can also interfere with the performance of RAM, displays, SATA cables... etc. Even if the reception of these signals can be done through Smartwatch, still someone must have compromised the Airgap isolated device.
- **In June 2025, CVE-2025-49113, a critical remote code execution (RCE) vulnerability in Roundcube Webmail** affecting all versions from 1.1.0 to 1.6.10, is disclosed. The flaw, caused by a lack of validation on the _from parameter in the upload.php script, allows an authenticated attacker to execute malicious code on the server, compromising the security of the platform. More than 84,000 Roundcube instances remain vulnerable.
- **In June 2025, two critical local privilege escalation vulnerabilities are discovered in Linux, CVE-2025-6018 and CVE-2025-6019,** which allow any user with access to a session (even via SSH) to gain root permissions on most modern distributions. The first flaw, in SUSE's PAM configuration, grants "allow_active" privileges to remote users; the second, in libblockdev and the udisks daemon (present by default in almost all Linux systems), allows these users to elevate themselves to root with very little effort.

- It is reported that more than **46,000 instances of Grafana exposed to the Internet remain vulnerable** to the critical flaw CVE-2025-4123, an open redirect and XSS issue that allows attackers to execute malicious plugins and hijack user accounts via a simple link. The exploit, which does not require elevated privileges and can run even with anonymous access enabled, allows session hijacking, credential switching and, in some cases, SSRF attacks against internal systems.

MOBILES

Apple iOS

The new iOS 18 security improvements

We will have to wait until after the summer to know the security improvements of the new version of iOS.

What we do know, so far, is that it will not be called iOS 19 but iOS 26. We are moving from linear numbering to naming the operating systems in a unified way according to the year; but even though it will be released in 2025, the name will be brought forward to the next one.

In the following report we will reveal the new security features that iOS 26 will bring us.

Vulnerabilities and versions released in the first half of 2025

We closed out 2024 with iOS 18.2. 2025 didn't take long to bring a new update, although in this case, it didn't include any security fixes. On January 6, iOS 18.2.1 was released.

It wasn't until January 27 that the first patches arrived. That day, Apple released iOS 18.3, which addressed 38 CVEs — some of them affecting the kernel, with the potential for arbitrary code execution or privilege escalation. Notably, CVE-2025-24085, a vulnerability in CoreMedia, was being actively exploited by attackers to escalate privileges on iOS versions prior to 17.2.

On February 10, Apple issued an emergency update: iOS 18.3.1. This patch addressed a very specific vulnerability, known to have been exploited in a targeted attack against a high-profile individual. Specifically, the update fixed an issue allowing access to a locked iPhone (CVE-2025-24200).

The alarms didn't stop there. On March 11, Apple released iOS 18.3.2 to patch another vulnerability that had been used in highly targeted attacks. This time, it was a WebKit buffer overflow (out-of-bounds write), registered as CVE-2025-24201.

Given the severity of these two vulnerabilities, Apple also decided to release updates for older devices. On March 31, versions 15.8.4 and 16.7.11 were published to address the same issues on legacy systems.

That same day, iOS 18.4 was released, patching a staggering 75 unique vulnerabilities. Despite the high number, only one of them — CVE-2025-24243, involving the Audio component — could be exploited to execute arbitrary code.

On April 16, Apple released another critical update: iOS 18.4.1. It patched two zero-day vulnerabilities that had been actively exploited in targeted operations — similar to previous cases. One was in the CoreAudio component, and the other in RPAC (Reconfigurable Preprocessing Architecture Core).

From that point until May 12 — the release date of iOS 18.5 — no further incidents were reported. This update addressed 33 vulnerabilities, including two that could allow arbitrary code execution.

Evolution of vulnerabilities in iOS during the first half of 2025

The second half of 2024 closed with 149 patched vulnerabilities, four considered high-risk, with the possibility of executing arbitrary code.

This number is similar to that of the first half of 2025, which is slightly higher with 153 patches.

VULNERABILITTIES IN IOS 2025-H1

Evolution of vulnerabilities by year

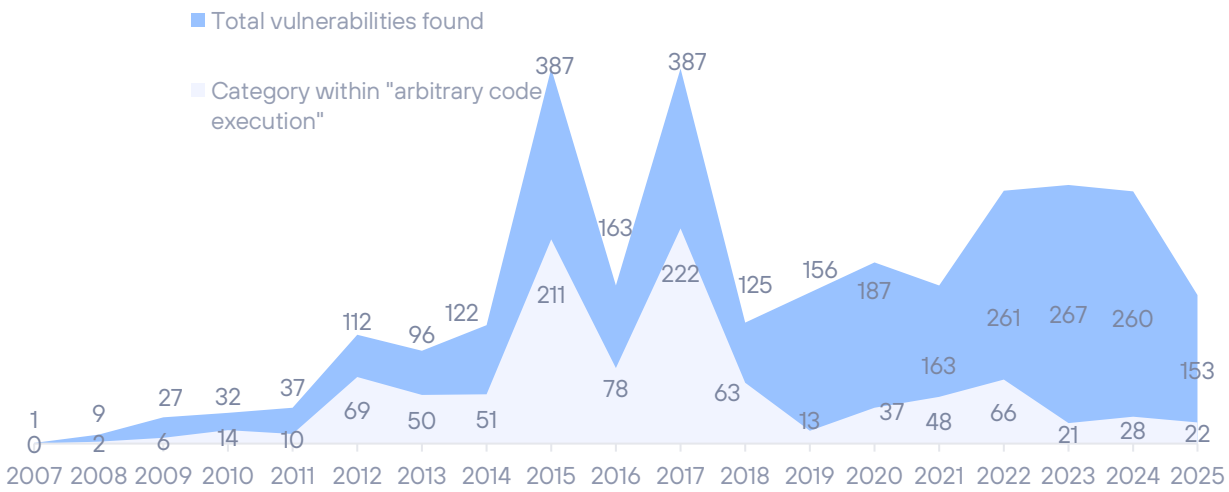

Fragmentation of versions during the first half of 2025

Fragmentation has never been traditionally an issue for iOS developers. The advantage of having a homogeneous platform is undisputed and continues to yield near-identical figures every time we review iPhone user adoption of a new version of the operating system.

And no wonder. The first four share positions are occupied by iOS 18 and its different versions.

