VSPAM.ORG · INDEPENDENT RESEARCH

Trends in Spam, Phishing, Spoofing, Malware & DNS Abuse

A monthly research brief on the email, identity, and DNS threat landscape. May 2026 edition — focused on the month three flagship datasets landed and disagreed with one another.

ISSUE	**Volume 2026, Issue 5**
ISSN	**2026-VSPAM-005**
REPORTING PERIOD	1 – 31 May 2026
RELEASED	June 2026
PUBLISHER	vSpam.org Independent Research
LICENSE	Creative Commons Attribution 4.0 (CC BY 4.0)

ABSTRACT

May 2026 delivered three flagship datasets in sixteen days, and they did not agree. Barracuda's **2026 Email Threats Report** (May 12) found **one in three messages malicious or unwanted** across 3.1 billion emails. The **Verizon DBIR** (May 20), the largest in its history at 22,000+ confirmed breaches, placed vulnerability exploitation ahead of social engineering as the dominant initial-access vector — directly contradicting Cisco Talos a month earlier. And **APWG's Q1 2026 report** (May 27) recorded **971,181 phishing attacks**, up 13.8%, with telecom rocketing from 5.9% to **33% of all attacks** in two quarters. This issue reconciles the divergence — it is a window-alignment artefact, not a change in the threat — reports the vSpam.org collector corpus for the month, including a **bulk-registration burst concentrated in four novelty gTLDs**, and closes with operator guidance.

Table of Contents

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

Editorial disclosure. This issue was produced with editorial support from DDMARC (ddmarc.com), a DMARC and anti-spoofing platform that is a vSpam.org research partner. DDMARC had no veto over data, citations, or conclusions; competitor products and services are cited on the basis of the underlying data quality. Readers can verify every numeric claim against the references at the back of this report.

Executive Summary

May 2026 was the month the industry's measurement apparatus visibly disagreed with itself. Three flagship datasets landed inside sixteen days, each methodologically sound, each pointing somewhere different. Barracuda's 2026 Email Threats Report (May 12), built on 3.1 billion messages of January telemetry, found that **one in three email messages is now malicious or unwanted spam**, and attributed the growth to AI-powered social engineering and phishing-as-a-service.[1] The Verizon Data Breach Investigations Report (May 20), the largest in its history at more than **22,000 confirmed breaches**, found the human element in **62% of breaches** but placed vulnerability exploitation — not social engineering — as the dominant initial-access vector.[2] And APWG's Q1 2026 Phishing Activity Trends Report (May 27) counted **971,181 phishing attacks**, a **13.8% quarter-on-quarter rise** from 853,244 in Q4 2025.[3]

Those findings are not actually in conflict; they measure different populations over different windows. Section 7 reconciles them in full. The short version: the DBIR's observation window closed on **31 October 2025**, months before the Q1 2026 phishing rebound that Talos and APWG both recorded, and its unit of analysis is the confirmed breach rather than the attempted intrusion. Reading the DBIR as evidence that phishing has receded is a window-alignment error, and in our assessment it is the most consequential misreading available to a defender this quarter.

Four signals stand out beyond the measurement story.

1. **Telecom became the primary phishing target, and it happened fast.** APWG recorded telecom rising from **5.9% of all attacks in Q3 2025 to 33% in Q1 2026**, displacing SaaS and webmail (20%) from the top slot. URL-phishing frequency inside the telecom category rose **75% since Q4 2025** alone.[3]
2. **Ransomware fell below its 2025 pace for the first time in 2026.** BreachSense recorded **646 victims across 61 leak sites in 73 countries** — a 16% month-on-month decline and the year's lowest — while the ecosystem splintered further, with three groups entering the top ten for the first time in a single month.[4]
3. **The Tycoon2FA question moved from "did the takedown work" to "did it hold".** March's Europol-coordinated seizure removed a platform with at least 2,000 operators and more than 30,000 phishing domains; CrowdStrike reported in the weeks following that the service **persists in degraded form** rather than having been eliminated.[5] [6] Public platform-level telemetry for May had not been released at the time of writing; our own corpus offers an early read in Section 6.
4. **Our collector corpus recorded a bulk-registration campaign in novelty gTLDs.** Of the domain indicators vSpam.org ingested in May, **3,459 sat in just four TLDs** — `.lat`, `.surf`, `.pics` and `.garden` — a concentration with no precedent in our data and, as the following issue will show, no successor.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

Phishing Volume and the Telecom Pivot — APWG Q1 2026

Attack count rose 13.8% quarter-on-quarter; sector concentration shifted far more sharply

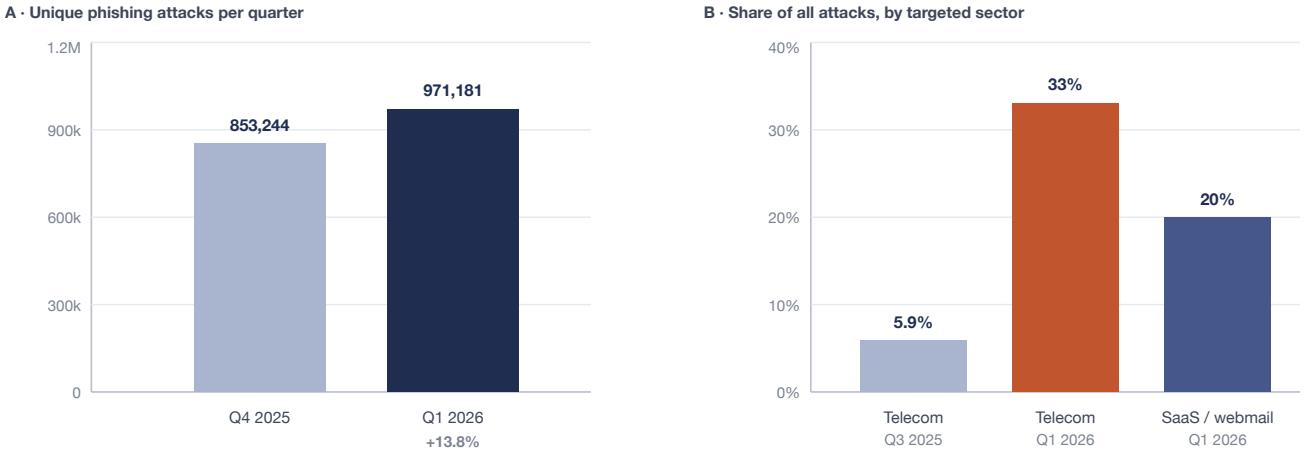

Source: APWG Phishing Activity Trends Report, 1st Quarter 2026 (published 27 May 2026).

FIGURE 1 *Quarterly attack volume and sector concentration. The 13.8% volume rise is the headline number, but the sector rotation is the more actionable finding: a category carrying 5.9% of attacks two quarters earlier now carries a third of them.*

A sector moving from 5.9% to 33% of observed attacks in two quarters is not organic drift. It is the signature of a small number of campaign operators re-pointing existing tooling at a new brand set — the same kits, the same hosting, a different logo in the template. That matters operationally because it means the defensive controls that worked against the previous target set are still valid; what has changed is which brands a filter should expect to see impersonated, and which customer base is receiving the lures.

