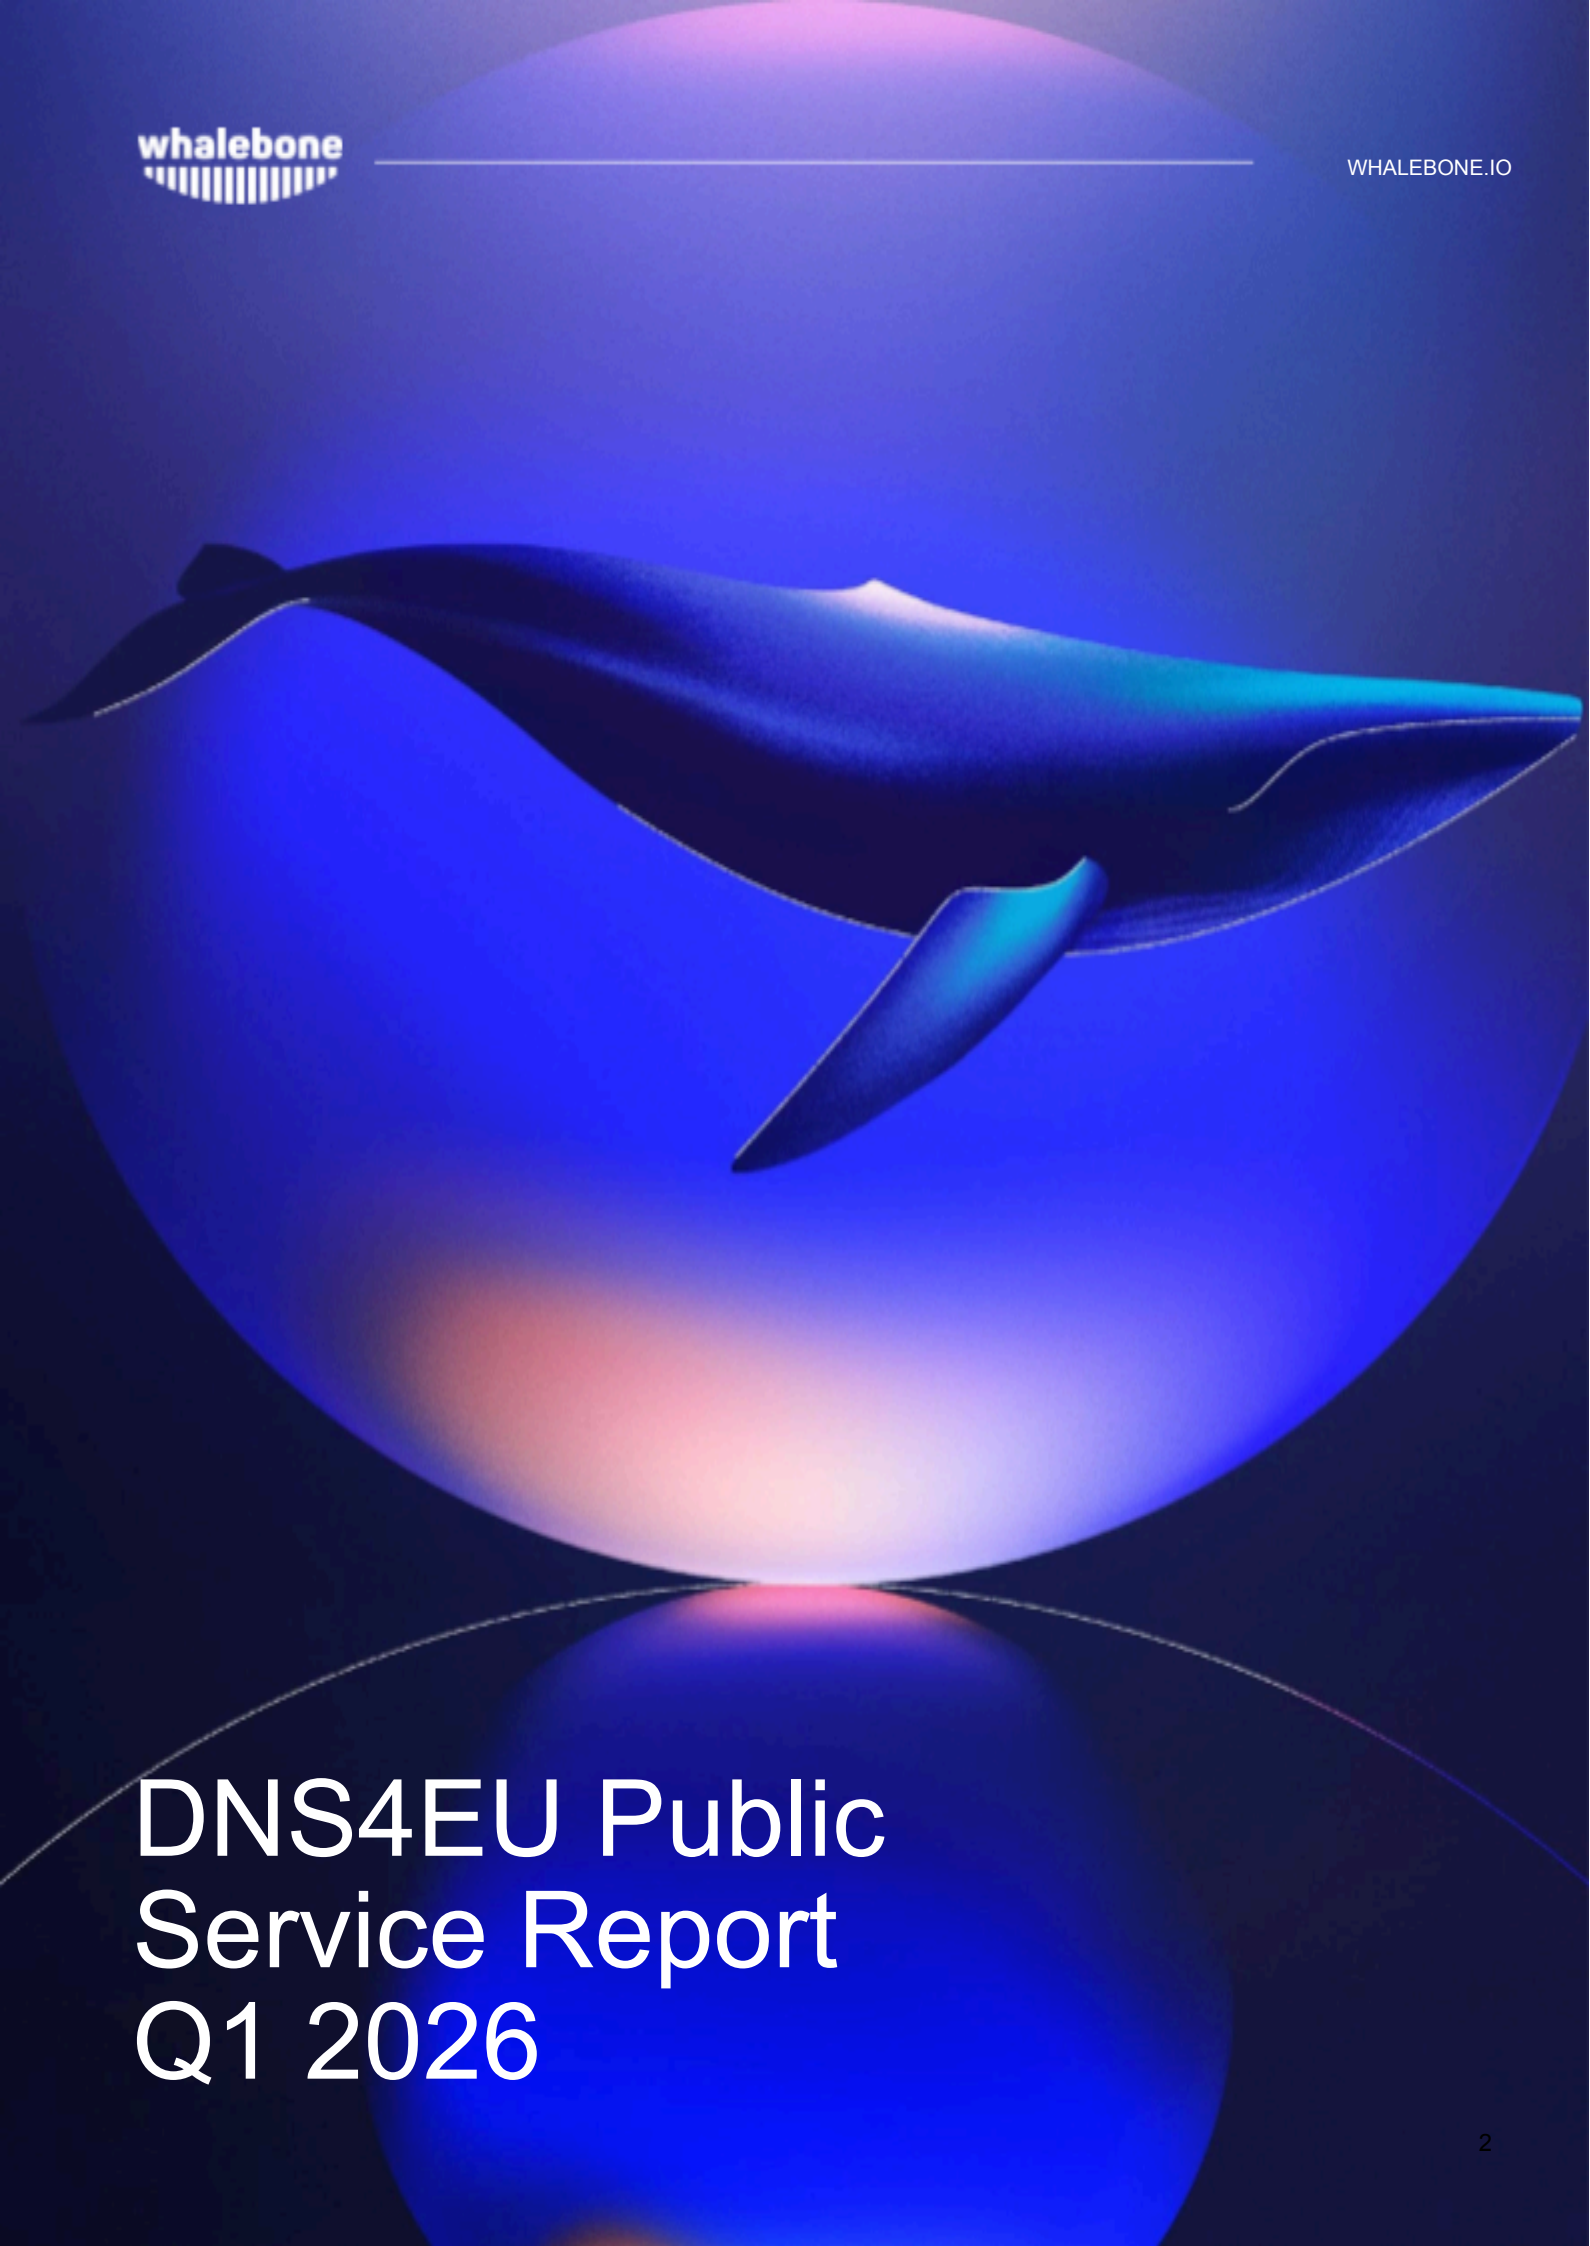

DNS4EU Public Service Report Q1 2026

Table of Contents

1. Executive Summary

During **Q1 2026**, the DNS4EU Public Service achieved significant growth in query volume and threat mitigation, validating its performance and reliability as a scalable tech export.

Key outcomes from the quarter include:

- Traffic Analysis: The infrastructure processed several billion queries and maintaining low-latency performance.
- Threat Analysis: The service intercepted and neutralized millions of security threats at the network layer. Telemetry from Q1 identified high volumes of malicious activity, specifically targeting automated phishing, malware distribution networks, spam infrastructure, and advanced Command and Control (C&C) communication channels.

2. DNS4EU Public Service Overview

The DNS4EU Public Service is a free, privacy-first European recursive DNS resolver operated by the Czech cybersecurity company Whalebone. Launched in June 2025, the service is designed **exclusively for individual citizens**, with adoption steadily growing among both the public and cybersecurity experts. It offers users five