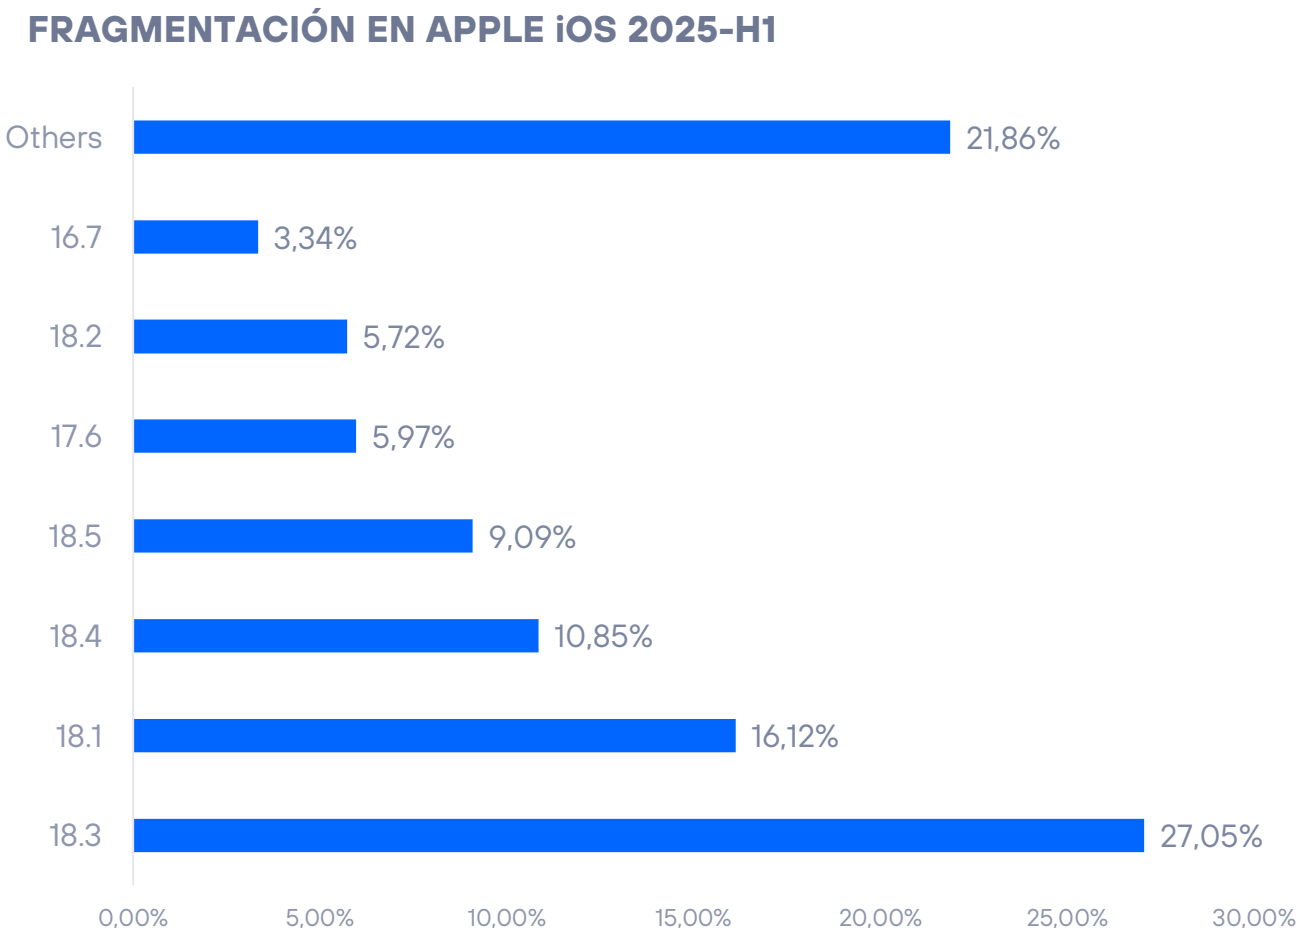

That said, we have a 21.86% share that does not specify version, so it is a bag of versions that we cannot identify and could be supported or not.

The latest version supported by Apple is iOS 15, released in September 2021.

Apple Transparency Report

Governments sometimes need to rely on large corporations to do their job. When a threat involves knowing the identity or having access to the data of a potential attacker or a victim in danger, the digital information stored by these companies can prove vital to the investigation and avert a catastrophe. Apple publishes a comprehensive report every six months on what data governments request from it, which data and to what extent the requests are fulfilled. We update here some data we have extracted from information published by Apple for **the second half of 2023 and first half of 2024 (the latest published by Apple as of the first half of 2025)** on government activities and requests to the company.

Device-based requests

It represents requests from government agencies requesting Apple device information, such as serial number or IMEI number. When law enforcement acts on behalf of customers who, for example, have had their device stolen or lost. It also receives requests related to fraud investigations: they typically request details of Apple customers associated with Apple devices or connections to Apple services.

Country	Requests 2023-H2	Accepted 2023-H2	Requests 2024-H1	Accepted 2024-H1	Addition	% Accepted
United States	6.410	5.233	12.043	10.377	18.543	85%
Germany	8.866	4.463	9.778	4.713	18.644	49%
China mainland	905	853	1.212	1.146	2.117	94%
Brazil	5.858	4.765	8.776	6.808	14.634	79%
United Kingdom	1.633	1.337	2.925	2.278	4.558	79%
Spain	620	290	540	186	1.160	41%

As usual, Germany leads the way in requests for device information, although in 2024 the United States made a strong comeback. Brazil is also a major player here. In Spain, there is a curious detail: the acceptance rate is quite low.

Requests based on financial data

These requests happen when law enforcement acts on behalf of customers who require assistance related to fraudulent credit card or gift card activity that has been used to purchase Apple products.

Country	Requests 2023-H2	Data Prov. 2023-H2	Requests 2024-H1	Data Prov. 2024-H1	Addition	% Accepted
Taiwan	4.415	4.345	4.968	4.819	9.383	98%
Japan	477	300	1.345	1.142	1.822	79%
China mainland	327	278	465	361	792	81%
United States	1.018	744	1.341	930	2.359	71%
South Korea	154	131	199	111	353	69%

Spain	593	197	640	224	1.233	34%

Taiwan again surpasses the United States in requests for fraud information during the second half of 2023 and first half of 2024. Spain occupies a prominent position, also, as in the previous case, with a very low acceptance rate.

Account-based requests

Requests are made from governments to Apple related to accounts that may have been used in violation of the law and Apple's terms of use. These are iCloud or iTunes accounts and their name, address and even content in the cloud (backup, photos, contacts...).

Country	Requests 2023-H2	Requests 2024-H1	Addition	% Accepted 23H2	% Accepted 24H1
United States	10.827	12.812	23.639	90%	90%
Germany	1.925	2.655	4.580	63%	65%
Brazil	3.327	3.664	6.991	62%	71%
United Kingdom	1.414	2.550	3.964	81%	82%
Japan	259	841	1.100	55%	77%
Spain	101	156	257	34%	36%

The United States again leads by a comfortable margin in account information requests sent to Apple during late 2023 and early 2024.

Requests Related to Account Preservation

Under the context of the U.S. Electronic Communications Privacy Act (ECPA), Apple may be requested to "freeze" an account's data and hold it for 90 to 180 days. This is a preliminary step to requesting access to the account, pending legal permission to request data and to prevent the account from being deleted by the respondent.

Country	Requests 2023-H2	Preserved 2023-H2	Requests 2024-H1	Preserved 2024-H1
United States	6.610	16.682	8.170	20.513
Brazil	242	466	264	407
United Kingdom	47	54	64	99
Germany	53	61	52	50
France	2	6	41	85
Spain	1	1	0	0

The United States is by far the country with the most requests, followed by Brazil. In this respect, Spain hardly ever makes this type of request.

Emergency requests

Also under the U.S. Electronic Communications Privacy Act (ECPA), it is possible to request Apple to provide private account data if in emergency situations it is believed that this could avert a danger of death or serious harm to individuals.

Country	Requests 2023-H2	Data Prov. 2023-H2	Requests 2024-H1	Data Prov. 2024-H1	Add. Data Prov.
United Kingdom	655	597	726	658	**1.255**
United States	636	451	793	601	**1.052**
Japan	259	215	288	239	**454**
Canada	147	125	169	141	**266**
Germany	99	73	99	74	**147**
Spain	5	3	1	1	**4**

The United Kingdom continues to be the country with the highest number of such requests to Apple, with most of them being fulfilled.

Requests related to the removal of apps from the market

This data is no longer provided in Apple's report. We will update the following with more information and new developments in Apple's transparency report.

> ***To clarify:*** *In this exercise we have graphed the tables published by Apple itself. It is important to specify that requests are made in batches that may include more than one account or device. Apple, for example, counts the number of requests for device information, and in turn each request can contain an undetermined number of devices in them. Same with account requests and the number of accounts in each request. When Apple talks about the percentage of fulfilled requests, it is talking about requests, but not about specific accounts. For example: Apple receives 10 requests, with 100 devices among all the requests and then says it has satisfied 90% of the requests, we don't know how many individual devices have been provided. So, this is an exercise that can give us a rough idea of the actual number of devices provided for the example given.*

Android

New Android 16

Android released version 16 of Google's operating system on June 10 this year. Let's see what's new in terms of security.

Advance Protection 2

In May 2025, Google announced a security feature called Advanced Protection for Android devices. This feature is aimed at users with a high risk profile, such as journalists, activists or public figures or other profiles of interest, but it can be activated by anyone who wants to strengthen the security of their mobile device. In a way, it is similar to the "Security LockDown" functionality available on iOS for some years now.

Advanced Protection brings together multiple protection measures that already existed on Android, but groups them under a single system that can be activated simply starting with Android 16. When activated, it locks various system settings to prevent their accidental or intentional disabling by malware or third parties with physical access to the device. This activation also automatically adjusts certain security settings in Google apps, such as Chrome, Messages or Phone, among others.