The remainder of this issue analyses May along the five tracks used throughout the series — phishing, malware and ransomware, email authentication and spoofing, DNS abuse, and major incidents — then adds a sixth track reporting the vSpam.org collector corpus directly, reconciles the month's conflicting datasets, and closes with recommendations for messaging operators, registrars, and enterprise defenders.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

Headline Numbers at a Glance

METRIC	VALUE	SOURCE
Unique phishing attacks, Q1 2026	**971,181 (+13.8% QoQ)**	APWG
Prior quarter baseline, Q4 2025	**853,244**	APWG
Telecom share of all phishing attacks	**33% (from 5.9% in Q3 2025)**	APWG
SaaS / webmail share	**20%**	APWG
URL-phishing frequency growth, telecom	**+75% since Q4 2025**	APWG
Impersonation share of social-media threats	**43.8%**	APWG
Scam share of social-media threats	**27.1%**	APWG
Email messages malicious or unwanted spam	**1 in 3**	Barracuda
Email corpus analysed (January telemetry)	**3.1 billion messages**	Barracuda
Confirmed breaches analysed, DBIR 2026	**22,000+ (record)**	Verizon DBIR
Breaches involving the human element	**62%**	Verizon DBIR
Social engineering share of breaches	**16%**	Verizon DBIR
Phishing share of AI-assisted initial access	**44% (largest category)**	Verizon DBIR
GenAI techniques per documented campaign	**15 (median)**	Verizon DBIR
Mobile-channel social-engineering engagement	**+40% vs email lures**	Verizon DBIR
Phishing share of observed email attacks	**58%**	Abnormal AI
Phishing attacks using redirect chains	**21.6%**	Abnormal AI
Vendor email compromise share of BEC	**61%**	Abnormal AI
Ransomware victims posted to leak sites	**646 (61 groups, 73 countries)**	BreachSense
Month-on-month ransomware change	**–16% (year's lowest)**	BreachSense
YTD ransomware victims (Jan–May)	**3,583 (annualises ≈8,600)**	BreachSense
Most active ransomware group	**Qilin, 101 (5th month at #1)**	BreachSense
Domains with any DMARC record	**52.1% (937,931 domains)**	EasyDMARC
Domains combining enforcement with reporting	**≈9%**	EasyDMARC
vSpam.org collector indicators ingested	**30,039**	vSpam.org
vSpam.org domain IOCs in four novelty gTLDs	**3,459**	vSpam.org

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

REFERENCE

Methodology and Data Sources

This report consolidates public, attributable data published during or directly referencing May 2026 and — new from this issue — reports directly from the vSpam.org collector corpus. The two are kept in separate sections and separately labelled throughout. Every numeric claim drawn from an external source links to a primary or secondary public reference at the back of this report.

Primary external sources include the Anti-Phishing Working Group (APWG), the Verizon Data Breach Investigations Report, CISA, ENISA, M3AAWG, NIST, ICANN, vendor incident-response and product telemetry (Cisco Talos, Microsoft Threat Intelligence, Barracuda, Abnormal AI, CrowdStrike, Trend Micro, KnowBe4), and breach-disclosure aggregators (BreachSense, PKWARE, UpGuard, SharkStriker, HIPAA Journal).

Publication window

This issue reports what was publicly knowable at the time of writing in early June 2026. Several major Q2 datasets — notably platform-level email telemetry covering April to June and the Q2 incident-response trend reports — were not yet published, and are deliberately not anticipated here. Where this issue makes a forward claim that later data can confirm or falsify, it is marked as such and will be revisited in a subsequent issue.

Handling of conflicting sources

Where two sources conflict, we present both figures and explain the methodological difference rather than selecting one. May 2026 contains the clearest example in the series to date — the Verizon DBIR and Cisco Talos reach opposite conclusions about the leading initial-access vector — and Section 7 is devoted to it. Where a vendor headline figure depends on a proprietary corpus that cannot be externally validated, we mark it explicitly.

vSpam.org corpus scope and limitations

Section 6 reports the **vSpam.org collector corpus**: indicators ingested, normalised, and scored by our automated collection pipeline from cooperating public feeds (URLhaus, ThreatFox, Feodo Tracker). Figures are counted by *first-seen* timestamp within the calendar month and de-duplicated per source and indicator.

Two limitations must be stated plainly. First, **community submission volume was low during this reporting window**; the platform's operator-submission channel was quiescent through Q2 2026 following the March–April bulk-ingestion phase. Section 6 therefore draws on the automated collector corpus, and no community-derived volume claims are made. Second, malware-family labels are inherited from upstream feed taxonomies — predominantly URLhaus — which mix family names, file architectures, and behavioural tags in a single field. We report labels as received and identify which are true family attributions wherever it affects the analysis.

Normalisation

All dates are normalised to UTC. Where a source uses a calendar-quarter label, we state the corresponding months explicitly to avoid ambiguity — "Q1 2026" means January–March 2026, and the datasets carrying that label were published in April and May.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

SECTION 1

Phishing Landscape — May 2026

1.1 Volume and Vector Mix

APWG's Q1 2026 Phishing Activity Trends Report, published 27 May, recorded **971,181 unique phishing attacks** for January–March 2026, up **13.8%** from 853,244 in Q4 2025.[3] That is a return to growth after two quarters of relative flatness, and it aligns with the initial-access picture Cisco Talos reported in April, in which phishing reclaimed the leading position for the first time in three quarters.[7]

Barracuda's 2026 Email Threats Report, released 12 May and built on **3.1 billion messages** of January 2026 telemetry, supplied the volumetric counterpart: **one in three email messages is now malicious or unwanted spam**. Barracuda attributes the increase to two accelerants operating together — AI-assisted social engineering, which raises the quality of each lure, and phishing-as-a-service, which lowers the skill floor required to run a campaign at all.[1] The report also documents a structural change in payload strategy: attackers are moving away from file-based payloads toward **URL-based delivery**, and are embedding QR codes inside document formats the recipient already trusts rather than attaching obviously suspicious files.

Abnormal AI's 2026 Attack Landscape Report, covering roughly 800,000 email attacks across more than 4,600 organisations, provides the composition view: **phishing accounts for 58% of all observed attacks**, making it comfortably the most common threat an employee encounters.[8] Two of its subsidiary findings deserve more attention than they received:

- **21.6% of phishing attacks now use redirect chains** — routing the victim through multiple intermediate URLs to obscure the final destination. Any filter that evaluates only the URL present in the message body is evaluating the wrong artefact in more than a fifth of cases.
- **Vendor email compromise now constitutes 61% of all business email compromise.** The majority of BEC no longer originates from a spoofed or lookalike sender at all; it arrives from a genuine, authenticated mailbox belonging to a real supplier whose account has been taken over. Authentication passes because the mail is authentic.

Taken together, the three corpora describe a threat that is simultaneously higher in volume, higher in quality, and structurally harder to catch with the controls that dominated the previous decade. The table on the following page sets the three side by side, because their differences in scope are the key to reading them correctly.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

DATASET	PUBLISHED	CORPUS	OBSERVATION WINDOW
Barracuda 2026 Email Threats	12 May 2026	3.1B email messages	January 2026
Verizon DBIR 2026	20 May 2026	22,000+ confirmed breaches	1 Nov 2024 – 31 Oct 2025
APWG Q1 2026 Trends	27 May 2026	971,181 unique attacks	January – March 2026
Abnormal AI Attack Landscape	22 Apr 2026	~800k attacks, 4,600+ orgs	July – December 2025

> **READING NOTE**
>
> No two of these four datasets share an observation window, and only one of them (APWG) actually observes the reporting month's own quarter. A defender comparing headline percentages across them is comparing measurements of different years. Section 7 works through the consequences.

1.2 The Telecom Pivot

The single most striking number in APWG's Q1 report is sectoral, not volumetric. The telecom category rose from **5.9% of all observed attacks in Q3 2025 to 33% in Q1 2026**, overtaking SaaS and webmail (20%) to become the most-frequently attacked category outright.[3] Within telecom, the frequency of URL-based phishing rose **75% since Q4 2025** alone — meaning the pivot was still accelerating through the end of the measured window.

A shift of that magnitude in two quarters has a small number of plausible explanations, and the most likely is the least exotic: telecom accounts have become the highest-value identity anchor available to an attacker. A compromised mobile carrier account enables SIM swap and number port-out, which in turn defeats SMS-based one-time passwords across every downstream service the victim uses — banking, email recovery, cryptocurrency exchanges, and enterprise MFA where SMS remains a fallback factor. Telecom is not the target; it is the master key.

That framing has a direct consequence for messaging operators. Telecom-brand impersonation should be scored as an *account-takeover precursor*, not merely as a credential-phishing attempt, and it warrants the same escalation posture normally reserved for financial-brand impersonation. The downstream loss event may occur weeks later, in an entirely different service, and will not be traceable to the original mail.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

APWG's social-media findings point the same direction. Threat volume increased on **every** measured social platform during Q1 2026, with the mix dominated by **impersonation (43.8% of all threats)** and **scams (27.1%)**.[3] Impersonation outweighing scams by better than three to two indicates an ecosystem oriented toward building durable false identities rather than executing one-shot frauds — the infrastructure layer of social engineering rather than its payoff layer.