distinct configuration options - ranging from standard security filtering to specialized configurations for child protection and ad blocking - ensuring robust, tailored network safety directly at the DNS layer.

3. Threat Intelligence: Covered Attack Vectors

Our threat intelligence capabilities have significantly matured, allowing DNS4EU to preemptively shield users from a wide array of cyber threats. Over the past reporting period, we successfully identified and blocked:

- **Malware:** any software designed to harm, exploit, or gain unauthorized access to devices, networks, or data. Common types include viruses, worms, trojans, ransomware, spyware, and adware. Malware can steal information, disrupt operations, or extort users. It often spreads through phishing, infected websites, or software vulnerabilities.
- **Phishing:** a type of cyberattack where attackers trick people into sharing sensitive information—such as passwords, credit card numbers, or personal details—by posing as a trustworthy source. These attacks often come through fake emails, messages, or websites that appear legitimate. The goal is to steal data or gain access to accounts and systems.
- **Blacklist/Deny List:** consists of websites, domains, or IP addresses that have been identified as malicious or unsafe and added to a security blocklist (deny list). These sites may host malware, phishing pages, or other harmful content. Security systems use Deny Lists to prevent users from reaching such destinations, helping to reduce the risk of infection or data theft.
- **Command & Control:** refers to a method attackers use to remotely control compromised devices or networks. After malware infects a system, it connects to an external server—the "command and control" center—to receive instructions or send stolen data. This connection allows attackers to move laterally, exfiltrate information, or deploy further attacks, often without being immediately detected.
- **Spam:** involves the mass distribution of unsolicited messages, usually by email, often used to spread malware, phishing links, or scams. While some spam is merely unwanted advertising, others pose serious risks by tricking users into clicking harmful links or downloading infected

attachments. Spam can also overwhelm inboxes and reduce productivity.

