Event Report

CloudSEK

# THE 5060 SIEGE

## Industrialized Attacks Against the SIP Telephony Ecosystem

**Category:** Honeypot Threat Study · Adversary Intelligence
**Threat Vector:** SIP / VoIP — UDP+TCP / 5060
**Reporting Period:** 2026-05-04 to 2026-05-22 (18 days)
**Telemetry:** 15,183,358 events · 323 source IPs
**Classification:** TLP: RED · Confidential

## Executive Summary

**Over an 18-day window, a single Internet-facing SIP service recorded 15,183,358 telemetry events — roughly 3,787,791 distinct SIP requests — from 323 unique source addresses. The traffic was not random noise. It was a sustained, automated assault on the telephony layer, dominated by industrial-scale credential theft and a parallel stream of international toll-fraud call attempts.**

Two activities account for almost all of the intent-bearing traffic. The first is **SIP registration brute force at industrial scale**: 1,869,521 authentication attempts carrying full Digest credentials, spread across 29,433 distinct extension identities. The second is **toll fraud**: 89,465 call-setup (INVITE) attempts, overwhelmingly aimed at United Kingdom revenue-share number ranges and executed through mechanical dial-plan probing. A smaller but strategically important slice of traffic replays authentication challenges harvested from other, real PBX systems — evidence that this sensor sits inside a much larger credential-harvesting economy.

Because the captured authentication material is complete, the actual plaintext password behind **96.09% of all 1,869,521 credential attempts** could be determined. The result is a recovered dictionary of 277,632 unique passwords and 1,499,846 unique extension/password pairs — a direct, unobstructed view into the wordlist an active VoIP-fraud operation is spraying across the Internet today.

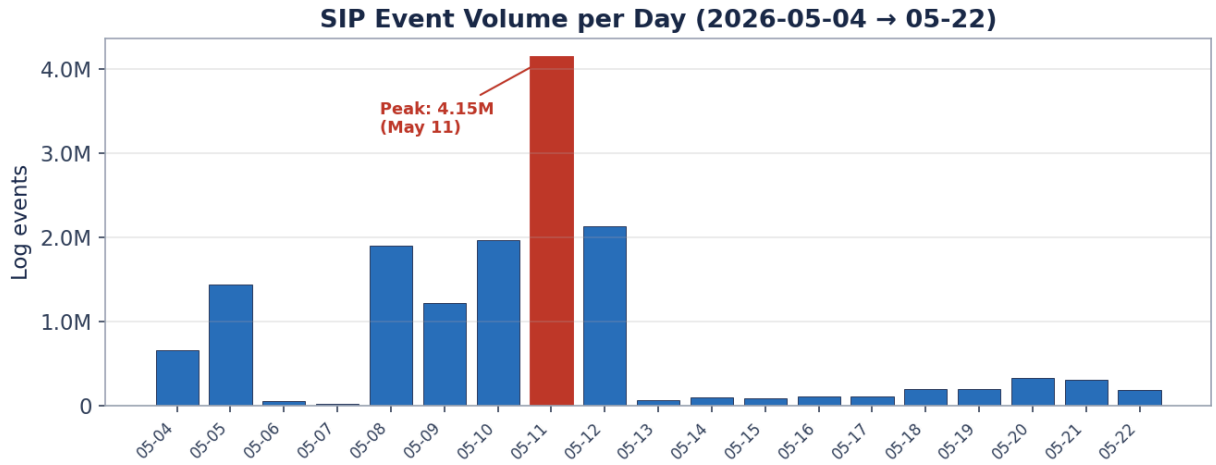


*Figure 1. Daily SIP event volume. A concentrated high-tempo campaign on 8–12 May dwarfs the surrounding baseline.*

### Key Judgments

- **The dominant threat is credential harvesting, not opportunistic guessing.** A handful of hosts sprayed a curated dictionary of 277,632 passwords — weighted toward medium- and high-complexity strings, not just 1234 — against every common PBX extension number.

- **Toll fraud is targeted and revenue-driven.** 47,273 of 89,465 call attempts targeted UK numbers, concentrated on a small set of rural and Northern Ireland ranges consistent with International Revenue Share Fraud (IRSF).
- **The attack infrastructure is cheap, European, and highly concentrated.** A single hosting network (OVH, AS16276) originated 6,559,589 events; one /24 (15.204.157.0/24) produced 5,178,084 on its own.
- **The sources are known-bad and server-hosted, not victims' home connections.** Cross-referenced against third-party IP intelligence, 93.5% of attacker addresses already appear on a known-abuser list and 99.8% of source-attributed traffic originated from datacenter/hosting ranges.
- **Operations run around the clock and rotate identity, not behavior.** Tooling impersonates FreePBX, Cisco, Polycom and Avaya user agents, yet leaves stable signatures — most notably a hardcoded registration Contact of `sip:123@1.1.1.1` present in 3,396,685 registration requests.
- **The honeypot never granted access.** Every authentication was rejected; the value here is intelligence — adversary wordlists, target numbers, tooling, and infrastructure — not breach impact.

## The Campaign at a Glance

| | |
|---|---|
| **Observation window** | 2026-05-04 → 2026-05-22 (18 days) |
| **Sensor** | Single Internet-facing SIP service (Kamailio), UDP+TCP/5060 |
| **Total telemetry events** | 15,183,358 |
| **Distinct SIP requests** | ~3,787,791 |
| **Unique source IPs** | 323 |
| **Credential attempts captured** | 1,869,521 |
| **Plaintext passwords recovered** | 1,796,362 (96.09%) |
| **Unique passwords in adversary dictionary** | 277,632 |
| **Toll-fraud call attempts (INVITE)** | 89,465 |
| **Busiest day** | 2026-05-11 (4,152,736 events) |
| **Dominant hosting network** | AS16276 OVH SAS — 6,559,589 events |
| **Sources on known-abuser lists** | 93.5% of IPs |
| **Traffic from datacenter ranges** | 99.8% of attributed events |

## Scope, Sensor & Methodology

The sensor is a purpose-built SIP service exposed directly to the public Internet on port 5060 over both UDP and TCP. It answers like a misconfigured private branch exchange (PBX): it challenges registration and call attempts for authentication and records every field of every request, but it never accepts a credential and never completes a call. This design draws the full sequence of an attack — discovery, registration, credential submission, and call setup — while keeping the system inert and safe.

### Reading the Numbers

A single SIP transaction produces several telemetry records: a raw inbound packet, a parsed request, the server's challenge or rejection, and, for credentialed attempts, the captured authentication header. The 15,183,358 total events therefore correspond to approximately 3,787,791 actual SIP requests. Throughout this report, volumetric claims are stated against the measure that matches them — requests for attack counts, events for tempo — and both are reported so the relationship is transparent.

**Telemetry composition by record type**

| Record type | Count | What it represents |
|---|---|---|
| Raw inbound packets | 3,813,171 | UDP datagrams received at the socket |
| Parsed SIP requests | 3,787,791 | Fully parsed inbound requests |
| Server responses | 3,787,728 | Challenges (401/407) and rejections (403) |
| REGISTER — probe (no auth) | 1,851,566 | Initial registration, no credentials |
| REGISTER — credentialed | 1,842,267 | Registration carrying a Digest response |
| INVITE — no auth | 62,211 | Call setup attempt |
| INVITE — credentialed | 27,254 | Call setup carrying proxy auth |
| OPTIONS / ACK / other | 930 | Service discovery and misc methods |

*Counts exclude synthetic pipeline-health events (see below). REGISTER requests total ~3,693,833; INVITE requests ~89,465.*

## Attack Surface & Traffic Overview

Activity falls into three operational categories. Reconnaissance and registration probing account for the largest share of raw events; credential attacks are the largest share of intent-bearing traffic; and toll fraud, though smallest by volume, carries the most direct financial motive. The chart below normalizes all categories against the full real-event population, including the protocol-state records that accompany every transaction.