1.3 Technique Composition

Three techniques defined May's observable campaign traffic, and all three share a structural property: they move the decisive step of the attack out of the email body, where the majority of filtering happens, and into somewhere the filter cannot easily follow.

Redirect chains

At **21.6% of phishing attacks**, redirect chains are now a mainstream technique rather than an evasion nicety.[8] The operational implication is that URL reputation must be evaluated at the *terminal* destination, following redirects at scan time, and that any intermediate hop on a high-reputation domain — a URL shortener, a marketing click-tracker, an open redirect on a legitimate SaaS host — must be treated as unresolved rather than as trustworthy.

QR codes inside trusted document formats

Barracuda documents the migration of QR-code lures from the message body into attached documents of formats recipients routinely open.[1] This defeats two controls at once: the QR code is not visible to body-content scanning, and the document format itself carries no macro or executable payload to detect. The victim's own phone camera completes the delivery on a device that typically sits outside the enterprise mail security perimeter entirely.

Adversary-in-the-middle proxying

Reverse-proxy kits that relay a genuine login page while harvesting the resulting session token remain the standard method for defeating multi-factor authentication. Barracuda notes that such kits **raise account-takeover risk even in fully MFA-enabled environments**, because the credential is never the object of the theft — the authenticated session is.[1]

> **KEY TAKEAWAY**
>
> All three dominant techniques of the month are filter-relocation strategies. None of them require a novel exploit, a fresh CVE, or a signature the defender has not seen. They require only that the decisive step happen somewhere the defender is not looking — a later redirect hop, a phone camera, or an authenticated session.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

1.4 After the Tycoon2FA Takedown

March 2026's coordinated seizure of the Tycoon2FA phishing-as-a-service platform was the largest anti-phishing enforcement action in several years, and May was the first full month in which its effects could in principle be measured. The operation, coordinated by Europol with Microsoft leading the technical disruption, obtained a US court order to seize **330 active domains** underpinning the platform's core infrastructure, while law enforcement in Latvia, Lithuania, Portugal, Poland, Spain, and the United Kingdom seized in-country infrastructure.[5]

The scale of what was removed explains why the action mattered. At the point of disruption, Trend Micro assessed that Tycoon2FA had accumulated a user base of at least **2,000 operators** and had leveraged more than **30,000 phishing domains**. Since its emergence in August 2023 the platform had supported campaigns delivering tens of millions of messages to more than **500,000 organisations per month**.[6] A single service, in other words, accounted for a material share of the world's adversary-in-the-middle credential phishing.

The unresolved question in May was durability. CrowdStrike reported in the period following the takedown that the platform **persists in degraded form** rather than having been eliminated — infrastructure was rebuilt, and a subset of operators continued to transact.[9] That is the historically normal outcome: takedowns of criminal service platforms typically produce a step change in volume followed by partial recovery, and the durable benefit lies in the size of the step and the slope of the recovery rather than in elimination.

> **OPEN QUESTION — TO BE REVISITED**
>
> Public platform-level telemetry quantifying the takedown's effect across April, May, and June had not been published at the time of writing. Our forward assessment is that the step change will prove large — a majority reduction in platform-linked volume — and that the displaced demand will surface in adjacent channels rather than disappearing. We will report against this claim when the Q2 telemetry lands.

One early indicator is available now, from our own corpus. If AiTM credential phishing had simply migrated wholesale to a successor platform, we would expect to see it in the phishing and credential-harvest indicators our collector ingests. We do not, at least not yet; what May shows instead is a fake-update and browser-lure cluster at its highest level of the year. Section 6 reports the detail.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

SECTION 2

Malware and Ransomware

2.1 Volume and Group Dynamics

BreachSense recorded **646 victims posted across 61 active leak sites** in May 2026, spanning 73 countries and 59 industries — a **16% decline** from April's 772 and the lowest monthly total of the year.[4] More significant than the drop itself is what it does to the annual trend line: at 3,583 victims through five months, May was **the first month of 2026 to fall below the corresponding 2025 pace**, ending four consecutive months in which the year ran ahead of its predecessor.

The year-to-date figure still annualises to roughly **8,600 victims**, an 18% increase over the 7,307 recorded across the whole of 2025 — so the direction of travel remains upward. A single month below trend is not an inflection, and it would be a mistake to read it as one.

Ransomware Leak-Site Postings — 2026 to Date

May broke a four-month run above the 2025 pace; group leadership did not change

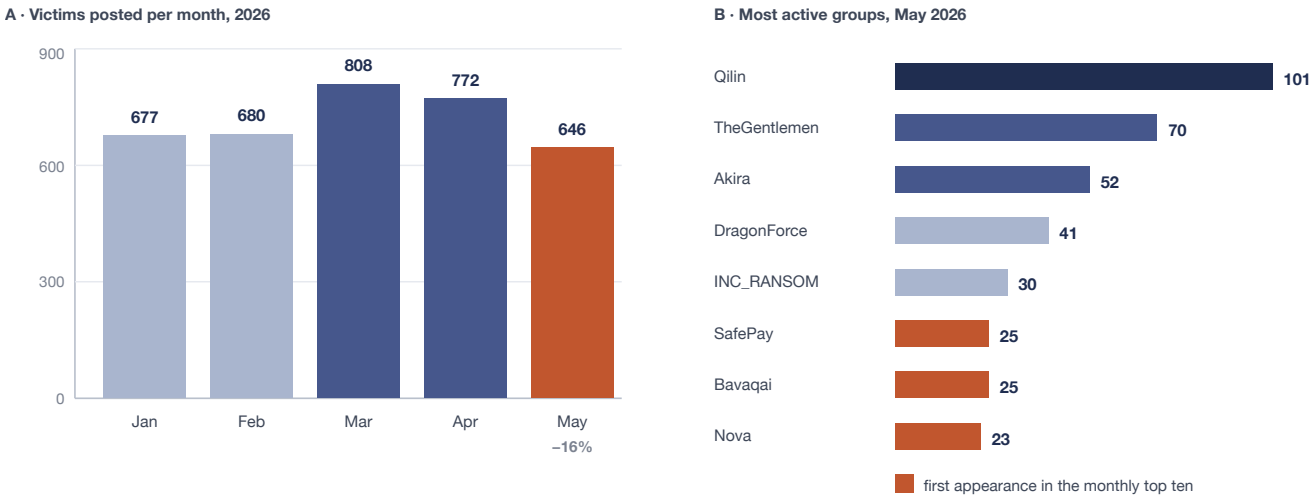

Source: BreachSense May 2026 Ransomware Report; BreachSense monthly series, January–May 2026.

FIGURE 2 *Monthly victim counts and May's group ranking. Three of the eight most active groups — SafePay, Bavaqai, and Nova — entered the top ten for the first time in the same month, while Qilin held first place for a fifth consecutive month.*

Panel B is the more informative half. A top-eight list in which three names are new arrivals, and in which the three of them together post 73 victims — more than the second-placed group — describes a market absorbing displaced affiliates rather than one consolidating. Section 2.2 takes up what that means for defenders.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

2.2 Ecosystem Fragmentation

The more durable story in May's data is structural. **Three groups — SafePay, Bavaqai, and Nova — debuted in the monthly top ten simultaneously**, which BreachSense characterises as an ecosystem that "keeps splintering instead of settling around a few big players".[4] Sixty-one distinct groups posted victims in a month when total volume fell 16%; the count of active operations is holding up considerably better than the count of victims.

That combination — fewer victims, similar numbers of groups — is what a market looks like after a period of successful law-enforcement pressure on its largest participants. Affiliates disperse rather than retire. The consequence for defenders is unwelcome: negotiation precedent, leak-site behaviour, and encryption tooling all become less predictable as the long tail lengthens, and threat-intelligence keyed to a handful of named brands covers a steadily smaller share of the actual risk.

Qilin's stability against that background is notable. It held first place for a **fifth consecutive month** at 101 victims — the only real constant in a ranking that otherwise reshuffled substantially. LockBit, by contrast, posted 18 victims, down from 39 in April, continuing its decline from pre-disruption levels.