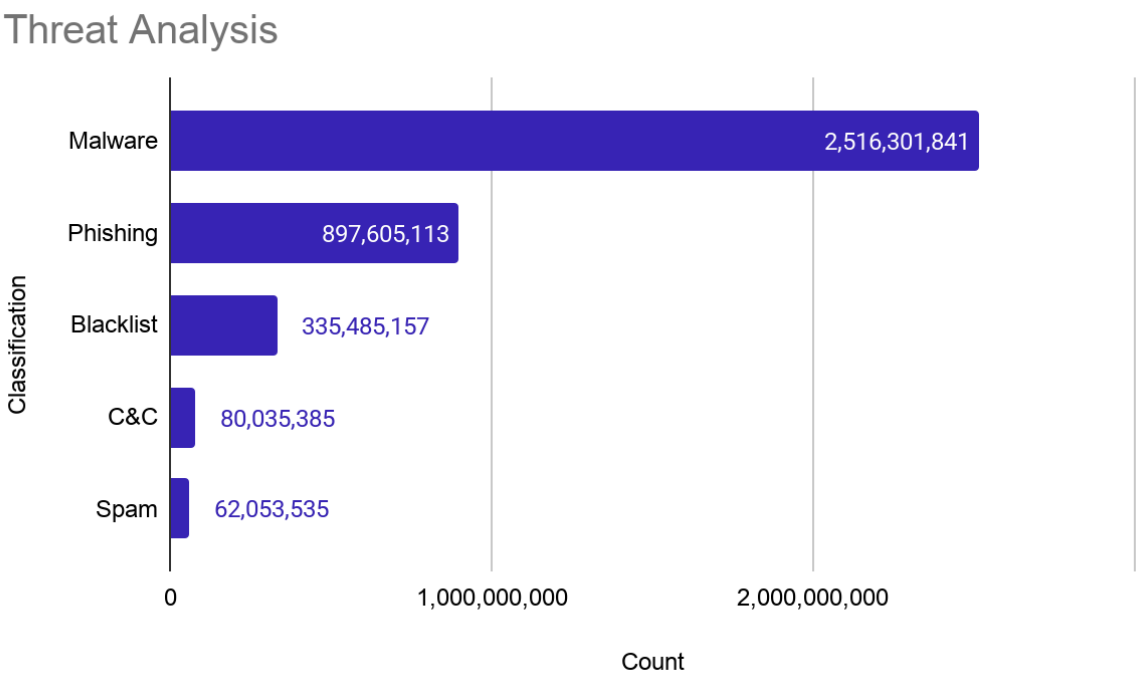

In addition to our standard security feeds, the DNS4EU Public Service operates in strict compliance with applicable legal and regulatory frameworks. Over the past reporting period, this included executing a number of binding court orders to implement specific domain blocks on our blocklists. Notably, these legal mandates extended enforcement to certain traffic within our otherwise 'unfiltered' category. In alignment with our commitment to accountability and open governance, DNS4EU maintains full public transparency regarding these regulatory actions; a complete, up-to-date registry of all court-ordered blocking activities is actively published and maintained on our website.

4. Case Studies: Concrete Campaigns Stopped

Campaign 1: Whatsapp Hijack

Throughout February 2026, Whalebone Threat Intelligence was monitoring a widespread **phishing campaign targeting WhatsApp users** in more than **15 countries world-wide**, including Czechia, Serbia, Spain, and Brazil. During February alone, Whalebone identified and blocked **over 200 domains** associated with the current phishing campaign.

The attackers used a previously hijacked account belonging to someone the victim knows and **exploited the existing trust** to direct victims toward a fraudulent voting website. Under the guise of a simple "authentication" for a contest, **the platform attempted to take complete control over the victim's WhatsApp account.**

Stolen accounts were then used for further dissemination of phishing messages to the victim's contact list and to carry out social engineering attacks by **requesting short-term loans** from the victim's close contacts.

The Bait

“Hi! Please, can you vote for Alexandra in this poll? She is my friend's daughter.
The main prize is a free scholarship for next year and it is really very important for her.
Thank you very much!”

Messages as such have been spreading across Czech WhatsApp users in the past weeks. The message arrives from a known contact, which has already been hijacked upstream and requests a seemingly benign action – to vote for a specific contestant in a children's dance competition. Its sole purpose is to direct the recipient to the attached link.

The Catch

Upon navigating to the malicious link, the victim is directed to a landing page which displays images of two child participants in gymnastic poses. Selecting the target "competitor," prompts the victim with a popup window to log-in and “verify” the vote through WhatsApp.

Once the victim inputs their phone number, they are presented with a pairing code and instructed to enter it into the "Linked Devices" section of their WhatsApp settings. As the victim inputs the pairing code into their WhatsApp settings, the attacker’s WhatsApp instance is authorized as a linked device, obtaining full access to the victim's contacts, chat history, and the ability to initiate new chat sessions.