Among the features included, there is a more rigorous local app scanning system using Google Play Protect, which improves detection of potentially harmful software. There is also an announced future feature called "Intrusion Logging", which will log relevant system events to facilitate forensic analysis in the event of a compromise. This log is stored in a protected location, inaccessible even to users with elevated privileges.

In addition, the program includes enhanced protection against potentially dangerous physical or network connections. In future updates, Advanced Protection is expected to automatically block USB connections to unauthorized devices, as well as disconnect the phone from open Wi-Fi networks that are considered risky. It will also be integrated with system functions to detect potential phone scams.

Advanced Protection is designed to evolve over time, adding new layers of security without the need for user intervention. Its activation does not require advanced technical knowledge and seeks

to balance ease of use with a higher level of protection than the standard Android default.

Fraud prevention during calls

Interesting measure adopted in Android. During calls, disabling Google Play Protect will be prevented to prevent social engineering from convincing the user to install a malicious app. The use of AI to verify whether a call is fraudulent is also being explored.

Wifi

Android 16 allows connecting to WiFi 6 or 802.11az networks with AES-256 encryption and protection against man-in-the-middle attacks.

Security evento logging

Another interesting measure. A security event log is stored in an area of the device that is difficult to access even for users with advanced privileges.

This makes it easy to obtain a log of events on devices that have been subject to attacks or infections. The idea is that it is not easily accessible but can be used by forensic teams to extract these events when performing an analysis of the device. A sort of virtual "black box".

Vulnerabilities

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

Month	CVEs	Criticals or RCE
January	34	6
February	46	1
March	41	10
April	57	4
May	46	1
June	34	0

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

In total, 258 patches in this half year (167 in the previous half year); 22 of them are considered critical (10 in the previous half year).

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

ANDROID VULNERABILITIES 2025-H1

Evolution of vulnerabilities by year

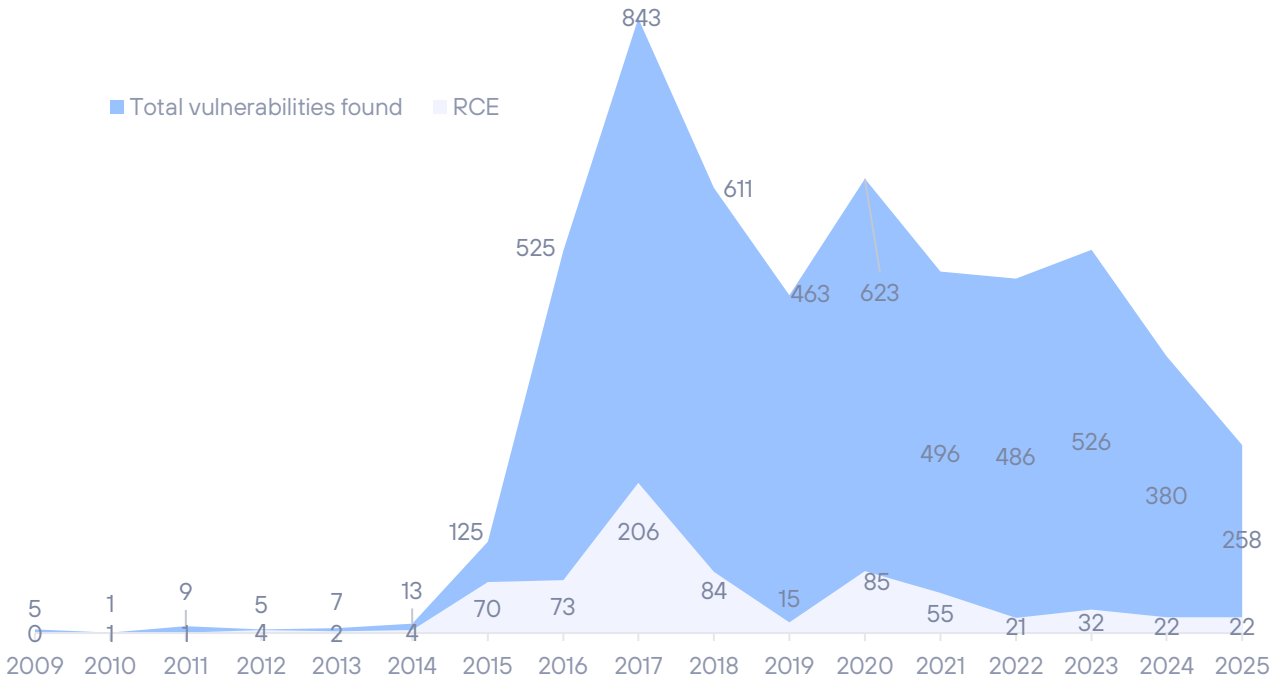

Fragmentation on Android systems

Statcounter's latest release at the time of publication, indicates that the most deployed version of Android is still Android 14, with a share of 33.31%, which is even higher than the 25.64% of the last semester. In other words, despite the release of Android 15 in October 2024, 14 continues to grow in its user base.

Android 15 has not yet had a noticeable effect on the number of Android 15 mobile devices. We will have to wait for the next edition of this report to assess the situation. Even the increase in figures of the previous version, version 14, which has even grown in penetration, is remarkable; something that we already said and happened in the previous report. In fact, we are surprised that the growth of 14 is even sharper than last semester. We will have to wait one more semester to see the share evolution of version 15.

Android versions prior to version 13 (version released in August 2022) are no longer supported for updates.

The ranking continues to get us used to seeing a worrying share of unsupported Android versions. Hence, they do not receive updates.

The ranking is in:

FRAGMENTACIÓN EN ANDROID 2025-H1

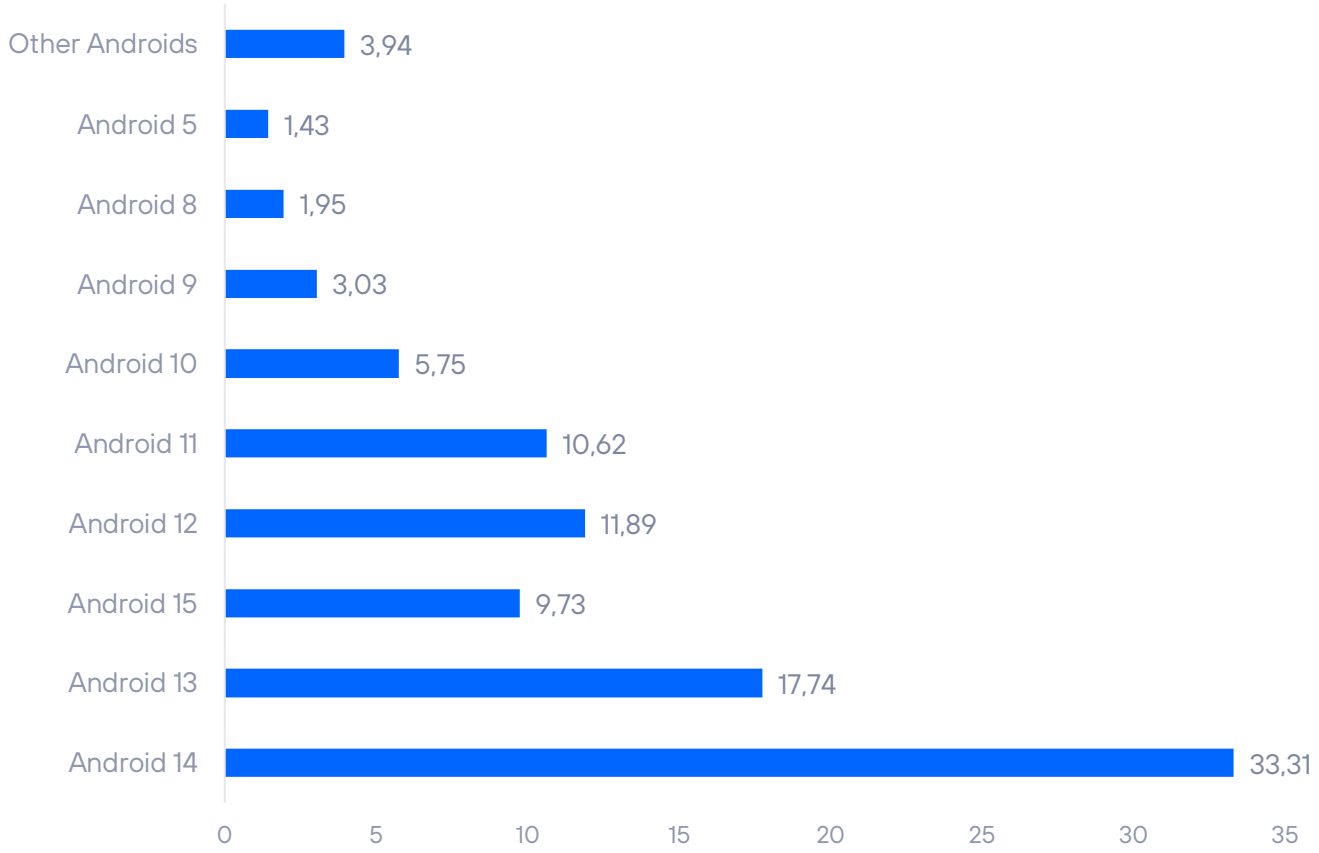

FRAGMENTATION IN ANDROID 2025-H1

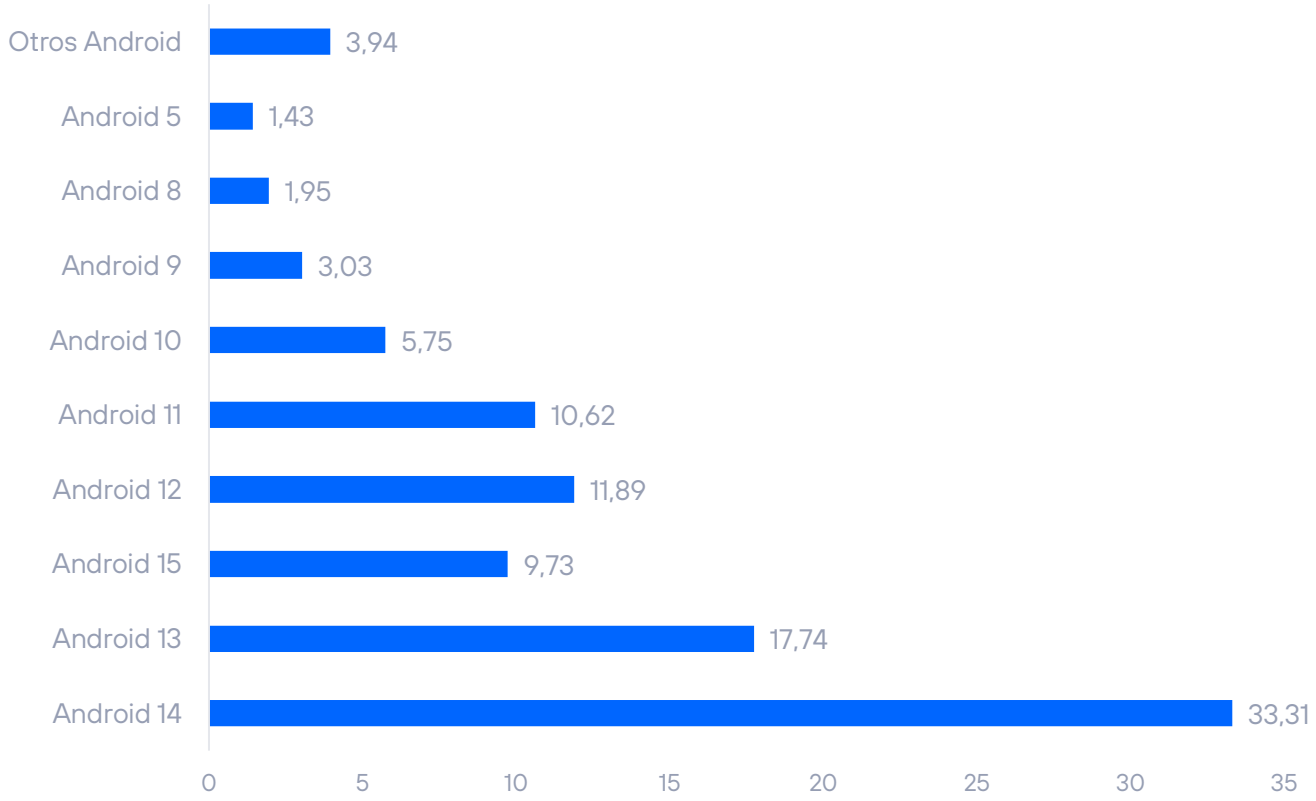

Supported versions account for 60.78%. The rest of the share, almost 40%, no longer have official update support:

We could legitimately ask ourselves: what are such old versions of Android still doing on the market? We should keep in mind that many Android terminals with a long life are still in operation in countries with less developed economies. They are cheap handsets with humble features but still serve a basic function for people in those regions.

Another large part of the terminals without official support are phones that still work well. Consider that Android 12, without support, is only a little more than three years old.

SIGNIFICANT VULNERABILITIES

In this section, we discuss some of the most notable vulnerabilities in our view from the first half of 2024.

CVE ID	TARGET	DESCRIPTION	SCORING
CVE-2025-21556	Oracle Agile PLM Framework	It allows full control of the software via HTTP, easily exploitable remotely.	9,9
CVE-2025-21298	Windows OLE	An attacker can execute code when sending a malicious email to Outlook, even in preview	9.8
CVE-2025-21311	Windows NTLMv1	Remotely exploitable with low complexity, it allows to obtain SYSTEM privileges.	9.8
CVE-2025-21524	Oracle JD Edwards EnterpriseOne Tools.	It allows code execution without authentication	9.8
CVE-2025-29966	Windows RDP (cliente)	It allows arbitrary code execution by connecting to a malicious RDP server	8.8
CVE-2025-20188	Cisco IOS XE	It allows unauthenticated remote attackers to upload arbitrary files and execute commands as root using hardcoded JWT. It requires the out-of-band image download feature to be enabled.	10
CVE-2025-20281	Cisco ISE e ISE-PIC	It allows unauthenticated remote attackers to execute commands as root by exploiting the API, completely compromising the system.	10
CVE-2025-22467	Ivanti	It allows authenticated attackers to execute arbitrary code and cause system memory corruption	9.9
CVE-2025-22457	Ivanti	It allows an attacker to take full control of the affected system. It is being actively exploited	9.8
CVE-2024-55591	Fortinet	It allows remote attackers to obtain super administrator privileges through manipulated requests. Active exploit confirmed (published on 01/14/2025)	9.6

CVE-2025-20188	Cisco IOS XE	It allows unauthenticated remote attackers to upload arbitrary files and execute commands as root via a hardcoded JWT	10
CVE-2025-1960	Schneider Electric SE	A vulnerability of initializing a resource with an insecure default value that could cause an attacker to execute unauthorized commands when a system's default password credentials have not been changed on first use. The default username is not displayed correctly in the WebHMI interface.	9.2
CVE-2024-54092	Industrial Edge Device Kit de Siemens	Although the CVE Code is dated 2024, the release date is April 2025. Successful exploitation of this vulnerability could allow an unauthenticated remote attacker to bypass authentication and impersonate a legitimate user.	9.3
CVE-2025-2567	XPort de Lantronix	An attacker could modify or disable the configuration, disrupt fuel monitoring and supply chain operations. This could create security risks in fuel storage and transportation.	9.3
CVE-2025-36535	ModBus Gateway de AutomationDirect	The embedded web server lacks authentication and access controls, allowing unrestricted remote access: configuration changes, operational interruptions or arbitrary code execution...	10
CVE-2025-40585	Energy Services	A vulnerability has been identified in Energy Services (all versions with G5DFR - a multifunctional recorder capable of recording and storing all electrical waveforms). Affected solutions contain default credentials.	9.5
CVE-2025-6029	Automotive Security Research Group (ASRG)	Use of fixed learning codes, one for locking and one for unlocking the vehicle, on the key fob transmitter of KIA's generic aftermarket smart keyless entry system.	9.4