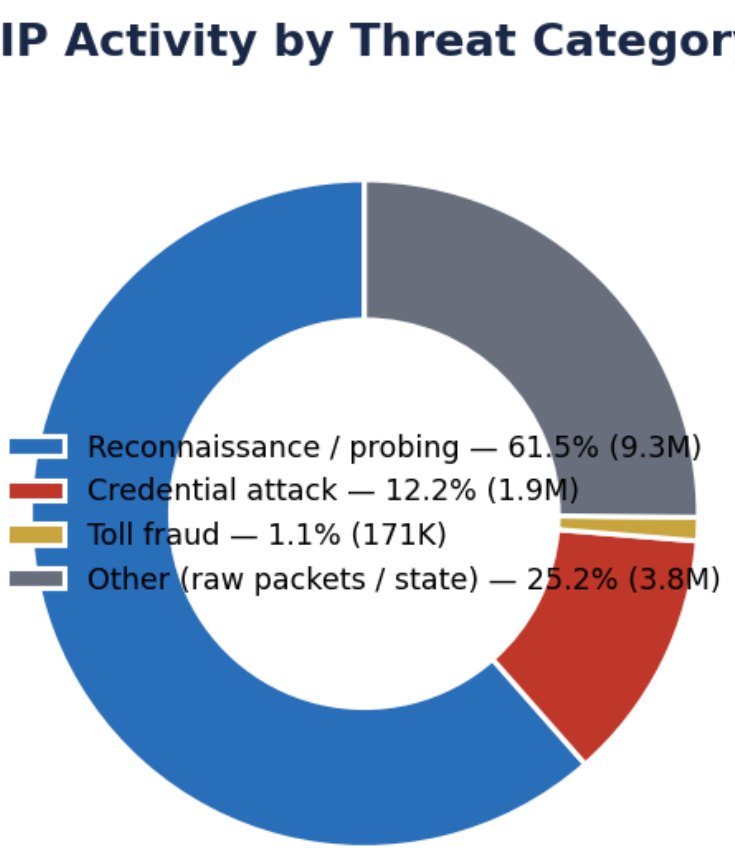


*Figure 2. Share of activity by threat category, over all real events.*

**Server responses issued (a proxy for request intent)**

| Response | Count | Triggered by |
|---|---|---|
| 401 Unauthorized | 1,851,570 | REGISTER probe — challenge issued |
| 403 Forbidden | 1,869,563 | Credentialed REGISTER/INVITE — rejected |
| 407 Proxy Auth Required | 62,211 | INVITE — proxy challenge issued |
| 200 OK | 851 | OPTIONS — service discovery answered |
| 400 Bad Request | 3,518 | Malformed input |

*Among parsed requests, transport resolved as 3,272,995 UDP vs 23,721 TCP — the campaign is almost entirely UDP.*

## Tempo

The campaign is bursty. A single high-intensity push between 8 and 12 May produced the bulk of all traffic, peaking at 4,152,736 events on 2026-05-11, before settling into a persistent lower-rate baseline. Activity is spread evenly across all 24 hours of the day, with no business-hours signature — the hallmark of unattended, automated operation rather than hands-on-keyboard work.

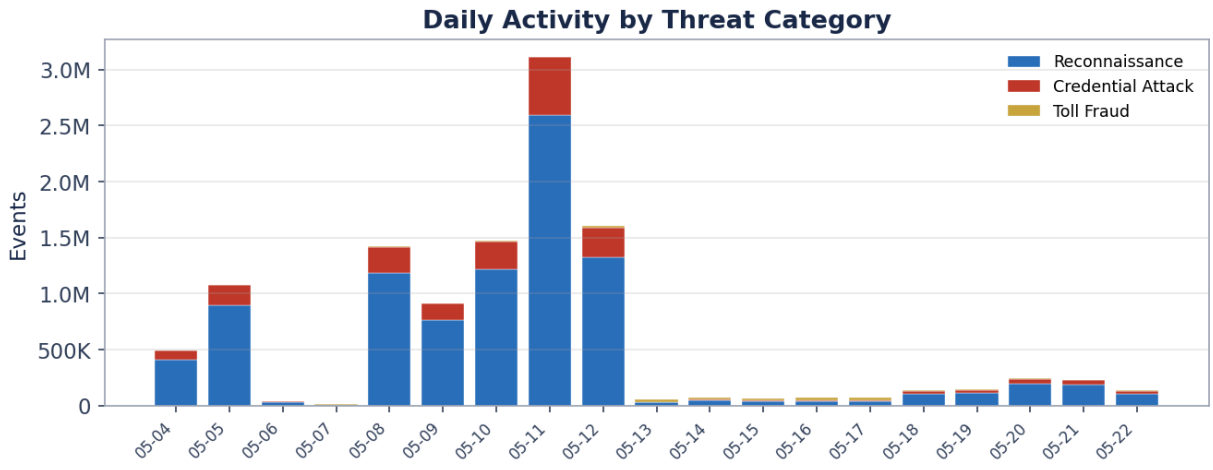


*Figure 3. Daily activity decomposed by threat category.*

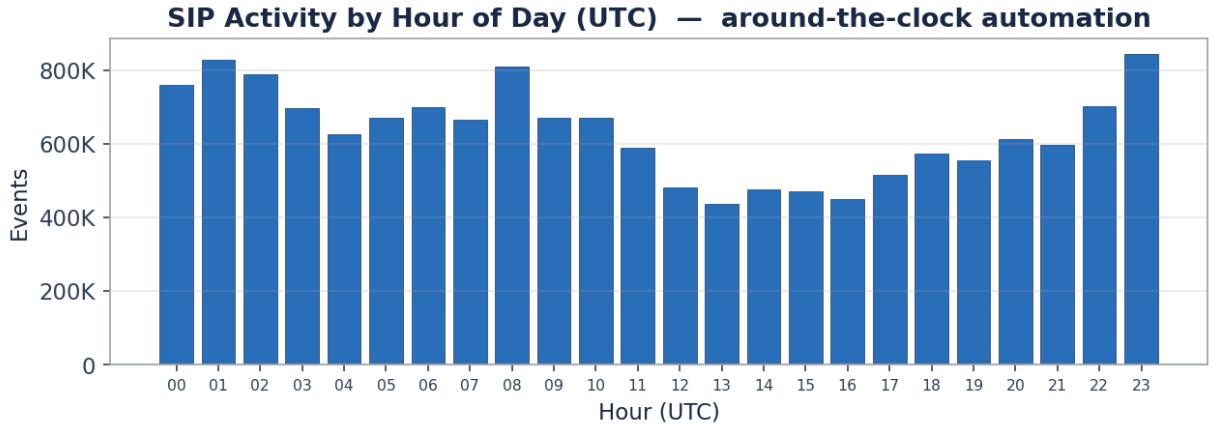


*Figure 4. Activity by hour of day (UTC). Continuous, around-the-clock automation.*

## The Attack Progression

The activity observed is not a single technique but a graduated campaign that escalates in sophistication. The sections that follow are ordered along this progression — from the lowest-effort, highest-volume reconnaissance through to the most complex, financially-motivated fraud. Each stage builds on the intelligence the previous one produced.

**Attack stages, ordered by complexity**

| Stage | Technique | Relative complexity | Observed volume |
|---|---|---|---|
| 1 | Reconnaissance & extension enumeration | Low | 9,335,528 events (61.5%) |
| 2 | Credential brute force & password spraying | Low–Medium | 1,869,521 credential attempts |
| 3 | Credential replay (harvested material) | Medium–High | 45,580 replayed attempts |
| 4 | Toll fraud & IRSF (monetization) | High | 89,465 call attempts |

*Complexity reflects operator effort and prerequisites, not traffic volume — the simplest stage is by far the loudest, while the most complex stage is the quietest but carries the direct financial payoff.*

## Stage 1 · Reconnaissance & Extension Enumeration

> **STAGE 1 · WHAT THIS ATTACK IS — SIP RECONNAISSANCE & ENUMERATION**
> Reconnaissance is the discovery phase that precedes everything else. Attackers send `OPTIONS` requests to confirm a SIP service is alive and fingerprint its software, and unauthenticated `REGISTER` probes to learn which extension numbers exist — a server often answers valid and invalid extensions differently, leaking its dial plan. It is low-effort, high-volume, and entirely automated; its only goal is to map the attack surface and build the target list that the credential and fraud stages then consume.