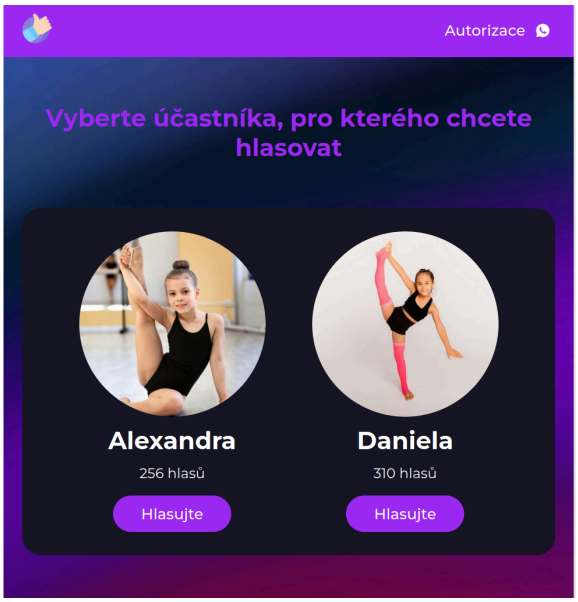

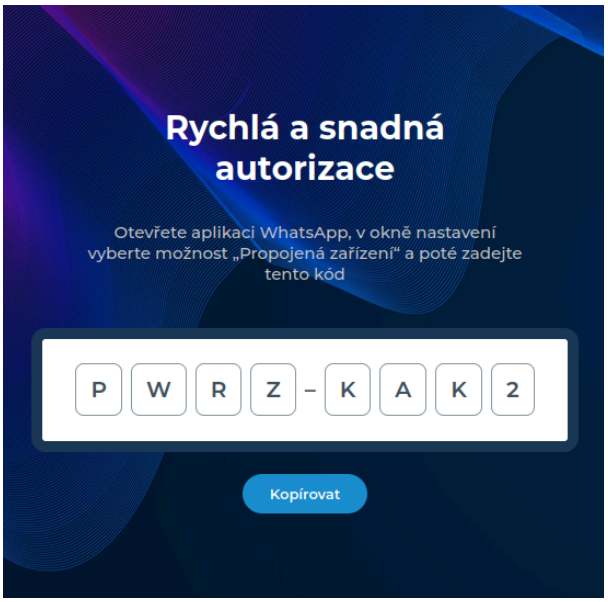

Infrastructure & Technical Observations

The adversary utilized **randomized domain names associated with dance or competitions**, such as denceegimcz[.]life, baletidancecz[.]run, or stardencer[.]fun, often incorporating typographical errors and the specific country code of the targeted region. The domains are **registered through multiple registrars and the entire operational infrastructure is masked behind Cloudflare** proxy services to obfuscate the true origin of the servers.

As of mid-February 2026, at least **15 localized variations** of the scam have been identified, including Czechia, Slovakia, Slovenia, Serbia, Romania, Bulgaria, Poland, and Croatia, as well as Western European and international targets like Spain, Italy, Greece, Brazil, Mexico, or Ireland. Based on the volume of unique domains allocated to each region, the threat actors appear to be prioritizing **Central and Southern European European countries**.

The campaign relied on a communication channel established via a **customized Socket.IO framework**, orchestrated by a specialized JavaScript file internally named number.js.

The client-side script initializes a connection to a Cloudflare-masked command and control (C2) server. Once the victim submits their phone number, a server-side request to WhatsApp's official pairing service is triggered, which is then sent back to the victim's browser to display the 6-digit linking code. Once the hijack is complete, the attacker's C2 server redirects the victim back to the main site with a popup message confirming that their vote was successfully counted.

The infrastructure furthermore exhibits specific artifacts that facilitate the attribution of the threat actor's origin. Multiple analyzed index.html files contain the <html lang="ru"> tag, suggesting the use of a Russian-language development environment or a phishing kit developed by Russian-speaking authors. This, combined with the geographical focus, indicates a **likely Russian-speaking threat actor or group.**

Campaign 2: Courier-themed Phishing

Sectors targeted by cybercriminals remain largely unchanged – services that make it easy to exploit unsuspecting victims, such as delivery providers, telecommunications companies or financial services. Their tools however allow them to launch new phishing domains faster and in greater numbers, create more convincing sites and target more potential victims.

Since early 2026, Whalebone Threat Intelligence has been intercepting a large number of phishing sites impersonating popular delivery services and targeting predominantly European consumers. Similar to campaigns observed in previous years, these phishing sites were distributed via smishing messages and aimed to steal payment-related information.

Thousands of domains observed shared the same pattern: a scrambled domain name incorporating the impersonated courier brand, phishing content delivered on a secondary URL path, and nearly identical message content repackaged for a specific delivery company and country, centered on rescheduling parcel delivery.