Vulnerabilities in figures

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

VULNERABILITIES RISK

Distribution of vulnerabilities by risk

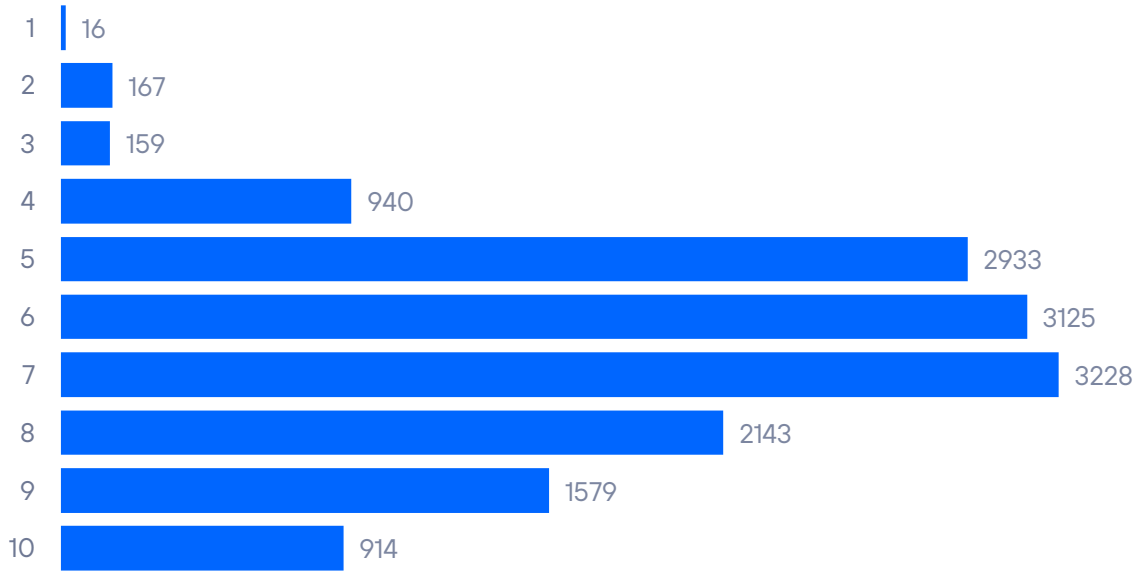

Top 25 companies with most accumulated CVEs

During the first half of 2025, Microsoft leads in the number of known vulnerabilities, followed by Adobe and Apple. In general, it is common for Microsoft, Adobe, Google, and Oracle to always be among the top in terms of the number of vulnerabilities. Linux, after being well above the rest last semester (the reason is explained in this blog post: Linux and the paradox of vulnerabilities: more reports, more security?), now occupies the fourth position.

VULNERABILITIES BY MANUFACTURER

Top 25 manufacturers by cumulative CVEs

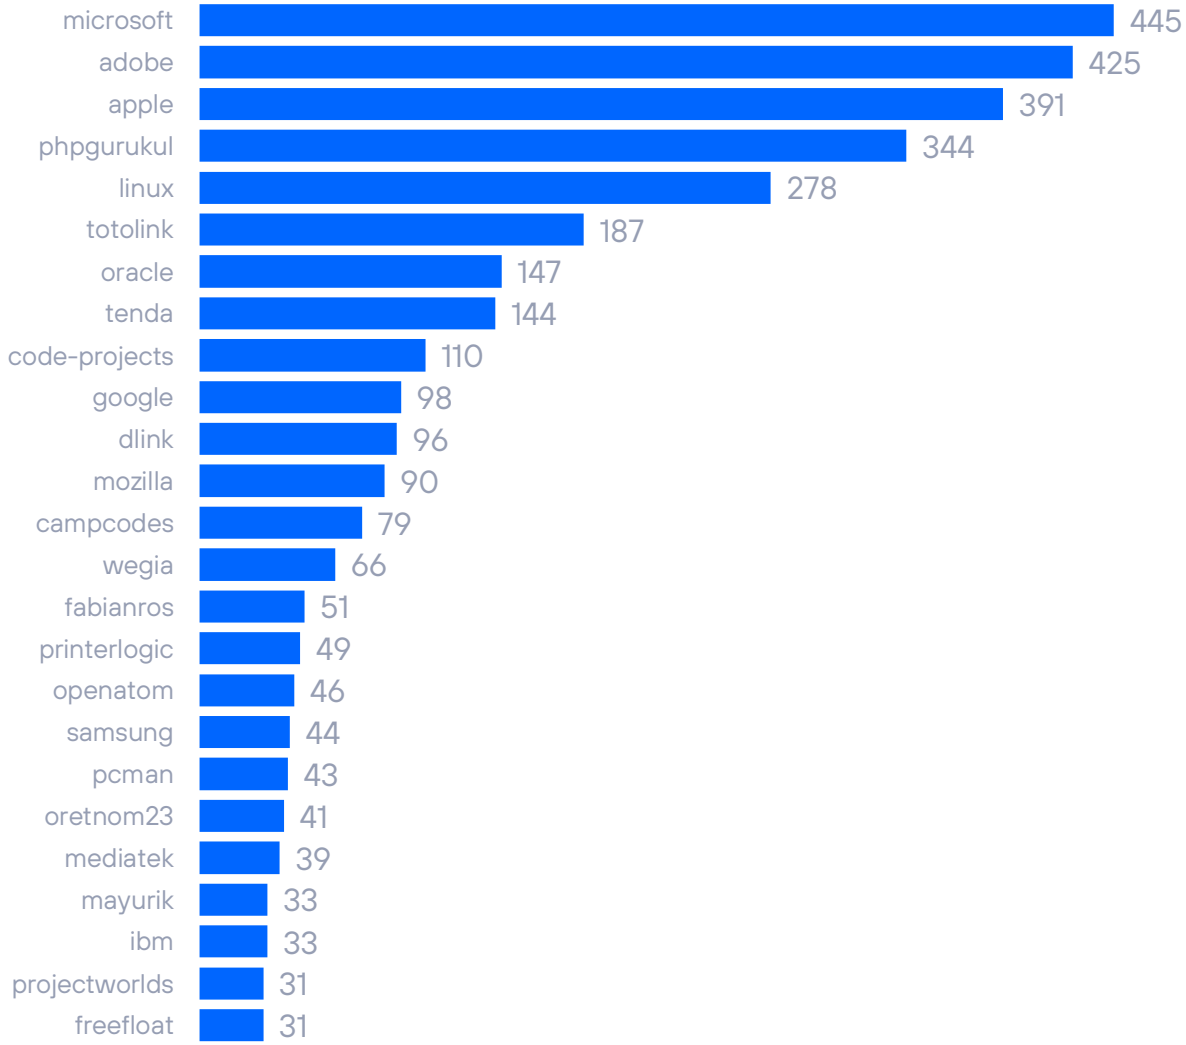

APT OPERATIONS, ORGANISED GROUPS, AND ASSOCIATED MALWARE

We reviewed the activity of the various groups attributed with responsibility for APT operations or notable campaigns.

We warn that the attribution of this type of operations, as well as the composition, origin and ideology of the organized groups is complex and, necessarily, cannot 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 in order to divert attention or harm them.

Notable APT activity detected during the first half of 2025.

APT28 – Fancy bear: The finger in every pie.

Lots of experience and lots of activity so far in 2025. They have been caught in intelligence operations using the "Clickfix" social engineering technique.

We've seen them accused by the French Ministry of Foreign Affairs of attacking at least a dozen French entities in recent years.

We have seen them running a cyber espionage campaign called "RoundPress", exploiting 0-day to access mail servers of government organizations in several countries.

We have seen them attacking defense, transportation, IT, air traffic, and maritime entities in several European countries and in the US with a common element: they were providing services to Ukraine.

And all this... in 6 months. What will the next six months bring?

More information: https://www.bleepingcomputer.com/tag/apt28/

Lazarus: Get up, walk... and succeed

Lazarus is one of the best-known APT groups in the cyber security ecosystem. They are in fact an important part of their sponsor's funding.

In this case, and despite the matter that it is a hectic six months, we are going to comment in particular on a communication from the FBI. It confirms that it was this group the peroetrator of the Bybit Exchange, where 1.5 billion dollars were stolen. Lazarus managed to breach one of the cold wallets, theoretically considered more secure, to connect it to a hot wallet, access the funds and split them into multiple wallets (and cryptoassets) to make them harder to trace.