Reconnaissance and probing is the single largest activity class in the dataset — 9,335,528 events, 61.5% of all real traffic. It is dominated by **1,851,566 unauthenticated REGISTER probes**, each of which drew a `401` challenge (1,851,570 issued in total) before any credential was offered — the classic two-step where a tool first checks whether an extension answers, then escalates. Alongside these sit 851 `OPTIONS` service-discovery scans used purely to confirm and fingerprint the service.

The enumeration is exhaustive and numeric. Across the campaign attackers probed deep into the common PBX extension space, and in the credentialed traffic that follows, 29,433 distinct extension identities were targeted — every one a number, with no named accounts. Mapped to ATT&CK, this stage is *Gather Victim Identity Information* (T1589, 5,546,049 events) and *Network Service Discovery* (T1046, 3,789,473 events).

**Reconnaissance signals observed**

| Signal | Count | Purpose |
|---|---|---|
| Unauthenticated REGISTER probes | 1,851,566 | Extension enumeration (elicits 401 challenge) |
| 401 challenges issued | 1,851,570 | Server response to each probe |
| OPTIONS service-discovery scans | 851 | Confirm/fingerprint the SIP service |
| Distinct extensions targeted | 29,433 | Size of the enumerated dial-plan space |

*Once a live service and an extension list are confirmed, the same tooling escalates directly to Stage 2.*

## Stage 2 · Credential Brute Force & Password Spraying

> **STAGE 2 · WHAT THIS ATTACK IS — SIP REGISTRATION BRUTE FORCE**
> SIP phones and softphones log in to a PBX by sending a REGISTER request that binds an extension — a short numeric account such as `100` or `2001` — to a shared secret. In a registration brute-force attack, an adversary floods the server with REGISTER attempts, cycling through likely extension numbers and candidate passwords to find any account whose credential is weak, default, or reused. Because a PBX's accounts are simply its extension numbers, the username space is small and predictable, so attackers

> enumerate it exhaustively rather than guessing names like `admin`. A single cracked extension is enough to place fraudulent calls, eavesdrop, or pivot further into the network.

The defining activity of the dataset is registration brute force. The sensor captured 1,869,521 authentication attempts carrying a full Digest response (1,842,267 on REGISTER, 27,254 on INVITE), targeting 29,433 distinct identities. Critically, **every target is a numeric extension** — 3- and 4-digit PBX extension numbers such as those below. There is no meaningful presence of `admin`, `root`, or named accounts. The adversary understands that on a SIP PBX, the username space *is* the extension plan, and enumerates it exhaustively.

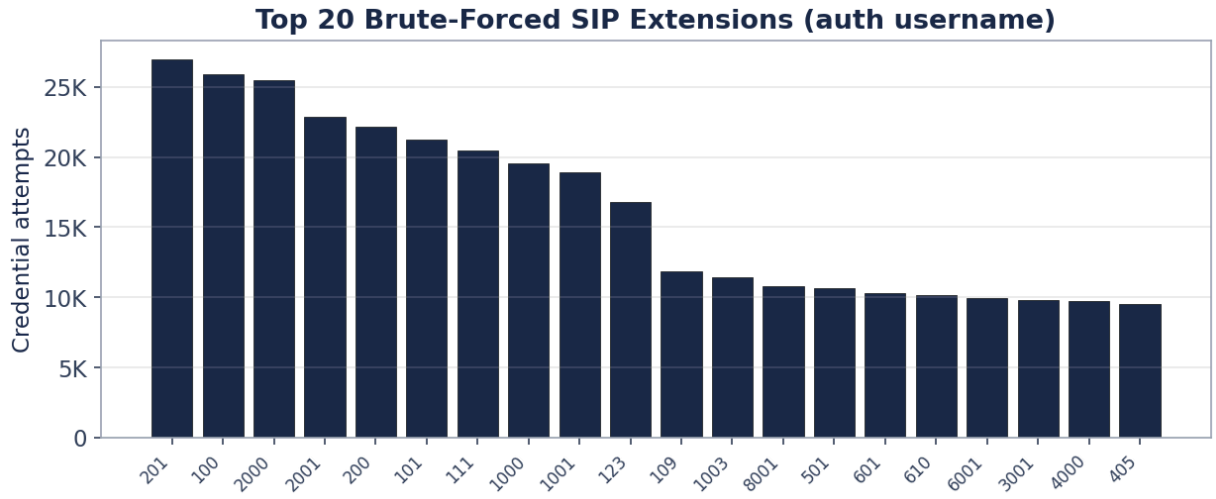


*Figure 5. The most-attacked extensions — pure numeric enumeration of common PBX dial plans.*

**Top targeted extensions**

| Extension | Attempts | Extension | Attempts |
|---|---|---|---|
| 201 | 26,959 | 100 | 25,897 |
| 2000 | 25,503 | 2001 | 22,897 |
| 200 | 22,129 | 101 | 21,223 |
| 111 | 20,434 | 1000 | 19,578 |
| 1001 | 18,932 | 123 | 16,825 |
| 109 | 11,838 | 1003 | 11,433 |
| 8001 | 10,822 | 501 | 10,686 |
| 601 | 10,303 | 610 | 10,161 |
| 6001 | 9,933 | 3001 | 9,775 |
| 4000 | 9,701 | 405 | 9,502 |

**A high-fidelity tool signature**

| Extension | Attempts | Extension | Attempts |
|---|---|---|---|

> Of all registration attempts, 3,396,685 carried the identical registration Contact `sip:123@1.1.1.1`. A live PBX would never see this. The hardcoded value is a stable fingerprint of the dominant credential-spraying tool and is one of the single highest-confidence indicators in this dataset.

## Stage 2 · Inside the Recovered Dictionary

> **STAGE 2 · WHAT THIS ATTACK IS — DIGEST AUTHENTICATION & PASSWORD SPRAYING**
> SIP secures logins with HTTP Digest authentication: instead of sending the password, the client returns an MD5 hash computed over the username, password, a server-issued realm and nonce, and the request line. On a properly configured server this keeps the password off the wire. 'Password spraying' is the practice of testing a large dictionary of candidate passwords across many accounts. Capturing these authentication exchanges is valuable because they reveal exactly which passwords an adversary is attempting — exposing the wordlists in active circulation rather than merely recording that an attempt occurred.

Because the captured authentication material is complete, the plaintext password behind **96.09% of the 1,869,521 credential attempts** was recovered — 1,693,562 confirmed where the tooling leaked the cleartext pair and the recovered value reproduced the captured hash, plus 102,800 more resolved by other means. The exposed wordlist contains 277,632 unique passwords across 1,499,846 unique extension/password pairs.

The composition matters. This is **not** a naive weak-password list. Weighted by attempt volume, recovered passwords skew toward medium and high complexity — a curated dictionary that mixes classic weak entries with device-default and breach-derived strings of real cryptographic strength. The implication for defenders is direct: a strong-looking SIP secret is not protection if it appears in a list like this one.