GEOGRAPHY	VICTIMS	SHARE	INDUSTRY	VICTIMS	Δ VS APR
United States	**254**	39.3%	Manufacturing	**58**	+8
United Kingdom	**36**	5.6%	Healthcare	**54**	−10
Canada	**31**	4.8%	Construction	**42**	+5
Germany	**29**	4.5%	Consumer goods	**38**	+1
Spain	**21**	3.3%	Finance	**31**	+3
Italy	**19**	2.9%	Technology	**31**	−25
Australia	**17**	2.6%	Legal	**27**	−13
Mexico	**17**	2.6%	IT services	**26**	−5
France	**16**	2.5%	Engineering	**23**	−6
India	**11**	1.7%	Government	**22**	−6

Manufacturing reclaimed the most-targeted position at 58 victims, displacing healthcare, which fell 10 to 54. The sharpest single move was technology, down 25 to 31 — a decline steep enough that it accounts for roughly a fifth of the month's total drop on its own. US share, at 39.3%, held approximately level with April's 39%, so the month's contraction was broadly distributed rather than concentrated in one geography.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

2.3 Loader and Delivery Observations

Three delivery patterns carried through May from the prior quarter's reporting, all of them consistent with the technique composition described in Section 1.3 — and all of them structured so that the decisive step happens outside the mail stream.

ClickFix and the self-inflicted payload

ClickFix-style social engineering induces the user to paste an attacker-supplied command into a Run dialog, terminal, or browser console — typically framed as a fix for a fabricated rendering error, a CAPTCHA that will not clear, or a document that will not open. The attack begins in email or on a compromised web property, but the payload is delivered by the victim's own hands. Attachment scanning has nothing to scan and URL reputation has nothing to resolve; the malicious artefact never traverses the mail path at all.

This is the technique behind the fake-update cluster that dominates our own May corpus, and it is the reason that cluster is worth watching as a proxy for the health of the wider social-engineering ecosystem. Section 6.3 reports the detail.

OAuth abuse for staging and exfiltration

Tokens granted to legitimate cloud applications continue to serve as both a data-exfiltration channel and a delivery channel for follow-on tooling. The defensive difficulty is that the traffic is genuinely authorised: it originates from a sanctioned application, carries a valid token, and terminates at a first-party API endpoint. Detection has to rest on behavioural baselines — volume, timing, and scope of access — rather than on any property of the request itself.

Trusted-platform hosting

Landing pages and second-stage payloads increasingly sit on developer and automation platforms rather than on attacker-registered domains. This shifts the takedown burden from registrars, who now operate under contractually specified response windows, to platform abuse teams, who largely do not — and it renders domain-level blocking substantially less effective, because the hostname a defender would block is shared with a large volume of legitimate traffic.

> **KEY TAKEAWAY**
>
> A 16% monthly decline in leak-site postings alongside a stable count of active groups is not a retreat — it is redistribution. Intelligence programmes organised around the top five ransomware brands are covering a shrinking fraction of the real exposure, and the delivery techniques feeding those operations are converging on methods that leave no artefact in the mail stream to inspect.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

SECTION 3

Email Authentication and Spoofing

3.1 The Enforcement Gap at Mid-Year

The benchmark dataset for 2026 remains EasyDMARC's adoption report, covering 1.8 million domains plus the Fortune 500 and Inc. 5000. Its headline is that DMARC adoption crossed a symbolic threshold — **52.1% of measured domains now publish a DMARC record**, 937,931 in absolute terms, up from 47.7% a year earlier.[12] Adoption, however, is the wrong metric, and the report is unusually direct about saying so.

The number that matters is the share of domains where DMARC is actually doing something. Of the 937,931 domains publishing a record, **411,935 are at an enforcing policy** (`p=quarantine` or `p=reject`) — and only **about 9% of domains combine an enforcement policy with reporting**, which is the configuration required both to block spoofed mail and to retain visibility into what is being sent under the domain's name.[12]

METRIC	VALUE	INTERPRETATION
Domains publishing any DMARC record	**52.1%**	Up from 47.7% year-on-year; the widely quoted headline
Domains at `p=quarantine` or `p=reject`	**411,935**	~44% of record-publishing domains; the real protection base
Domains with enforcement *and* reporting	**≈9%**	The only configuration that both blocks and shows you what it blocked
Fortune 500 with a DMARC record	**475 / 500**	95% adoption; more than 80% at enforcement
Inc. 5000 at monitor-only	**> half**	The gap is a company-size gap, not an awareness gap

The Fortune 500 and Inc. 5000 comparison is the most useful pair of numbers in the report. Enterprises with dedicated messaging teams have substantially solved DMARC: **475 of 500 publish a record and more than 80% enforce**. Mid-market companies — the Inc. 5000 — remain majority monitor-only. The bottleneck is neither awareness nor tooling cost; it is the labour of enumerating every legitimate sending source before a policy can safely be tightened, and that labour does not scale down to organisations without a dedicated owner.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

3.2 Where the Gap Actually Lives

DMARC: Records Published vs. Protection Delivered

Adoption is a majority; effective, observable enforcement is a small minority

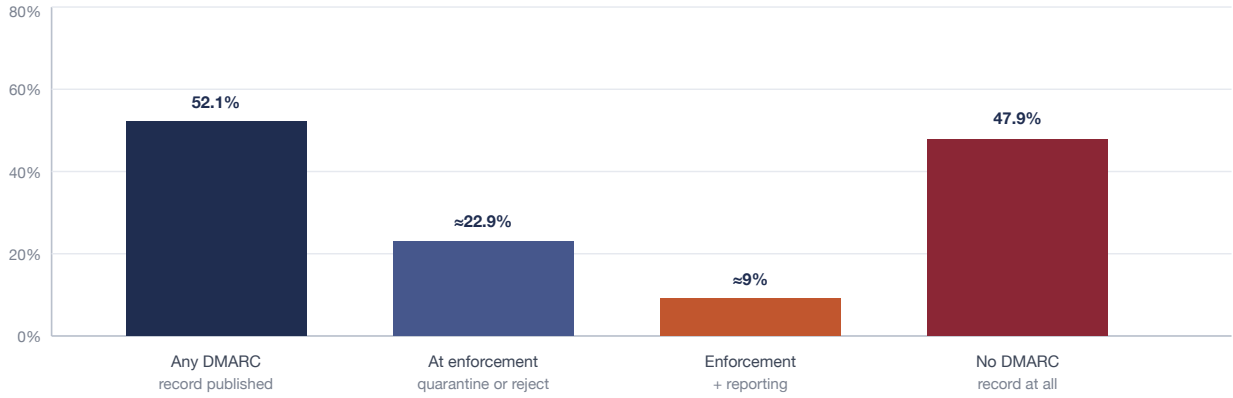

All bars expressed as a share of the full 1.8M-domain corpus.

Source: EasyDMARC 2026 DMARC Adoption & Enforcement Report (1.8M domains, Fortune 500, Inc. 5000).

FIGURE 3 *The adoption headline (52.1%) and the protection reality diverge by roughly a factor of six. Fewer than one domain in ten is configured so that spoofed mail is both blocked and visible to the domain owner.*

The practical reading is that the industry has spent a decade successfully persuading domain owners to publish a record, and has not yet solved the harder problem of getting them to finish. A record at `p=none` provides no protection against spoofing whatsoever; it provides telemetry, which is valuable only if somebody reads it. Roughly nine domains in ten sit in one of the two failure states — no record at all, or a record that does not enforce.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

The transition from monitor to enforce is a project with a predictable shape: enumerate every legitimate sending source, align SPF and DKIM for each, verify against aggregate reports, then tighten policy in stages. Managed platforms exist precisely to compress that work — DDMARC reports that most teams using its policy simulator and ARC-chain monitoring move from `p=none` to `p=reject` in **four to six weeks**, against the twelve-plus months typical of unaided transitions.[13] Valimail, dmarcian, EasyDMARC, and PowerDMARC operate in the same category. We note that this supplier-reported figure has not been independently audited.