According to the BBC, part of the North Korean missile program would have been funded through the theft of cryptoassets.

More information: https://www.ic3.gov/PSA/2025/PSA250226

Primitive Bear: primitive, but very smart

This primitive bear has evolved and has set up an infrastructure to deploy malware using tunnels using the TryCloudflare testing service. This tactic is not new and they were already detected using it in September 2024.

However, the goat is the goat and the bear is the bear in its primitive state. Researchers describe this bear as unsophisticated, as it leaves a lot of traces and performs redundant actions, such as deploying several backdoors or downloaders.

However, they make up for being unsophisticated with their work. They tend to update and change the obfuscation methods of their tools to avoid being neutralized.

More information: Hackers Leveraging Cloudflare Tunnels, DNS Fast-Flux to Hide GammaDrop Malware

Mythic Leopard: The cat that handles rats

This is about RATS.

The ones deployed by this group attacking companies in various sectors in India, particularly: Xeno RAT, Spark RAT and a new detection, CurlBack RAT.

Railroads, oil and gas... it's a big step for someone who was mainly targeting universities and research groups. Once inside, they capture all information that might be of value.

APT36, as they are also known, focuses primarily on Linux systems, while other "sister" groups focus on Windows systems.

More information: https://thehackernews.com/2025/04/pakistan-linked-hackers-expand-targets.html

OT THREAT ANALYSIS

The following information comes from the OT threat capture and analysis system, Aristeo. **Aristeo** incorporates a **network of decoys, made of real industrial hardware,** that look and behave like real industrial systems in production, but are extracting all the information about threats accessing the system. Aristeo applies relationships and intelligence to go beyond the data, being able to proactively detect campaigns, targeted or sectorized attacks, 0-day vulnerabilities, etc., thanks to the information from all the devices deployed in the different node-signature.

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 client 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

The most significant case of this semester has been (and still is) an intrusion of which we made an initial analysis in this article in March's blog post. To sum it up briefly, at the beginning of the year, an attacker managed to breach one of our decoys (oil&gas sector) through the engineering bay. A strong password and typical security measures in real industrial environments (because that's what Aristeo is all about, let's not forget) were not enough to stop them. We will not comment on the entry vector for the moment.

After this initial intrusion, the attacker changed his behavior. Let's say that he became more "novice", so we clearly saw that it was another attacker. The activity inside the machine was limited to minor and unsophisticated actions against specific targets. They even left a document with a list of targets and port scanning software in the recycle garbage can....

After the first intrusion, we did what anyone would do in our case: we let them play. We kept some maintenance, which they could see, but we didn't want to touch too much. Until, a month ago, when we did a deep cleaning and even changed the password (17 random characters with the whole pack of symbols, numbers...). Finally, they would think, "the systems guy has figured something out". We were not interested in the second attacker, although he gave us good information. We were interested in the high-capacity attacker. So, we took the opportunity to deploy even more information capture capabilities and see what happened.

A week later when the less knowledgeable attackers were unable to gain access, the professional attacker reappeared and... got back in. We were still holding on to the entry vector until we could do further checks and studies. This attacker persisted the access again, and someone on computer B (we suspect) changed the RDP password, so we "lost" access to the machine.

At that point, we disconnected the machine and went back into action. We had captured enough information and could not allow a machine to be held hostage and at the mercy of an attacker.

By the way, regarding the origins of the attackers, it is curious to see how they were IP addresses with little activity and that, in fact, they were 99% unregistered in other intelligence engines and information sources of the most consulted on the Internet. This, together with other information, led us to deduce that the attacker is an APT-Group and was not renting access to third parties. It was an "A-Team" that has very high skills and knowledge and a "B-Team" that was engaged in "retail", so to speak.

One last thing, even though we had captured enough information and locked down access to the machine, we have plans. Stay tuned.

We move on to the general statistics of the recorded information. In the first half of 2025, almost 82 million cyber security events were detected. At this point, it is worth remembering that these are complex events, and that, thanks to Aristeo 2.0, the events are now associated with each other, which turns the more than 369 million "simple" events that we have had into this figure of 82. Comparing the figures obtained with the first half of 2024, this represents an increase of almost 18%.

The distribution by country is as follows:

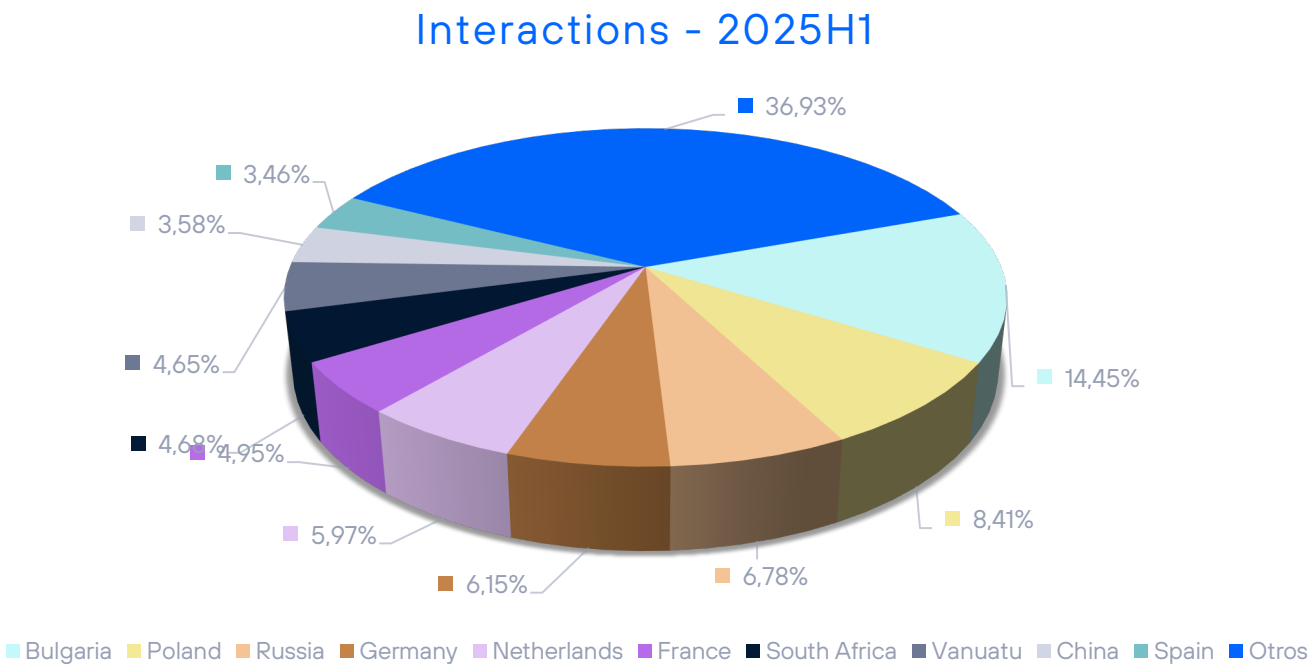

This semester, the top 10 is headed by Bulgaria, an old acquaintance regarding the origins that always visit us in large numbers. The dispersion, on the other hand, is very similar to last semester.

Now let's take a look at the top ten IP addresses with the most interaction with the Aristeo system. In this semester, the countries with the most visits to Aristeo are the ones that position their IPs in the list. In this case, one could speak of a certain stability.

As in the previous semester, the vast majority of IP addresses in this Top 10 come from central-north-eastern Europe. The difference compared to H2 2024 is that, in this case, 100% of the addresses in the Top 10 are from this region of the world (85% in the previous half-year). In addition, the Top 10 represents almost 19% of the total number of origins detected by Aristeo.

To see the magnitude, we can see the specific context: almost 2 million events out of 10 million registered belong to the Top 10 IP, and they are all from Europe. And so on up to 82 million complex events.