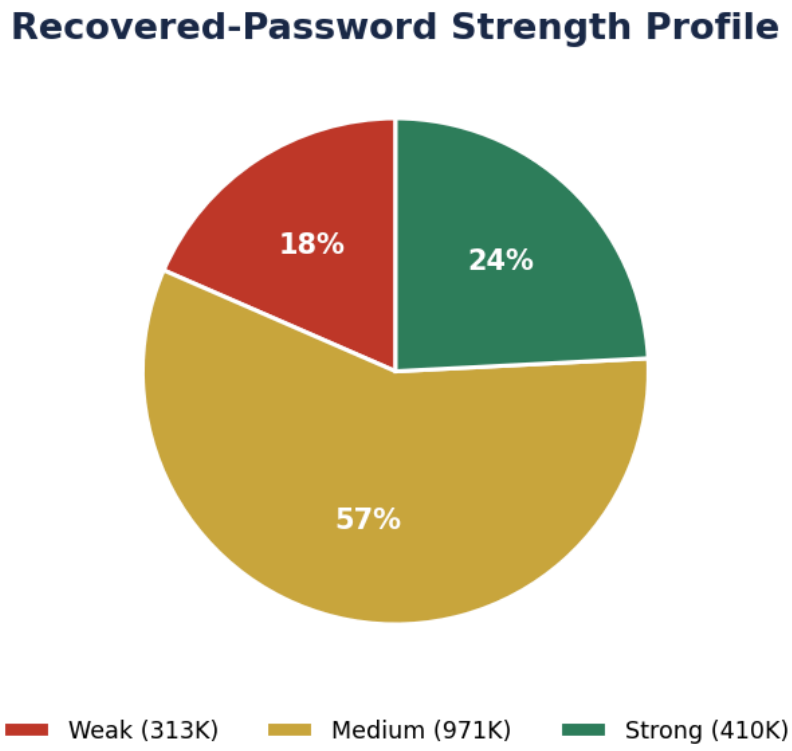


*Figure 6. Complexity profile of recovered passwords, weighted by attempts.*

### Most frequently attempted passwords (recovered plaintext)

| Password | Attempts | Password | Attempts |
|---|---|---|---|
| a%2508%23r8qH20HN | 806 | abc123 | 528 |
| abc123456 | 523 | 1234567 | 523 |
| 1234 | 523 | abc@123 | 521 |
| abc1234 | 521 | 12345678 | 521 |
| 12345 | 520 | abc12345 | 520 |
| 123 | 519 | 123456 | 517 |
| abc12345678 | 510 | abc1234567 | 508 |
| abc123456789 | 508 | Abc123 | 507 |
| abc@1234 | 507 | abc@123456 | 505 |
| admin123 | 503 | abc@12345 | 503 |
| Abc12345 | 502 | admin1234 | 501 |
| admin@123 | 499 | Abc1234 | 498 |

*Values shown as transmitted (some are percent-encoded). The distribution is strikingly flat — the dictionary is broad rather than top-heavy.*

### Passwords sprayed across the most distinct extensions

| Password | Distinct extensions | Total attempts |
|---|---|---|
| 987654321 | 320 | 457 |
| 1234567 | 313 | 523 |

| Password | Distinct extensions | Total attempts |
|---|---|---|
| abc123456 | 311 | 523 |
| 12345678 | 311 | 521 |
| abc12345678 | 311 | 510 |
| abc123 | 310 | 528 |
| 12345 | 310 | 520 |
| abc12345 | 310 | 520 |
| abc1234 | 310 | 521 |
| 1234 | 310 | 523 |
| abc1234567 | 310 | 508 |
| 123 | 309 | 519 |

The flat frequency distribution — the single most common password was tried only 806 times out of 1,796,362 recoveries — confirms a breadth-first strategy: the operator pushes a very large dictionary across the entire extension space rather than hammering a few guesses, maximizing the chance of catching any reused or default secret anywhere it lands.

> **Published dataset**
> The full recovered dictionary — 277,632 unique passwords and 1,499,846 extension/password pairs — is released publicly so operators can screen their own SIP secrets against this live attacker wordlist: **https://github.com/<your-org>/whisperpot-sip-credential-dataset**

## Stage 3 · Credential Replay & Harvested Infrastructure

> **STAGE 3 · WHAT THIS ATTACK IS — CREDENTIAL REPLAY & HARVESTING**
> In a credential-replay attack the adversary does not guess passwords against the target at all. Instead it submits authentication material — usernames, realms, and pre-computed response hashes — that it harvested from other systems it has already scanned or compromised, betting that the same credentials or configuration are reused elsewhere. The realm and server nonce embedded in a replayed attempt act as fingerprints of where that material originally came from, which makes this traffic a window into the operation's wider portfolio of victims.

Most attempts answered this sensor's own challenge. But 45,580 credential attempts (2.44%) carried an authentication realm that the sensor never issued, and 1,721 carried a nonce the sensor never generated. These attempts are **replays** — Digest responses computed against challenges harvested from other systems and fired blind at this one.

The foreign realms are a window into the operation's other targets. They include generic platform defaults (`asterisk`, `sipproxy`), specific PBX and device vendors (`Intelbras`, `grandstream`, `STARFACE`), and — most tellingly — the IP addresses of other PBX systems, one of them an **RFC 1918 private address**

(`10.13.0.16`) that can only have come from inside a victim network. This is direct evidence that the actor maintains a portfolio of compromised or probed PBXs and recycles harvested material across them.

**Authentication realms not issued by this sensor (replay indicators)**

| Realm presented | Attempts |
|---|---|
| asterisk | 24,751 |
| Intelbras | 3,454 |
| sipproxy | 3,273 |
| 168.121.253.140 | 2,552 |
| 10.13.0.16 | 1,989 |
| 51.75.37.4 | 1,974 |
| 41.76.197.42 | 1,786 |
| grandstream | 1,711 |
| STARFACE | 1,696 |
| 84.49.129.185 | 960 |
| 188.113.135.160 | 680 |
| 83.101.48.102 | 482 |
| 67.219.100.7 | 272 |

*Realms naming bare IP addresses correspond to other PBX systems the operation has interacted with.*

## Stage 4 · Toll Fraud & International Revenue-Share Fraud

> **STAGE 4 · WHAT THIS ATTACK IS — TOLL FRAUD & IRSF**
> Toll fraud is the abuse of a victim's phone system to place calls the victim ultimately pays for. Its most lucrative form is International Revenue Share Fraud (IRSF): criminals acquire premium-rate or international 'revenue-share' numbers that pay them a cut of every call's termination charge, then drive a compromised or misconfigured PBX to dial those numbers at volume — typically automated and outside business hours — turning someone else's phone bill into direct profit. Before the calls can connect, attackers probe the PBX dial plan (the rules governing which numbers it will route, and in what format) to find a dialing pattern the system accepts.

Running in parallel with the credential assault is a stream of 89,465 call-setup attempts. Their destinations are not random: 47,273 attempts targeted the United Kingdom, far ahead of any other country, and concentrated on a small number of rural and Northern Ireland landline ranges. This pattern is the signature of **International Revenue Share Fraud (IRSF)** — the attacker controls or rents premium terminating numbers, then tries to make a victim PBX dial them so that the resulting call charges convert into shared revenue.

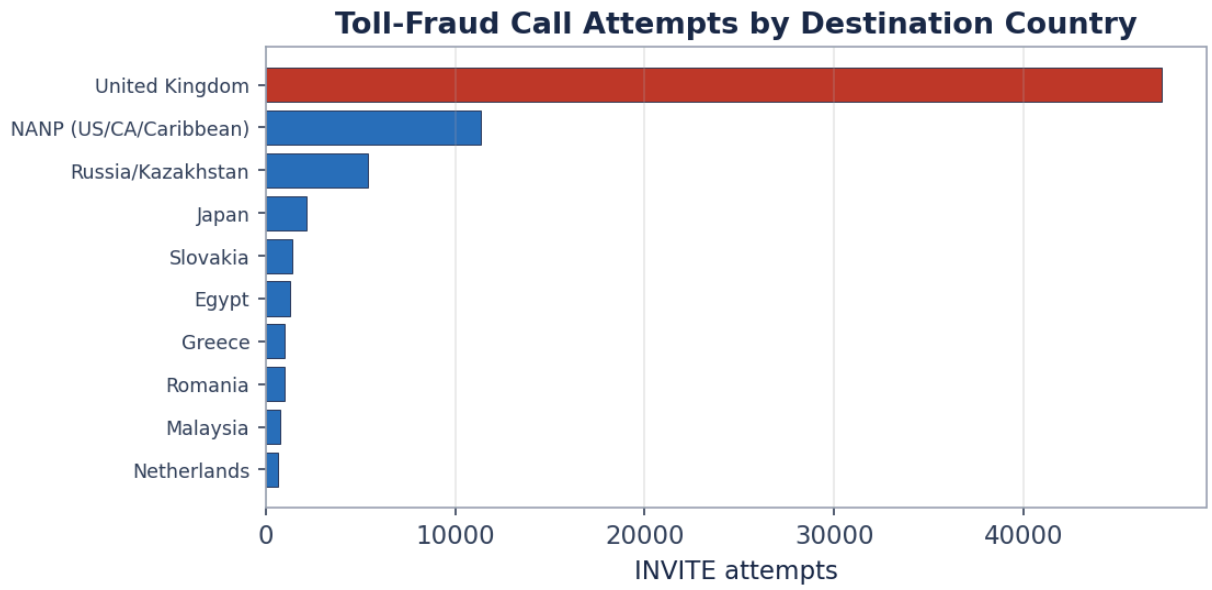


*Figure 7. Toll-fraud call attempts by destination country. UK targeting dominates.*