3.3 Spoofing Beyond Header Forgery

As DMARC enforcement closes the header-from spoofing channel for the domains that do finish the transition, attackers move to the channels DMARC does not cover. Three are prominent:

- **Display-name impersonation.** The `From:` display name is unauthenticated free text. A message from an unremarkable domain, correctly signed and correctly aligned, can present as any brand or colleague the attacker chooses, and will pass every authentication check cleanly.
- **Lookalike and homoglyph domains.** Freshly registered domains — frequently less than 72 hours old at time of use — are too new for reputation systems to have rated, and the attacker controls them outright, so SPF, DKIM, and DMARC all pass on a domain the recipient has simply never seen before.
- **Authenticated compromise.** The dominant case, and the one no authentication control can address: mail sent from a genuine, authenticated mailbox whose owner has been compromised. This is the mechanism behind supplier email compromise reaching **61% of all BEC**.[8]

> **KEY TAKEAWAY**
>
> Email authentication solves sender-domain forgery, and it solves it well. It does not solve identity deception, and the majority of today's high-value email fraud is identity deception conducted over correctly authenticated channels. Both controls are necessary; only one of them is widely deployed.

The operational implication for messaging operators is that authentication results should feed reputation scoring as one input among several rather than function as a gate. A message that passes DMARC from a domain registered yesterday, or one that passes DMARC while presenting a display name matching a well-known brand the sending domain has no relationship with, should score *worse* — not better — than an unauthenticated message from a long-established correspondent.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

SECTION 4

DNS Abuse and Infrastructure

4.1 Two Years of Contractual Enforcement

ICANN's DNS Abuse mitigation requirements reached their two-year enforcement anniversary in April 2026, and the compliance data published alongside it is the most concrete measure yet of whether contractual obligation changes registrar behaviour. Between 5 April 2024 and 5 April 2026, ICANN Contractual Compliance launched **nearly 530 investigations** under the DNS Abuse requirements, of which **approximately 66% resulted in the contracted party taking action to stop the abuse**.[14]

A two-thirds action rate is a genuine result, and worth stating plainly against the prevailing scepticism about ICANN enforcement. It is also a ceiling worth understanding: the remaining third represents cases where the registrar contested the evidence, the domain was already gone, or the report did not meet the actionability threshold. The mechanism works on well-evidenced reports against responsive parties — which is precisely the population that was least of a problem to begin with.

4.2 The 2026 gTLD Round Raises the Floor

The contractual environment tightens further with the next round of new gTLD applications. Registry agreements on offer in the 2026 round require operators to deploy **automated abuse-detection tooling**, to participate in cross-industry threat-intelligence sharing including feeds from CISA-aligned bodies and APWG, and to respond to verified abuse reports **within contractually specified windows**. Registries missing those thresholds face escalating compliance review up to contract termination.[15]

Under the revised Registrar Accreditation Agreement, registrars must maintain accessible abuse reporting mechanisms and ensure reports are effectively processed — with the substantive change being an explicit **duty to act** when presented with actionable evidence of DNS abuse involving a sponsored domain.[15] The shift from "must receive reports" to "must act on evidence" is the most consequential wording change in the registrar-facing contract in a decade.

4.3 Measurement After DAAR

ICANN Domain Metrica, which replaced the retired DAAR platform in September 2025, now supplies registries with daily feeds through MoSAPI alongside service-level performance data.[16] The practical effect for researchers is that long-run trend lines crossing the September 2025 boundary remain discontinuous, and comparisons across it should still be treated with caution — a caveat that will persist until Metrica has accumulated several more quarters of history.

> **WHAT OUR OWN DATA ADDS**
>
> Contractual response windows govern how fast a registrar acts once a domain is reported. They say nothing about how fast domains are created. Our May corpus captured a bulk-registration event concentrated in four novelty gTLDs — exactly the pattern that registry-level automated detection is meant to catch at creation time rather than at report time. Section 6.2 reports it in full.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

SECTION 5

Major Incidents — May 2026

DISCLOSED	ORGANISATION	NATURE	REPORTED SCALE
1 May	**Instructure (Canvas)**	Learning-platform data exposure	275M users / ~8,809 institutions *(actor claim)*
4 May	**Trellix**	Source-code repository access	Partial repository; scope not quantified
May	**Carnival Corporation**	Unauthorised access (incident 10 Apr)	5,995,277 people

5.1 Instructure / Canvas — and How to Cite It

The largest claimed exposure of the month was against Instructure, operator of the Canvas learning management system, disclosed on 1 May. The attackers claimed **275 million users across approximately 8,809 institutions**. Instructure has **not** independently confirmed that 275 million unique people were affected.[10]

We report the figure as a claim because that is what it is, and because actor-supplied record counts are systematically inflated. Leak-site totals routinely include duplicate rows, historical archives long since deleted, service accounts, and non-personal records. Where a breach touches an education platform, per-institution rosters frequently overlap heavily across academic years. A number sourced from the party with the strongest incentive to maximise it is not a measurement, and the practice of laundering such numbers into industry statistics through repetition is among the more persistent hygiene problems in this field.

5.2 Carnival — the Disclosure Lag as an Attack Window

The Carnival timeline is worth internalising precisely because the response was competent: the incident occurred on **10 April**, was identified on **14 April**, had its scope confirmed on **22 April**, and reached public disclosure in **May**, affecting **5,995,277 people**.[11] Four days to detection and eight more to scoping is a good result by prevailing standards.

Even so, roughly a month elapsed between compromise and public knowledge. That interval is the period of maximum value to an attacker holding mailbox access, because every correspondent still trusts the sender and no downstream party has any reason to apply additional scrutiny. It is the operational basis for supplier email compromise now constituting **61% of all BEC**.[8]

> **THE COMMON THREAD**
>
> May's incidents share a property with April's: the compromised party and the exposed party are not the same organisation. A platform breach exposes thousands of institutions; a supplier breach exposes every customer that supplier serves. Inbound trust scoring must therefore account for third-party identity, and disclosure lag must be treated as an active exposure window rather than an administrative delay.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

SECTION 6

vSpam Researcher Team Analysis

This section reports directly from the vSpam.org collector corpus. It is new to the series from this issue, and its purpose is to contribute original observation rather than to restate the month's published research. Where our data agrees with the public datasets we say so; where it diverges, the divergence is the finding.

6.1 Corpus and Scope

30,039
INDICATORS FIRST SEEN IN MAY

9,075
DOMAIN-TYPE INDICATORS

450
DISTINCT FAMILY / TAG LABELS

38.1%
DOMAIN IOCS IN FOUR GTLDS

May 2026 was the first full month at steady state for the vSpam.org collector after its April scale-up, which makes it the first month whose composition can be read as a measurement rather than as a ramp artefact. The pipeline ingested **30,039 indicators** by first-seen timestamp — 29,856 from URLhaus and 183 from ThreatFox — carrying 450 distinct family and tag labels between them.[20]

INDICATOR TYPE	CLASSIFICATION	MAY 2026	SHARE
URL	Malware distribution	**20,801**	69.2%
Domain	Malware distribution	**9,055**	30.1%
IP	Botnet command & control	**155**	0.5%
Domain	Botnet command & control	**11**	<0.1%
Domain	Payload delivery	**9**	<0.1%
URL	Botnet command & control	**8**	<0.1%

The distribution is worth stating explicitly because it bounds what this corpus can and cannot say. **More than 99% of May's indicators describe malware distribution**, split roughly seven-to-three between URL-level and domain-level observations. Command-and-control infrastructure is present but marginal at 174 indicators. This is a corpus about *where malicious content is served from*, not about who is operating the botnets that consume it — and the analysis below is scoped accordingly.

> **SCOPE STATEMENT**
>
> Community submission volume was low through this reporting window; the operator-submission channel was quiescent following the March–April bulk-ingestion phase. Nothing in this section derives from community reporting, and no claim is made about community-reported volume. Family labels are inherited from upstream feed taxonomies which mix true family names with file architectures and behavioural tags; we identify which is which below.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

6.2 The Novelty-gTLD Burst

The most distinctive feature of May's corpus is a concentration in the domain namespace that we have not observed before or since. Of **9,075 domain-type indicators**, **3,459 — 38.1% — sat in just four top-level domains**: `.lat`, `.surf`, `.pics`, and `.garden`. Legacy `.com` placed sixth, with 172 indicators.