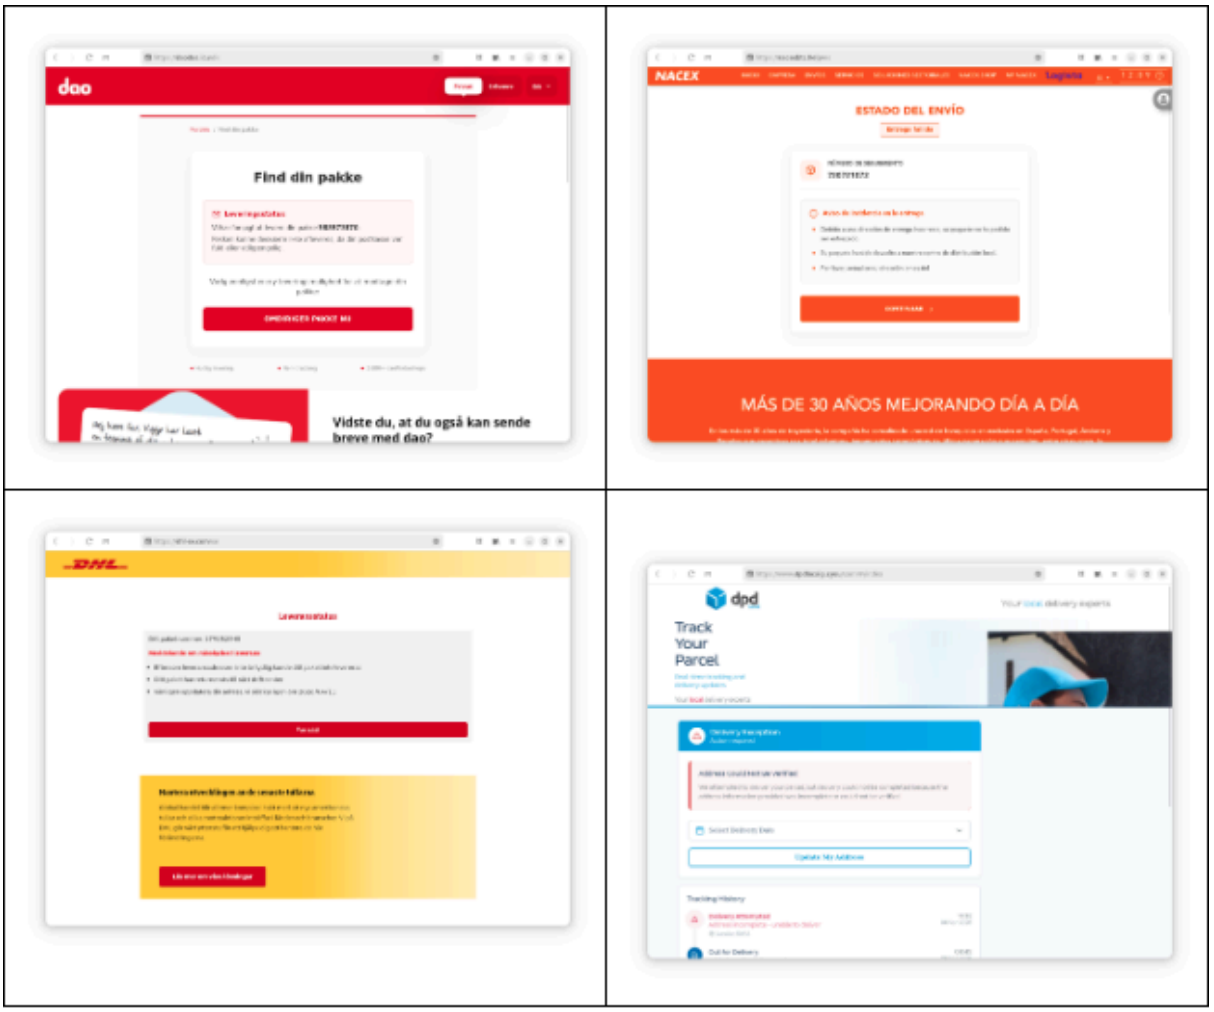

Phishing Flow

The phishing starts by a familiar vector - a smishing message, often shared via iMessage or RCS (Rich Communication Services), notifying the recipients of a fraudulent “failed delivery,” and prompting them to “reschedule.”

Clicking the link in the smishing message leads the recipient to a website closely mimicking the targeted delivery company. Upon proceeding with “rescheduling,” the victim is guided through steps that include extensive collection of personal information, such as phone number, full name, and full address. This information is immediately sent field by field to the attacker - even an incomplete or unsent form is therefore recorded by the operators. At this stage, the victim is also prompted to select a new delivery date for the supposedly missed parcel.

In the final stage of the phishing flow, the victim is prompted to pay a small “re-delivery fee,” by inputting payment card details. Once submitted, the payment data is recorded. Based on the authentication flows observed in the code, the user may then be prompted to complete MFA (multi factor authentication).

Infrastructure

While the flow may seem rather generic, what was more interesting was the infrastructure of the operation. As of March 2026, we have observed thousands of domains impersonating predominantly

European courier services and business entities. Given the observed extent, we assumed that we were looking at a large-scale, automated operation.

Many of these domains were registered on AS132203 (Tencent Building, Kejizhongyi Avenue), with entire IP addresses apparently dedicated to this phishing activity. Several IP addresses observed have hundreds of domains assigned to them, all of which were associated with current or earlier waves of courier-themed phishing. Some domains were also observed behind Cloudflare or hosted by smaller ISPs.

The domain names are intentionally chosen to resemble the targeted courier service. For instance, in the case of DAO Denmark, domains often incorporate both “dao” and “as,” or slight variations such as “daoas[.]icu,” “daoaso[.]sbs,” or “dao-pakke[.]shop.” In other cases, the courier brand appears in a subdomain, while the registered domain itself begins with a TLD-like string, as in dao.as-nfd[.]top. In most cases, generic, low trust TLDs are used, such as cyou, icu or sbs. Likely in order to hinder analysis, phishing content is furthermore delivered not on the holding page, but on secondary paths, such as /[country code], /com or /pay. Similar naming patterns can be observed across the vast majority of the observed phishing domains.

We have also observed that the analyzed websites had a feature that allowed the attackers to display a dynamic, country-specific validation page in order to capture MFA (multi factor authentication): in the example of sites impersonating DAO, this included language-specific messages and local authentication methods, such as email/phone OTP, MitID, Nordea Bank or Danske Bank custom MFA verification pages. This phishing modus operandi is consistent with public reporting on the Darcula-family / Magic Cat Phishing-as-a-Service (PhaaS) platform – a pre-packaged set of software tools, scripts, and website templates that allows operators, even those with limited technical skills, to quickly deploy and localize phishing websites at scale.