**Most-targeted destination numbers**

| Destination (normalized) | Attempts | Country | Prefix variants |
|---|---|---|---|
| 442820539014 | 11,049 | United Kingdom | 19 |
| 441863614031 | 5,211 | United Kingdom | 7 |
| 441863614012 | 4,774 | United Kingdom | 80 |
| 441863614013 | 3,807 | United Kingdom | 9 |
| 441980774302 | 3,173 | United Kingdom | 7 |
| 442893587015 | 2,961 | United Kingdom | 58 |
| 441330562015 | 2,952 | United Kingdom | 7 |
| 441833542050 | 1,856 | United Kingdom | 8 |
| 441873901207 | 1,457 | United Kingdom | 5 |
| 441887593451 | 1,424 | United Kingdom | 5 |
| 7233751167 | 1,301 | Russia/Kazakhstan | 169 |
| 442820539018 | 1,256 | United Kingdom | 8 |

## Dial-Plan Bypass by Prefix Rotation

The tooling does not know the victim PBX's dial plan, so it brute-forces it. For each target number, the same source re-dials under a rotating set of outbound and international access prefixes — bare, `00`, `000`, `900`, `+`, `011`, `0011` — probing for any permutation the dial plan will accept. One UK number was attempted under 80+ distinct prefix variants. The table below shows the most active source/destination pairs and the prefixes each cycled through.

**Prefix-rotation behavior (dial-plan probing)**

| Source IP | Target number | Attempts | Prefixes cycled |
|---|---|---|---|
| 217.160.24.49 | 441863614031 | 5,207 | 00441863614031, 000441863614031, 441863614031, 900441863614031, +441863614031 |
| 87.106.105.253 | 441863614012 | 4,625 | 00441863614012, 000441863614012, 441863614012, +441863614012, 900441863614012 |
| 87.106.206.249 | 442820539014 | 3,894 | 000442820539014, 00442820539014, +442820539014, 442820539014, 900442820539014 |
| 87.106.141.148 | 441863614013 | 3,799 | 00441863614013, 000441863614013, 900441863614013, +441863614013, 441863614013 |
| 87.106.98.147 | 441980774302 | 3,171 | 00441980774302, 000441980774302, +441980774302, 900441980774302, 441980774302 |
| 87.106.78.3 | 441330562015 | 2,950 | 00441330562015, 000441330562015, +441330562015, 441330562015, 900441330562015 |
| 217.160.24.45 | 442893587015 | 2,903 | 00442893587015, 000442893587015, 442893587015, 900442893587015, +442893587015 |
| 87.106.189.67 | 442820539014 | 2,864 | 900442820539014, 00442820539014, 442820539014, +442820539014, 000442820539014 |
| 87.106.223.228 | 442820539014 | 2,607 | 000442820539014, 00442820539014, 442820539014, +442820539014, 900442820539014 |
| 31.70.66.9 | 441833542050 | 1,818 | 00441833542050, +441833542050, 000441833542050, 011441833542050, 900441833542050, 441833542050 |

### Spoofed Identities and Injection Probes

Caller identities on these attempts are fabricated and varied: internal-looking extensions (`10001`, `100`), device and trunk labels (`cisco`, `atcom`, `trunk_1`), and spoofed full E.164 caller IDs. Notably, the caller field also carried SQL injection probes such as `'or''='` — the same tooling reflexively tests whether the SIP front end passes caller data into a backend database, blending telephony fraud with conventional web-style application attacks.

## Tooling & Tradecraft

The User-Agent strings are an exercise in camouflage. The most common identities — `FPBX` (FreePBX), generic `PBX`, and a range of `Cisco`, `Polycom`, `Avaya` and `Linksys` device strings — are designed to blend into legitimate VoIP traffic. They are trivially spoofed and should be read as costumes, not identities. Alongside the impersonators sit openly recognizable scanning tools, chiefly `pplsip` and the SIPVicious-derived `friendly-scanner`.

**Most common User-Agent strings**

| User-Agent (claimed) | Events | Interpretation |
|---|---|---|
| FPBX | 4,419,822 | FreePBX impersonation |
| PBX | 1,961,611 | Generic PBX impersonation |
| pplsip | 439,610 | Known SIP scanner |
| Cisco | 411,629 | IP-phone / device impersonation |
| <null> | 98,713 | No User-Agent header (unidentified) |
| Avaya | 70,311 | IP-phone / device impersonation |
| SNAPmobile | 39,716 | Mobile softphone client |
| PolycomSoundPointIP-SPIP_450-UA/3.3.4.0085 | 33,000 | IP-phone / device impersonation |
| PolycomSoundPointIP-SPIP_450-UA | 25,180 | IP-phone / device impersonation |
| Cisco/SPA504G-7.4.9c | 16,444 | IP-phone / device impersonation |
| Linksys-SPA942 | 15,034 | IP-phone / device impersonation |
| PolycomVVX-VVX_450-UA/6.4.3.5059 | 12,025 | IP-phone / device impersonation |

Behaviorally, the actors rotate identity far more readily than method. The same hardcoded registration Contact, the same prefix-rotation routine, and the same From-tag credential encoding recur across many source addresses and many claimed User-Agents — which is exactly what makes those behaviors, rather than the spoofed banners, the durable detections.

## Attacker Infrastructure

The infrastructure is cheap, rented, and remarkably concentrated. A single network — AS16276 OVH SAS — originated 6,559,589 events from 44 addresses, more than every other network combined. The next tier

is a familiar roster of low-cost European hosting providers. This concentration is itself an opportunity: network- and ASN-level controls are disproportionately effective against this kind of campaign.