Top-Level Domain Concentration — vSpam.org Collector, May 2026

Four novelty gTLDs carried 38.1% of the month's domain indicators; .com placed sixth

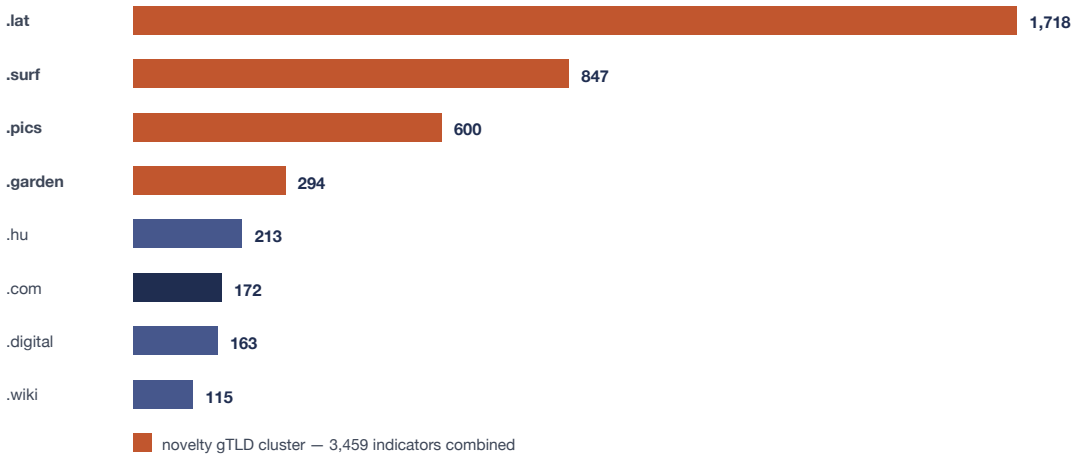

Source: vSpam.org collector corpus, domain-type indicators by first-seen timestamp, 1–31 May 2026.

FIGURE 4 *Domain indicators by top-level domain. A four-TLD cluster carrying more than a third of the month's domain observations, with legacy `.com` an order of magnitude below the leader, is the signature of bulk registration rather than of opportunistic acquisition.*

A distribution this skewed does not arise from attackers independently choosing registrars. It is what a small number of bulk-registration events look like from downstream: one actor, or a small group, acquiring domains in volume where they are cheapest and least scrutinised, then burning through them over a period of weeks.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

The finding has a direct bearing on Section 4. ICANN's tightened contracts govern how quickly a registrar acts once a domain has been reported. They do not govern the rate at which domains enter the zone. An actor able to register several thousand domains across four TLDs inside a month is operating entirely upstream of the abuse-report mechanism, and will remain so until registry-side automated detection acts on *registration patterns* — velocity, shared registrant fingerprints, sequential nameserver assignment — rather than on downstream reports. That capability is written into the 2026 registry agreements. This is precisely the case it needs to catch.

6.3 The Fake-Update Cluster Peaks

May's family-label distribution was dominated by a single behavioural cluster. **ClearFake accounted for 9,632 indicators** — nearly a third of the month's entire corpus, and more than the next four true family labels combined. ClearFake is the fake-browser-update lure family that underpins much of the ClickFix ecosystem described in Section 2.3: a compromised or attacker-controlled page presents a fabricated update prompt or verification step, and the victim is walked into executing the payload themselves.

Weekly Ingestion — Fake-Update Cluster vs. IoT Botnet Baseline

vSpam.org collector, weeks commencing 27 April to 25 May 2026

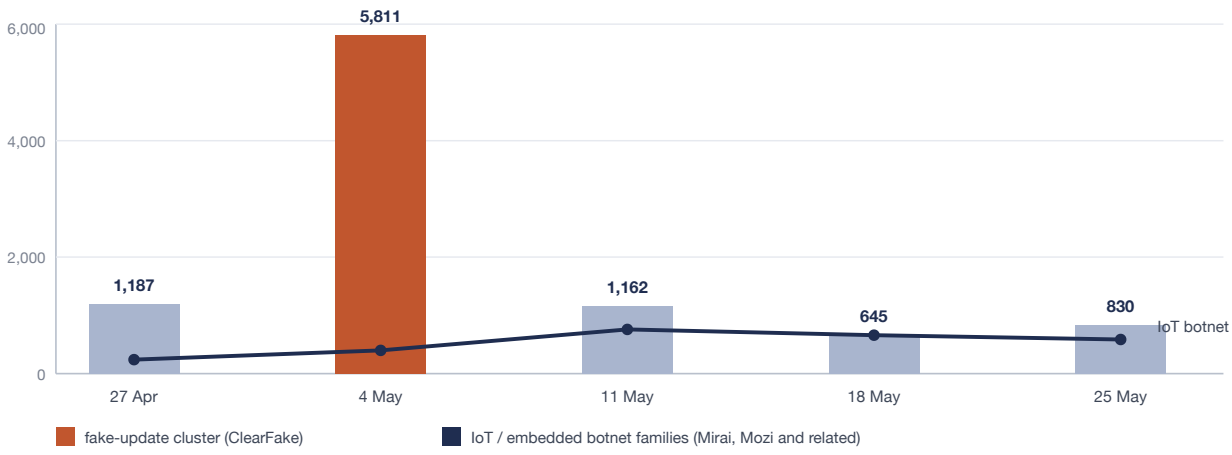

Source: vSpam.org collector corpus, indicators grouped by upstream family label, weeks commencing 27 April – 25 May 2026.

FIGURE 5 *The fake-update cluster spiked to 5,811 indicators in the week commencing 4 May, then fell 89% over the following fortnight. The IoT botnet baseline, by contrast, varied within a narrow band throughout — the difference between campaign traffic and infrastructure traffic.*

Note the relationship between the two series in the week commencing 11 May: the fake-update cluster had already fallen 80% from its peak while the IoT baseline reached its own monthly high of 758. The two are uncorrelated, which is what one would expect if they represent genuinely separate operations sharing a collection pipeline rather than a single actor's output.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

Figure 5 shows two fundamentally different kinds of traffic sharing one pipeline. The fake-update cluster is **campaign traffic**: it spikes to 5,811 indicators in the week commencing 4 May, then collapses 89% to 645 within a fortnight. The IoT and embedded botnet families — Mirai, Mozi, and their relatives — are **infrastructure traffic**, moving between 239 and 758 all month with no comparable excursion.

The distinction matters for anyone building alerting on feed volume. A threefold week-on-week move in campaign families is normal and carries little information on its own; the same move in the infrastructure baseline would be highly significant. Treating aggregate feed volume as a single signal conflates the two and produces an alert stream dominated by ordinary campaign churn.

6.4 Reading Our Corpus Against the Public Datasets

Three points of contact are worth drawing out between what we observe and what the month's published research reports.

On the Tycoon2FA question

Section 1.4 left open whether demand displaced by the March takedown had migrated to a successor platform. Our corpus offers a partial negative: had AiTM credential phishing simply relocated, we would expect elevated credential-harvest and phishing indicators. What we see instead is the fake-update and browser-lure cluster at its annual peak. That is consistent with displacement *across technique* rather than across platform — operators moving from a proxied-login model to a user-executes-the-payload model, which requires no phishing infrastructure at all. We flag this as an interpretation, not a finding: our corpus observes malware distribution, not credential-harvest campaigns, so its silence on the latter is weak evidence.

On the file-to-URL shift

Barracuda's finding that attackers are migrating from file-based payloads to URL-based delivery is strongly corroborated here.[1] URL-type indicators outnumber domain-type indicators better than two to one in our May corpus, and account for **69.2% of everything ingested**. The unit of malicious infrastructure is now the individual URL, not the domain — which has direct consequences for any blocklist built at domain granularity.

On sector targeting

We cannot corroborate APWG's telecom finding, and we say so explicitly. Our corpus records malware-distribution infrastructure, which carries no reliable brand-target attribution; APWG's records phishing attacks against identified brands. The two measure different objects. A reader should not treat the absence of a telecom signal here as evidence against APWG's.

> **RESEARCHER TEAM ASSESSMENT**
>
> May's corpus describes an ecosystem in which the domain is becoming a disposable, bulk- acquired commodity and the URL is becoming the meaningful unit of observation. Defensive architectures that treat the domain as the atomic reputation object — including most DNS- layer blocking — are measuring a unit the attacker has already stopped optimising for.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

SECTION 7

Cross-Cutting Analysis — Reconciling the Divergence

The defining analytical problem of May 2026 is that its two most authoritative datasets appear to contradict each other on the most basic question a defender asks: what is the leading way attackers get in?