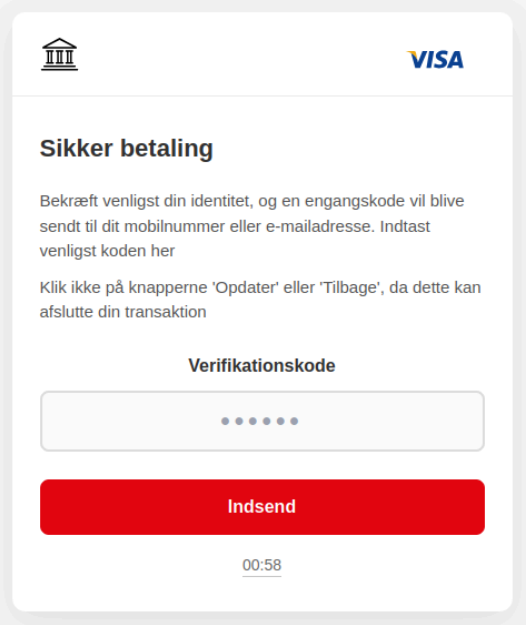

It remains unclear whether this campaign is operated by a single organized criminal group, or by multiple groups using the same PhaaS kit simultaneously, due to minor differences in infrastructure choices and content delivery structure, including the ASN used, phishing domain naming patterns, brand name placement (for example in a subdomain or second-level domain), and secondary URL path structure.

Campaign 3: JS Keylogger

During our day-to-day operations, we regularly uncover threats that are less common and often significantly harder to detect. There may be no obvious sign of malicious activity or misuse: no file is downloaded, no phishing page is displayed, and no forwarding chain is initiated. In the case below, a legitimate website was compromised with a malicious script that recorded every keystroke entered by the user and sent the captured data to a command-and-control (C2) server.

Initially, the WordPress (WP) website we investigated showed no obvious signs of malicious behavior. Yet, examining communication initiated by the domain, one particular item stood out. Among the many

`.js` files commonly found in the WP ecosystem, a file named `sc-top.js` was identified. The primary feature of this script was the fact that it was being loaded from the IP address `94.156.79.8`, which had no apparent connection to the domain itself.

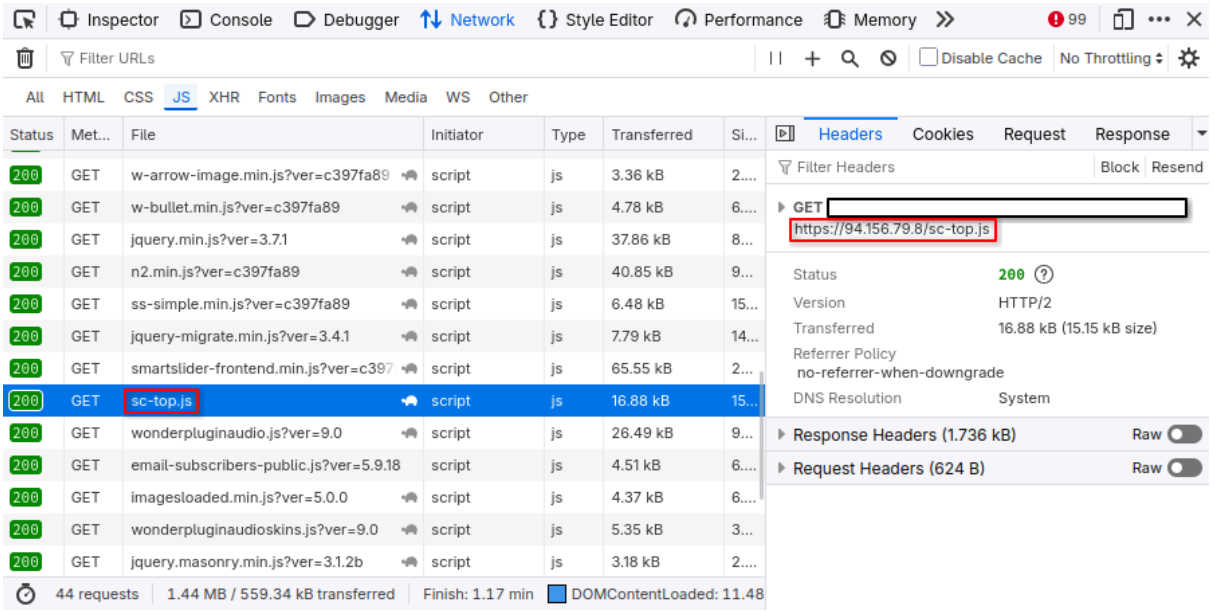

Analysis of the page source revealed that this external script was being loaded via an embedded obfuscated script. In this case, the purpose of the obfuscation was clearly to conceal the presence of malicious code.