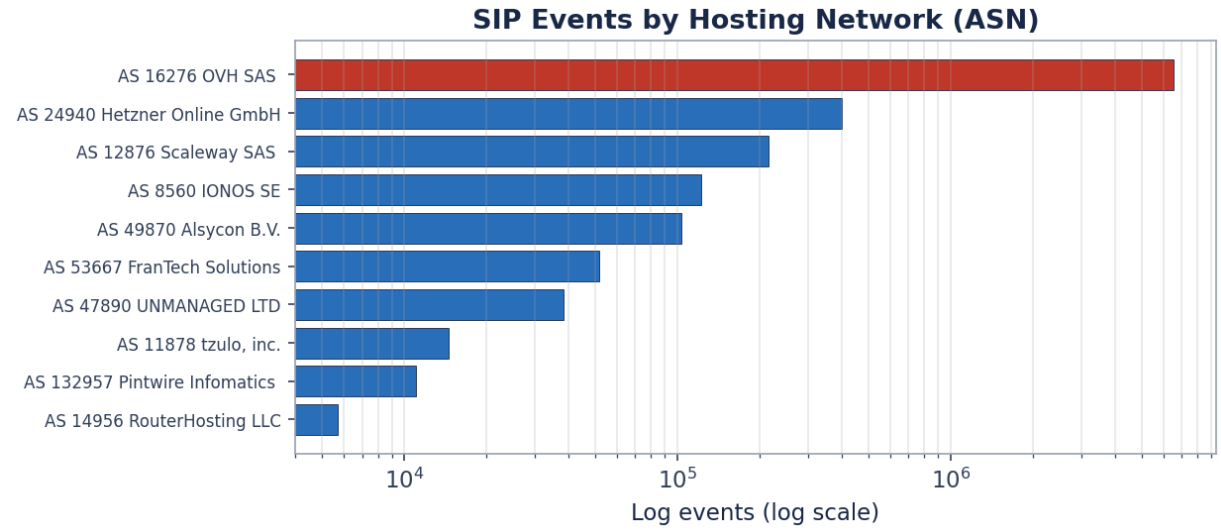


*Figure 8. SIP events by originating network (log scale).*

**Top originating networks (ASN)**

| Network | Events | Unique IPs |
|---|---|---|
| AS16276 OVH SAS | 6,559,589 | 44 |
| AS24940 Hetzner Online GmbH | 400,625 | 2 |
| AS12876 Scaleway SAS | 215,637 | 4 |
| AS8560 IONOS SE | 122,446 | 27 |
| AS49870 Alsycon B.V. | 103,144 | 9 |
| AS53667 FranTech Solutions | 51,759 | 11 |
| AS47890 UNMANAGED LTD | 38,253 | 1 |
| AS11878 tzulo, inc. | 14,570 | 2 |
| AS132957 Pintwire Infomatics Private Limited | 10,986 | 1 |
| AS14956 RouterHosting LLC | 5,684 | 3 |

## Adjacent-IP Fleets

Several actors operate consecutive addresses inside a single /24 — rented in blocks and burned together. The clearest example, 15.204.157.0/24, accounts for 5,178,084 events across 3 addresses behaving identically.

**Multi-IP subnet clusters**

| Subnet | IPs | Events | Network | Claimed UA |
|---|---|---|---|---|
| 15.204.157.0/24 | 3 | 5,178,084 | AS16276 OVH SAS | PBX, FPBX |

| Subnet | IPs | Events | Network | Claimed UA |
|---|---|---|---|---|
| 89.190.156.0/24 | 3 | 103,104 | AS49870 Alsycon B.V. | PBX, Cisco-CP7960G/8.0 |
| 172.110.223.0/24 | 13 | 38,349 | — | friendly-scanner, PBX |
| 45.61.184.0/24 | 2 | 19,480 | AS53667 FranTech Solutions | SNAPmobile |
| 217.160.24.0/24 | 2 | 19,190 | AS8560 IONOS SE | pplsip, Avaya |
| 31.70.75.0/24 | 6 | 15,375 | AS8560 IONOS SE | Avaya, pplsip |
| 45.61.186.0/24 | 2 | 12,025 | AS53667 FranTech Solutions | PolycomVVX-VVX_450-UA/6.4.3.5059, PortSIP |
| 5.135.71.0/24 | 2 | 9,326 | AS16276 OVH SAS | Avaya, friendly-scanner |

## Actor Archetypes

Classifying each source by its dominant tooling reveals a steep concentration of impact. A small group of PBX-impersonating hosts drives the overwhelming majority of credential attempts, while a long tail of scanner-tool and device-impersonation hosts contributes the breadth of coverage.

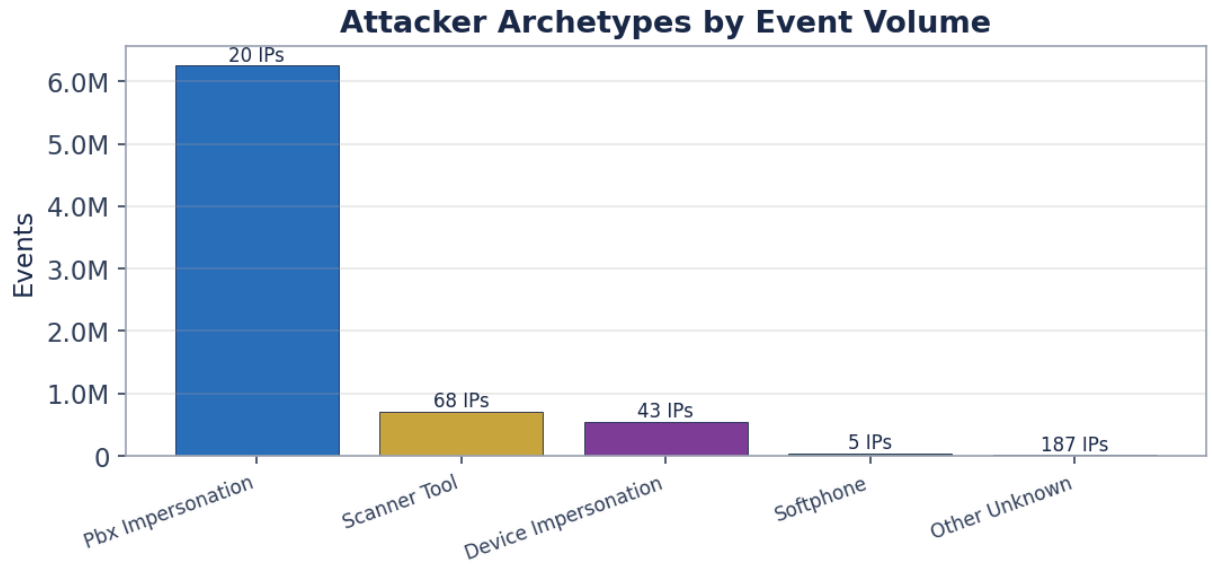


*Figure 9. Attacker archetypes by event volume (IP counts annotated).*

### Source breakdown by archetype

| Archetype | Source IPs | Events | Credential attempts |
|---|---|---|---|
| Pbx Impersonation | 20 | 6,252,999 | 1,556,855 |
| Scanner Tool | 68 | 712,555 | 166,520 |
| Device Impersonation | 43 | 551,478 | 134,671 |
| Softphone | 5 | 39,738 | 9,920 |
| Other Unknown | 187 | 15,264 | 1,555 |

## Infrastructure Reputation & Anonymization

To corroborate the picture above with an independent source, all 323 attacker addresses were cross-referenced against a commercial IP-intelligence dataset. The result removes any ambiguity about what this infrastructure is. **99.8% of source-attributed attack traffic came from datacenter and hosting ranges**, and **93.5% of the source addresses (302 of 323) were already on a known-abuser reputation list** before they ever reached this sensor. This is not background Internet noise from compromised home machines; it is purpose-rented server infrastructure with a pre-existing abuse history.

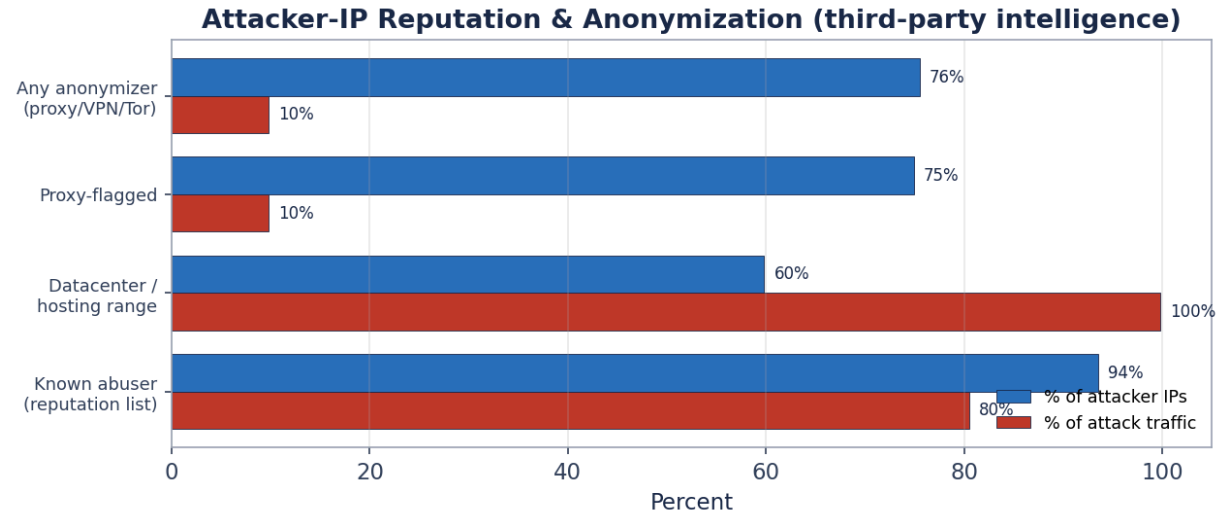


*Figure 11. Reputation and anonymization profile of attacker IPs, by share of IPs and of attack traffic.*