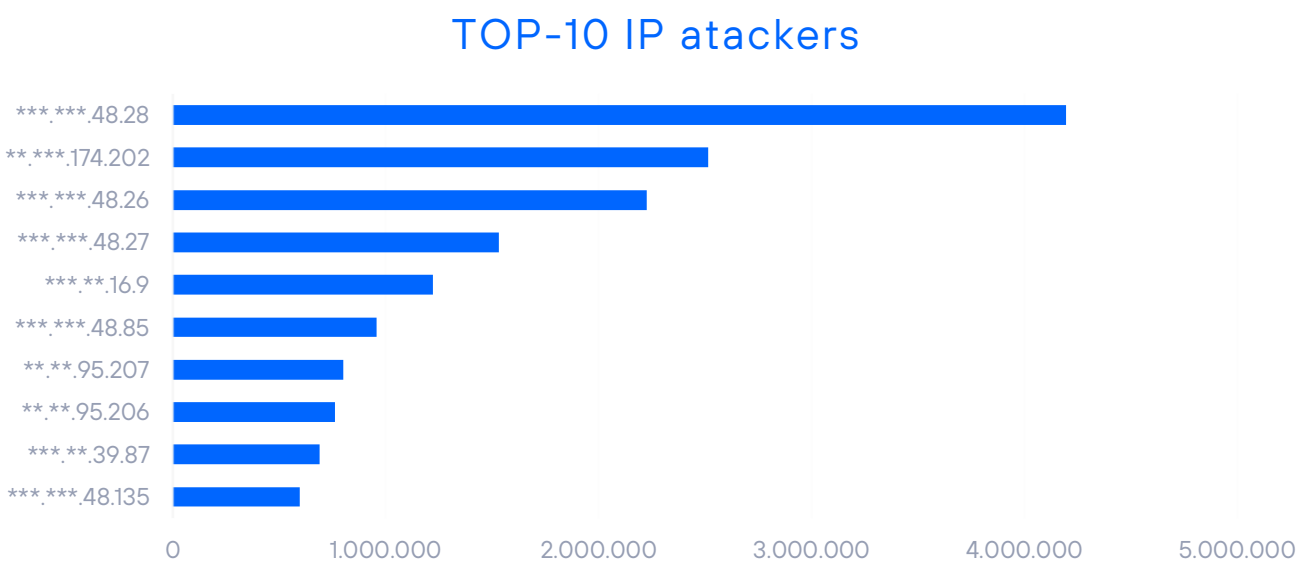

We'll find below a breakdown of the top 10 registered countries. In this semester, Spanish origins return and two new stars appear: South Africa and Vanuatu. Surely, we all know how to locate South Africa on a map (if only by the clues given by the name). But Vanuatu... the archipelago of Vanuatu, east of Australia, is made up of 12 large islands and 70 islets. It has a population similar to Lugo and Orense, and 3 Spaniards live there (according to the Spanish Ministry of Foreign Affairs).

After the geography lesson (you learn something new every day) we have not found a clear pattern why this country has slipped as eighth in the list of countries of origin with more interactions against the Aristeo network. We can only justify it through two arguments.

- In the previous half-year, we mentioned that we had opened Aristeo sites outside Europe and perhaps this relocation has to do with the interest of other attackers who take into account the region of the world in which they are active. In fact, this is something we already saw and discussed in the previous report:

Again, further offshoring Aristeo implies les focus on a specific point of the world geography and more leverage with respect to the rest.

- Another reason that may play a role is that Vanuatu is becoming fashionable. It is a country with no personal taxes and with a citizenship program in which “little” money and 3 months is granted to any foreigner who meets the conditions. This aspect, which may seem commonplace at first glance, provides 30% of the country's GDP for the amount of people who are “moving in”. These types of situations affect the economic landscape and also in the cyber landscape. This is not the first relatively small country we've seen pass through Aristeo's Top 10. Previously Belize (several times), Panama (several times) have passed through Aristeo

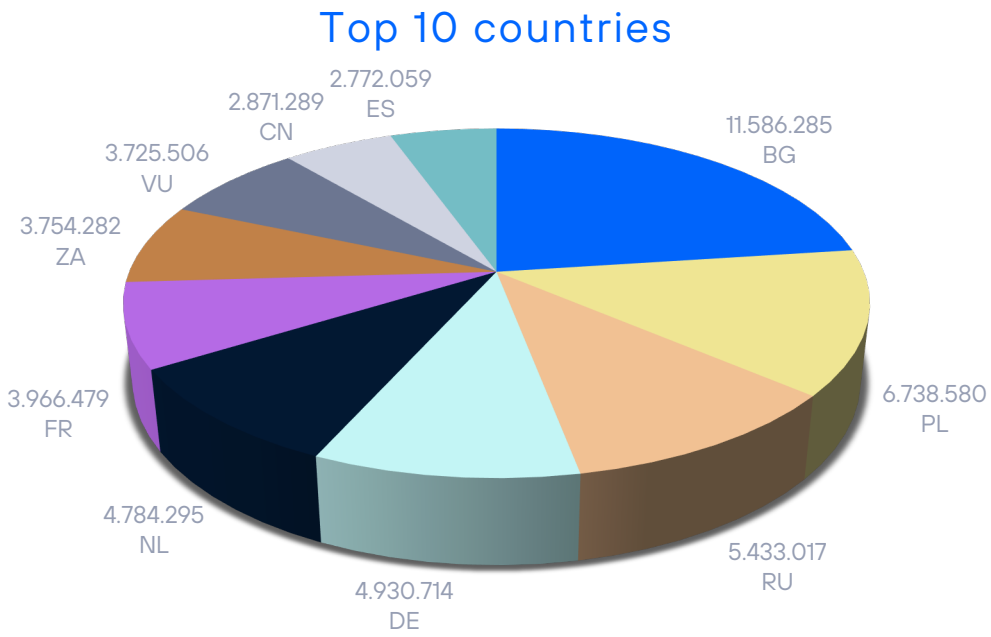

This semester, thanks to Aristeo 2.0, we have incorporated a new graphic with the TTPs most exploited by attackers.