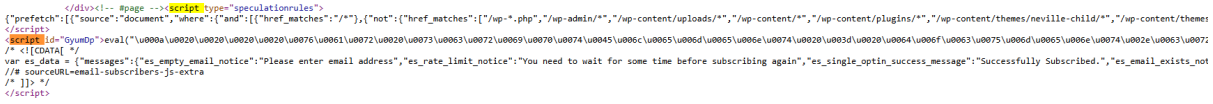

The `sc-top.js` script itself injects into the domain and records text entered into web forms on the infected site, effectively functioning as a keylogger. The script treats each individual keystroke or each piece of text pasted from the clipboard as a separate record. It then forwards the captured information to a PHP script called `pinche.php`, located on the domain `hostpdf[.]co`. The data is sent via a POST request containing three fields: `uid`, a unique identifier used to identify a specific user, `i_name`, and `b`, the actual text entered by the user, encoded using Base64.

What does this mean for a visitor to an infected website? Any information they enter on the site may be captured and sent to the attacker. Depending on the nature of the data, it may then be abused accordingly - credit card details may be used for unauthorized payments, submitted passwords may be

exploited for account theft, and personal information may be sold on the dark web. As for the specific website on which Whalebone Threat Intelligence uncovered the compromise, its owner was informed of the situation and took remedial action – the site no longer poses a threat to visitors.

This particular attack forms part of a broader campaign uncovered in June 2024 by Wordfence. As part of the campaign, five WordPress plugins across various versions were infected: Social Warfare, Blaze Widget, Wrapper Link Element, Contact Form 7 Multi-Step Addon, and Simply Show Hooks. The implanted code enabled the attackers to create a new user account with administrator privileges in the WordPress administration panel, effectively taking control of the website. The code was then used to inject malicious scripts of the attacker's choice.

Although the vulnerabilities in the targeted plugins have since been addressed, the example above clearly shows that not all websites using them have kept up with security updates. As a result, many sites may remain vulnerable to exploitation or, worse, may already be compromised, as seen in the example above.

5. Traffic Insights & Data Analytics

Our traffic analysis reveals a steady growth in adoption and a highly clear picture of the regional threat landscape.

Metric	**Reporting Period Value**
Traffic analysis	4 billion requests
Average Latency	13.19 ms
Peak Throughput	~70k QPS
Geographical distribution (most utilized locations)	Germany Netherlands France Czech Republic Spain

DNS4EU Public Service offers 5 resolver variants. In the graph below, you can see what is the protocol distribution between encrypted DNS over HTTPS (DoH), DNS over TLS (DoT), UDP and TCP.

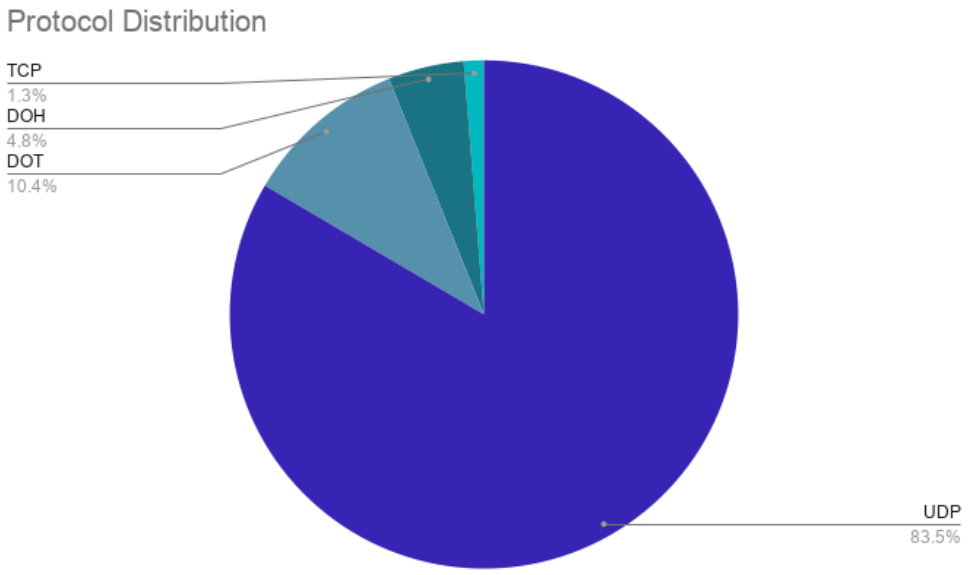

6. Conclusion

The performance and telemetry data from Q1 2026 demonstrate that the DNS4EU Public Service has successfully transitioned from an initial launch phase into a mature, stable infrastructure capable of operating at a multi-billion query scale. By successfully intercepting millions of network threats - ranging from routine malware and spam to sophisticated C&C channels - the service has validated its security efficacy under real-world conditions.

Moving forward, maintaining this level of resilience requires clear alignment of use cases: the **DNS4EU Public Service is strictly engineered and dedicated to protecting individual citizens**. Because corporate and public sector networks face vastly different threat profiles and regulatory requirements, standard consumer protection is not enough. If you are an enterprise, institution, or organization looking to secure your infrastructure with tailored, industrial-grade threat intelligence, we invite you to evaluate your defenses by trying a dedicated solution - **Whalebone Immunity**.

Contributors:
Tadeáš Hájek
Roman Šilhan
Magdalena Krucká

immunity@whalebone.io
contact@joindns4.eu

Whalebone, s.r.o., Jezuitská 14/13
602 00 Brno, Czech Republic