**Attacker-IP reputation summary (independent IP intelligence)**

| Classification | Attacker IPs | % of IPs | % of attack traffic |
|---|---|---|---|
| On known-abuser list | 302 | 93.5% | 80.5% |
| Datacenter / hosting range | 193 | 59.8% | 99.8% |
| Proxy-flagged | 242 | 74.9% | 9.8% |
| Any anonymizer (proxy/VPN/Tor) | 244 | 75.5% | 9.8% |

*Anonymizers are common among low-volume IPs but rare among the high-volume credential-spray hosts — VPN, Tor and mobile sources are effectively absent from the heavy traffic.*

Operator-type classification reinforces the same conclusion: 89.8% of source-attributed traffic is attributed to `hosting`-type companies. The independently-derived list of datacenters below re-confirms the OVH-led concentration seen earlier — arrived at from a completely separate dataset.

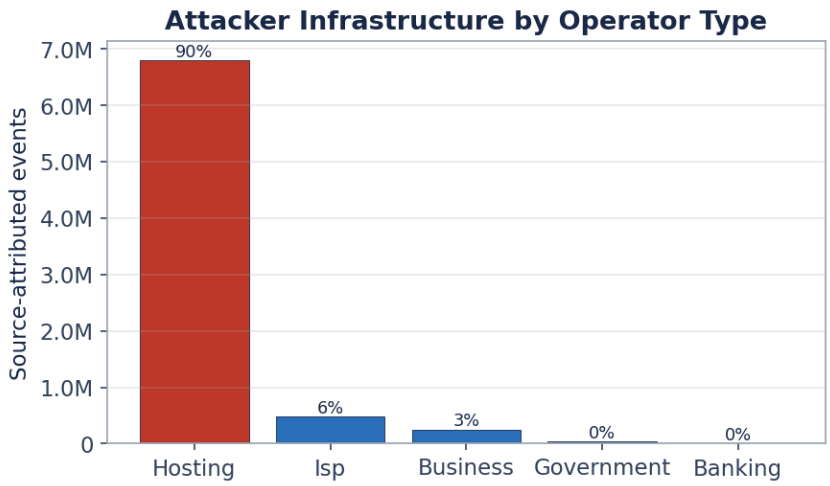


*Figure 12. Attacker infrastructure by operator type (share of attack traffic labeled).*

**Top datacenters hosting the attack traffic (independent intelligence)**

| Datacenter / hosting provider | Attacker IPs | Events |
|---|---|---|
| OVH US LLC | 6 | 5,178,226 |
| OVH Ltd | 3 | 884,064 |
| Hetzner Online GmbH | 1 | 400,623 |
| OVH BE | 2 | 329,349 |
| Scaleway SAS | 1 | 161,053 |
| IONOS SE | 26 | 115,886 |
| HostUS Solutions LLC | 9 | 103,144 |
| OVH SAS | 8 | 88,609 |
| Scaleway | 3 | 54,584 |
| frantech.ca | 8 | 52,207 |

**Why this matters for defenders**

When 93.5% of attackers are already known-bad and 99.8% of the traffic comes from hosting ranges, reputation- and hosting-aware filtering at the SIP edge becomes one of the highest-leverage controls available: a typical enterprise PBX has little legitimate reason to accept registrations or calls originating from bulk datacenter space.

## Threat Actor Profiles

A small number of operators define the dataset. The pair 15.204.157.10 / 15.204.157.11 (consecutive addresses on AS16276 OVH SAS) ran the central credential-spray campaign, together producing over 5,177,956 events in five days. 5.39.101.60 (User-Agent `pplsip`) is the most persistent single actor, active across the full observation window. A separate cluster on IONOS infrastructure (the `87.106.*` and `217.160.24.*` ranges) specialized in the UK toll-fraud calling described earlier.

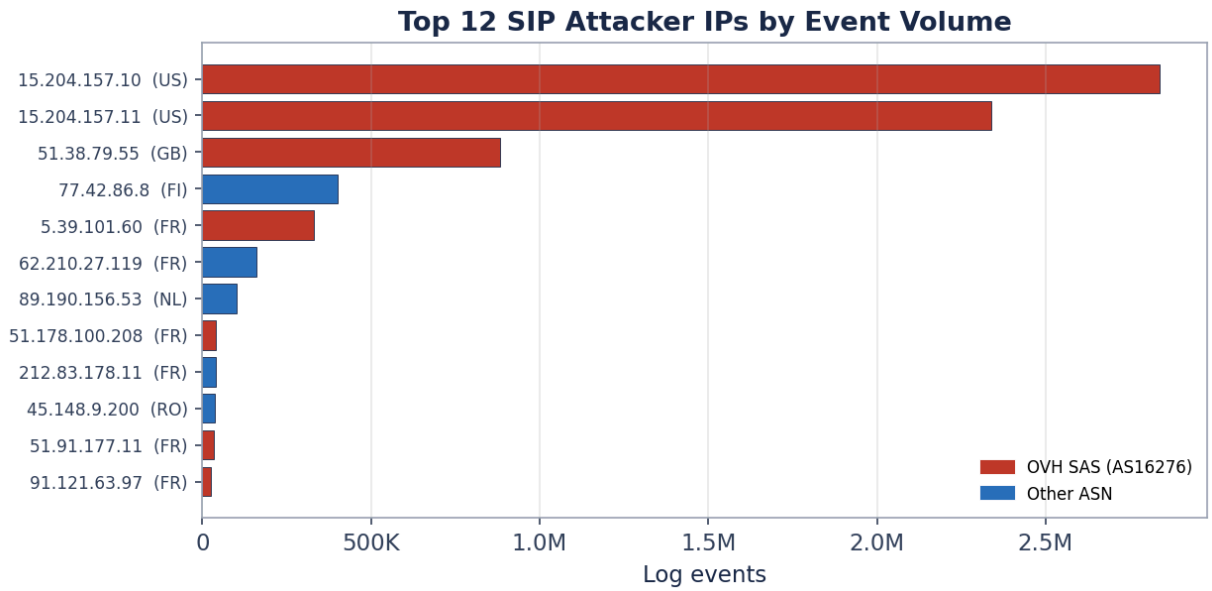


*Figure 10. Top source addresses by event volume; OVH-hosted hosts highlighted.*

**Top source addresses (with independent reputation classification)**

| Source IP | Events | Cred. attempts | Country | Network | Reputation |
|---|---|---|---|---|---|
| 15.204.157.10 | 2,838,149 | 707,961 | US | AS16276 OVH SAS | abuser(Elevated), datacenter |
| 15.204.157.11 | 2,339,807 | 583,214 | US | AS16276 OVH SAS | abuser(Elevated), datacenter |
| 51.38.79.55 | 883,944 | 220,216 | GB | AS16276 OVH SAS | datacenter |
| 77.42.86.8 | 400,623 | 99,943 | FI | AS24940 Hetzner Online | datacenter |
| 5.39.101.60 | 329,347 | 82,295 | FR | AS16276 OVH SAS | abuser(High), datacenter, proxy |
| 62.210.27.119 | 161,053 | 40,065 | FR | AS12876 Scaleway SAS | datacenter |
| 89.190.156.53 | 101,433 | 25,299 | NL | AS49870 Alsycon B.V. | abuser(Very Low), datacenter, proxy |
| 51.178.100.208 | 40,178 | 10,020 | FR | AS16276 OVH SAS | abuser(Low), datacenter, proxy |

| Source IP | Events | Cred. attempts | Country | Network | Reputation |
|---|---|---|---|---|---|
| 212.83.178.11 | 38,668 | 9,637 | FR | AS12876 Scaleway SAS | abuser(High), datacenter |
| 45.148.9.200 | 38,253 | 9,533 | RO | AS47890 UNMANAGED LTD | abuser(Elevated), datacenter, proxy |
| 51.91.177.11 | 33,673 | 8,393 | FR | AS16276 OVH SAS | abuser(High), datacenter, proxy |
| 91.121.63.97 | 25,152 | 6,282 | FR | AS16276 OVH SAS | abuser(Elevated), datacenter |