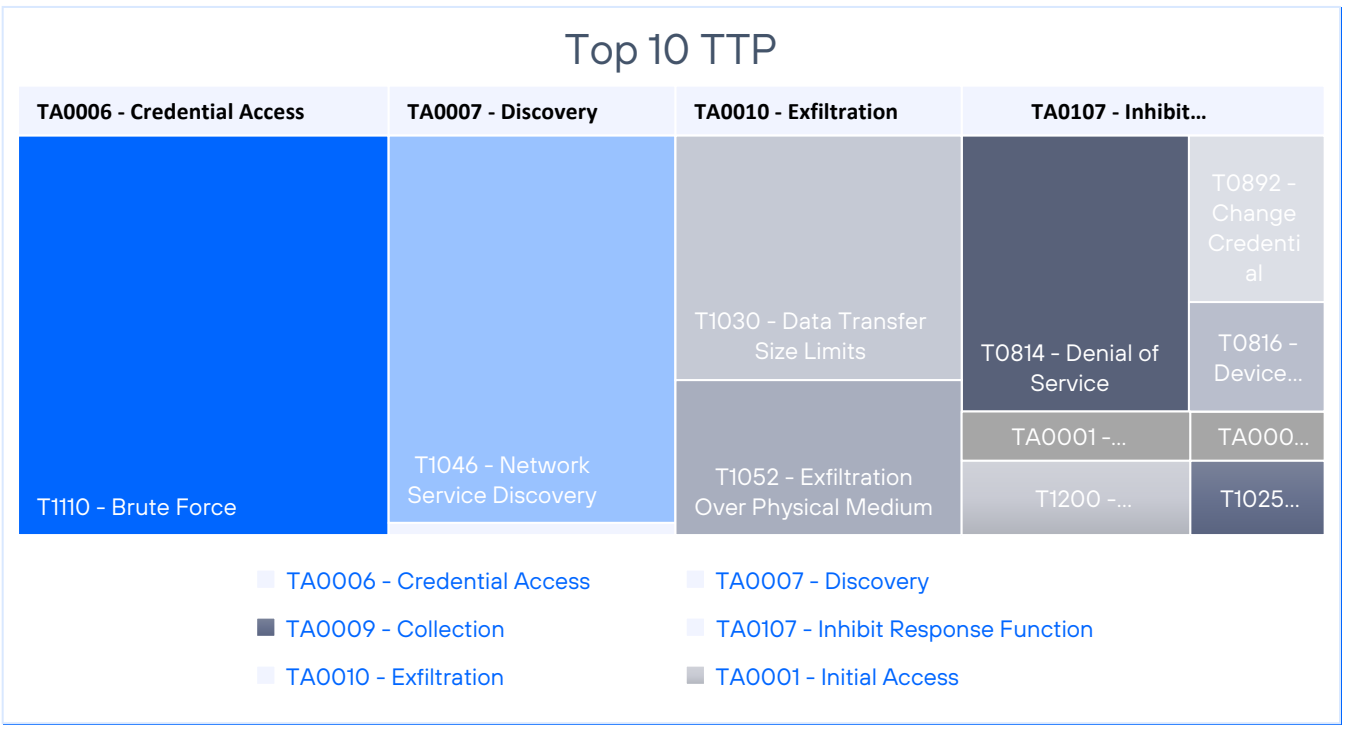

It can be seen that most of the activity is focused on initial actions, such as attempting to gain access through brute force. Other actions such as discovery, exfiltration, information gathering... are minor because Aristeo decoys are not gratis. As a good Deception environment, the decoys are properly configured and only high level attackers can access and still demonstrate their TTPs. Regarding references to physical media, Aristeo environments are usually not physically accessible to anyone, but they can be used as learning environments for your employees, or put as bait if the customer wants to detect potential insiders.

THREAT STUDY 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.

In total, for the different IOCs involved we have studied: 335.637 IP addresses, 146.329 domains and 297.615 URLs.

What kind 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 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 in the Maltiverse database, and the following results have been found:

Malware Download	185119	62,20%
Phishing	95977	32,25%
Lumma Stealer	5613	1,89%
Formbook	4559	1,53%
Clearfake	851	0,29%

FAKEUPDATES	695	0,23%
DCRat	537	0,18%
Stealc	291	0,10%
Vidar	290	0,10%
Coper	272	0,09%

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

In the last edition we had Stela Stealer as the star malware. It has disappeared, but in its place comes Lumma Stealer, hitting hard and with the same functionality: stealing data. It affects Microsoft Windows systems and mainly uses the phishing vector via e-mails.

It is followed in the ranking by FormBook, another stealer disguised as a trojan, which also has a MacOS version in addition to Windows.

The rest are distributed, as we can see, among the most widespread malware families. Infrastructure that serves as a download point, to capture orders and even to temporarily deposit stolen information. A long line of families with different DNA but with a common malicious payload.

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

We consulted Maltiverse this year to find out which domains appear most frequently in the URLs studied.

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

In the end, a URL will have a host or redirect 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 used (illegitimately, for example).

vercel.app	10281	3,45%
webflow.io	9480	3,19%

github.io	6892	2,32%
pages.dev	4940	1,66%
github.com	3448	1,16%
weebly.com	2325	0,78%
godaddysites.com	1913	0,64%
duckdns.org	1519	0,51%
r2.dev	1319	0,44%
glitch.me	1154	0,39%

As usual, the top spots belong to online services that allow you to host web content for free: vercel.app, weflow.io, github.io.

It's a common pattern: Why take a risk on private hosting or compromised servers when they offer free and anonymous hosting?

There are also domains associated with these malicious URLs that use dynamic domain resolvers: duckdns.org. In essence, 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 will continue to resolve to the new location.

As we can see, in one type of service or another, the keynote is always: free and anonymous. Two characteristics that are sought after and used zealously by cybercriminals.

Which countries are the IP addresses on which malicious activity has been detected?

Before answering the question, it should be clarified that just because a country appears in this ranking does not mean that there is 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 greater fraudulent use. A server can be hosted in a country and the criminal organization that uses it can come from another nationality.

India	68265	20,34%
United States of America	52462	15,63%

China	30877	9,20%
Singapur	17491	5,21%
Vietnam	16084	4,79%
Russia	13246	3,95%
Germany	11178	3,33%
Brazil	8914	2,66%
United Kingdom	7091	2,11%
Pakistan	6855	2,04%

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?

We could conclude that certain governments request access to data "too often," but we could also argue that justice may operate more quickly there, or that there may be more fraud in these locations the interpretation is free. Below are some conclusions based on our analysis:

Suspicious host	158322	47,17%
Mail Spammer	143953	42,89%
HTTP Spammer	117679	35,06%
Malicious host	81556	24,30%
Bruteforce	59340	17,68%
Hacking	57975	17,27%

Port Scanner	55216	16,45%
SSH Attacker	54335	16,19%
Proxy	51675	15,40%
HTTP Attacker	42955	12,80%

Later, when a label is added with the explanation 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 more than 80% of the tags. Remember that tags overlap, so the same IP can contain several of them. A generalist of "suspicious" and "HTTP Spammer", for example, or even that the same IP is 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 feature of the Internet for decades (as was Rlogin or telnet in its early days). Almost 16.19% of IP addresses have been observed performing attacks on SSH (mostly dictionary attacks on the login).

Likewise, "Bruteforce" refers to the continuous attempt to perform brute-force authentication (actually, again: dictionaries of common usernames and passwords). This category adds up to almost 17.68%.

In another subcategory, indiscriminate scans, we find: Port and Host scanner. IP addresses that have been detected by performing mass scans of entire ranges or multiple ports on specific hosts. That is, horizontal scans looking for certain ports or vertical scans (in depth) on a group of hosts.

The "Proxy" category with almost 15.4% are systems that, either deliberately or unsuspectedly, serve as a gateway or hop to other machines to hide the origin of certain attacks or unauthorized access.

Overall, we find the "hacking" category with 17.27% closing the ranking. These are nodes that have been observed performing attacks in general, either trying to find SQL vulnerabilities or launching exploits. Often, these are vulnerability scanners used indiscriminately and, of course, without authorization.

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

As we know, a domain resolves to an IP address. In the world of cybercrime, domains are of paramount importance because they allow you to make use of it and change the IP address if the currently active server ceases its malicious activity.

A domain is composed of several levels. If you look at them, they are stretches of strings separated by dots. If we get these groups from right to left they form a hierarchy. The rightmost one is the highest level domain.

Hence, we can group the domains categorized as malicious by their highest level domain. The result of the top 10 is this:

xyz	43193	29,52%
com	34609	23,65%
io	9768	6,68%
top	5170	3,53%
org	4705	3,22%
app	4188	2,86%
net	3532	2,41%
dev	2896	1,98%
shop	2450	1,67%
gg	1831	1,25%

“xyz” tops our TLD ranking this semester. Although it disputes in alternation the leadership with “com”, it rises strongly and snatches the baton. “xyz” was born in 2014 and was pushed by Google for matching its parent: “abc.xyz”.

Why is “xyz” the “favorite” of malware creators? Because of its competitive prices: from $0.99 per year it is a more than attractive figure to use domains of this TLD.

Regarding “.app” it is especially curious as it is a TLD for which Google paid more than $25 million to ICANN in February 2015 to take control of it. Moreover, it is a TLD for which HTTPS traffic is mandatory.

“gg” which sneaks into our top 10 is a geographic TLD belonging to Guernsey. An island in the English Channel, belonging to the United Kingdom. It is a TLD lately associated with video games and e-sports sites.

What malicious categorization do the studied domains possess?

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

Finally, let us see how the top 10 of these domains have been classified over the last six months.

Phishing	44483	30,40%
Metastealer	39659	27,10%
Formbook	14389	9,83%
Lumma Stealer	8623	5,89%
Virut	6708	4,58%
Malware Download	5077	3,47%
Joker	2394	1,64%
Xworm	1734	1,19%
BumbleBee	1608	1,10%
Bankbot	1406	0,96%

As we have already mentioned, there is a very close relationship between domains and URLs and 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

The truth about the 320 seconds to hack Bitcoin: a technical analysis

Analysis of an intrusion on the Aristeo platform as a demo of Its predictive capabilities

Linux and the vulnerability paradox: More reports, more security?

ARTIFICIAL INTELLIGENCE

https://telefonicatech.com/en/blog/the-incredible-inner-world-of-llms-i

The incredible inner world of LLMs (II)

La tokenización y el caballero andante Don Quijote