Cisco Talos, reporting on Q1 2026 incident-response engagements, found phishing back at the top of the initial-access ranking for the first time in three quarters, accounting for more than a third of engagements where the vector could be determined.[7] The Verizon DBIR, published a month later on a corpus six times larger, placed **vulnerability exploitation** as the dominant initial-access vector, with social engineering at 16% of breaches.[2] Both are correct. The reconciliation lies in four methodological differences, and understanding them is more valuable than either headline.

Why the Month's Two Flagship Datasets Disagree

Observation windows, plotted against the period each purports to describe

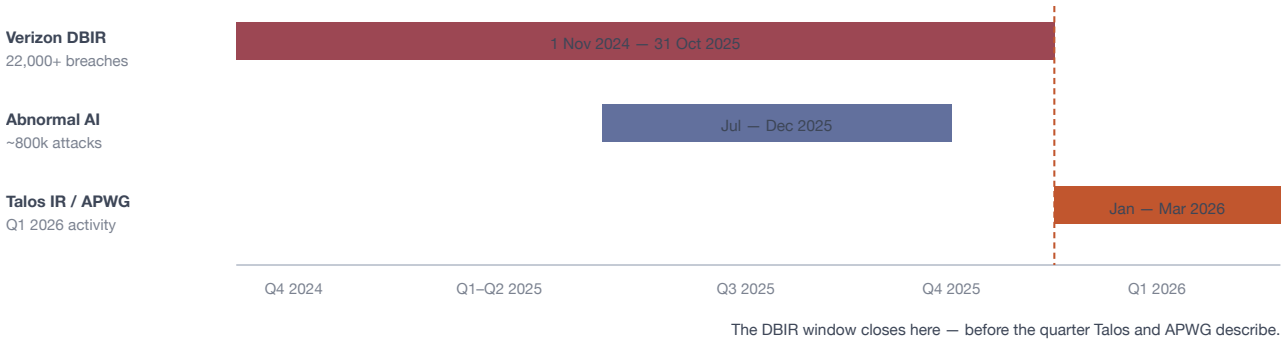

Sources: Verizon DBIR 2026; Cisco Talos IR Trends Q1 2026; APWG Q1 2026; Abnormal AI 2026 Attack Landscape Report.

FIGURE 6 *No overlap exists between the DBIR's observation window and the quarter that Talos and APWG report on. The datasets are not contradicting each other about the same period; they are describing different periods.*

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

7.1 The Four Differences

1. The observation windows do not overlap

The DBIR's incidents occurred between **1 November 2024 and 31 October 2025**. The Talos and APWG findings describe **January to March 2026**. The DBIR's window closes three months before the quarter in which phishing rebounded. It cannot report a trend that had not yet happened when its data collection stopped, and it does not claim to.

2. The units of analysis differ

The DBIR counts **confirmed breaches** — incidents with verified data compromise. APWG counts **attacks**, whether or not they succeeded. Talos counts **engagements** its responders were retained for. A vector that generates enormous attempt volume but converts poorly will dominate APWG and under-index in the DBIR. Phishing is exactly such a vector.

3. Selection bias runs in opposite directions

Incident-response firms are engaged for incidents an organisation could not handle alone, which over-samples sophisticated intrusions. Breach-disclosure corpora over-sample incidents that triggered a regulatory notification threshold, which skews toward large volumes of regulated personal data. Neither is a random sample of reality, and the biases do not point the same way.

4. The vector taxonomies are not aligned

A campaign that begins with a phishing email, harvests a session token through an adversary-in-the-middle proxy, and proceeds through a legitimate authenticated session can reasonably be classified as phishing, as credential abuse, or as valid-account use — and different datasets classify it differently. The 2025–2026 rise of AiTM tooling has made this ambiguity structural rather than occasional.

> **THE SYNTHESIS**
>
> Across the 2024–2025 breach population, exploitation converted to confirmed breaches more efficiently than phishing did. Across Q1 2026 attempt volume and IR engagements, phishing leads. Both statements can be true simultaneously, and a defender should act on both: patch the internet-facing estate on the exploitation timeline, and treat the human layer as the highest-volume attack surface. The datasets are not in conflict. The summaries of them are.

7.2 What This Means for Reading Q2 Data

The same reconciliation discipline will be needed when Q2 2026 datasets publish over the coming weeks. Platform-level email telemetry covering April to June, and the Q2 incident-response trend reports, will arrive with windows that again do not align with one another. We commit to three practices in reporting them: state each dataset's window before quoting its headline, never compare percentages across datasets with different units of analysis, and revisit the forward claims made in this issue — specifically the Tycoon2FA displacement assessment in Section 1.4 and the technique-displacement interpretation in Section 6.4 — against whatever the new data shows, including where it contradicts us.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

SECTION 8

Standards and Regulatory Developments

8.1 ICANN

The two-year enforcement review of the DNS Abuse mitigation requirements — nearly 530 investigations, roughly 66% resulting in remedial action — is the substantive regulatory datapoint of the period, and is treated in Section 4.1.[14] The 2026 gTLD round's registry agreements and the revised Registrar Accreditation Agreement raise the contractual floor further, with the duty to act on actionable evidence being the material change.[15]

8.2 M3AAWG

M3AAWG's 67th General Meeting was scheduled for **8–11 June 2026 in Montréal**, immediately following this reporting period. All four Priority Committees were slated to meet in person to set roadmaps against AI-driven messaging abuse, identity, and online safety. We will report outcomes in the next issue.[17]

8.3 NIST

NIST's email-security framework — SP 800-177 Rev. 1 (*Trustworthy Email*) and SP 800-45 Ver. 2 (*Guidelines on Electronic Mail Security*) — remains the federal anchor, recommending the SPF / DKIM / DMARC / TLS / S-MIME stack.[18] No revision was issued in May 2026. Federal enforcement of the binding operational directives on DMARC `p=reject` and transport security continues to expand the addressable enforcement surface, and remains the single largest driver of enforcement-rate improvement in the US domain population.

8.4 CISA

CISA's operational tempo continued through May with regular additions to the Known Exploited Vulnerabilities catalogue under Binding Operational Directive 26-04, which sets the minimum remediation standard for federal civilian agencies.[19] No May 2026 advisory was dedicated exclusively to phishing or email security at alert level.

8.5 The Regulatory Pattern Worth Noting

Across ICANN, NIST, and CISA, the direction of travel in 2026 is consistent: from *disclosure* obligations toward *action* obligations with defined windows. Registrars must act on evidence, not merely receive reports. Federal agencies must remediate catalogued vulnerabilities on a clock, not merely track them. This is a meaningful shift, and its limitation is equally clear — every one of these regimes acts on threats after they have been identified and reported. None of them changes the economics of creating the threat in the first place, which is where the bulk-registration pattern documented in Section 6.2 lives.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

SECTION 9

Outlook and Recommendations

9.1 Forecast (June – August 2026)

1. **Q2 2026 phishing volume will exceed Q1's 971,181.** The Q1 rise was recorded against a quarter in which the Tycoon2FA platform was still fully operational for two of three months; the displacement of its operators into other tooling is more likely to redistribute volume than to remove it. We expect APWG's Q2 report to land in the **1.0M–1.15M** range.
2. **The telecom concentration will partially revert.** A single sector holding 33% of observed attacks is an unstable equilibrium; campaign operators rotate targets as detection catches up. We expect telecom to remain the leading category in Q2 but to fall back toward the low-to-mid twenties as a percentage.
3. **Ransomware will resume its upward trend after May's dip.** One month below the prior year's pace, against a stable count of active groups and a lengthening long tail, reads as reporting noise rather than as a turn. We expect June and July to return above the 2025 comparison.
4. **Post-takedown telemetry will show a large step change with partial recovery.** Stated in Section 1.4 and repeated here so it can be scored: we expect published Q2 telemetry to show a majority reduction in Tycoon2FA-linked volume, with displaced activity appearing in adjacent channels rather than disappearing.