*Reputation/datacenter/proxy flags are drawn from third-party IP intelligence, independent of this sensor's telemetry.*

## MITRE ATT&CK Mapping

Observed behaviors map to the following ATT&CK techniques, ranked by event volume.

| Technique | Name / context in this dataset | Events |
|---|---|---|
| T1589 | Gather Victim Identity Information (extension enumeration) | 5,546,049 |
| T1046 | Network Service Discovery / scanning | 3,789,473 |
| T1110 | Brute Force (credential spraying) | 1,842,267 |
| T1498 | Network Denial of Service / resource abuse (toll fraud) | 154,561 |
| T1078 | Valid Accounts (credentialed call attempts) | 27,254 |
| T1071 | Application Layer Protocol (SIP session activity) | 146 |
| T1087 | Account Discovery | 6 |
| T1563 | Remote Service Session Hijacking | 6 |

## Detection & Mitigation Recommendations

The behaviors in this dataset are noisy and repetitive, which makes them highly detectable and largely preventable with controls aimed at the specific tradecraft observed.

### Detection

- Alert on REGISTER traffic whose Contact header is `sip:123@1.1.1.1` — in this dataset it is exclusively adversarial.
- Flag any single source attempting many distinct extensions in a short window (extension enumeration), and any REGISTER carrying an authentication realm or nonce your platform never issued (replayed/harvested credentials).
- Treat repeated INVITEs to the same destination under rotating dial prefixes (`00`/`000`/`900`/`+`/`011`) as dial-plan probing, regardless of caller identity.

- Monitor for spoofed or recognizable scanner User-Agents (`pplsip`, `friendly-scanner`) and for non-numeric or injection-like values (`'or''='`) in caller and username fields.

### Mitigation

- Do not expose SIP registration to the open Internet; restrict 5060/UDP+TCP to known peers, VPN, or session-border-controller front ends with rate limiting and fail2ban-style lockout.
- Enforce high-entropy, unique SIP secrets and screen them against breach/credential lists — recovered passwords here include strong-looking strings, so complexity alone is insufficient.
- Lock down the outbound dial plan: deny international and premium ranges by default, allow-list only required destinations, and require explicit authorization for high-cost prefixes.
- Consider ASN/geo-aware filtering: a small set of hosting networks (led by AS16276) produced most of this traffic and rarely originates legitimate calls for a typical enterprise PBX.
- Set per-account and per-trunk call-spend ceilings and real-time anomaly alerts so that any successful fraud is contained financially.

## Indicators of Compromise

The following indicators are derived directly from observed activity. IP and User-Agent indicators reflect rented, rotating infrastructure and spoofable banners — prioritize the behavioral indicators (the Contact string, prefix rotation, replayed realms) for durable detection.

**High-confidence behavioral indicators**

| Indicator | Value | Type |
|---|---|---|
| Registration Contact | sip:123@1.1.1.1 | Tool signature |
| Fixed challenge realm abused | honeypot.local | Context |
| SQL-injection probe in caller | 'or"=' | Tradecraft |
| Scanner User-Agents | pplsip, friendly-scanner | Tooling |
| Dominant hosting ASN | AS16276 OVH SAS | Infrastructure |
| Sources on known-abuser lists | 93.5% of attacker IPs | Reputation |
| Sources in datacenter ranges | 99.8% of traffic | Reputation |

**Top source IP indicators**

| IP | Network | Country |
|---|---|---|
| 15.204.157.10 | AS16276 OVH SAS | US |
| 15.204.157.11 | AS16276 OVH SAS | US |
| 51.38.79.55 | AS16276 OVH SAS | GB |
| 77.42.86.8 | AS24940 Hetzner Online GmbH | FI |
| 5.39.101.60 | AS16276 OVH SAS | FR |

| IP | Network | Country |
|---|---|---|
| 62.210.27.119 | AS12876 Scaleway SAS | FR |
| 89.190.156.53 | AS49870 Alsycon B.V. | NL |
| 51.178.100.208 | AS16276 OVH SAS | FR |
| 212.83.178.11 | AS12876 Scaleway SAS | FR |
| 45.148.9.200 | AS47890 UNMANAGED LTD | RO |
| 51.91.177.11 | AS16276 OVH SAS | FR |
| 91.121.63.97 | AS16276 OVH SAS | FR |

**Toll-fraud destination indicators (normalized)**

| Number | Country | Number | Country |
|---|---|---|---|
| 442820539014 | United Kingdom | 441863614031 | United Kingdom |
| 441863614012 | United Kingdom | 441863614013 | United Kingdom |
| 441980774302 | United Kingdom | 442893587015 | United Kingdom |
| 441330562015 | United Kingdom | 441833542050 | United Kingdom |
| 441873901207 | United Kingdom | 441887593451 | United Kingdom |
| 7233751167 | Russia/Kazakhstan | 442820539018 | United Kingdom |

**Most-targeted extension identities**

| Extensions targeted (sample) |
|---|
| 201, 100, 2000, 2001, 200, 101, 111, 1000, 1001, 123, 109, 1003, 8001, 501, 601, 610, 6001, 3001, 4000, 405 |

## Appendix: Analytical Notes

- All figures exclude 2,460 synthetic pipeline-health events; analysis covers 15,180,898 real events.
- Volumetric claims distinguish telemetry events from SIP requests; attack counts use the request or credentialed-event measure that matches the claim.
- Credential recovery was validated by agreement between two independent methods on 99.9944% of decoded attempts.
- Geolocation and network ownership reflect hosting, not operator location; User-Agent and caller fields are attacker-controlled.
- No authentication succeeded and no call completed; there is no breach impact, only intelligence value.
- The recovered credential dictionary is published for defensive screening at https://github.com/<your-org>/whisperpot-sip-credential-dataset.

## Disclaimer

This report is derived from application-layer telemetry collected by a controlled SIP honeypot and from independent third-party IP-intelligence enrichment. No production system was attacked and no real victim credentials are contained herein; all credential material consists of attacker-supplied inputs to the honeypot.

Indicators of Compromise (IOCs) — IP addresses, hosting networks, User-Agent strings, dialed numbers and the like — reflect rented, frequently-rotated infrastructure and spoofable identifiers. They should be operationalized with the behavioral indicators prioritized in this report and validated against the consumer's own environment before blocking.

The information is provided "as is" for defensive and research purposes. CloudSEK accepts no liability for actions taken on the basis of this report.

CloudSEK

Predict ◆ Protect ◆ Quantify

# We **Predict** Cyber Threats **Before They Strike**


### Registered Office:

CloudSEK Research Pte. Ltd.
160 Robinson Road, #20-03, Singapore Business Federation Center, Singapore - 068914

### Regional Office : India

CloudSEK Information Security Pvt Ltd.
16/1, Cambridge Rd, Halasuru, Cambridge Layout, Jogupalya, Bengaluru, Karnataka - 560008

### Regional Office : United States

CloudSEK Inc.
8 The Green, Ste A, Dover, DE - 19901, United States

### Regional Office : United Kingdom

CloudSEK, 2 Kingdom Street,
6th Floor London, W2 6BD - United Kingdom