9.2 Recommendations — Messaging Operators and ISPs

- **Score telecom-brand impersonation as an account-takeover precursor.** With telecom at 33% of observed attacks and carrier accounts functioning as the master key to SMS-based recovery across every downstream service, these lures warrant the escalation posture normally reserved for financial-brand impersonation.
- **Resolve URLs to their terminal destination at scan time.** With 21.6% of phishing using redirect chains, evaluating only the URL present in the message body means evaluating the wrong artefact in more than one message in five. Treat intermediate hops on high-reputation hosts as unresolved, not as trusted.
- **Move blocklist granularity from domain to URL.** URL-type indicators outnumber domain-type better than two to one in our corpus. Where the malicious content sits on a shared or trusted host, the domain is not a usable blocking unit.
- **Weight `p=none` domains down in reputation scoring.** With roughly nine domains in ten either publishing no record or publishing one that does not enforce, reputation pressure is the only lever the receiving side controls.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

9.3 Recommendations — Enterprise Defenders

— **Treat third-party mailboxes as the primary BEC channel.** At 61% of all BEC, supplier email compromise is no longer an edge case, and no authentication control will stop it — the mail is genuinely authenticated. Out-of-band verification for payment-detail changes is the only control that addresses the actual mechanism.

— **Adopt phishing-resistant MFA for high-value accounts.** Adversary-in-the-middle proxies defeat TOTP and push-based factors by design, because the session token is the target rather than the credential. Device-bound credentials — FIDO2, passkeys — are the control that changes the outcome.

— **Inventory and review OAuth grants.** Tokens issued to third-party applications function as both an exfiltration channel and a delivery channel, and the traffic is genuinely authorised. Scope review is the only point at which this is cheap to control.

— **Treat ClickFix as a phishing-class threat, not an endpoint-class one.** The attack begins in email or on the web and is executed by the user; defences that wait for endpoint detection are operating after the decisive step has already happened.

— **Plan on a month of exposure after a supplier breach.** The Carnival timeline — competent by prevailing standards — still ran roughly a month from compromise to public disclosure. Assume that window exists for every supplier you correspond with.

9.4 Recommendations — Registrars and Registries

— **Detect on registration velocity, not only on abuse reports.** Section 6.2 documents several thousand domains concentrated in four TLDs inside one month. That pattern is visible at creation time — shared registrant fingerprints, sequential nameserver assignment, acquisition velocity — and invisible to any process that begins when the first report arrives.

— **Front-load abuse automation ahead of the 2026 round obligations.** The new registry agreements mandate automated detection tooling and contractual response windows. Registries building that capability reactively will be doing so under compliance review.

— **Apply heightened scrutiny to low-cost novelty gTLDs.** Our May data shows legacy `.com` an order of magnitude below the leading novelty TLD in malicious-domain count. Price and scrutiny are the variables attackers optimise; registry operators control both.

> **WHERE WE MAY BE WRONG**
>
> The forecasts in Section 9.1 are falsifiable and we intend them to be scored. The interpretation most likely to prove wrong is the technique-displacement reading in Section 6.4: our corpus observes malware distribution rather than credential harvesting, so its silence on AiTM migration is weak evidence, and Q2 platform telemetry may well show that displaced Tycoon2FA demand went to a successor platform after all.

ISSN 2026-VSPAM-005 · VOL. 2026 · ISSUE 5

APPENDIX A

Tools and Suppliers Referenced

This report mentions multiple commercial suppliers as illustrative examples of categories of tooling. Inclusion is not an endorsement, and the list is non-exhaustive. Capability, pricing, and regional availability vary substantially; readers should evaluate against their own requirements.

CATEGORY	EXAMPLES CITED IN THIS ISSUE
Managed DMARC and email authentication	DDMARC (research partner — see editorial disclosure), Valimail, dmarcian, EasyDMARC, PowerDMARC
Email security and threat telemetry	Barracuda, Abnormal AI, Microsoft Threat Intelligence, Proofpoint
Threat intelligence and incident response	Cisco Talos, CrowdStrike, Trend Micro, Verizon Threat Research Advisory Center
Phishing simulation and training	KnowBe4, Hoxhunt
Breach and ransomware tracking	BreachSense, PKWARE, UpGuard, SharkStriker
Blocklist and reputation	Spamhaus, SURBL, URIBL, vSpam.org
Open threat-intelligence feeds	URLhaus, ThreatFox, Feodo Tracker (abuse.ch)
DNS abuse measurement	ICANN Domain Metrica, NetBeacon Institute, DomainTools

Data availability

The vSpam.org figures reported in Section 6 are derived from the public collector corpus. Aggregate statistics are available through the vSpam.org public API, and researchers seeking the underlying per-indicator data for replication may request access at **research@vspam.org**. We will supply the exact queries used to produce every figure in Section 6 on request.

REFERENCES

References

1. Barracuda Networks (12 May 2026). *2026 Email Threats Report.* barracuda.com/reports/2026-email-threats-report
2. Verizon Business (20 May 2026). *2026 Data Breach Investigations Report.* verizon.com/business/resources/reports/dbir/
3. APWG (27 May 2026). *Phishing Activity Trends Report, 1st Quarter 2026.* docs.apwg.org/reports/apwg_trends_report_q1_2026.pdf
4. BreachSense (June 2026). *May 2026 Ransomware Report: 646 Victims, 61 Groups.* breachsense.com/ransomware-reports/may-2026/
5. Europol (2026). *Global phishing-as-a-service platform taken down in coordinated public-private action.* europol.europa.eu
6. Trend Micro (March 2026). *Europol, Microsoft, TrendAI and Collaborators Halt Tycoon 2FA Operations.* trendmicro.com/en_us/research/26/c/tycoon2fa-takedown.html
7. Cisco Talos (22 April 2026). *IR Trends Q1 2026: Phishing reemerges as top initial access vector.* blog.talosintelligence.com/ir-trends-q1-2026/
8. Abnormal AI (22 April 2026). *2026 Attack Landscape Report.* abnormal.ai/blog/2026-attack-landscape-report-phishing
9. CrowdStrike (2026). *Tycoon2FA Phishing-as-a-Service Platform Persists Following Takedown.* crowdstrike.com
10. Wikipedia / contemporaneous reporting (May 2026). *2026 Canvas data breach.* See also Instructure public statements, 1 May 2026.
11. Carnival Corporation breach notification (May 2026), as compiled by BrightDefense, *List of Recent Data Breaches in 2026.* brightdefense.com/resources/recent-data-breaches/
12. EasyDMARC (April 2026). *2026 DMARC Adoption & Enforcement Report: Insights from 1.8M Domains, the Fortune 500, and Inc. 5000.* easydmarc.com/blog/ebook/dmarc-adoption-report-2026/
13. DDMARC. *Platform documentation — policy simulator, ARC-chain monitoring, managed MTA-STS and BIMI hosting.* ddmarc.com — supplier-reported; not independently audited at time of publication.
14. ICANN (6 February 2026). *Two Years of Enforcing DNS Abuse Mitigation Requirements: Progress & Next Steps.* icann.org/en/blogs
15. Active Domain (2026). *The Crackdown on DNS Abuse: What New gTLD Rules Mean For Everyone Online.* active-domain.com
16. ICANN. *Domain Metrica: A Measurement Platform* and *Domain Metrica FAQs.* icann.org/octo-ssr/metrica-en
17. M3AAWG. *67th General Meeting, Montréal, 8–11 June 2026.* m3aawg.org/events/67th-general-meeting
18. NIST. *SP 800-177 Rev. 1 — Trustworthy Email; SP 800-45 Ver. 2 — Guidelines on Electronic Mail Security.* csrc.nist.gov
19. CISA. *Known Exploited Vulnerabilities Catalog* and *Binding Operational Directive 26-04.* cisa.gov
20. vSpam.org collector corpus, 1–31 May 2026. Aggregate figures reproducible via the vSpam.org public API; per-indicator data available on request to research@vspam.org.

This report is published by vSpam.org as part of the ongoing independent research series on email and DNS abuse. ISSN 2026-VSPAM-005. Volume 2026, Issue 5. Released under the Creative Commons Attribution 4.0 International License (CC BY 4.0). Citation: vSpam.org Independent Research. (June 2026). "Trends in Spam, Phishing, Spoofing, Malware, and DNS Abuse — May 2026." Volume 2026, Issue 5. ISSN 2026-VSPAM-005. Corrections, additional data, or feedback: research@vspam.org